2023

PRESS KIT

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.

ISSCC 2023
VISION STATEMENT

The International Solid-State Circuits Conference is the foremost global forum for presentation of advances in solid-state circuits and systems-on-a-chip. The Conference offers a unique opportunity for engineers working at the cutting edge of IC design and use to maintain technical currency, and to network with leading experts.

Table of Contents

Table of Contents **4**

Preamble **7**

FAQ on ISSCC 7

Plenary Session (Session 1) 9

Plenary Session — Invited Papers **11**

Plenary Session — Invited Papers **12**

Evening Events (EE) **14**

Sunday, February 19 14

Mentoring Session/Networking Bingo Event 14

EE1: Student Research Preview 14

EE2: Integrated Circuits in an Interconnected World 14

EE3: The Path to Sustainable IC Ecosystems 15

EE4: The Smartest Designer in the Universe, Post-Pandemic! 15

EE5: What will be the Essential Skills for IC Designers in the Next Decade? 15

Session Overviews and Highlights **16**

Conditions of Publication **17**

PREAMBLE 17

FOOTNOTE 17

Session 2 Overview: Digital Processors **18**

Session 2: Digital Processors **19**

Session 3 Overview: Amplifiers and Oscillators **21**

Session 3 Highlights: Amplifiers and Oscillators **22**

Session 4 Overview: Frequency Synthesizers **24**

Session 4 Highlights: Frequency Synthesizers **25**

Session 5 Overview: Image Sensors **26**

Session 5 Highlights: Image Sensors **27**

Session 6 Overview: Advanced Wireline Links and Techniques **28**

Session 6 Highlights: Advanced Wireline Links and Techniques **29**

Session 7 Overview: SRAM Compute-In-Memory **30**

Session 7 Highlights: SRAM Compute-In-Memory **31**

Session 8 Overview: GHz-to-Millimeter Wave Frequency Generation **33**

Session 8 Highlights: GHz-to-millimeter Wave Frequency Generation **34**

Session 9 Overview: Highlighted Chip Releases: Digital and Machine Learning Processors **35**

Session 9 Highlights: Highlighted Chip Releases: Digital and Machine Learning Processors **36**

Session 9 Highlights: Highlighted Chip Releases: Digital and Machine Learning Processors **37**

Session 9 Highlights: Highlighted Chip Releases: Digital and Machine Learning Processors **38**

Session 9 Highlights: Highlighted Chip Releases: Digital and Machine Learning Processors **39**

Session 10 Overview: Pipelined and Noise-Shaping ADCs **40**

Session 10 Highlights: Pipelined and Noise-Shaping ADCs **41**
Session 11 Overview: USB and Compute Power Delivery **42**
Session 11 Highlights: USB and Compute Power Delivery **43**
Session 12 Overview: High Performance Optical Receivers **44**
Session 12 Highlights: High Performance Optical Receivers **45**
Session 13 Overview: Ideas for the Future **47**
Session 13 Highlights: Ideas for the Future **48**
Session 14 Overview: Digital Techniques for Clocking and Power Management **49**
Session 14 Highlights: Digital Techniques for Clocking and Power Management **50**
Session 15 Overview: IoT & Security **51**
Session 15 Highlights: IoT & Security **52**
Session 16 Overview: Efficient Compute-In-Memory-Based Processors for ML **54**
Session 16 Highlights: Efficient Compute-In-Memory-Based Processors for ML **55**
Session 17 Overview: High-Speed Data Converters **57**
Session 17 Highlights: High-Speed Data Converters **58**
Session 18 Overview: mm-Wave & sub-THz for Wireless and Sensing **59**
Session 18 Highlights: mm-Wave & sub-THz for Wireless and Sensing **60**
Session 19 Overview: 5G and Satcom: Receivers and Transmitters **61**
Session 19 Highlights: 5G and Satcom - Receivers and Transmitters **62**
Session 20 Overview: GaN Power Conversion **63**
Session 20 Highlights: GaN Power Conversion **64**
Session 21 Overview: Emerging Sensing Systems and IOT **65**
Session 21 Highlights: Emerging Sensing Systems and IOT **66**
Session 22 Overview: Heterogeneous ML Accelerators **67**
Session 22 Highlights: Heterogeneous ML Accelerators **68**
Session 23 Overview: Analog Sensor Interfaces **70**
Session 23 Highlights: Analog Sensor Interfaces **71**
Session 24 Overview: THz Signal Generation **73**
Session 24 Highlights: THz Signal Generation **74**
Session 25 Overview: RF Transceiver Building Blocks **75**
Session 25 Highlights: RF Transceiver Building Blocks **76**
Session 26 Overview: Display and User Interaction Technologies **77**
Session 26 Highlights: Display and User Interaction Technologies **78**
Session 27 Overview: Innovations from Outside the (ISSCC) Box **79**
Session 28 Overview: High-Density Memories and High-Speed Interfaces **80**
Session 28 Highlights: High-Density Memories and High-Speed Interfaces **81**
Session 29 Overview: Digital Accelerators and Circuit Techniques **83**
Session 29 Highlights: Digital Accelerators and Circuit Techniques **84**
Session 29 Highlights: Digital Accelerators and Circuit Techniques **85**
Session 30 Overview: Power Management Techniques **87**
Session 30 Highlights: Power Management Techniques **88**

Session 31 Overview: Energy-Efficient Radios for UWB, BMI, and IoT Systems **90**
Session 31 Highlights: Energy-Efficient Radios for UWB, BMI, and IoT Systems **91**
Session 32 Overview: Intelligent Biomedical Circuits and Systems **92**
Session 32 Highlights: Intelligent Biomedical Circuits and Systems **93**
Session 33 Overview: Non-Volatile Memory and Compute-In-Memory **95**
Session 33 Highlights: Non-Volatile Memory and Compute-In-Memory **96**
Session 34 Overview: Cryo-CMOS for Quantum Computing **97**
Session 34 Highlights: Cryo-CMOS for Quantum Computing **98**
Trends **99**
Conditions of Publication **100**
PREAMBLE 100
FOOTNOTE 100
Analog – 2023 Trends **102**
Power Management – 2023 Trends **103**
Data Converters – 2023 Trends **104**
RF Subcommittee – 2023 Trends **109**
Wireless- 2023 Trends **111**
Wireline – 2023 Trends **113**
Digital Architectures & Systems (DAS) – 2023 Trends **118**
Digital Circuits – 2023 Trends **124**
Machine Learning (ML) & AI – 2023 Trends **126**
Memory – 2023 Trends **130**
IMMD – 2023 Trends (Medical) **135**
IMMD – 2022 Trends (Imagers) **136**
Technology Directions – 2023 Trends **137**
INDEX **138**
Technical Topics Mapped to Papers 139
Selected Presenting Companies/Institution Mapped to Papers 139
Contact Information **143**

Preamble

FAQ on ISSCC

What is ISSCC?

ISSCC (International Solid-State Circuits Conference) is the **flagship** conference of the IEEE Solid-State Circuits Society. According to the SIA, the Semiconductor industry generated US$555.9 billion in sales in 2021 and ISSCC continues to be the premier technical forum for presenting advances in solid-state circuits and systems. According to the SIA, in 2022, semiconductor sales are expected to exceed U.S$600 billion worldwide. Semiconductors are crucial components of electronics devices and the industry is highly competitive. The year-on-year growth rate in 2022 is expected to see a growth rate of 8.8% percent.

Who Attends ISSCC?

Attendance at ISSCC 2023 is expected to be around **2700**. Corporate attendees from the semiconductor and system industries typically represent around **60%**.

Where is ISSCC?

The **70th ISSCC** will be held in-person with some online offering from February 19th through February 23rd, 2023.

Are there Keynote Speakers?

After a day devoted to educational events, ISSCC 2023 begins formally on Monday, February 20, 2023, with four exciting plenary talks:

- **Lisa Su,** *Chair and Chief Executive Officer, AMD, Austin, TX*
- **Akira Matsuzawa,** *Professor Emeritus of Tokyo Institute of Technology and CEO of Tech Idea, Kawasaki, Japan*
- **Jo De Boeck,** *Executive Vice President and Chief Strategy Officer, imec & KU Leuven, Leuven, Belgium*
- **Erik Ekudden,** *Senior Vice President & Chief Technology Officer, Ericsson, Kista, Sweden*
-

What is the Technical Coverage at ISSCC?

ISSCC covers a full spectrum of design approaches in advanced technical areas broadly categorized as: (1) Communication Systems, (2) Analog Systems, (3) Digital Systems, and (4) Innovations including micro-machines and MEMS, imagers, sensors, biomedical devices, as well as forward-looking developments that may take three or more years for commercialization.

How are ISSCC Papers Selected?
Currently around 650 submissions are received each year across the broad spectrum of specified topics. Review is by a team of over 150 scientific and industry experts from the Far-East, Europe, and North America. These experts are organized into 12 Sub-Committees that cover the 4 broad areas described earlier:

- **Communication Systems** includes Wireless, RF, and Wireline Subcommittees
- **Analog Systems** includes Analog, Power Management, and Data Converter Subcommittees
- **Digital Systems** includes Memory, Digital Circuits, Digital Architectures and Systems, and Machine Leaning and AI Subcommittees
- **Innovative Topics** includes Imagers/MEMS/Medical Devices/Displays and Technology Directions Subcommittees

What Companies are Presenting this year?

Companies presenting papers at ISSCC 2023 include

AMD, Apple, Arm, Asahi Kasei Microdevices, Autosilicon, Axelspace, Broadcom, Cisco Systems, Everactive, Google Quantum AI, Graphcore, Intel, MediaTek, Qualcomm, Realtek Semiconductor, Renesas Electronics, Samsung, SiTime, SKHynix Semiconductor, Sony, STMicroelectronics, Syntiant, Tesla, Texas Instruments, TSMC, just to name a few.

A more complete list can be found in the Index.

Are there educational sessions?

ISSCC features a variety of educational events which include:

- Twelve Tutorials (targeted toward participants looking to broaden their horizon)
- Seven Forums (targeted toward experts in an information sharing context)
- One Short Course (targeted toward in-depth appreciation of a current hot topic)

Are There Other Events?

A more complete list of all activities at ISSCC 2023:

- Four Plenary Presentations
- Eight Invited Talks, consisting of four Industry Talks on Highlighted Chip Releases and four Talks on Innovative Out-of-the-box Ideas.
- Technical Sessions (34 distinct sessions)
- Five Special Events and Panels, including:
    - Integrated Circuits in an Interconnected World The Path to Sustainable IC Ecosystems The Smartest Designer in The Universe, Post-Pandemic! What will be the Essential Skills for IC Designers in the Next Decade?
    - Student Research Preview (for the introduction of graduate-student research-in-progress)
- Educational Sessions Featuring:
    - Twelve Tutorials
    - Seven Forums
    - One Short Course
- Demonstration Sessions from Academia and Industry
- Networking Events
- Author Interview Sessions
- A Number of University Alumni Events

How Do I Use this Press Kit?

The Press Kit provides a PREAMBLE section that features this FAQ and other general information. The kit also includes SESSION OVERVIEWS AND HIGHLIGHTS of all 35 technical sessions into which the 208 papers are grouped, together with brief descriptions and context for each. As well, there is an abstract for each of the Plenary talks. For your convenience, the Kit includes two structural charts in the INDEX section: (a) a list of the 4 Technical Topics and their associated Subcommittees and Sessions; (b) a list of contributing companies and institutions with their associated papers. Thus, to locate information of interest you can access Chart 4.1 to identify sessions of interest, after which you might logically access its Session's Overview or Highlight section. Alternatively, if your interest is in particular organization then Chart 4.1 will direct you immediately to papers of interest each of which is detailed in its corresponding Session Overview and possibly in the Highlights section. For anyone's interest it is useful to use Chart 4.1 to access the appropriate Trends information which provides a broad historical view of the context of your interest and often includes reference to current ISSCC 2022 papers.

Anything New This Year?

ISSCC will hold an invited Industry Track (Session 9) which will highlight recent hot-product releases from Tesla, Syntiant, Nvidia, and Intel, and will discuss innovative ways they solved product-level challenges. ISSCC will also hold an invited session (Session 27) that will highlight some of the new and important developments outside of the circuit community that can have profound impact on the solid-state circuits society either through emergence of new applications or via offering of new platforms for processing, communications, and sensing.

Overview: ISSCC 2023 – "Building on 70 Years of Innovation in Solid-State Circuit Design"

Now in its 70th year, ISSCC is a flagship conference in solid-state circuit design. ISSCC reports innovative ideas that will fuel the continued, rapid evolution of integrated circuits (ICs). The IC design community has proven to be agile and resilient in a world scarred by the worldwide pandemic, and we are ready to commemorate the 70th anniversary of ISSCC with an edition focused more than ever on innovation. To enhance our lives and support an increasingly more sustainable development of our society, ISSCC promotes and shares new circuit ideas with potential to advance the state-of-the-art for IC design and to enable systems with ever-improving functionality and efficiency.

.

Plenary Session (Session 1)

The Plenary Session on the mornings of February 20, 2023, will feature four renowned speakers:

- **Lisa Su,** *Chair and Chief Executive Officer, AMD, Austin, TX*
- **Akira Matsuzawa,** *Professor Emeritus of Tokyo Institute of Technology and CEO of Tech Idea, Kawasaki, Japan*
- **Jo De Boeck,** *Executive Vice President and Chief Strategy Officer, imec & KU Leuven, Leuven, Belgium*
- **Erik Ekudden,** *Senior Vice President & Chief Technology Officer, Ericsson, Kista, Sweden*

Highlights of these Plenary talks are provided in the following section.

ISSCC 2023

PLENARY SESSION – INVITED PAPERS

Plenary Session — Invited Papers

Chair: ***Eugenio Cantatore,*** *Eindhoven University of Technology, Eindhoven, The Netherlands*

ISSCC Conference Chair

Associate Chair: ***Piet Wambacq,*** *imec, Heverlee, Belgium*

ISSCC International Technical-Program Chair

1.1 Innovation For the Next Decade of Compute Efficiency

Lisa Su, *Chair and Chief Executive Officer, AMD, Austin, TX*

Although traditional scaling has slowed over the past decade, we have made tremendous progress as an industry with new approaches including chiplet-based architectures, domain-specific accelerators, and advanced packaging technologies which have enabled major milestones including the first exascale supercomputers. As we look into the future, we need to accelerate the pace of innovation to drive the next decade of advancement in high-performance computing. By far, the largest limiting factor to delivering continued compounded growth in computation power is energy efficiency. In this paper, we highlight a holistic strategy for accelerating innovation in energy efficiency required for next-generation high-performance computing and ultimately achieving zettascale performance. These approaches will be built on continued innovation in process technologies, modular chiplet architectures, and advanced packaging. Fully meeting the challenge will require new dimensions of improvement through extending domain-specific architectures to accelerate core algorithms in combination with wide-scale deployment of AI across all aspects of the system from transistors to software.

1.2 Shape the World with Mixed-Signal and RF Integrated Circuits - Past, Present, and Future

Akira Matsuzawa, *Professor Emeritus of Tokyo Institute of Technology and CEO of Tech Idea, Kawasaki, Japan*

The past 50 years has been an era in which analog equipment has been replaced by digital counterparts. Audio, TV, video, camcorder, camera, recording, wired connection, and wireless communication have been subject to digitization. The digitization of these devices and systems was due to the technological shift from bipolar to CMOS, and to the development of logic and memory circuits supported by scaling laws. In addition, design innovation in mixed-signal integrated circuits such as ADCs and DACs has shown to be indispensable. This talk will look back on the digitization of equipment and the mixed-signal integrated circuit technology that contributed to it. Further, we will look forward to future applications and developments.

1.3 EU Chips Act Drives Pan-European Full-Stack Innovation Partnerships

Jo De Boeck, *Executive Vice President and Chief Strategy Officer, imec & KU Leuven, Leuven, Belgium*

In every aspect of our life and society, semiconductors play a major role. The pandemic in conjunction with supply chain hiccups and geopolitical tensions made all regions realize that they need to revisit their presence in the semiconductor value chain. The European Commission projected the ambition of achieving a 20% share of the global semiconductor production by 2030.

Europe can leverage existing strengths such as, among others, the unique position of equipment companies and leadership positions in 300mm semiconductor technology R&D. The Chips-for-Europe initiative will invest in pilot lines and ecosystems for chip manufacturing, embracing leading-edge and first-of-a-kind technologies. The pilot lines will allow early exploration of the potential impact of new technology features in advanced chip and system architectures. This will trigger increased demand and accelerate industrial uptake of the novel technologies. This type of innovation loop is also essential for deep-tech start-ups building their unique value proposition. The full-stack, networked model of industry collaboration is at the core of the EU Chips Act ambition and will impact different application domains such as heterogeneous cloud and distributed computing, connectivity, automotive, and health.

It is crucial for all this innovation potential that we, as an industry, consider that semiconductor manufacturing is resource-intensive with respect to energy, water, chemicals, and raw materials. Design-technology co-optimization (DTCO) and System-Technology co-optimization (STCO) methodologies can develop a framework for early sustainability assessments of logic technologies. Finally, we urgently need to get the message across that climate, health, safety, and human connectedness all require complex digital backbones, if we want to stand a chance of attracting the right talent.

1.4 5G Drives Exponential Increase in Processing Needs Across all Industries

Erik Ekudden, *Senior Vice President & Chief Technology Officer, Ericsson, Kista, Sweden*

Across essentially all industrial sectors, advanced semiconductor technology is the key enabler for innovations in customer offerings and internal efficiencies. The increase in the value of data and the related push for AI are examples of forces that increase the demand for compute power, which translates to more complex and powerful silicon. Moore's Law, supported by rapidly evolving semiconductor technology and ever more advanced building practices and assembly technologies, has met the need for decades.

But what is driving 5G today? If we look at the processing requirements, it is the digital front-end, physical layer processing, and beam forming. Back in 2010, LTE/4G was a 20MHz carrier with two receive and two transmit branches, and there was a transmission time interval of one millisecond. Fast forward to where we are today on 5G with massive MIMO, we typically have 100MHz carrier bandwidth. That is a factor of five increase. We have 64 transmitter and 64 receiver radios, which is an increase by a factor of 32, and the transmission time is down to 0.5 milliseconds. In other words, which is an increase by a factor of 32, and to do 160 times more processing. This is driving an exponential increase in processing needs across the telecom business today and will continue to do so as we race towards 6G. This talk will address whether the semiconductor industry is ready to tackle these challenges.

ISSCC 2023

EVENING EVENTS

Evening Events (EE)

ISSCC 2023 will continue the popular tradition of evening sessions where experts often of opposing views, discuss topics which range from the lighthearted to the controversial (but always informative and entertaining!). This year's are "*Integrated Circuits in an Interconnected World*", "*The Path to Sustainable IC Ecosystems*", "*The Smartest Designer in the Universe, Post-Pandemic!*", and "*What will be the Essential Skills for IC Designers in the Next Decade?*".
In addition, ISSCC 2023 will include additional special events including "*Student Research Preview'* and a networking and mentoring session.

Sunday, February 19

Mentoring Session/Networking Bingo Event

Women in Circuits (WiC) together with ISSCC will be holding a networking and mentoring session on Sunday afternoon. Distinguished panelists from the "Integrated Circuits in an Interconnected World" panel, WiC members, and other participants will play getting-to-know-you bingo to promote engagement between various members of the community. This will give participants the chance to network and mingle with people across a spectrum of seniority in the field in a casual setting. This event is open to all ISSCC attendees and the public.

EE1: Student Research Preview

The Student Research Preview (SRP) will highlight selected student research projects in progress. The SRP consists of 90 second presentations followed by a Poster Session, by graduate students from around the world, which have been selected on the basis of a short submission concerning their on-going research.

The SRP will start with an inspirational lecture by Professor Mark Horowitz (Stanford University).

Monday, February 20

EE2: Integrated Circuits in an Interconnected World

Data-centric applications such as automotive, IoT, and machine learning primarily depend on IC connectivity to meet the demands of high-bandwidth and energy-efficient communication. Product requirement diversity continues to grow, and IC components now face a bottleneck in system performance and power, particularly due to their interconnects. To ensure communication capabilities that meet the demands of emerging applications with ever-increasing features, additional research and focus is needed. This panel focuses on connectivity for the next generation of communication systems and brings together expert panelists to share their perspectives on topics in IC connectivity across wireless, wireline, chip-to-chip, and optical link communications.

Monday, February 20

EE3: The Path to Sustainable IC Ecosystems

Sustainability has become a major concern in our lives in general, and all the way to IC design. It widened from energy management to include greenhouse gas emissions and pressure on natural resources all along the product lifecycle. These challenges must be taken into account early in product design phases, to rethink system architecture and circuit techniques to favor frugality and reuse and minimize the impact of manufacturing. A deep restructuration of the IC ecosystem to integrate eco-design and reuse will need to arise, either from top-down political intervention or company-driven initiatives. How do we facilitate an economically viable path to sustainable IC ecosystem? Will this come from market incentives or are government regulations required? With a growing awareness of this challenge, several initiatives have already emerged for sustainable electronics, coming from research labs, companies, citizens, or governments. This panel confronts these approaches and explores their potential to create this economy viable path via market incentive or government regulation, for a sustainable value chain all over the product lifecycle.

Tuesday, February 22

EE4: The Smartest Designer in the Universe, Post-Pandemic!

At ISSCC 2020 there was a battle of epic proportions between industry, academia and students to determine who was the smartest designer in the universe. Industry came out victorious. Now, at ISSCC 2023, as we return to an in-person conference, academia and students have their chance to get their revenge and set the record straight. In this interactive quiz show, three teams representing industry, academia and students will compete for the honor and the prestigious title: “The Smartest Designer in the Universe”. In several rounds, the contestants will solve questions and puzzles covering all parts of electrical engineering. They will baffle you with their knowledge, surprise you with their wit and entertain you with their to the point remarks. This all topped with a gentle sauce of irony since the smartest designer in the universe should be smart enough to appreciate the special relativity of it all. Join this session not only to support your own team but enroll in the game. Everybody will be able to actively participate using an app.

Tuesday, February 22

EE5: What will be the Essential Skills for IC Designers in the Next Decade?

Based on emerging trends in design methodology such as AI for IC design and verification, this session of academic and industry leaders will predict and discuss how future design automation will change the way IC designers work in the next decade. Will more and more of IC design be automated by then? Is our field shrinking? Are we attracting and training enough students to learn IC design to meet potential industry needs? Join this special evening topic session to get the perspective of industry and academic leaders in IC design.

ISSCC 2023

SESSION OVERVIEWS AND HIGHLIGHTS

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:

- 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.
- 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 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.

Session 2 Overview: Digital Processors

Digital Architectures and Systems Subcommittee

Session Chair: Shidhartha Das, AMD, Cambridge, UK

Session Co-Chair: Ji-Hoon Kim, Ewha Woman's University, Korea

In this year's conference, mainstream processors designed in advanced process technologies share the stage with domain-specific processors to address a range of applications from high-performance general-purpose computing through to genomics. The session leads off with AMD's latest-generation "Zen 4" core, MediaTek's flagship 5G mobile SoC and features university researchers demonstrating simulated-annealing processors and data-flow computing SoCs for AR/VR, robotics and next-generation genomic sequencing.

- In Paper 2.1, AMD highlights a 5.7GHz "Zen 4" 8-core complex fabricated in a 5nm FinFET process, occupying 55mm^2. 13% IPC improvement is accomplished through architectural enhancements including increased structure sizes and conflict reduction. This coupled with physical design innovations including V_{th} management, selective path tuning and intelligent power grid optimization drives a 6% reduction in switching capacitance, 16% increase in frequency, and an F_{MAX} of 5.7GHz. Over 30% increase in iso-power performance vs. the prior generation in desktop products is demonstrated.
- In Paper 2.2, MediaTek demonstrates a high-performance thermal management system for a 110mm^2 5G mobile gaming SoC featuring a tri-gear CPU with GPU, designed in 4nm. With CPUs running up to 3.35GHz, a power-predictor and calculator system fed by multiple on-chip sensors combined with energy-thermal-aware task reallocation facilitates an average 10°C increase in throttling threshold enabling a record AnTuTu score of 114.6M.
- In Paper 2.3, Tokyo Institute of Technology presents a 40nm programmable multi-policy simulated-annealing processor integrating 4-replica 512 fully-connected spins, extensible across 4-chips. The 9mm^2 die operates at 336MHz consuming 151-474mW at 1.1V.
- In Paper 2.4, National Taiwan University researchers present a 28nm processor for next-generation genomic sequencing supporting end-to-end workflow from short-read mapping through to genotyping. The 16mm^2 die operates at 400MHz @ 0.9V and is designed in a TSMC 28nm CMOS process. The chip delivers 59× higher throughput and 935-to-4910× higher energy-efficiency compared to state-of-the-art cloud-based solutions.
- In Paper 2.5, National Taiwan University researchers present a 28mW 5mm^2 200MHz SoC for autonomous robot control that incorporates sampling-based motor control that enables high parallelization. Optimizations such as trajectory pruning and use of an acceleration-based model facilitate a ~5KHz maximum rate control, with less than 1.6% tracking error, and over 350× improvement in energy efficiency as compared to prior state of the art.
- In Paper 2.6, National Tsing Hua University presents a video CNN chip for 4K-UHD imaging/display applications, providing peak throughput of 60/50fps for spatial/temporal-interpolation with 704mW power dissipation. The 40nm 12.6mm^2 chip achieves comparable energy efficiency to prior work ranging from 3.8-6.4TOPS/W and area efficiency of 222.2GOPS/mm^2, while supporting the advanced feature of multi-image processing.
- In Paper 2.7, KAIST presents a real-time hyper-realistic-3D-NeRF processor, MetaVRain, for metaverse on mobile devices, which can create 3D models by training a DNN to memorize 3D scene geometry from a few photos. The 28nm chip, integrating 5K FP8-FP16 configurable MACs with 2MB of SRAM, demonstrates a maximum of 118fps, and consumes at least 99.95% lower power compared with modern GPUs and a TPU.

Session 2: Digital Processors

[2.1] "Zen 4": The AMD 5nm 5.7GHz x86-64 Microprocessor Core

[2.4] A Fully Integrated End-to-End Genome Analysis Accelerator for Next-Generation Sequencing

[2.7] MetaVRain: A 133mW Real-Time Hyper-Realistic 3D-NeRF Processor with 1D-2D Hybrid-Neural Engines for Metaverse on Mobile Devices

Paper 2.1 Authors: Benjamin Munger[1], Kathy Wilcox[1], Jeshuah Sniderman[1], Chuck Tung[1], Brett Johnson[2], Russell Schreiber[3], Carson Henrion[2], Kevin Gillespie[1], Harry Fair[1], Thomas Burd[4], Jonathan White[1], Scott McLelland[1], Steven Bakke[1], Javin Olson[1], Ryan McCracken[1], Matthew Pickett[2], Aaron Horiuchi[2]

Paper 2.1 Affiliation: [1]AMD, Boxborough, MA, [1]AMD, Fort Collins, CO, [3]AMD, Austin, TX, [4]AMD, Santa Clara, CA

Paper 2.4 Authors: Yen-Lung Chen[1], Chung-Hsuan Yang[1], Yi-Chung Wu[1], Chao-Hsi Lee[2], Wen-Ching Chen[3], Liang-Yi Lin[3], Nian-Shyang Chang[3], Chun-Pin Lin[3], Chi-Shi Chen[3], Jui-Hung Hung[2,4], Chia-Hsiang Yang[1,2]

Paper 2.4 Affiliation: [1]National Taiwan University, Taipei, Taiwan, [2]GeneASIC Technologies Corp., Hsinchu, Taiwan, [3]Taiwan Semiconductor Research Institute, Hsinchu, Taiwan, [4]National Yang Ming Chiao Tung University, Hsinchu, Taiwan

Paper 2.7 Authors: Donghyeon Han, Junha Ryu, Sangyeob Kim, Sangjin Kim, Hoi-Jun Yoo

Paper 2.7 Affiliation: Korea Advanced Institute of Science and Technology, Korea

Subcommittee Chair: Thomas Burd, AMD, Santa Clara, CA, Digital Architectures and Systems Subcommittee

CONTEXT AND STATE OF THE ART

- Architectural innovations are tailored to a range of applications from high-performance computing through to domain-specific applications, such as next-generation genomic sequencing and video-rendering engines for the emerging metaverse.
- Process technology continues to scale, enabling generational improvements in transistor integration, IPC and operating frequencies, however reliable power-delivery and thermal-management continue to present key technical challenges.
- Emerging applications continue to benefit from energy-efficient processing delivered through a combination of technology scaling and algorithmic/architecture co-optimization.

TECHNICAL HIGHLIGHTS

- **AMD presents "Zen 4", their next-generation 5.7GHz 8-core complex occupying 55mm^2 in a 5nm technology.**
  - A 13% core IPC improvement is achieved through architectural enhancements, along with a 16% frequency boost enables a generational performance gain of over 30% performance at an equivalent power level vs. the prior generation in desktop products
- **National Taiwan University researchers demonstrate a 400MHz 16mm^2 genome accelerator chip in a 28nm technology.**
  - A 59× higher throughput and 935-to-4910× higher energy-efficiency is demonstrated through end-to-end workflow optimization and inclusion of genotyping.
- **KAIST presents a real-time hyper-realistic 3D-NeRF processor, MetaVRain, integrating 5K FP8-FP16 configurable MACs with 2MB of SRAM in 28nm CMOS, for metaverse on mobile devices.**
  - At least 99.95% lower power consumption is achieved compared with modern GPUs and a TPU, while achieving a maximum of 118fps.

APPLICATIONS AND ECONOMIC IMPACT

- AMD continues to provide major advancements in mainstream CPU performance for desktop and server through architectural improvements combined with physical design optimization and design resource scaling.
- Researchers deliver domain-specific processors for a variety of emerging applications, such as next-generation genomic sequencing, simulated annealing, robotics and AR/VR, which provide more energy efficiency and higher performance compared with the conventional general-purpose processors.

Session 3 Overview: Amplifiers and Oscillators

Analog Subcommittee

Session Chair: Jens Anders, University of Stuttgart, Germany

Session Co-Chair: Shon-Hang Wen, MediaTek, Taiwan

This session highlights advances in state-of-the-art amplifiers and oscillators for several applications. The first three papers focus on improving DR in a Class-D amplifier, lowering the noise and input current in a chopper-stabilized opamp and generating a single tone with the lowest THD for ADC testing. RC oscillators achieve high accuracy with a high-resolution trimming and aging compensation scheme, and crystal oscillators with reduced startup time and low energy, a wide acceptable injection clock frequency error, and low power sensitivity to temperature are demonstrated. A PLL-less BAW-based oscillator with digital calibration achieves high frequency stability and low jitter.

- In Paper 3.1, Delft University of Technology and Goodix Technology present a digital-input capacitively-coupled Class-D amplifier. It achieves 120.9dB DR and -111.2dB peak THD+N, and can deliver 13W/23W at 10% THD into an 8Ω/4Ω load with 90%/86% efficiency.
- In Paper 3.2, Delft University of Technology presents a chopper-stabilized amplifier with a relaxed fill-in technique. By introducing a duty-cycled non-chopped fill-in OTA and a ripple reduction loop, this work achieves a 25× reduction in input current, while achieving a flat noise floor of 12nV/√Hz.
- In Paper 3.3, Texas Instruments and IIT Madras present bandpass filter and oscillator ICs with a new-benchmark THD performance at $10V_{ppd}$ swing for testing high-resolution ADCs. It is achieved by incorporating capacitor nonlinearity cancellation and opamp output conductance nonlinearity suppression techniques into an active-RC bandpass filter and oscillator.
- In Paper 3.4, Delft University of Technology, Silicon Integrated B.V. and Tsinghua University present a $0.01mm^2$ 10MHz RC frequency reference with high-resolution on-chip trimming of both its temperature coefficient and absolute frequency. This work achieves a ±0.28% inaccuracy from −45°C to 125°C after 1-point trim.
- In Paper 3.5, the University of Illinois at Urbana-Champaign presents a 100MHz RC oscillator in 65nm CMOS that achieves ±1030ppm frequency inaccuracy from -40°C to 85°C after accelerated aging for 500 hours at 125°C. This performance is achieved with aging compensation by periodically locking the oscillator to a less-aged reference oscillator.
- In Paper 3.6, the University of Macau and Instituto Superior Tecnico present a 12/13.56MHz crystal oscillator with 45.8µs startup time and 5nJ startup energy. It is achieved by introducing binary-search-assisted frequency locking to a 2-step injection method, implemented using a resettable fast-settling auxiliary DCO, edge aligner, frequency comparator, and control logic.
- In Paper 3.7, Nanjing University of Posts and Telecommunications and Hefei University present a 16MHz crystal oscillator in 40nm CMOS, which uses an automatic phase-error correction technique to achieve a 17.5µs startup time with an injection clock frequency error of up to 10^4 ppm. The startup energy is 9.2nJ, and the startup time variation over temperature is ±4.5%.
- In Paper 3.8, Peking University and the Advanced Institute of Information Technology of Peking University present a crystal oscillator with a Gm-C-based regulated current-injection technique, achieving very low power sensitivity to temperature of 0.017nW/°C.
- In Paper 3.9, Texas Instruments and IIT Madras present a 0.5-to-400MHz programmable BAW-based oscillator employing a new temperature/supply-insensitive Dual-Slope Fractional Output Divider architecture. It achieves ±4ppm frequency stability over -40°C to 85°C with <95fs rms phase jitter.

Session 3 Highlights: Amplifiers and Oscillators

[3.1] A 120.9dB DR, -111.2dB THD+N Digital-Input Capacitively-Coupled Chopper Class-D Audio Amplifier

[3.5] A 1.4µW/MHz 100MHz RC Oscillator with ±1030ppm Inaccuracy from -40°C to 85°C After Accelerated Aging for 500 Hours at 125°C

[3.9] A 0.5-to-400MHz Programmable BAW Oscillator with Fractional Output Divider Achieving 4ppm Frequency Stability over Temperature and <95fs Jitter

Paper 3.1 Authors: Huajun Zhang[1], Marco Berkhout[2], Kofi A. A. Makinwa[1], Qinwen Fan[1]

Paper 3.1 Affiliation: [1]Delft University of Technology, Delft, The Netherlands, [2]Goodix Technology, Nijmegen, The Netherlands

Paper 3.5 Authors: Kyu-Sang Park, Nilanjan Pal, Yongxin Li, Ruhao Xia, Tianyu Wang, Ahmed Abdelrahman, Pavan Kumar Hanumolu

Paper 3.5 Affiliation: University of Illinois at Urbana-Champaign, Urbana, IL

Paper 3.9 Authors: Subhashish Mukherjee[1], Yogesh Darwhekar[1], Jayawardan Janardhanan[1], Peeyoosh Mirajkar[1], Raghavendra Reddy[1], Harish Ramesh[1], Bichoy Bahr[2], Jagdish Chand[1], Uday Meda[1], Baher Haroun[2], Shankar Karantha[1], Ernest Yen[3], Keegan Martin[2], Daniel Gan[4], Amin Sijelmassi[2], Sankaran Aniruddhan[5]

Paper 3.9 Affiliation: [1]Texas Instruments, Bangalore, India, [2]Texas Instruments, Dallas, [3]Texas Instruments, Santa Clara, [4]Texas Instruments, Melaka, Malaysia, [5]IIT Madras, Chennai, India

Subcommittee Chair: Maurits Ortmanns, University of Ulm, Germany

CONTEXT AND STATE OF THE ART

- Conventional digital-input Class-D amplifiers suffer limited DR due to either high jitter sensitivity in open-loop topologies or thermal noise of feedback resistors in closed-loop topologies.
- Until today, research on CMOS time-references has mostly focused on achieving sub-10ppm frequency inaccuracy vs. temperature, which is required by many applications. The next step is to investigate the performance over the full lifetime of the chip and find ways to preserve the excellent initial performance.
- Today quartz oscillators are the first choice as references in phase-noise-critical applications. However, today's quartz oscillators do not provide frequencies beyond a few hundred MHz. Here, BAW oscillators present an interesting alternative for future 5G and 6G applications with fundamental frequencies in the GHz range and competitive phase-noise performance.

TECHNICAL HIGHLIGHTS

- **Delft University of Technology presents the first digital-input capacitively-coupled Class-D amplifier. It achieves 120.9dB DR and -111.2dB peak THD+N, and can deliver 13W/23W at 10% THD into an 8Ω/4Ω load with 90%/86% efficiency.**
    - This is the first digital-input capacitively-coupled Class-D amplifier (CDA). Compared to other high-voltage (>10V) digital-input CDAs, it achieves the best peak THD+N (14dB better), the highest dynamic range (5.4dB higher), and the lowest A-weighted integrated output noise (2× lower).

- **the University of Illinois at Urbana-Champaign presents a 100MHz RC oscillator in 65nm CMOS, which achieves ±1030ppm frequency inaccuracy from -40°C to 85°C after accelerated aging for 500 hours at 125°C. This performance is achieved with aging compensation by periodically locking the oscillator to a less-aged reference oscillator.**
    - This paper presents a temperature- and aging-compensated RC oscillator in which the long-term drift of the main oscillator is compensated by periodically locking its frequency to that of the less-aged reference oscillator. A prototype 100MHz RC oscillator fabricated in a 65nm CMOS process achieves a good power efficiency of 1.4μW/MHz and better than ±1030ppm frequency inaccuracy from -40°C to 85°C after accelerated aging for 500 hours at 125°C.
- **Texas Instruments presents a 0.5-to-400MHz programmable BAW-based oscillator employing a new temperature/supply-insensitive Dual-Slope Fractional Output Divider architecture. It achieves ±4ppm frequency stability over -40°C to 85°C with <95fs rms phase jitter.**
    - A BAW-based programmable oscillator is integrated as a multichip module (MCM) with tight frequency control. New fractional output divider architecture and calibration schemes are introduced to achieve best-in-class jitter and power (1.2× lower) compared to existing programmable oscillators in the industry.

APPLICATIONS AND ECONOMIC IMPACT

- Higher dynamic range and better linearity result in better user experience, especially for emerging 3D reality audio (spatial audio) applications, such as virtual reality.
- Combining capacitive DACs (CDACs) into closed-loop capacitively-coupled Class-D amplifers results in higher levels of integration and lower cost, achieving higher DR than the conventional topology, which consists of current-steering DAC (IDAC) and resistive-feedback Class-D amplifers.
- Aging-resilient high-accuracy CMOS time references are the door opener for such systems to enter a much broader market.
- Temperature-stabilized high-performance BAW oscillators can greatly reduce the footprint of today's quartz-based systems. Moreover, they represent ideal candidates for 5G and 6G applications with operating frequencies beyond 100 GHz.

Session 4 Overview: Frequency Synthesizers

RF Subcommittee

Session Chair: Dmytro Cherniak, Infineon Technologies, Villach, Austria

Session Co-Chair: Wanghua Wu, Samsung Semiconductor, San Jose, CA

Low-phase-noise frequency synthesizers are key building blocks of various wireless systems ranging from IoT radios to high-performance 5G transceivers and radars. This session presents several state-of-the-art frequency synthesizers including a wide-bandwidth FMCW synthesizer, a novel wide-range LO generator for 5G, digital PLLs with superior fractional spur performance, an ultra-low-jitter PLL in sub-THz bands, and an ultra-low-voltage PLL for IoT.

- In Paper 4.1, imec and Vrije Universiteit Brussel present a 16GHz charge-pump PLL-based FMCW synthesizer with robust duty-cycled operation achieving 1.5GHz modulation bandwidth and 41kHz$_{rms}$ error with below 1µs re-lock time in 28nm CMOS.
- In Paper 4.2, KAIST presents a low-jitter and wide-range LO generator for 5G FR1 in 65nm CMOS achieving 0.6-to-7.7GHz locking range and 135fs$_{rms}$ jitter using a single LC-VCO cascaded with a ring-oscillator-based sub-integer-N frequency multiplier.
- In Paper 4.3, Politecnico di Milano and Infineon Technologies present a 9-to-11GHz fractional-N digital PLL with an inverse-constant-slope DTC and FCW subtractive dithering technique achieving 76.7fs$_{rms}$ jitter and -71.9dBc in-band fractional spur in 28nm CMOS.
- In Paper 4.4, Tokyo Institute of Technology shows a 2.4GHz fractional-N PLL, which operates from a 32kHz reference clock and employs nonuniform-timing reference oversampling to avoid noisy sampling points. The PLL achieves a 3.94ps$_{rms}$ jitter and -43dBc fractional spur at 3.8mW power in 65nm CMOS.
- In Paper 4.5, Politecnico di Milano and Infineon Technologies demonstrate a 9.25GHz bang-bang digital PLL exploiting a multi-DTC topology with phase-shifted quantization-error sequences achieving -60.3dBc in-band fractional spur and 77fs$_{rms}$ jitter at 17.9mW power in 28nm CMOS.
- In Paper 4.6, KAIST presents a 103.5GHz PLL for ultra-low-jitter LO generation in sub-THz bands employing a power-gating injection-locked frequency-multiplier-based phase detector and achieving 47fs$_{rms}$ jitter and -253dB FoM at 26.6mW power in 65nm CMOS.
- In Paper 4.7, the Chinese Academy of Sciences and the University of Macau introduce a 2.4GHz ultra-low-voltage subsampling PLL for IoT radios operating from a 0.4V supply and achieving 236fs$_{rms}$ jitter, -76.1dB reference spur, and -253dB FoM in 40nm CMOS.

Session 4 Highlights: Frequency Synthesizers

[4.2] A 135fs$_{rms}$-Jitter 0.6-to-7.7GHz LO Generator Using a Single LC-VCO-Based Subsampling PLL and a Ring-Oscillator-Based Sub-Integer-N Frequency Multiplier

[4.3] A 76.7fs-Integrated-Jitter and -71.9dBc In-Band Fractional-Spur Bang-Bang Digital PLL Based on an Inverse-Constant-Slope DTC and FCW Subtractive Dithering

Paper 4.2 Authors: Yongwoo Jo*, Juyeop Kim*, Yuhwan Shin, Chanwoong Hwang, Hangi Park, Jaehyouk Choi

Paper 4.2 Affiliation: KAIST, Daejeon, Korea

Paper 4.3 Authors: Mattia Dartizio[1], Francesco Tesolin[1], Giacomo Castoro[1], Francesco Buccoleri[1], Luca Lanzoni[1], Michele Rossoni[1], Dmytro Cherniak[2], Luca Bertulessi[1], Carlo Samori[1], Andrea Leonardo Lacaita[1], Salvatore Levantino[1]

Paper 4.3 Affiliation: [1]Politecnico di Milano, Milano, Italy, [2]Infineon Technologies, Villach, Austria

Subcommittee Chair: Jan Craninckx, imec, Leuven, Belgium

CONTEXT AND STATE OF THE ART

- Low-phase-noise and low-power frequency synthesizers are key building blocks of wireless transceivers for applications ranging from IoT radios to 5G and radars.
- Circuit non-idealities and variability with technology process, supply, and temperature have been a fundamental limitation to achieving high purity of frequency synthesis.
- Innovations in circuit techniques and digital algorithms for frequency synthesizers allows to break the trade-offs between spectral purity and power consumption.

TECHNICAL HIGHLIGHTS

- **KAIST introduces a low-jitter and wide-range LO generation for 5G FR1 band in 65nm CMOS based on cascaded-PLL architecture achieving better FoM than previous state-of-the-art LO generators.**
    - A low-jitter and wide-range PLL for 5G FR1 in 65nm CMOS achieving 0.6-to-7.7GHz locking range, 135fs$_{rms}$ jitter and -244.9dB FoM using a single LC-VCO cascaded with a ring-oscillator-based sub-integer-N frequency multiplier.
- **Politecnico di Milano and Infineon Technologies present an ultra-low-jitter digital PLL in 28nm CMOS with a novel DTC topology and FCW dithering technique, which achieves much better fractional spur comparing to state-of-the-art PLLs.**
    - A 9-to-11GHz fractional-N digital PLL with a novel inverse-constant-slope DTC and FCW subtractive dithering technique achieving 76.7fs$_{rms}$ jitter, -249dB FoM, and -71.9dBc in-band fractional spur in 28nm CMOS

APPLICATIONS AND ECONOMIC IMPACT

- Advanced circuit techniques in combination with digital algorithms for high-performance frequency synthesizers enable new applications.
- High-performance and low-power frequency synthesizers with advanced calibration techniques require no factory trimming, thereby reducing manufacturing cost significantly.

Session 5 Overview: Image Sensors

Imagers, MEMS, Medical & Displays Subcommittee

Session Chair: Kazuko Nishimura, Panasonic Holdings, Moriguchi, Japan

Session Co-Chair: Masaki Sakakibara, Sony Semiconductor Solutions, Kanagawa, Japan

Imagers continue to diversify for an ever-wider range of applications. The first three papers describe high-performance hybrid Event-based Vision Sensors (EVS)/CIS sensors, followed by a 1.28μm quad-pixel structure implemented in a 50Mpixel CIS. This is followed by a 3THz, 73dB dynamic range image sensor, a SPAD X-ray sensor with high frame rate and HDR. The final two papers showcase an ultra-low-power energy-harvested sensor, operating at just 55pW/pixel, and a HDR imager with an ADC-free flux-readout scheme.

- In Paper 5.1, OmniVision Technologies describes the high-resolution, hybrid EVS/CIS SoC with 15 MPixel CIS and 1 MPixel EVS resolution and a record peak event rate of 4.6 GEvent/s at only 55pJ/event. This results in FOMs of 4.4 MP × GEvent/s and 0.08 MP × GEvent/s/pJ.
- In Paper 5.2, Sony Semiconductor Solutions presents a 1.22μm-pixel 35.6Mpixel RGB hybrid event-based vision sensor. The sensor demonstrates high-quality RGB output with random noise of 1.57e- and high-speed event output up to 10Kfps using adaptive frame rate control based on the event sparsity.
- In Paper 5.3, Sony Semiconductor Solutions shows a 2.97μm-pixel back-illuminated stacked event-based vision sensor with front-end circuit sharing and time division detection technique. The sensor demonstrates both an event-detection mode with 79.8mW power consumption at 1412Meps and an intensity-acquisition mode with 10,600e- full-well capacity and 2.6e- readout noise.
- In Paper 5.4, Samsung Electronics presents 1.28μm quad pixel structure for 50Mpixel CMOS image sensor. Full-well capacity of 20,000e- is enabled by newly quad pixel and dual overflow paths with intra-pixel DTI, and a temporal noise of 0.98e- is achieved by quarter-ring source-follower and floating diffusion sharing technology.
- In Paper 5.5, Chinese Academy of Sciences introduces a 16.4kpixel 3.08-to-3.86THz image sensor based on a column-parallel architecture. The pixel with a step-covered patch antenna and a defected ground structure achieves a 753V/W sensitivity. The sensor exhibits a 73dB dynamic range and realizes real-time imaging with QCLs.
- In Paper 5.6, Yonsei University describes a SPAD X-ray detector with seamless global shutter, which is enabled by using two time-encoded extrapolation counters. It has a pixel pitch of 49.5μm and a 400×200 pixel resolution, achieves 600fps and a DR of 117.7dB while consuming only 127.2mW.
- In Paper 5.7, National University of Singapore reveals a μW-range low-cost imager with on-chip detection of Regions of Interest (ROI) with novel content and power adaptation to purely harvested sources. Always-on power of 55-pW/pixel is achieved in a 180nm test chip powered by a 3.3×3.3-mm^2 solar cell down to 100lux.
- In Paper 5.8, University of Toronto reveals an image sensor with an ADC-free flux-readout scheme that uses regression on sinusoidal-reference binary comparator outputs read out at 26KHz during exposure to produce one or more digital HDR flux samples per frame. The sensor is demonstrated for HDR imaging and exhibits a >300× faster pixel-wise exposure coding rate than previous designs.

Session 5 Highlights: Image Sensors

[5.1] A 3-Wafer-Stacked Hybrid 15MPixel CIS + 1 MPixel EVS with 4.6GEvent/s Readout, In-Pixel TDC and On-Chip ISP and ESP Function

[5.4] A 0.64µm 4-Photodiode 1.28µm 50Mpixel CMOS Image Sensor with 0.98e- Temporal Noise and 20Ke- Full-Well Capacity Employing Quarter-Ring Source-Follower

Paper 5.1 Authors: Menghan Guo[1], Shoushun Chen[1], Zhe Gao[2], Wenlei Yang[1], Peter Bartkovjak[2], Qing Qin[2], Xiaoqin Hu[1], Dahei Zhou[1], Masayuki Uchiyama[2], Yoshiharu Kudo[3], Shimpei Fukuoka[3], Chengcheng Xu[2], Hiroaki Ebihara[2], Andy Wang[2], Peiwen Jiang[2], Bo Jiang[2], Bo Mu[2], Huan Chen[1], Jason Yang[2], TJ Dai[2], Andreas Suess[2]
Paper 5.1 Affiliation: [1]OmniVision Technologies, Shanghai, China, [2]OmniVision Technologies, Santa Clara, CA, [3]OmniVision Technologies, Yokohama, Japan

Paper 5.4 Authors: Hyuncheol Kim, Yun Hyeok Kim, Sanghyuck Moon, Hwanwoong Kim, Byeongjun Yoo, Jueun Park, Seyoung Kim, June-Mo Koo, Sewon Seo, Hye Ji Shin, Jinwoo Kim, Kyungil Kim, Jae-Hoon Seo, Seunghyun Lim, Taesub Jung, Howoo Park, Juhyun Ko, Kyungho Lee, JungChak Ahn, JoonSeo Yim
Paper 5.4 Affiliation: Samsung Electronics, Hwasung, Korea

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

CONTEXT AND STATE OF THE ART

- State-of-the-art devices are combined with imaging technology to expand the functionality of image sensors.
- Three-wafer-stacked technology combining a CIS and event-based sensing with on-chip signal processing enable a high-speed image sensor.
- Dynamic range enhanced by innovations in pixel read-out circuitry with quarter-ring source-follower and a quad-pixel structure with dual overflow paths.

TECHNICAL HIGHLIGHTS

- **Omni Vision introduces a triple-wafer-stacking CMOS Image Sensor (CIS) and Event-Based Vision Sensor (EVS) hybrid sensor.**
    - A hybrid EVS/CIS SoC with 15 MPixel CIS and 1 MPixel EVS resolution demonstrates a record peak event rate of 4.6 GEvent/s at 55pJ/event. This results in record FOMs of 4.4 MP × GEvent/s and 0.08 MP × GEvent/s/pJ.
- **Samsung Electronics presents a 1.28µm quad-pixel structure for 50Mpixel CMOS Image Sensor (CIS) with full-well capacity of 20,000ke-.**
    - A full-well capacity of 20,000ke- is achieved by using quad-pixel and dual overflow paths with intra-pixel DTI. A temporal noise of 0.98e- can be achieved by a quarter-ring source-follower with floating sharing technology resulting in a dynamic range of 86.2dB.

APPLICATIONS AND ECONOMIC IMPACT

- Three-tier wafer stack technology for CIS and EVS hybrid sensors enable new imager applications with improved efficiency.
- Low-noise and wide-dynamic-range CMOS image sensors enable low-noise image capture under a wide-range of conditions and camera applications.

Session 6 Overview:
Advanced Wireline Links and Techniques

WLN Subcommittee

Session Chair: *Friedel Gerfers, TU Berlin, Berlin, Germany*

Session Co-Chair: *Takashi Takemoto, Hitachi Ltd., Sapporo, Japan*

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

Increasing demand for bandwidth in networking and computing drives wireline links to push data rate limits, while at the same time improving the energy-efficiency, bit-error-rate, and throughput per millimeter of chip edge and silicon area. The first paper in the session describes the design of a 112Gb/s PAM-4 transceiver system dealing with 43dB channel loss by utilizing a 3-tap FFE and 18-tap DFE. The second paper demonstrates current state of the art of 112Gb/s ADC-DSP based transceivers, achieving transmission over a 48dB loss channel while consuming only 4.63pJ/b. The following paper exhibits how an on-chip coplanar waveguide is used to realize a 5-tap FFE receiver for 200Gb/s transmission compensating 17.2dB channel loss with only 0.43pJ/b. Paper 6.4 demonstrates a 32Gb/s die-to-die chiplet transceiver, achieving 8Tb/s/mm beach-front bandwidth while consuming 0.44pJ/b. The next two papers showcase advanced clock-and-data recovery (CDR) techniques. One paper applies a network of autocorrelators for flash frequency acquisition, and the other achieves low-power operation by adopting a pattern-based PD with a simple baud-rate operation for short-link operation. The last two short papers describe advanced techniques to implement >100Gb/s PAM-4 and PAM-8 transmitters in 28nm and 40nm processes.

- In Paper 6.1, Broadcom demonstrates a flexible transceiver for a wide range of data rates up to 112Gb/s. The RX features a 3-tap FFE and an 18-tap DFE equalizer while the TX applies a 7b DAC driver with a 6-tap FFE achieving an exceptional RLM of 0.999. The PLL has a tuning range of 40-to-60GHz with only 0.12ps$_{rms}$ jitter measured at 56GHz. The transceiver is able to compensate up to 43.9dB channel loss at 112Gb/s. The 7nm CMOS TX/RX consumes 690mW and occupies a silicon area of 0.63mm per TX/RX.
- In Paper 6.2, MediaTek presents a high-performance DSP-based transceiver achieving 48dB and 49dB loss compensation at 112.5Gb/s and 106.25Gb/s data rates, respectively. The system has two PLLs per lane for fully independent TX and RX data rate programming. Robustness bottlenecks arising from PVT corner variation and EM coupling are handled by an adaptive biasing technique and a wideband PLL.
- In Paper 6.3, Peking University describes a 200Gb/s 5-tap delay-line-based receiver FFE in 28nm CMOS. The FFE employs on-chip grounded coplanar waveguides as delay lines and a one-stage topology for higher bandwidth and lower power. The h0 and h1 taps are implemented using distributed amplifiers to alleviate reflection. The RX FFE achieves a power efficiency of 0.43pJ/b and compensates for a 17.2dB-loss channel.
- In Paper 6.4, Samsung Electronics demonstrates a 32Gb/s per-lane die-to-die chiplet transceiver in 4nm CMOS. The TX features a reflection cancellation driver for suppressing the effects of impedance mismatch in the interposer. The RX has a 1-tap direct DFE with timing relaxation scheme for compensating ISI. The transceiver achieves a power efficiency of 0.44pJ/b and beach-front bandwidth density of 8Tb/s/mm.
- In Paper 6.5, MediaTek solves the difficult problem of a CDR to fast lock input data with significant frequency offset using 37.8dB channel loss. The proposed CDR uses a network of autocorrelators for flash frequency acquisition, adopting an open-loop frequency acquisition to avoid noise accumulation. The CDR achieves 0.6µs lock time independent of frequency offset under 37.8dB channel loss.
- In Paper 6.6, Korea University presents a 52Gb/s PAM-4 baud-rate CDR in 28nm CMOS for short-reach application. To attain a power-efficient RX by reducing the number of comparators, the proposed CDR employs a sampling method under a single-phase clock owing to a shared path of clock and data recoveries. The CDR achieves a power efficiency of 0.83pJ/b under 7.2dB channel loss while occupying only 0.011mm^2.
- In Paper 6.7, Peking University describes a 128Gb/s PAM-4 transmitter, which improves signal degradations due to transitions between non-adjacent levels and data dependent jitter. The transmitter has programmable-width pulse generation for edge-optimization and a pattern-dependent pre-emphasis for enlarging eye opening. A 28nm CMOS transmitter IC achieves a power efficiency of 1.4pJ/b and silicon area of 0.137mm^2.
- In Paper 6.8, Hanyang University demonstrates a 100Gb/s PAM-8 high-swing transmitter with 3-tap FFE in 40nm CMOS. The proposed transmitter adopts a current-mode driver with current bleeders and high-voltage protection cascode to obtain a 1.6V$_{ppd}$ output swing. The transmitter achieves a power efficiency of 3.35pJ/b with sufficient eye opening under 9.4dB channel loss and compact silicon area of 0.36mm^2.

Session 6 Highlights:
Advanced Wireline Links and Techniques

[6.1] A 112Gb/s Serial Link Transceiver With 3-tap FFE and 18-tap DFE Receiver for up to 43dB Insertion Loss Channel in 7nm FinFET Technology

Paper Authors: Bo Zhang[1], Anand Vasani[1], Ashutosh Sinha[2], Alireza Nilchi[1], Haitao Tong[1], Lakshmi Rao[1], Karapet Khanoyan[1], Hamid Hatamkhani[1], Alex Yang[1], Xin Meng[1], Alexander Wong[2], Jun Kim[2], Ping Jing[2], Yehui Sun[2], Ali Nazemi[1], Dean Liu[2], Anthony Brewster[1], Jun Cao[1], Afshin Momtaz[1]

Paper Affiliation: [1]Broadcom, Irvine, CA, [2]Broadcom, San Jose, CA

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

CONTEXT AND STATE OF THE ART

- The emergence of cloud computing, machine learning, and artificial intelligence drives rapid growth in data center bandwidth which approximately doubles every 4 years.
- This paper describes the industry's first non-ADC/DSP-based PAM-4 wireline transceiver supporting a wide range of operating speeds up to 112Gb/s and capable of operating under long reach channel loss over 40dB. This architecture comes with improved area and power efficiencies compared to other state-of-the-art designs with an ADC/DSP RX architecture.

TECHNICAL HIGHLIGHTS

- 112Gb/s analog-based (non-ADC/DSP) long reach (LR) transceiver in 7nm FinFET technology, featuring a fully adaptive 2-stage CTLE, 3-tap FFE (one pre, one post), and 18-tap DFE receiver, along with a 6-tap FFE 7b transmitter. In its RX, the number of samplers is reduced by a 1+D technique (also called Duo-PAM-4). The DFE feedback timing is relaxed to 2 UIs by a look-ahead technique. DFE taps 12-18 are floating and can be configured up to tap 38.
- The design achieves a high CDR bandwidth (thus better jitter tolerance) compared to the alternative ADC-based solutions with long CDR latencies coming from data parallelization and processing. The CDR bandwidth achieved is 16MHz, or 2× higher than similar published solutions. The clock path runs at quarter-rate, and supports sampling-based baud-rate and non-sampling-based operations.
- The PLL is shared among 8 RX/TX pairs and utilizes two LC VCOs to cover dual frequency bands with 5GHz overlap. The measured rms jitters are 0.12ps and 0.15ps for integer and fractional-N modes, respectively.
- Favorable area and power consumptions are achieved compared to other ADC/DSP-based 7nm transceivers published.

APPLICATIONS AND ECONOMIC IMPACT

- The long-reach 112Gb/s transceiver described in this paper targets switch, router, and other similar large-scale ASIC chips at the cores of most metro networks and data centers, which in turn deliver to ever-increasing demands of bandwidths for social media, video streaming and working-from-home applications. The power and chip area reduction enabled by innovation in the paper will be reflected in the cost-effectiveness of the end systems.

Session 7 Overview: SRAM Compute-In-Memory

Memory Subcommittee

Session Chair: Kyu-Hyoun (KH) Kim, IBM, Yorktown Heights, NY

Session Co-Chair: Violante Moschiano, Intel, Rome, Italia

Memory access has been the major bottleneck to system performance and energy consumption in both traditional (von Neumann) and nontraditional architectures (such as machine learning). A compute-in-memory (CIM) architecture eliminates this bottleneck by bringing the compute operations into the memory array. Innovations in CIM continue to improve energy and area efficiencies while maintaining overall AI network accuracy. This session showcases the 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 a hybrid-domain floating-point SRAM CIM macro. UC Santa Barbara introduces a continuous-time latch-based Ising machine implementation.

- In Paper 7.1, National Tsing Hua University introduces the first true floating-point hybrid-domain (analog & digital) SRAM CIM macro with 16.2 - 70.2TFLOPS/W in 22nm.
- In Paper 7.2, Southeast University presents a digital-domain floating-point computing-unit and a double-bit 6T-SRAM CIM macro with 31.6TFLOPS/W.
- In Paper 7.3, Tsinghua University presents an 8b multiply-less approximate digital SRAM CIM macro with 38 - 102TOPS/W.
- In Paper 7.4, TSMC presents a SRAM-based digital CIM macro with 6163TOPS/W/b and 4790TOPS/mm^2 in 4nm.
- In Paper 7.5, Southeast University presents a horizontal-weight and vertical-feature-shift-based 6T-SRAM CIM macro for depth-wise neural networks.
- In Paper 7.6, MediaTek presents an 8b word-wise analog CIM macro with 70.9 - 86.3TOPS/W.
- In Paper 7.7, Tsinghua University presents a XOR-derived similarity-aware CIM macro for cost volume construction.
- In Paper 7.8, Peking University presents a delta-sigma ($\Delta\Sigma$) SRAM CIM macro with 21.38TOPS/W with 8b precision.
- In Paper 7.9, UC Santa Barbara introduces the first continuous-time latch-based Ising machine with a 1440 spins.

Session 7 Highlights: SRAM Compute-In-Memory

[7.1] 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

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

Paper 7.1 Authors: Ping-Chun Wu[1], Jian-Wei Su[1,2], Li-Yang Hong[1], Jin-Sheng Ren[1], Chih-Han Chien[1], Ho-Yu Chen[1], Chao-En Ke[1], Hsu-Ming Hsiao[2], Sih-Han Li[2], Shyh-Shyuan Sheu[2], Wei-Chung Lo[2], Shih-Chieh Chang[2], Chung-Chuan Lo[1], Ren-Shuo Liu[1], Chih-Cheng Hsieh[1], Kea-Tiong Tang[1], Mengfan Chang

Paper 7.1 Affiliation: [1]National Tsing Hua University, Hsinchu, Taiwan, [2]Industrial Technology Research Institute, Hsinchu, Taiwan

Paper 7.9 Authors: Jooyoung Bae, Wonsik Oh, Jahyun Koo, Bongjin Kim

Paper 7.9 Affiliation: University of California, Santa Barbara, Santa Barbara, CA

Subcommittee Chair: Memory

CONTEXT AND STATE OF THE ART

- Memory access has been the major bottleneck to system performance and energy consumption in both traditional (von Neumann) and nontraditional architectures (such as machine learning). A compute-in-memory (CIM) architecture can eliminate this bottleneck by bringing the compute operations into the memory array.
- Thus far, mostly analog CIMs have been pursued as analog implementations, in principle, have a much higher energy efficiency compared to digital implementations. However, analog CIMs suffer from transistor variation and mismatch imposing limitations on bit precision, which eventually undermine the overall model accuracy.
- It is difficult to implement floating point operations in CIM by via the structural nature of repeating simple cells. Although research toward utilizing lower bit precision has progressed a lot of accuracy improvement, still some models and applications show a higher accuracy when using floating point operations.
- Machine learning (ML) and deep learning (DL) are the most popular and successful nontraditional computing algorithms to date. However, there are several other more-[radical non-conventional computing paradigms that are being researched. One such non-conventional method is the Ising compute (or machine) approach. An Ising machine can find the optimum solution in a similar way that quantum computers do, but rather using analog spins instead of qubits. There have been attempts at implementing an Ising Machine using the most well-known cost-efficient way – using CMOS. However, prior work used CMOS ring oscillators to model analog spin, which consume excessive energy.

TECHNICAL HIGHLIGHTS

- **National Tsing Hua University introduces the first true floating-point hybrid-domain (analog & digital) SRAM CIM macro with 16.2 - 70.2TFLOPS/W in 22nm.**
    - A 22-nm 832-kb Floating-Point SRAM CIM macro achieves 70.2TFLOPS/W at BF16-IN and BF16-W with FP32-OUT and up to 128 accumulations
- **UC Santa Barbara introduces the first continuous-time latch-based Ising Machine with a 1440 spins.**
    - The proposed continuous-time Ising machine with 40 × 36 (1440) spins enables fully-parallel operations via continuous-time operations and achieves a 1051× faster time-to-solution (< 20ns) than prior discrete-time solutions, while consuming 0.2 - 3nJ.

APPLICATIONS AND ECONOMIC IMPACT

- Highly-energy efficient CIM enables many new AI applications (image to speed) in various markets (edge to cloud).
- Energy-efficient CMOS Ising Machine with fast time-to-solution has a significant impact and the potential to create new applications in various areas and may fill the gap until quantum computers become practical.

Session 8 Overview:
GHz-to-Millimeter Wave Frequency Generation

RF Subcommittee

Session Chair: Swaminathan Sankaran, Texas Instruments Inc., TX.

Session Co-Chair: Mona Hella, Rensselaer Polytechnic Institute, NY.

High-quality frequency references are fundamental building blocks for enabling enhanced fidelity and throughput in 5G/emerging-xG communications and mm-wave radar/imaging. This session presents the latest advancements in high-quality frequency generation in bulk CMOS processes with potential impact on maturing 5G, emerging-xG, connectivity, and sensing applications while allowing for superior levels of integration and low-cost adoption.

- In Paper 8.1, Tsinghua University presents a Class-F VCO with inherent common-mode-noise self-cancelation and isolation achieving 192.8dBc/Hz FoM @ 1MHz offset at ~12GHz in a 65nm bulk CMOS process. With the introduction of common-mode drain-to-source coupling and subsequent impedance increase, the authors present reduction to noise-current injection resulting in excellent performance without any additional area overhead.
- In Paper 8.2, the University of Macau and the University of Lisboa demonstrate a dual-path synchronized quad-core oscillator in a 65nm bulk CMOS process. The oscillator achieves 193.3dBc/Hz FoM @ 10MHz offset at ~26GHz. Low phase-noise and high FoM are concurrently achieved through simultaneous synchronization of all cores that enables a true dB-per-dB improvement in phase noise with the number of cores.
- In Paper 8.3, the University of Electronic Science and Technology of China reports a uniquely scalable PMOS-NMOS topology with concurrent common- and differential-mode enhancement using capacitive and inductive coupling to freely scale oscillator cores while enabling compromise-free phase-noise improvement for mm-wave applications. The ~25-to-30GHz 4- and 20-core VCOs adopting the technique achieve 193.3dBc/Hz FoM @ 1MHz offset using 40nm CMOS.
- In Paper 8.4, the University of Macau presents a W-band 3rd harmonic extraction VCO employing a multi-resonance/core/mode technique in a 65nm bulk CMOS process. With the multi-pronged approach, the work demonstrates an excellent phase-noise FoM balance with a 21% tuning range.

Session 8 Highlights:
GHz-to-millimeter Wave Frequency Generation

[8.3] A 28GHz Scalable Inter-Core-Shaping Multi-Core Oscillator with DM/CM-Configured Coupling Achieving 193.3dBc/Hz FoM and 205.5dBc/Hz FoM_A at 1MHz Offset

[8.4] An 83.3-to-104.7GHz Harmonic-Extraction VCO Incorporating Multi-Resonance, Multi-Core, and Multi-Mode (3M) Techniques Achieving -124dBc/Hz Absolute PN and 190.7dBc/Hz FoM_T

Paper 8.3 Authors: Yiyang Shu, Huizhen Jenny Qian, and Xun Luo

Paper 8.3 Affiliation: University of Electronic Science and Technology of China, Chengdu, China

Paper 8.4 Authors: Hao Guo, Yong Chen, Yunbo Huang, Pui-In Mak, and Rui P. Martins

Paper 8.4 Affiliation: University of Macau, Macao, China

Subcommittee Chair: Jan Craninckx, imec, Leuven, Belgium

CONTEXT AND STATE OF THE ART

- High-quality VCOs are essential building blocks for high-throughput-communication and high-fidelity-sensing systems.
- A true dB-per-dB phase-noise scaling with the number of cores is challenging due to mismatch and limited synchronization, thereby degrading FoM.
- Wide tuning-range VCOs can save silicon area; however, such VCOs typically need to compromise phase-noise performance and/or power consumption.

TECHNICAL HIGHLIGHTS

- **The University of Electronic Science and Technology of China presents a unique, scalable oscillator topology with concurrent enhancement of differential- and common-mode passive quality factor demonstrating best-in-class FoM.**
    - Quality factor degradation is circumvented with innovative construction and optimal modal configuration of the LC tank allowing the 40nm bulk-CMOS VCOs to achieve >205dB area-normalized FoM.
- **The University of Macau reports a wide-tuning-range W-band VCO simultaneously manipulating the operating mode, resonance locations, and the number of cores to achieve an excellent phase-noise/FoM balance.**
    - The CMOS VCO in 65nm bulk CMOS achieves an FoM_T of ~190dBc/Hz @ 1MHz offset with a ~21% tuning range.

APPLICATIONS AND ECONOMIC IMPACT

- The reported VCOs achieve excellent phase-noise performance with compact areas. These techniques are desirable and could be used in communication systems such as 5/xG/mm-wave and sensing applications including mm-wave radar and imaging.

Session 9 Overview: Highlighted Chip Releases: Digital and Machine Learning Processors

Invited Industry Session

Session Chair: Alicia Klinefelter, NVIDIA, Durham, NC

Session Co-Chair: Vivek De, Intel, Hillsboro, OR

This session highlights four recent digital processor and system products spanning several applications. In addition to machine learning accelerators for both training and inference, papers describing design of an ultra-high bandwidth chip-to-chip link and an Infrastructure Processing Unit (IPU) for offloading datacenter management are presented. A range of techniques used in leading-edge systems, including testability and power reduction, are described. The papers also discuss the critical role of software support in hardware adoption. Integration, clocking, thermal, and power considerations are discussed in each paper in the context of its unique deployment and application.

- In Invited Paper 9.1, Tesla describes a 7nm ML training processor, D1, used in the DOJO exa-scale computer system. D1 contains 354 compute nodes clocked in a mesochronous wave distribution, and 576 lanes of 112Gbps SerDes for die-die communication in a system-on-wafer package providing a total of 362 TFlops of BFP16/CFP8 performance
- In Invited Paper 9.2, Syntiant presents their neural decision processor, the NDP200, designed using at-memory compute to minimize the movement of data for machine learning inference. When running the MobileNetV1 network on NDP200 at 5 frames per second, the device consumes only 0.83mW of average power when supplied from a 0.9V core supply
- In Invited Paper 9.3, NVIDIA presents a 5nm chip-to-chip link (NVLink C2C) that enables its Grace-Hopper and Grace Superchip systems, with 900GB/s link between Grace and Hopper, or between two Grace chips. It achieves 552Gbps/mm^2 bandwidth with 40Gbps/pin at 1E-15 or lower BER over a 12dB loss off-package channel.
- In Invited Paper 9.4, Intel describes an Infrastructure Processing Unit (IPU) that features a rich packet processing pipeline, significant RDMA and storage capability including NVMe offload, and an ARM Neoverse based compute complex enabling customer software to execute features ranging from complex soft dataplanes to storage transport to device management and telemetry.

Session 9 Highlights: Highlighted Chip Releases: Digital and Machine Learning Processors

[9.1] D1: A 7nm ML Training Processor with Wave Clock Distribution

Paper Authors: Tim C Fischer[1], Anantha Kumar Nivarti[1], Raghuvir Ramachandran[1], Ram Bharti[1], Derek Carson[1], Anton Lawrendra[1], Vineet Mudgal[1], Vivek Santhosh[1], Sunil Shukla[2], Te-Chen Tsai[1]

Paper Affiliation: [1]Tesla, Palo Alto, CA, [2]Tesla, Austin, TX

CONTEXT AND STATE OF THE ART

- Automotive is an emerging target application for IC designers and continues to gain momentum. Machine learning accelerators (especially for training) with hundreds of compute nodes tend to be very power-hungry and require extensive system integration effort.
- Clocking power comprises a large share of chip power budgets and di/dt mitigation is important for reducing power supply noise. Alternatives to globally synchronous clocking such as locally synchronous and globally mesochronous scheme used on DOJO ameliorate this noise.

TECHNICAL HIGHLIGHTS

- **Tesla describes their D1 ML training processor used in the DOJO exa-scale computer system.**
    - The 7nm ASIC has a wide operating range from 500MHz-2.3GHz and 576 lanes of high-speed SerDes. The clocking scheme provides frequency-dependent filtering of switching-induced current noise on the global supply rail.
    - The 50B transistors and 645mm^2 chip is modularly constructed, with a 20x18 array of compute and sensor nodes connected by abutment and without top-level routing.
    - The D1 node is a fully programmable 4-way multi-threaded 64b superscalar processor core achieving 362 BF16/CFP8 TFlops of compute at 2GHz.

APPLICATIONS AND ECONOMIC IMPACT

- Automotive applications such as full self-driving computers are incorporating more machine learning hardware for both inference and online training. Training has traditionally been accelerated using GPUs but is now starting to migrate to ASIC implementations for additional acceleration and power efficiency.
- The automotive electronics market is forecast to expand at an annual rate of 7.9% and in 2020, was valued at USD 217.86 billion.

Session 9 Highlights: Highlighted Chip Releases: Digital and Machine Learning Processors

[9.2] A 1mW Always-on Computer Vision Deep Learning Neural Decision Processor

Paper Authors: David Garrett, Youn-sung Park, Josh Kim, Jay Sharma, Wenbin Huang, Majid Shaghaghi, Vinay Parthasarathy, Stephen Gibellini, Stephen Bailey, Mallik Moturi, Kurt Busch, Pieter Vorenkamp Jeremy Holleman

Paper Affiliation: Syntiant, Irvine, CA

CONTEXT AND STATE OF THE ART

- Machine learning at the edge is gaining steam for low-latency and power-constrained applications such as computer vision, home automation, and mobile electronics. These devices are designed to operate without needing to send data to the cloud, and this helps address privacy concerns.
- Data movement is often cited as the largest source of power for machine learning hardware. Syntiant coordinates memory accesses versus focusing on logic to reduce power.

TECHNICAL HIGHLIGHTS

- **Syntiant presents an extremely efficient Neural Decision Processor designed to meet the high-speed inference needs of computer vision on the edge. When running the MobileNetV1 network at 5 frames per second, the device consumes <1mW of average power when supplied from a 0.9V core supply**
    - The 40nm ASIC implements the Syntiant Core 2 (SC2) architecture that achieves 740k cycles per inference and coordinates with an embedded Tensilica HiFi3 DSP and ARM Cortex M0 to provide additional compute.
    - The NDP200 is supported by Syntiant's Software Development Kit (SDK), which allows for pre-trained models to be deployed on the device, and Syntiant's Training Development Kit (TDK) which enables customers to train and package their own networks with quick feedback on predicted energy consumption.

APPLICATIONS AND ECONOMIC IMPACT

- Ultra-low power ML accelerators are considered TinyML, an emerging field of research. According to the Emerging Spaces review of Pitchbook, $26 million has been invested in TinyML since January 2020 and this area is expected to continue to grow.
- Modern CPUs consume between 65-85W and standard consumer GPUs consume anywhere between 200-400W. This is overkill for many applications looking to deploy machine learning and novel solutions are needed to meet upcoming demand and more complex networks.

Session 9 Highlights: Highlighted Chip Releases: Digital and Machine Learning Processors

[9.3] NVLink-C2C: A Coherent Off Package Chip-to-Chip Interconnect with 40Gbps/pin Single-ended Signaling

Paper Authors: [1]Ying Wei, [2]Yi Chieh Huang, [1]Haiming Tang, [1]Nithya Sankaran, [1]Ish Chadha, [1]Dai Dai, [1]Olakanmi Oluwole, [1]Vishnu Balan, [1]Edward Lee

Paper Affiliation: [1]Nvidia, Santa Clara, CA, [2]Nvidia, Hsinchu, Taiwan

CONTEXT AND STATE OF THE ART

- Chiplets and heterogeneous computing are predicted to offset the losses from a slowing Moore's Law. NVLink-C2C allows custom die and smaller chiplets to coherently interconnect to GPUs, SoCs, and NICs.
- Complex workloads such as AI and HPC require high-bandwidth, low-latency, and low-power connectivity. NVLink C2C provides a CPU-to-GPU coherent memory model for such applications.

TECHNICAL HIGHLIGHTS

- **NVIDIA presents a 5nm chip-to-chip link (NVLink C2C) that is the enabler for Nvidia's Grace-Hopper and Grace Superchip systems, with 900GB/s link between Grace and Hopper, or between two Grace chips.**
    - Grace and Hopper each incorporate 10 NVLink-C2C links to achieve 900GB/s of off-package bandwidth.
    - NVLink C2C pushes ground referenced single-ended link speed to 40Gbps with enough margin to withstand high-volume manufacturing (HVM) variations for off-package connections.
    - NVLink-C2C uses a clock-forwarded architecture with an innovative super wideband phase rotating phase locked loop in the receiver for per-lane skew adjustment and wideband jitter tracking between data and clock paths.

APPLICATIONS AND ECONOMIC IMPACT

- With the recent development of the Universal Chiplet Interconnect Express (UCIe) standard proposed by Intel, we're seeing a focus on heterogeneous compute including chiplet-based design.
- Fueled by the slowing of Moore's Law and certain designs being better suited to different technology nodes, the chiplets market is expected to cross US$47.2 billion by 2031.

Session 9 Highlights: Highlighted Chip Releases: Digital and Machine Learning Processors

[9.4] An In-depth Look at the Intel IPU E2000

Paper Authors: Brad Burres*[1], Naru Sundar*[2], Yadong Li*[3]

Paper Affiliation: [1]Intel, Hudson, MA, [2]Intel, Santa Clara, CA, [3]Intel, Portland, OR

CONTEXT AND STATE OF THE ART

- Infrastructure services such as virtual switching, security, and storage can consume a significant number of CPU cycles. Offloading these capabilities onto a dedicated processor frees up CPU resources.
- Server architectures today include disk with every server and since capacity is difficult to predict, disk storage is overprovisioned resulting in unused capacity. With an IPU, you can move to a diskless server architecture.
- Infrastructure accelerators exist today that are implemented using FPGA-based acceleration or the use SmartNICs, but an ASIC implementation maximizes performance and energy-efficiency.

TECHNICAL HIGHLIGHTS

- **Intel presents their Infrastructure Processing Unit (IPU) that was co-designed with Google and enables a rich variety of services to be provided to CPU-attached clients such as GPUs or FPGAs. The IPU consists of both a data and networking complex interconnected by a high-performance system level cache.**
    - The networking complex is a streaming architecture consisting of client facing PCIe logic and protocol engines, a flexible and high-performance packet processing engine, an inline crypto engine and a traffic shaping engine.
    - The compute complex includes the 16 ARM Neoverse N1 cores, a dedicated lookaside crypto and compression engine, a high bandwidth memory interface consisting of 3 channels of either DDR4 or LPDDR4, and a dedicated system management engine.

APPLICATIONS AND ECONOMIC IMPACT

- The data center market size is estimated to grow by USD 615.96 billion from 2021 to 2026, growing at a CAGR of 21.98%. The IPU supports the rapid innovation necessary for modern and future data centers.

Session 10 Overview: Pipelined and Noise-Shaping ADCs

Data Converters Subcommittee

Session Chair: Nima Maghari, University of Florida, Gainesville, FL

Session Co-Chair: Ping Gui, Southern Methodist University, TX

Architectural and implementation improvements in both Nyquist and oversampling converters continue to push the performance envelope in terms of speed, resolution, and power consumption. Innovative circuits such as ring-amplifiers and floating inverter amplifiers combined with other scaling friendly approaches such as time-domain quantizers and noise-shaping SARs are some of the driving forces in these works. The first three papers target beyond GS/s medium-to-high resolution ADCs for modern direct wireless and wireline receivers. These are followed by four papers with very high resolution and in a variety of bandwidths targeting a broad range of applications from sensor interface to wireless receivers.

- In Paper 10.1, MediaTek presents a 2× time-interleaved 12b ADC leveraging a pre-sampling technique with robust PVT radix utilizing residue voltage stacking. This chip achieves 60dB SNDR in 800MHz BW in 7nm CMOS with 7.55mW power consumption.
- In Paper 10.2, University of Macau presents a PVT-robust single-channel 2.6GS/s ADC leveraging a robust time-domain quantization and early release of CDAC resulting in relaxed timing of the frontend sampling. Fabricated in 28nm CMOS, this ADC achieves 51.4dB SNDR with 13.9mW power consumption.
- In Paper 10.3, University of Macau shows single-channel 2GS/s ADCs with a critically damped ring-amplifier, which employs two additional diode-tied inverters in the chain to remove the undesired ringing and enhance settling performance. This ADC achieves an SNDR of 60dB in Nyquist bandwidth with 27mW power consumption in 28nm CMOS.
- In Paper 10.4, Oregon State University presents a predictive level-shifting technique to move the input signal to the linear range of the input buffer, therefore enhancing the front-end linearity and swing. The fabricated ADC operates at 12MS/s and achieves 94dB SNDR with 30.4mW in 180nm.
- In Paper 10.5, the University of Macau introduces a gain-error shaping combined with NS-SAR in an oversampled pipelined ADC, which relaxes the constraints on the loading of the residue amplifier. This prototype achieves 77dB SNDR in 25MHz bandwidth with 2.03mW power consumption in 28nm CMOS.
- In Paper 10.6, Peking University describes a fully dynamic incremental zoom ADC with noise-shaping SAR quantizer, requiring only a single dynamic buffer to achieve N^{th}-order loop filter. The ADC achieves peak SNDR of 92.5dB in 150kHz bandwidth and power consumption of 160µW in 28nm.
- In Paper 10.7, University of Macau reveals a 100MHz 4^{th}-order noise-shaping hybrid pipelined-SAR with residue amplifier error shaping robust across voltage and temperature variations. This work obtains 70dB SNDR and 90dB SFDR with 4.5mW power consumption in 28nm CMOS.

Session 10 Highlights: Pipelined and Noise-Shaping ADCs

[10.1] A 1.8GHz 12b Pre-Sampling Pipelined ADC with Reference Buffer and OP Power Relaxations

[10.3] A Single-Channel 12b 2GS/s PVT-Robust Pipelined ADC with Critically Damped Ring Amplifier and Time-Domain Quantizer

Paper 10.1 Authors: Sung-En Hsieh, Tzu-Chien Wu, Chun-Chih Hou

Paper 10.1 Affiliation: MediaTek, Hsinchu, Taiwan

Paper 10.3 Authors: Yuefeng Cao[1], Minglei Zhang[1], Yan Zhu[1], Chi-Hang Chan[1], Rui P. Martins[1,2]

Paper 10.3 Affiliation: [1]University of Macau, Macau, China, [2]University of Lisboa, Lisbon, Portugal

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

CONTEXT AND STATE OF THE ART

- High-speed and wide dynamic range enable direct RF sampling, which can simplify a broad range of wireless applications such as 5G and beyond and WiFi 6 communications.
- Leveraging time-domain conversion and quantization in nano-scale CMOS nodes brings forth the digital scaling advantages into the mixed-signal and ADC domains.

TECHNICAL HIGHLIGHTS

- **MediaTek introduces a gigahertz ADC for direct RF sampling in 7nm CMOS achieving excellent SNDR with minimal power consumption.**
    - A 1.8GHz 12b Pre-Sampling Pipelined ADC with Reference Buffer and OP Power Relaxations.
- **The University of Macau presents a process, voltage, temperature (PVT)-robust pipelined ADC in 28nm CMOS leveraging a new critically damped ring amplifier and time-domain quantization.**
    - A Single-Channel 12b 2GS/s PVT-Robust Pipelined ADC with Critically Damped Ring Amplifier and Time-Domain Quantizer.

APPLICATIONS AND ECONOMIC IMPACT

- New generations of high-speed low-power ADCs enable direct RF sampling thereby minimizing the cost and area overhead of traditional receiver chains.
- Time-domain processing and conversion leverages process scaling in analog and mixed-signal domains.
- Innovative robust circuit architectures enable stable circuit operation in a variety of external conditions.

Session 11 Overview: USB and Compute Power Delivery

Power Management Subcommittee

Session Chair: Harish Krishnamurthy, Intel, Hillsboro, OR

Session Co-Chairs: Chan-Hong Chern, TSMC, Hsinchu, Taiwan

Improved power density, efficiency, and transient response have always been the key performance metrics that modern day DC-DC converters are measured by, while powering various applications such as multi-core microprocessors, energy harvesting and direct battery-attached systems, automotive electronics, and LED drivers. Novel topologies that interface higher input large conversion ratios while minimizing inductor current, maximizing inductor slew rate for quicker transient response while maintaining the high-power conversion efficiency, fully integrated power converters for maximum power density and alternative approaches to sense inductor current in traditional buck converters are some of the techniques presented in this session demonstrating the latest in DC-DC converters at both system and circuit levels.

- In Paper 11.1, University of California San Diego and Murata present a heterogeneously integrated vertical power delivery system that combines a 0.18μm HV BCD 20V:4V hybrid converter, a 65nm 4V:1V switched-capacitor converter, and an interposer filled with 3D nano-porous deep-trench flying and decoupling capacitors with a total efficiency of 86% from 20V:0.9V.
- In Paper 11.2, University of Macau and Instituto Superior Tecnico/University of Lisboa present a 12V to 1V quad-output switched-capacitor buck converter with shared DC capacitors achieving a peak efficiency of 90.4%, and 48mA/mm3 power density at 85% efficiency.
- In Paper 11.3, Intel presents a high-power (1.8W) high-frequency SIMO converter in 16nm FinFET CMOS that operates in DCM at 10MHz with an ultra-small 5~10 nH inductor, delivering power to 4 outputs (1.4V/1.2V/1V/0.8V), featuring a digital "sensor-less" computational zero-current detector with low overheads to achieve zero-current operation on a per cycle basis.
- In Paper 11.4, National Yang Ming Chiao Tung University presents a double step-down dual-output converter with hybrid sum and deviation technique can achieve 25mV of output ripple with a cross regulation performance of 0.025mV/mA. A pulse-based level shifter is proposed improving transition speed and robustness with low voltage stress transistors.
- In Paper 11.5, University of California San Diego presents a reconfigurable step-down hybrid converter that couples two synchronous switched-capacitor stages with a single inductor to efficiently support 1-cell battery charging from a 5V to 24V supply. The design delivers 5A maximum current and achieves peak efficiencies of 94.8% and 92.4% from 5V/24V respectively
- In Paper 11.6, University of Macau and Instituto Superior Tecnico/University of Lisboa present a bidirectional power delivery solution, integrating the charging IC and its passives onto the Type-C connector, thereby eliminating VR and its associated thermals from the platform. The converter achieves a peak efficiency of 95.1% at 22.8W without the cable loss at a power density of 3.17W/mm^2 due to fewer off-chip components with a maximum power of 42W.
- In Paper 11.7, National Yang Ming Chiao Tung University, Chip-GaN Power Semiconductor, and Realtek Semiconductor propose a symmetric hybrid buck-boost converter with a wide-range conversion ratio (CR=0.1-to-10) to support USB PD 3.1 (5-to-48V) requirement. The design achieves a peak efficiency of 95.4% at I_{LOAD}=1.2A in the buck mode with a fast transient detector that achieves a recovery time of 4.5μs for a load step from 1 to 3A.
- In Paper 11.8, Zhejiang University proposes a 9~16V-to-1V dual-path series-capacitor converter, with the lowest VA metric. The design achieves a peak efficiency of 94.5% at 5A load with a die-area current density of 1.30A/mm^2.
- In Paper 11.9, University of Macau and Instituto Superior Tecnico/University of Lisboa propose a hybrid resonant SC parallel inductor buck converter that effectively reduces inductor currents to 0.5-to-0.67 of the output current for high VCRs ranging from 10 to 20. The design achieves a peak efficiency of 91.8% for a 12-to-1.2V conversion with a power inductor as small as 2.5×2×1.2mm^3.
- In Paper 11.10, University of Macau, Zhuhai UM Science & Technology Research Institute and Instituto Superior Tecnico/University of Lisboa propose a dual-inductor quad-path hybrid buck converter that can reduce inductor DC current by 30% and inductor current ripple by 18% at VCR=0.15 (12-to-1.8V). Measurement results show a peak efficiency of 93.7% at VCR=0.15, and a maximum current density of 0.19A/mm^2.

Session 11 Highlights: USB and Compute Power Delivery

[11.1] A Scalable Heterogeneous Integrated Two-Stage Vertical Power Delivery Architecture for High Performance Computing

Paper Authors: Casey Hardy[1], Hieu Pham[1], Mohamed Mehdi Jatlaoui[2], Frederic Voiron[1], Tianshi Xie[1], Po-Han Chen[1], Saket Jha[1], Patrick Mercier[1], Hanh-Phuc Le[1]

Paper Affiliation: [1]University of California, San Diego, CA, [2]Murata, Caen, France

[11.6] A 42W Reconfigurable Bidirectional Power Delivery Voltage-Regulating Cable

Paper Authors: Zhiguo Tong, Junwei Huang, Yan Lu, Rui P. Martins

Paper Affiliation: University of Macau, Macau, China

Subcommittee Chair: Bernhard Wicht, Microelectronic Systems, University of Hannover, Germany

CONTEXT AND STATE OF THE ART

- **Compute Power Delivery**: Emerging high-performance computing microprocessors demand increasingly large peak currents and bringing the voltage regulators (VRs) closer to the point of load (POL) is becoming critical to avoid power distribution losses. High-voltage regulators are currently too large and bulky to even come remotely close to the microprocessor resulting in significant efficiency penalty (~5-8%).
- **USB Power Delivery**: Fast charging our mobile devices is becoming common practice but thermal management is increasingly becoming a bottleneck to make the charging even faster. USB cables that help us charge our devices are becoming part and parcel of our daily lives and VR integration in the cable itself is becoming an attractive approach.

TECHNICAL HIGHLIGHTS

- **Compute Power Delivery: A 2-stage approach, using a high-voltage regulator module (HVRM) followed by a switched-capacitor voltage regulator (SCVR), is proposed that leverages heterogeneous integration to confine the highest currents near the POL while simultaneously supporting larger input voltages.**
    - The converter achieves a peak efficiency of 93.1% for the HVRM (1st stage) and 89.9% for the SCVR (2nd stage), resulting in the 2-stage solution achieving 86% when converting a 20V input to a 0.9V output.
- **USB Power Delivery: A bidirectional power delivery solution is proposed that integrates the charging IC and its passives onto the USB Type-C connector, thereby eliminating the VR and its associated thermals from the platform.**
    - The converter achieves a peak efficiency of 95.1% at 22.8W without the cable loss at a power density of 3.17W/mm^2 due to fewer off-chip components with a maximum power of 42W.

APPLICATIONS AND ECONOMIC IMPACT

- Power delivery technologies are a key enabler of sustainable high-power data centers as well as improved-form-factor battery-powered mobile and handheld systems of the future. With breakthroughs in advanced packaging, and topologies that can exploit devices and passives in unconventional ways, the contributions of the two papers highlighted above provide excitement that power delivery and management can shape the future of both the edge devices and the data centers.

Session 12 Overview: High Performance Optical Receivers

WLN Subcommittee

Session Chair: *Byungsub Kim, POSTECH, Pohang, Korea,*

Session Co-Chair: *Thomas Toifl, Cisco Systems, Thalwil, Switzerland*

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

Data rate and power efficiency of optical links have surged impressively in recent years, making them the work horse communication medium for data center networking and high-performance computing. The papers in this session reflect key developments on how to increase data rates as well as how to enable low-power coherent optical solutions. The first paper in the session demonstrates a low-power ring-resonator-based silicon photonic WDM receiver module at 350Gb/s aggregate data rate. The second paper describes the state of the art for an optical RX frontend, where -14dBm sensitivity was achieved for a TIA operating with a 106.25Gb/s PAM-4 signal. Analog implementations of coherent optical links are attractive since they enable lower power consumption and smaller area than DSP-based designs. The third paper in this session demonstrates how to achieve a low-latency, high-bandwidth carrier-phase recovery (CPR) loop for a coherent optical RX using a purely analog approach.

- In Paper 12.1, AMD demonstrates a 7×50Gb/s NRZ WDM receiver module, which achieves 350Gb/s aggregate data rate at <1e-12 BER without forward error correction. The RX incorporates stacked 7nm CMOS and 45nm silicon photonic dies and uses an array of cascaded optical ring resonators to receive optical data on a 1.5nm-spaced laser grid. High sensitivity (-11.1dBm median) was measured at 0.96pJ/b energy efficiency.
- In Paper 12.2, Cisco Systems and University of Illinois at Urbana-Champaign describe a linear PAM-4 CMOS differential TIA with asymmetric signal paths utilizing signal currents from both terminals of the PD to improve the SNR by 3dB. Designed in 16nm FinFET, the TIA has a gain of 77dBΩ, a bandwidth of 18.4GHz, and an input-referred noise density of 7pA/√Hz to achieve a -14dBm optical sensitivity at 106.25Gb/s.
- In Paper 12.3, University of Illinois at Urbana-Champaign describes a wide bandwidth analog carrier-phase recovery (CPR) method for short-reach QPSK coherent optical links. Using 16-phase switched-inverter-based harmonic-rejection complex mixers (HRMs) and low-latency phase detection circuits, the prototype QPSK receiver fabricated in 28nm CMOS achieves 100MHz bandwidth, 600MHz tracking range, and recovers 24Gb/s QPSK data without errors. The power efficiency of the QPSK coherent receiver is 3.2pJ/b.

Session 12 Highlights:
High Performance Optical Receivers

[12.1] A 0.96pJ/b 7×50Gb/s-per-fiber WDM Receiver with Stacked 7nm CMOS and 45nm Silicon Photonic Dies

Paper Authors: Mayank Raj[1], Chuan Xie[1], Ade Bekele[1], Adam Chou[1], Wenfeng Zhang[1], Ying Cao[1], Jae Wook Kim[1], Nakul Narang[2], Hongyuan Zhao[2], Yipeng Wang[2], Kee Hian Tan[2], Winson Lin[1], Jay Im[1], David Mahashin[1], Santiago Asuncion[1], Parag Upadhyaya[1], Yohan Frans[1]
Paper Affiliation: [1]AMD, San Jose, CA, [2]AMD, Singapore, Singapore
Subcommittee Chair: *Yohan Frans, AMD, San Jose, CA*

CONTEXT AND STATE OF THE ART

- Applications such as machine learning, high-performance computing, and cloud storage continue to drive demand in data transmission throughputs in and between data centers. Servers in data centers are connected through fiber optics to avoid bandwidth limitations of long reach wireline links.
- The electrical-to-optical conversions (and the other way around) happen in optical modules that are mounted on the PCB some distance away from the logic chips that generate and use the data. A significant amount of system power is consumed in long-reach transceivers to move the data between the logic chips and the optical modules across the PCB.
- To eliminate inefficiencies in existing solutions, future systems will be made of chips with co-packaged optics, which perform the electrical-to-optical conversion inside the logic chip package.

TECHNICAL HIGHLIGHTS

- This paper describes a high-performance receiver on a 45nm silicon-photonics die with a 7nm electrical die stacked on top. This design uses wavelength division multiplexing (WDM) to pack 7 wavelengths of laser light in a fiber. Each wavelength carries an NRZ data stream at 50Gb/s.
- The receiver achieves low-power (0.96pJ/b), high sensitivity (-11.1dBm median), and less than 1e-12 BER without forward-error-correction. With its aggregated throughput of 350Gb/s per fiber, the performance of this design is state of the art.
- The optical die has V-grooves to couple in light from the fiber. It implements wavelength selection with an array of cascaded ring resonators, collectively thermal-tuned, to provide second-order filtering with a flat-top characteristic centered around the target wavelength. This achieves crosstalk less than -15dB from the adjacent wavelength spaced 1.5nm away. Ring resonator output drives a SiGe photo detector.
- The electrical die RX has 3-stage inverter-based TIA with T-coil peaking to improve power efficiency while meeting a transimpedance of >4kΩ, bandwidth >35GHz and 3μA integrated input-referred noise. TIA bandwidth and peaking are programmable. A feedback loop stabilizes the photo detector bias. An on-chip voltage regulator improves power supply rejection of the TIA to greater than 25dB.
- A phase-interpolator-based CDR loop provides quarter-rate clocks to slicers and de-serializers. The phase-interpolator is placed between two injection-locked local oscillators in the clock generation. An area-efficient OTA-free scheme implements quadrature phase correction. Robustness of the design is proven by measurement showing 0.14UI eye opening at 1e-12 BER over all 7 wavelength receivers.

APPLICATIONS AND ECONOMIC IMPACT

- The work in this paper achieves the highest aggregate data-rate and the highest power efficiency compared to previously reported silicon photonics designs. Overall, this paper has demonstrated feasibility of merging the electrical dies with fiber optics communication dies directly in the same package, which brings the system integration to the next level. The silicon photonics solution as described in this paper will be key to enable much higher efficiency for future computing and communication systems.

Session 13 Overview: Ideas for the Future

Technology Directions Subcommittee

Session Chair: Sudip Shekhar, University of British Columbia, Vancouver, Canada

Session Co-Chair: Daniel Morris, Meta, Menlo Park, CA

This session illustrates the scope of what integrated circuits may become in the future. The back-end-of-line (BEOL) hosts 3D memory in X.1 and dielectric waveguides in X.5. Signal processing is shown with low-power mixed-signal circuits in X.2 and with reconfigurable photonics in X.6. Energy harvesting circuit operation is enabled in 0.1mm-class systems in X.3 and in wafer scribe lines in X.4.

- In Paper 13.1, SEL and Fukuoka University present a 3D AI SoC with one tier of CMOS logic, and two tiers of oxide semiconductor memory provide retention, enabling intermittent operation at 25.15uW power.
- In Paper 13.2, University of Macau and Instituto Superior Tecnico/University of Lisboa describe a mixed-signal voice-activity detector consuming 47nW of power and achieving >92% accuracy rate. The chip employs RNN-based classifier and ROM/capacitor-based multiply-accumulate function.
- In Paper 13.3, Osaka University and Kobe University describe a system consisting of clusters of 0.1mm-class battery, power and data link ICs in close proximity, leveraging Inductive-Coupling Power-Line Communication. A temperature sensor application is demonstrated as a use case.
- In Paper 13.4, Osaka University leverages the high-energy plasma into antennas built in a dicing street to power up circuits during a semiconductor fabrication process. An oxide-breakdown PUF implementation is demonstrated.
- In Paper 13.5, Caltech introduces a technique to add photonic waveguides to a standard bulk CMOS process. BEOL doped glass waveguides are coupled to photodiodes and TIAs to operate over a broad range of wavelengths including visible and IR.
- In Paper 13.6, University of Delaware presents a reconfigurable optical analog processor based on a large-scale optical mesh fabricated in CMOS-compatible silicon photonics process.

Session 13 Highlights: Ideas for the Future

[13.1] Crystalline Oxide Semiconductor-based 3D Bank Memory System for Endpoint Artificial Intelligence with Multiple Neural Networks Facilitating Context Switching and Power Gating

Paper Authors: Yuto Yakubo[1], Kazuma Furutani[1], Kouhei Toyotaka[1], Haruki Katagiri[1], Masashi Fujita[1], Munehiro Kozuma[1], Yoshinori Ando[1], Yoshiyuki Kurokawa[1], Toru Nakura[2], Shunpei Yamazaki[1]

Paper Affiliation: [1]Semiconductor Energy Laboratory, Atsugi, Japan, [2]Fukuoka University, Fukuoka, Japan

Subcommittee Chair: Ali Hajimiri, California Institute of Technology, CA, Technology Directions

CONTEXT AND STATE OF THE ART

- Endpoint devices often need to support multiple neural networks but still consume very low power in a small die area. Context switching must be done with low latency.
- However, the requirement of large-scale SRAM memory to support multiple neural networks leads to power and area overhead.

TECHNICAL HIGHLIGHTS

- **The test chip uses two layers of crystalline oxide FET memory stacking on top of the CMOS layer, with a layer selection device fabricated in the OSFET layer.**
    - Inference of two neural networks can be achieved without rewriting weight data; power gating time is extended as well.
    - An average power of 25.15uW is achieved, which is 79% lower than a chip using SRAM.

APPLICATIONS AND ECONOMIC IMPACT

- Stacking two layers of crystalline oxide memory promises additional scalability to future technologies with even more layers.
- Improving power, area and latency for neural networks promise agile endpoint devices.

Session 14 Overview:
Digital Techniques for Clocking and Power Management

Digital Circuits Subcommittee

Session Chair: Arijit Raychowdhury, Georgia Institute of Technology, Atlanta, GA

Session Co-Chair: Eric Jia-Wei Fang, MediaTek, Hsinchu, Taiwan

Subcommittee Chair: Keith Bowman, Qualcomm, Raleigh, NC

In this session, four papers highlight developments in digital clocking and power management. The session opens with a paper on the design of a clock generator with an injection-error scrambler utilizing two cascaded DTCs to achieve state-of-the-art fractional and reference spurs. The next paper describes a fractional output divider with an auxiliary PLL which subtracts the carrier to eliminate the spur from the DTC output efficiently. The last two papers on power management demonstrate: 1) a 10nm digital LDO to achieve a current density over one hundred amperes per mm^2 and a smaller V_{IN}-induced V_{MIN} variation, and 2) a DC-DC converter with phase-merging turbo to improve the current delivery capability of continuously scalable conversion-ratio switched-capacitor voltage regulators.

- In Paper 14.1, the University of Southern California presents a fractional-N MDLL with an injection error scrambler utilizing two cascaded DTCs to break the injection error periodicity. The approach achieves 800fs jitter and -67dBc fractional spur at 1.5GHz output frequency, consuming 13.56mW in an area of 0.23mm^2 in 65nm CMOS.

- In Paper 14.2, Tsinghua University introduces a fractional output divider with an auxiliary PLL which subtracts a carrier to eliminate a spur from the DTC output efficiently. The divider achieves a worst-case fractional spur of -80dBc, corresponding to >40dB improvement from no calibration at 10-300MHz. The chip consumes 0.084mm^2 in 28nm CMOS.

- In Paper 14.3, Intel researchers present a 10nm digital LDO to improve the current density and to reduce V_{MIN} variability over V_{IN} variations. The test chip integrates the LDO and a multi-domain CPU, demonstrating a current density of 125A/mm^2 in 0.297mm^2, approximately 3ns loop delay, a maximum droop of 198mV for a 30A load transient, and reduced V_{IN}-induced V_{MIN} variation of less than $\pm$10mV.

- In Paper 14.4, Intel researchers present a technique to improve the current capability of continuously scalable conversion ratio switched-capacitor converters and a method to gang multiple regulators onto a shared output domain without communication. The methods are demonstrated on a switched-capacitor voltage regulator of 0.0093mm^2 active area fabricated in a 4nm class CMOS process with up to 4 cores, resulting in a record 26A/mm^2 peak current density and up to 88.5% efficiency, while sustaining V_{IN} droops of 180mV or load transients of more than 200mA without observable impact on V_{OUT}.

Session 14 Highlights:
Digital Techniques for Clocking and Power Management

[14.1] A Fractional-N Digital MDLL with Injection Error Scrambling and Background Third-Order DTC Delay Equalizer Achieving –67dBc Fractional Spur

Paper Authors: Qiaochu Zhang, Mike Chen

Paper Affiliation: University of Southern California, Los Angeles, CA

Subcommittee Chair: Keith Bowman, Qualcomm, Raleigh, NC, Digital Circuits Subcommittee

CONTEXT AND STATE OF THE ART

- Ring oscillator (RO)-based on-chip clock generators enable more flexible integration and operation for SoC building blocks, such as processors, memories, I/O interfaces and power management.
- Multiplying delay-locked loops (MDLLs) allow high-frequency non-integer multiplication of a low frequency input reference clock. Digital-to-time converters (DTCs) effectively suppress the phase noise of the RO with reference injection. These circuits, however, often suffer from unwanted spurious tones (spurs) and jitter due to the nonlinearity of DTCs.

TECHNICAL HIGHLIGHTS

- **The University of Southern California presents a fractional-N MDLL with an injection-error scrambling and background DTC nonlinearity calibration, which achieves 800fs RMS jitter and -67dBc fractional spur at 1.5GHz output frequency in 65nm CMOS.**
    - The MDLL utilizes two cascaded DTCs to completely break the periodicity of injection error and realizes low reference spur and fractional spur of -58dB and -67dB, respectively.
    - A 3rd-order calibration technique for DTC offset, gain and nonlinearity is proposed to minimize noise floor elevation due to proposed injection error scrambling and to preserve low jitter.

APPLICATIONS AND ECONOMIC IMPACT

- Low-jitter, minimally spurious and low-cost clock generators boost both the design flexibility and power efficiency of various SoCs, enabling wider dynamic circuit operation for higher performance and lower power.
- The proposed scrambler significantly reduces the jitter and spurs for frequency generation schemes, including DTC-based all-digital phase-locked loops (ADPLLs) and injection-locked PLLs.

Session 15 Overview: IoT & Security

Digital Architectures and Systems Subcommittee

Session Chair: Chiraag Juvekar, Apple, Cupertino, CA

Session Co-Chair: Ingrid Verbauwhede, KU Leuven, Leuven, Belgium

This session highlights two important applications of modern digital architectures: IoT and security. The first three papers address IoT topics. The first paper describes an ultra-low-power SoC that can be distributed along a fiber with harvested energy. The second paper presents a self-powered SoC with integrated multimodal energy harvesting. The third paper describes a low-power vehicle communication gateway. The last two papers address security topics: the first presents an accelerator for a Paillier homomorphic encryption scheme; the second presents several fault-attack countermeasures applied on an AES engine in a 4nm process.

- In Paper 15.1, University of Virginia and MIT Lincoln Laboratory present a 65nm system-in-fiber SoC featuring distributed cooperative energy harvesting and multi-chip power management. The SoC draws 33nW power for the whole chip under 92lux light and a 2.7nW minimum control power for the energy harvesting and power management unit.
- In Paper 15.2, Everactive describes a 2.19μW self-powered SoC in 55nm that can simultaneously harvest from multiple sources, and cold starts from poor harvesting conditions. The -92dBm wake-up receiver sensitivity allows for up to 2km in an industrial environment, while allowing the system to maintain a small form factor of 51cm^3.
- In Paper 15.3, Renesas presents a vehicle communication gateway SoC in 12nm that achieves energy-efficient Gb Ethernet connectivity, <50ms CAN bus startup for car control, <2mW stand-by power for battery savings and 33DMIPS for application processing. The SoC features a dedicated network accelerator to handle three 2.5Gbps Ethernet streams, with energy-efficiency of 10Gbps/W and tagging to identify sensitive packets for secure processing.
- In Paper 15.4, Tsinghua University, Xi'an JiaoTong University and Polar Bear Tech introduce a cryptographic processor to accelerate the Paillier homomorphic encryption scheme enabling efficient computation on encrypted data. The 43mm^2 test-chip in 28nm runs at 500MHz and supports a throughput of 68MOPS at 0.18μJ/Op.
- In Paper 15.5, Intel presents a 4nm CMOS fault-attack resistant AES 256 engine. It combines several fault-attack countermeasures, including an AES-specific arithmetic and parity checker, a byte-interleaved register placement and a distributed all-digital laser attack detection circuit. It results in a 99.1% fault coverage against raster- and box-scan laser injection attacks and a 99.99% fault coverage against undervoltage attacks.

Session 15 Highlights: IoT & Security

[15.1] A Self-Powered SoC with Distributed Cooperative Energy Harvesting and Multi-Chip Power Management for System-in-Fiber

[15.3] A 33kDMIPS 6.4W Vehicle Communication Gateway Processor Achieving 10Gbps/W Network Routing, 40ms CAN Bus Start-Up and 1.4mW Standby Power

Paper 15.1 Authors: Xinjian Liu[1], Daniel S. Truesdell[1], Omar Faruqe[1], Lalitha Parameswaran[2], Michael Rickley[2], Andrew Kopanski[2], Lauren Cantley[2], Austin Coon[2], Matthew Bernasconi[2], Tairan Wang[2], Benton H. Calhoun[1]

Paper 15.1 Affiliation: [1]University of Virginia, Charlottesville, VA, [2]MIT Lincoln Laboratory, Lexington, MA

Paper 15.3 Authors: Kenichi Shimada, Keiichiro Sano, Kazuki Fukuoka, Masayuki Daito, Tatsuya Kamei, Hiroyuki Hamasaki, Yasuhisa Shimazaki

Paper 15.3 Affiliation: Renesas Electronics, Tokyo, Japan

Subcommittee Chair: Thomas Burd, AMD, Santa Clara, CA, Digital Architectures and Systems Subcommittee

CONTEXT AND STATE OF THE ART

- Ultra-low power is a strong requirement for smart wearables and edge IoT devices. In an automotive context, low-power enables passive cooling and enables operation across a wide temperature range.
- Novel automotive applications such as in-car infotainment and autonomous driving require an increasing amount of electronic content. There is a need for energy-efficient computing and networking to support these emerging applications.
- With the proliferation of electronics in mobile devices, cyber-physical systems, IoT and other devices, the need for hardware security only grows. This requires highly efficient implementations of current and next-generation cryptographic algorithms, i.e. symmetric/public-key algorithms, as well as new algorithms that enable computing on encrypted data.
- At the same time, these devices need protection against physical manipulation, especially passive side-channel and active-fault attacks. In the context of side-channel attacks. As the attacker capabilities increase, the countermeasures need to become more robust. Countermeasures exist at the circuit level, e.g. intelligent circuits for detecting intrusive attacks, as well as at the mathematical level, e.g., masking countermeasures.

TECHNICAL HIGHLIGHTS

- **A team of researchers from the University of Virginia and MIT Lincoln Labs present an ultra-low-power system-in-fiber SoC that can be distributed along a fiber with harvested energy and cooperative power management with other distributed SoCs.**
    - This work implements a self-powered fully autonomous SoC that is capable of simultaneously harvesting energy, cooperatively scaling performance, sharing power and booting-up with other distributed SoCs in-fiber.
    - A 65nm SoC achieves 33nW full-chip power from the harvesting input and can maintain continuous sensing under 92lux light and a 2.7nW minimum control power for the energy harvesting and power management unit.

- **Renesas Electronics presents a vehicle communications gateway SoC with a dedicated security-aware network accelerator to achieve energy-efficient secure Gb connectivity.**
    - The work uses a dedicated network hardware accelerator to process three 2.5Gbps Ethernet streams with packet tagging to identify sensitive data for secure processing.
    - The 61mm^2 SoC is fabricated in 12nm and achieves <50ms CAN bus startup time for responsive control, <2mW stand-by power to optimize battery life and 33DMIPS for application processing.

APPLICATIONS AND ECONOMIC IMPACT

- Intelligent nano-watt ICs and new form factors can open up a wide range of new applications, such as wearable electronics and healthcare monitoring.
- To scale to a trillion IoT nodes, devices must untether from batteries. Such intelligent nano-watt ICs, when augmented with energy harvesting circuits, can help usher in the next generation of innovative, battery-less IoT sensor nodes
- Energy-efficient communication and security are key aspects of automotive system-on-chip design. With an increasing focus on connectivity, effective countermeasures to avoid exploits of security vulnerabilities are a prerequisite to protect data and preserve safety.
- Resilience to physical and cryptanalytic attacks has become a major driver in hardware security applications, and is being pursued through the adoption of specific techniques that increase the attack effort by several orders of magnitude.

Session 16 Overview:
Efficient Compute-In-Memory-Based Processors for ML

Machine Learning Subcommittee

Session Chair: Jae-sun Seo, Arizona State University, Tempe, AZ

Session Co-Chair: Yongpan Liu, Tsinghua University, Beijing, China

Compute-in-memory (CIM) immerses computation into memory to reduce memory access and data movement, and prior CIM works demonstrated high energy efficiency at the macro level. For end-to-end processors, application-specific circuits and architectures are required to maintain the energy benefits at the system-level. This session includes seven papers, each representing advances in integration of SRAM/eDRAM/ReRAM-based CIM macros for system-level processors with new techniques, such as hybrid/reconfigurable computing architectures, sparsity-aware hardware design, CIM utilization improvement and AI model-specific optimization.

- In Paper 16.1, Tsinghua University presents a 28nm 14.36mm^2 digital CIM-based multimodal transformer accelerator (MulTCIM) with attention-token-bit hybrid sparsity exploitation. Long-reuse attention elimination, token pruning with symmetry modal overlapping, and bit-balanced CIM are incorporated for improving CIM utilization and reducing latency. MulTCIM consumes 2.24μJ/token at 0.7V and 160MHz frequency for the ViLBERT-base model with INT8/INT16 mixed precision.
- In Paper 16.2, Fudan University, Birentech and Peng Cheng Laboratory demonstrate a 28nm 3.93mm^2 CIM-based sparse transformer accelerator, featuring a butterfly-network-based sparsity-aware feed-forward computing architecture and a digital CIM-based local attention reusable engine. Peak system energy efficiency of 53.83TOPS/W is achieved with INT8 precision at 0.67V and 40MHz frequency.
- In Paper 16.3, the Institute of Microelectronics of the Chinese Academy of Sciences and Tsinghua University describe a 28nm 4.54mm^2 CIM processor for integer and floating-point (FP) inference/training, which integrates a RISC-V CPU, a CIM core for dense integer computations and a digital core for sparse FP computations. For ResNet-50 model inference (ImageNet dataset), system energy efficiencies are 41.8TOPS/W with INT8 precision and 8.87TOPS/W with FP16 precision at 0.485V and 50MHz.
- In Paper 16.4, Tsinghua University demonstrates a 28nm 12.42mm^2 processor chiplet called TensorCIM for a CIM-based multi-chip module (MCM) system, featuring redundancy elimination for sparse gathering and inter-/intra-CIM utilization improvement for sparse algebra. With FP32 precision, TensorCIM achieves 7.6TFLOPS/W and 3.7nJ/gather for a graph convolutional network (Pubmed dataset) and 6.4TFLOPS/W and 2.3nJ/gather for a deep learning recommendation model (MovieLens dataset) at 0.6V and 115MHz.
- In Paper 16.5, Korea Advanced Institute of Science and Technology presents a 28nm 20.25mm^2 eDRAM-based CIM processor called DynaPlasia with a novel triple-mode 3T2C cell and a dynamic reconfigurable core architecture that enables high system efficiency for ML workloads. For ResNet-18 (ImageNet dataset), DynaPlasia achieves system energy efficiency of 37.2TOPS/W and compute density of 2.03TOPS/mm^2 at 1.0V and 250MHz for INT4/INT5 activation/weight precision.
- In Paper 16.6, National Tsing Hua University and TSMC report a 22nm 24.48mm^2 non-volatile AI processor that integrates 4MB of ReRAM CIM macros, featuring a hybrid computing architecture (in-memory and near-memory) and ReRAM devices of mixed precision (SLC/MLC memory states) to balance NVM capacity, energy and accuracy. For MobileNet-V2 (CIFAR-10 dataset), 251TOPS/W for INT4 precision and 68.9TOPS/W for INT8 precision have been achieved at 0.7V and 50MHz.
- In Paper 16.7, STMicroelectronics presents a 18nm FD-SOI 4.2mm^2 chip that integrates an Arm Cortex-M CPU and eight neural processing unit (NPU) clusters. Each NPU instantiates 256Kb of push-rule 8T SRAM-based digital CIM macros, which can be configured as compute modules or standard memory for storage. System-level peak energy efficiency of 77TOPS/W was achieved for INT4 precision at 0.525V, 600MHz with a forward body bias of up to 1.5V

Session 16 Highlights:
Efficient Compute-In-Memory-Based Processors for ML

[16.1] MulTCIM: A 28nm 2.24μJ/Token Attention-Token-Bit Hybrid Sparse Digital CIM-based Accelerator for Multimodal Transformers

[16.5] DynaPlasia: An eDRAM In-Memory-Computing-Based Reconfigurable Spatial Accelerator with Triple-Mode Cell for Dynamic Resource Switching

Paper 16.1 Authors: Fengbin Tu, Zihan Wu, Yiqi Wang, Weiwei Wu, Leibo Liu, Yang Hu, Shaojun Wei, Shouyi Yin

Paper 16.1 Affiliation: Tsinghua University, Beijing, China

Paper 16.5 Authors: Sangjin Kim, Zhiyong Li, Soyeon Um, Wooyoung Jo, Sangwoo Ha, Juhyoung Lee, Sangyeob Kim, Donghyeon Han, Hoi-Jun Yoo

Paper 16.5 Affiliation: Korea Advanced Institute of Science and Technology, Daejeon, Korea

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

CONTEXT AND STATE OF THE ART

- Memory access and data movement often dominate the energy and throughput of ML processors. To address the bottleneck, compute-in-memory (CIM) macros based on various memory technologies, e.g. SRAM, eDRAM, ReRAM, etc. have been presented. While such CIM macros demonstrate high energy efficiency at the macro level, system-level energy benefits could be diminished without appropriate end-to-end optimizations.
- Previous transformer accelerators have shown that a CIM accelerator with attention sparsity can efficiently accelerate vanilla transformers, but for emerging workloads such as multimodal AI tasks, irregular attention patterns and different token sparsity from multiple modalities pose challenges to optimally integrating CIM macros for efficient hardware acceleration.
- Previous eDRAM-based CIM designs showed lower cell density than SRAM-based CIMs, because they occupy considerable area to realize a large cell capacitor for long retention time. A higher density eDRAM-based CIM processor is required for greater system-level efficiency and throughput for various ML models.

TECHNICAL HIGHLIGHTS

A multimodal transformer accelerator incorporates algorithm and hardware optimizations for improving the CIM utilization and reducing the end-to-end latency.

- Paper 16.1 presents MulTCIM, a digital CIM-based accelerator for multimodal transformers fabricated in 28nm CMOS technology, featuring aggressive exploitation of dynamic sparsity to reduce power consumption, coupled with several novel techniques to improve CIM utilization, thereby achieving outstanding efficiency: 2.24μJ/token for the ViLBERT-base model.

KAIST's eDRAM-based CIM processor exhibits reconfigurability at the CIM cell level, as well as the CIM macro and dataflow architecture level, enabling dynamic optimization of the system resource utilization according to the ML workloads, while maintaining high storage/compute density.

- Paper 16.5 presents DynaPlasia, a high density eDRAM-IMC processor in 28nm CMOS based on a dynamically reconfigurable triple-mode cell. A hierarchical in-memory ADC, signed-input signed-weight IMC organization and a leakage-tolerant architecture are incorporated to boost density (2× over SOTA). The novel reconfigurable core architecture improves utilization and efficiency (2.5× over SOTA on ResNet50).

APPLICATIONS AND ECONOMIC IMPACT

- Emerging multimodal applications such as video question answering, multilingual image retrieval, and action prediction can be efficiently processed and deployed in cloud servers and mobile devices.
- High system-level energy efficiency and compute density can improve wider adoption of practical ML inference processors for energy-/storage-constrained edge devices.

Session 17 Overview: High-Speed Data Converters

Data Converters Subcommittee

Session Chair: John Keane, Keysight Technologies, Santa Clara, CA

Session Co-Chair: Ying-Zu Lin, MediaTek, Hsinchu, Taiwan

Data converters continue to increase in speed with all papers in this session achieving conversion rates of at least 1GS/s. Achieving high speed while maintaining power efficiency is driven by time-interleaving, pipelining, and time-domain quantization techniques. Time-interleaving is enabled by increased reliance on background calibration. Pipelining is aided by ring-amp-based residue amplification. DAC performance levels at very high conversion rates continue to improve, driven by advanced process technology and dynamic element matching (DEM) techniques.

- In Paper 17.1, University of Macau describes a 2× interleaved 9b 2.8GS/s 5b/cycle SAR ADC. The ADC shows a 51.8dB SNDR and 72.4dB SFDR with a Nyquist input. The achieved Walden FoM is 20.3fJ/conv.-step and the Schreier FoM is 160.7dB.
- In Paper 17.2, IBM Zurich Research Laboratory presents an 8b time-domain ADC designed in 5nm process that achieves 16.6fJ/conv.-step Walden FoM with 313μm^2 area. The digitally intensive architecture runs at 1.25GS/s with 0.8V supply and at 1GS/s with 0.7V supply.
- In Paper 17.3, MediaTek shows a 14b 16GS/s DAC in 7nm process for a multi-band software-defined-radio RF transceiver using time-interleaving DEM. It achieves IM3 of -70dB up to 7.8GHz and 64fs jitter at 8GHz DAC output frequency.
- In Paper 17.4, Socionext Europe GmBH presents a 24GS/s 12b time-interleaved ADC for direct RF sampling in a 7nm process. It achieves 7.2GHz bandwidth, 76dB SFDR and -147.5dBFS/Hz NSD in 750mW with an area of 0.9mm^2.
- In Paper 17.5, Tsinghua University introduces a 1GS/s ring-amp-based time-interleaved pipelined SAR ADC that achieves 10 ENOB. Total power of 10mW including the reference buffer results in a Schreier FoM of 169.2dB.
- In Paper 17.6, National Cheng Kung University describes a 7b 4.5GS/s 4× interleaved SAR ADC with fully on-chip background timing skew calibration. The ADC achieves 6.31 ENOB in 6.56mW and a Walden FoM of 18.5fJ/conv.-step.
- In Paper 17.7, National Cheng Kung University presents a 2.7GS/s 8b sub-ranging ADC with multiple-reference-embedded comparators. It achieves 45.9dB SNDR in 3mW resulting in a 6.9fJ/conv.-step Walden FoM.
- In Paper 17.8, Nanyang Technological University presents a single-channel 10GS/s 8b time-domain ADC. The prototype achieves 36.4d SNDR, 58.9fJ/conv.-step Walden FoM and a 0.009mm^2 active area.

Session 17 Highlights: High-Speed Data Converters

[17.1] A 2×-Interleaved 9b 2.8GS/s 5b/cycle SAR ADC with Linearized Configurable V2T Buffer Achieving >50dB SNDR at 3GHz Input

[17.3] A 14b 16GS/s Time-Interleaving Direct-RF Synthesis DAC with T-DEM Achieving -70dBc IM3 up to 7.8GHz in 7nm

Paper 17.1 Authors: Hongzhi Zhao[1], Minglei Zhang[1], Yan Zhu[1], Chi-Hang Chan[1], Rui Paulo da Silva Martins[1,2]

Paper 17.1 Affiliation: [1]University of Macau, Macao, Macau, [2]Instituto Superior Tecnico/University of Lisboa, Lisbon, Portugal

Paper 17.3 Authors: Wei-Hsin Tseng*[1], Willy Lin*[1], Chung-Wei Hsu[1], Chang-Yang Huang[1], Yu-Sian Lin[1], Hung-Yi Huang[1], HsinWei Chen[1], Sheng-Hui Liao[1], Kuan-Dar Chen[1], Jon Strange[2], Gabriele Manganaro[3]

Paper 17.3 Affiliation: [1]MediaTek, HsinChu, Taiwan, [2]MediaTek, Kent, United Kingdom, [3]MediaTek, Woburn, MA

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

CONTEXT AND STATE OF THE ART

- Power, performance and speed of conventional SAR ADCs are limited by device non-idealities. Innovations in hybrid architectures enable high-speed, low-power and area-efficient ADCs with high linearity and dynamic range.
- In high-performance transmission systems, the linearity and noise floor of DACs directly affect the in-band EVM and out-band mask performance. For base station and software-defined radio applications, DACs with sampling speeds over 16GS/s with very low noise floor are necessary.

TECHNICAL HIGHLIGHTS

- **University of Macau presents a 5b/cycle multi-bit SAR ADC enabled by a linearized dynamic integrator-based voltage-to-time (V2T) buffer. At 2.8GS/s, the proposed 9b SAR ADC shows a 51.8dB SNDR and 72.4dB SFDR with a Nyquist input owing to the ADC architecture with only one signal DAC.**
- **MediaTek presents a 14b 16GS/s DAC composed of two time-interleaved (TI) RZ DACs, achieving a -70dBc IM3 up to 7.8GHz. The dynamic skew mismatch is randomized by a new time-interleaving DEM (T-DEM) that randomizes and spreads the signal power associated to the TI artifacts across frequency by 20dB.**

APPLICATIONS AND ECONOMIC IMPACT

- Multi-bit/cycle, pipelining and residue amplifying further increase the medium resolution ADC sampling rate to break the performance barriers.
- Circuit techniques like dynamic element matching and static weight calibration improve the non-idealities of transistors to make high-speed DACs achieving state-of-the-art performance.

Session 18 Overview:
mm-Wave & sub-THz for Wireless and Sensing

Wireless Subcommittee

Session Chair: Jane Gu, University of California, Davis, CA

Session Co-Chair: Giuseppe Gramegna, imec, Leuven, Belgium

mm-Wave/sub-THz technologies advance existing and enable emerging applications. The first paper presents a W-band FMCW radar transceiver array enabling great scalability. The second paper features a fully integrated D-band receiver with ultra-high data-rates and energy efficiency, followed by an E-band power-DAC for high output power and efficiency. The final paper presents a CMOS THz receiver for high-resolution imaging applications.

- In Paper 18.1, Stanford University presents a W-band FMCW radar transceiver array with 2.4GHz LO synchronization to facilitate large array scalability in 40nm CMOS. The 4-receiver chip and 1-transmitter chip are demonstrated.
- In Paper 18.2, Intel features a D-band receiver with integrated low-noise quadrature PLL and energy-efficient high-speed time-interleaved ADC in 22nm FinFET. The complete receiver demonstrates a data-rate of 128Gb/s and energy efficiency of 1.95pJ/b.
- In Paper 18.3, the University of Electronic Science and Technology shows a high-efficiency 71-to-89GHz transmitter in 40nm CMOS. The full transmitter achieves 20.5dBm peak output power and 20.4% system efficiency.
- In Paper 18.4, KU Leuven presents a 4×4 607GHz harmonic injection-locking receiver in 28nm CMOS for imaging applications. The achieved average NEP of the array is 4.4pW/$\sqrt{\text{Hz}}$.

Session 18 Highlights:
mm-Wave & sub-THz for Wireless and Sensing

[18.2] A 128Gb/s 1.95pJ/b D-Band Receiver with Integrated PLL and ADC in 22nm FinFET

Paper Authors: Abhishek Agrawal[1], Amy Whitcombe[2], Ritesh Bhat[1], Somnath Kundu[1], Peter Sagazio[1], Hariprasad Chandrakumar[1], Thomas Brown[1], Brent Carlton[1], Steven Callender[1], Stefano Pellerano[1]

Paper Affiliation: [1]Intel, Hillsboro, OR, [2]Intel, Santa Clara, CA

Subcommittee Chairs: Jane Gu,UC Davis, CA, Giuseppe Gramegna, imec, Belgium

CONTEXT AND STATE OF THE ART

- mm-wave (30 to 100GHz) and sub-THz frequencies have been attracting great interest due to their inherently huge bandwidth, the most essential resource for high-throughput wireless communications and connectivity systems.

  One of the primary challenges in sub-THz transmitters (TX) and receivers (RX) is achieving good energy efficiency for both wideband RF and mixed-signal/digital circuits concurrently; a very limited number of highly integrated TXs/RXs have been reported so far.

TECHNICAL HIGHLIGHTS

- **A D-Band (140GHz) fully integrated RX with 128Gb/s at 1.95pJ/b total efficiency in 22nm FinFET technology.**
  - The RX with highest data-rate and level of integration published: integrated 20GHz RX RF front-end, PLL, frequency tripler for LO generations, 16GHz ADC Nyquist BW (at 32GS/s) and 5.1 ENOB.
  - The full D-band RX (246mW) and the stand-alone RX front-end (166mW) achieve 160Gb/s and 128Gb/s with -16.4dB and -15.2dB EVM respectively.

APPLICATIONS AND ECONOMIC IMPACT

- Advanced circuit techniques for both sub-THz and wide-bandwidth operation enable high data-rate applications such as short-range high-resolution video streaming and fast data transfer.
- Competitive energy efficiency while retaining high performance enables the use of such technology in battery-powered devices while minimizing thermal challenges.

Session 19 Overview:
5G and Satcom: Receivers and Transmitters

Wireless Subcommittee

Session Chair: Venumadhav Bhagavatula, Samsung Semiconductor, San Jose, CA

Session Co-Chair: Alireza Zolfagari, Broadcom, Irvine, CA

This session includes four papers presenting key advances in the fields of 5G and satellite communication systems. The first two papers describe wide bandwidth, high linearity front-end receiver architectures for 5G systems. In the next two papers, advances in the field of CMOS transceivers for satellite applications are discussed.

- In Paper 19.1, Deft University of Technology presents a 0.4-to-7.3GHz receiver achieving 300MHz RF bandwidth and up to 38dBm in-band OIP3 for base-station application. The proposed receiver utilizes a Rauch TIA with an extra compensation feedback between TIA and RF input to improve the linearity and enhance the receiver bandwidth.

- In Paper 19.2, Massachusetts Institute of Technology introduces a passive voltage-mode harmonic-resilient N-path mixer implemented in a receiver operating from 0.25 to 2.5GHz. Fabricated in 45nm SOI, the prototype achieves HB1-dB of greater than 10 and 4dBm at the 3rd and 5th harmonics, respectively.

- In Paper 19.3, Tokyo Institute of Technology and Axelspace present a small-satellite-mounted single/dual circular polarization 256-element Ka-band phased-array transmitter using 64 4-element ICs. Each front-end uses a reconfigurable coupler to enable efficient generation of single circularly polarized waveforms while also supporting single and dual polarizations.

- In Paper 19.4, Tokyo Institute of Technology and Axelspace present a 2.95mW/element, Ka-band phased-array receiver for small satellite constellations. The prototype device, fabricated in 65nm CMOS, introduces an on-chip total-ionizing-dose (TID) sensor-based compensation technique to limit the main lobe degradation to 0.22dB/Mrad.

Session 19 Highlights:
5G and Satcom - Receivers and Transmitters

[19.2] An Interferer-Tolerant Harmonic-Resilient Receiver with +10dBm 3rd-Harmonic Blocker P_{1dB} for 5G NR Applications

Paper Authors: Soroush Araei, Shahabeddin Mohin, Negar Reiskarimian

Paper Affiliation: Massachusetts Institute of Technology, Cambridge, MA

Subcommittee Chair: Chih-Ming Hung, MediaTek, Taipei, Taiwan

CONTEXT AND STATE OF THE ART

- The frequency spectrum below 6GHz is extremely crowded with the existing 2G/3G/4G and unlicensed WiFi bands now coexisting with sub-6GHz 5G new radio (NR) bands. In such a blocker-rich environment, novel techniques to suppress interferers near the received carrier frequency and its harmonics are important to maintain a high data-rate down-link stream.

TECHNICAL HIGHLIGHTS

- **Massachusetts Institute of Technology introduces a high-linearity mixer-first receiver operating from 0.25 to 2.5GHz.**
  - The receiver can handle 3rd/5th-harmonic blockers as large as 10dBm/4dBm with less than 1dB compression.
  - Introduces a fully passive, low-loss harmonic-reject N-path-filter/mixer topology at the RX front-end for blocker rejection.

APPLICATIONS AND ECONOMIC IMPACT

- 5G provides a path for 10 to 100× increase in data-rates while reducing the latency by a factor of 10×. This technology will be the key enabler for bringing the next generation of augmented/virtual reality (AR/VR) devices and pervasive connected devices to life.

Session 20 Overview: GaN Power Conversion

Power Management Subcommittee

Session Chair: Patrik Arno, STMicrolelectronics, Grenoble, France

Session Co-Chair: Saurav Bandyopadhyay, Texas Instruments, Dallas, TX,

Power conversion is ubiquitous and a variety of power-converter topologies have found their ways into both established and emerging applications. This session aims at showcasing some of the best works in the field of application-specific GaN power conversion, such as monolithically integrated GaN-on-Si and GaN-on-SOI gate drivers, an autonomous and self-adaptive EMI control for automotive GaN power converter and a bidirectional dual GaN control rectifier for automotive applications.

- In Paper 20.1, National Chiao Tung University, Chip-GaN Power Semiconductor, and Realtek Semiconductor present a High Common-Mode Transient Immunity GaN-on-SOI Gate Driver for High dV/dt SiC Power Switch, with on-chip capacitive isolator, high data rate and low propagation delay. Peak efficiency is 98.6% and > 90% at V_{IN}=800V and 1700V, respectively.
- In Paper 20.2, University of Texas at Dallas and Texas Instruments introduce a Self-Adaptive EMI Control GaN Switching Regulator with Modulation Frequency Envelope Tracking for Full-Spectrum Automotive CISPR 25 Compliance from 150KHz to 108MHz. Envelope tracking adaptively shields audible range noise along EMI optimization
- In Paper 20.3, National Chiao Tung University, Chip-GaN Power Semiconductor, and Realtek Semiconductor present a GaN gate driver with on-chip adaptive on-time controller to limit the maximum allowable switching frequency for high efficiency and negative-current slope detector to suppress shoot-through problems. The best efficiency of this design is 96.2% and its maximum driving current is 10A.
- In Paper 20.4, National Chiao Tung University, Chip-GaN Power Semiconductor, and Realtek Semiconductor introduce a bidirectional dual GaN control rectifier can convert energy between HV (48V) and LV (12V) batteries in a multiple-phase accelerated current control technique. Peak efficiency is 95.5% and 94.2% at buck 4A and boost 1A load current, respectively, under 20MHz switching frequency.

Session 20 Highlights: GaN Power Conversion

[20.1] A High Common-Mode Transient Immunity GaN-on-SOI Gate Driver for High dV/dt SiC Power Switch

Paper Authors: Si-Yi Li[1], K.-H. Chen[1], Y.-H. Lin[2], S.-R. Lin[2], T.-Y. Tsai[2]

Paper Affiliation: [1]National Chiao Tung University, Hsinchu, Taiwan, [2]Realtek Semiconductor, Hsinchu, Taiwan

Subcommittee Chair: Bernhard Wicht, University of Hannover, Hannover, Germany

CONTEXT AND STATE OF THE ART

- SiC Power FETs with superior FoMs allow high density and high efficiency power converters for 1.2-to-1.7kV applications.
- However, SiC drivers pose interesting challenges requiring high CMTI, negative voltage generation slew rate control for high dV/dt operation.

TECHNICAL HIGHLIGHTS

- **National Chiao Tung University and Realtek Semiconductor present a GaN-on-Soi Gate Driver for SiC Power FET**
    - A robust driver employing a capacitive level shifter, a TX modulating a carrier frequency with the PWM and an RX with a TIA is proposed enabling >200V/ns V_{ds} slew rates
    - Challenges with driving SiC Power FETs are addressed with a novel bootstrap circuit for generating a negative rail for reliably turning off the power FET and a quad drive control technique that senses the miller plateau for slew rate control.

APPLICATIONS AND ECONOMIC IMPACT

- 1.2-to-1.7kV SiC FET driver with GaN-on-SOI gate driver with capacitively coupled level shifter and a Si transmitter IC.
- Addressing driver challenges for SiC and enabling high-density high-efficiency power conversion.

Session 21 Overview: Emerging Sensing Systems and IOT

Technology Directions Subcommittee

Session Chair: Kaushik Sengupta, Princeton University, Princeton, NJ

Session Co-Chair: Milin Zhang, Tsinghua University, Beijing, China

This session covers a wide range of emerging sensing and ultra-low power IOT systems. The first paper presents a fully integrated fluorescence-based biosensor in 65nm CMOS with integrated optical filters. An ultra-low power interface IC for bioreactor monitoring comes next, followed by a CMOS biosensor array with on-chip dielectrophoresis (DEP) and sensing capabilities. The next two papers present a 263GHz IC for electron paramagnetic resonance sensing, and a battery monitoring and communication IC for electrical vehicles. The final paper showcases an ultra-low-power backscattering IoT tag that harvests energy from LTE, and backscatters tone-like BLE packets to WiFi channels.

- In Paper 21.1, California Institute of Technology presents a fully-integrated fluorescence sensor in 65nm CMOS technology for static and dynamic measurements of living bacterial cells. The fluorescence sensor integrates plasmonic filters at 600/700nm in the metal layers, on-chip photodiodes and readout, and achieves a sensitivity of 1.05fA with greater than 18dB SNR.
- In Paper 21.2, imec presents an electro-chemical sensor readout IC supporting temperature, current, voltage, and ISFET for wireless bioreactor monitoring while consuming 22μW of peak power. The system is validated against commercial setups for sensing glucose, pH, and dissolved oxygen.
- In Paper 21.3, ETH Zurich presents a dielectrophoresis (DEP)-assisted multi-functional biosensor array in 130nm CMOS technology achieving improved limit-of-detection and detection speed by DEP-based analyte enrichment. The system integrates electro-chemical potentiostat, complex impedance sensors, photodiode arrays, and DEP controller with on-chip electrodes.
- In Paper 21.4, University of Stuttgart, Helmholtz-Zentrum Berlin and IMS-CHIPS present a 263GHz 32-channel injection -locked VCO array for electron paramagnetic resonance and dynamic nuclear polarization. A proof-of-concept EPR signal is measured inside a 9.4T magnet to validate the proposed approach in the target application.
- In Paper 21.5, University of California, San Diego presents an ultra-low-power IoT tag that can backscatter a tone-like BLE packet to an adjacent WiFi channel. The power consumption during active mode of operation is as low as 25μW that allows it to be powered by harvesting energy from LTE.
- In Paper 21.6, Korea Advanced Institute of Science and Technology and Autosilicon present a battery monitoring IC in electrical vehicles with high measurement accuracy and communication reliability. The design demonstrates a measurement inaccuracy of 2mV with less than 3nA of current input.

Session 21 Highlights:
Emerging Sensing Systems and IOT

[21.1] A 65nm CMOS Living-Cell Dynamic Fluorescence Sensor with 1.05fA Sensitivity at 600/700nm Wavelengths

Paper 21.1 Authors: Fatemeh Aghlmand[1], Chelsea Hu[2], Saransh Sharma[1], Krishna Pochana[1], Richard Murray[1], Azita Emami[1]

Paper 21.1 Affiliation: [1]California Institute of Technology, Pasadena, CA, [2]Texas A&M University, College Station, TX.

Subcommittee Chair: Ali Hajimiri, California Institute of Technology, Pasadena, CA, Technology Directions

CONTEXT AND STATE OF THE ART

- Integrated fluorescence and luminescence sensors, eliminating large and bulky external optical elements, can enable new biomedical applications including point-of-care diagnostics, drug screening, and drug delivery.
- Metal-based integrated nano-optical filters can allow such integration, but designing band-pass filters for different fluorescence labels to allow color multiplexing has been a challenge.
- Innovations on optical and electronic co-design can address several of these bottlenecks for high-performance and chip-scale fluorescence sensing systems.

TECHNICAL HIGHLIGHTS

- **A fully integrated fluorescence sensor with on-chip bandpass filters allows multiplexed optical biosensing.**
    - The presented chip demonstrates integrated metal-based on-chip bandpass filters in the 600/700nm wavelength range using low-loss cavity modes.
    - The sensor achieves 1.05fA sensitivity with greater than 18dB SNR, allowing measurement of static and dynamic behavior of living E. coli bacteria cells with fluorescent proteins.

APPLICATIONS AND ECONOMIC IMPACT

- Eliminating large and bulky optical elements for integrated optical structures can allow miniaturization of complex optical instrumentation into chip-scale systems.

- Low-cost and miniaturized fluorescence systems that can detect clinically relevant biomarkers are crucial for the growing field of precision medicine, point-of-care diagnosis, continuous health monitoring, and closed-loop drug delivery.

Session 22 Overview: Heterogeneous ML Accelerators

Machine Learning Subcommittee

Session Chair: Rangharajan Venkatesan, NVIDIA, Santa Clara, CA

Session Co-Chair: Sophia Shao, University of California, Berkeley, Berkeley, CA

Machine learning has been adopted across a wide range of application domains, from computer vision, natural language processing, point-cloud processing and medical interfaces. The growing diversity in machine learning workloads exhibit quite distinct compute and data movement behaviors, leading to the development of heterogeneous SoCs to support different components of the applications. This session includes nine papers, covering a diversity of architectural and circuit innovations and heterogeneous integration to support a wide range of ML-based applications.

- In Paper 22.1, the University of Bologna, ETH Zürich and Dolphin Design present a heterogeneous system-on-chip (SoC) with 16 RISC-V cores and precision-scalable DNN accelerators supporting AI IoT applications. The chip is fabricated in 22nm technology and includes a 16-core RISC-V-based DSP cluster with 2-to-8b ISA extensions and a reconfigurable DNN engine to support mixed-precision DL acceleration. The SoC achieves energy efficiency of 1.64TOPS/W at 8b precision and 12.4TOPS/W at 2b precision.
- In Paper 22.2, Tsinghua University describes a 2D/3D unified sparse convolution accelerator for voxel-based point-cloud processing. The chip, fabricated in a 28nm technology, features a configurable core to provide unified support for both 2D and 3D convolutions, together with sparsity support and irregular memory-access handling mechanisms. The accelerator achieves efficiency of 1.55-4.14TOPS/W for hybrid 2D/3D point-cloud networks.
- In Paper 22.3, Pohang University of Science and Technology presents a 1-to-8b scalable-precision deep-learning accelerator supporting MAC operations with arbitrarily quantized two vectors. The accelerator features flexible quantization compression support to handle different precisions and data formats, together with a dynamic-precision bit-serial processing to improve energy efficiency. The chip is fabricated in a 28nm technology and delivers up to 127.8TOPS/W with 90% sparsity.
- In Paper 22.4, National Taiwan University demonstrates a hardware-utilization-aware neural-network accelerator with dynamic dataflow in 28nm CMOS. The chip supports various levels of parallelism along multiple dimensions to maximize the hardware utilization across all convolutional layers. It achieves >97.3% hardware utilization for commonly used neural networks, while achieving an energy efficiency of 11.2TOPS/W.
- In Paper 22.5, Korea Advanced Institute of Science and Technology presents a complementary DNN accelerator with heterogeneous convolutional neural network (CNN) and spiking neural network (SNN) support. The chip, fabricated in a 28nm technology, includes dedicated CNN and SNN cores with dynamic arbitration to map layers to the heterogeneous cores during training and inference. The chip achieves 24.5TOPS/W for CNN inference with ImageNet.
- In Paper 22.6, Tsinghua University describes a 28nm asynchronous spiking neural network (SNN) accelerator with on-chip learning support for edge AI. The SNN accelerator supports 8b/10b weight precision and consumes sub-0.1μJ energy/sample for edge AI tasks, including gesture recognition, keyword spotting, and image classification, while maintaining >92% accuracy.
- In Paper 22.7, University of Electronic Science and Technology of China presents a deep-learning-based visual object processor for mobile visual object detection and tracking. It exploits domain-specific features to achieve high energy efficiency and supports state-of-the art frameworks with multi-scale semantic feature extraction. Fabricated in 40nm technology, the design achieves 163.13fps and 24.08TOPS/W for combined detection and tracking tasks.
- In Paper 22.8, Seoul National University demonstrates a real-time speech-enhancement processor for hearing aids in 28nm CMOS. The accelerator features importance-aware pruning, together with reconfigurable PEs and weight quantization, to realize real-time speech enhancement with 740μW power consumption.
- In Paper 22.9, Harvard University, University of California at Berkeley, Peking University, Columbia University and Cornell University present a sparse transformer processor featuring fine-grained latency and power management tailored to the complexity of the input sentences. The processor implements a FP4 MAC datapath with per-vector exponent bias scaling, sentence-level voltage/frequency scaling, and entropy-based early exit. The chip is fabricated in a 12nm technology and achieves energy efficiency of 3.0-8.24TFLOPS/W for FP8 and 6.61-18.1TFLOPS/W for FP4.

Session 22 Highlights: Heterogeneous ML Accelerators

[22.1] A 12.4TOPS/W @ 136GOPS AI-IoT System-on-Chip with 16 RISC-V, 2-to-8b Precision-Scalable DNN Acceleration and 30%-Boost Adaptive Body Biasing

[22.5] C-DNN: A 24.5-to-85.8TOPS/W Complementary-Deep-Neural-Network Processor with Heterogeneous CNN/SNN Core Architecture and Forward-Gradient-Based Sparsity Generation

Paper 22.1 Authors: Francesco Conti[1], Davide Rossi[1], Gianna Paulin[2], Angelo Garofalo[1], Alfio Di Mauro[2], Georg Ruetishauer[2], Gianmarco Ottavi[1], Manuel Eggimann[2], Hayate Okuhara[1], Vincent Huard[3], Olivier Montfort[3], Lionel Jure[3], Nils Exibard[3], Pascal Gouedo[3], Mathieu Louvat[3], Emmanuel Botte[3], Luca Benini[1,2]

Paper 22.1 Affiliation: [1]University of Bologna, Bologna, Italy, [2]ETH Zürich, Zürich, Switzerland, [3]Dolphin Design, Meylan, France

Paper 22.5 Authors: Sangyeob Kim, Soyeon Kim, Seongyon Hong, Sangjin Kim, Donghyeon Han, Hoi-Jun Yoo

Paper 22.5 Affiliation: Korea Advanced Institute of Science and Technology, Daejeon, Korea

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

CONTEXT AND STATE OF THE ART

- Deep neural networks continue to evolve and find applications across a wide range of applications. Today, they consist of different kinds of layers such as convolutional layers, attention layers, depthwise-separable layers, etc., that differ significantly in their compute and memory demands. This motivates the need for heterogeneous ML accelerators with specialized cores for different tasks.

- 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 become the new bottleneck, necessitating the need to handle various types of operations well. A heterogeneous architecture can take the “best of both” approaches and operate at a point that is more optimal than each type of core.

- Emerging techniques continue in the direction of workload reduction via skipping unnecessary computations, where the techniques applied include sparse convolution, sparsity generation using forward gradient, entropy-based early exiting and local attention use. Such techniques exploit sparsity in convolution kernels, forward gradients, execution paths and attention spans, respectively. Dynamic adaptation of data precision, dataflow in the pipeline, resource switching, as well as body biasing also help push the boundary of efficiency and performance in ML processors.

TECHNICAL HIGHLIGHTS

Paper 22.1 presents an AI-IoT heterogeneous SoC in 22nm FD-SOI CMOS technology, with a 16-core RISC-V cluster (5.1GOPS, 440GOPS/W).

- The cores support wide-range (2GOPS, 1.64TOPS/W@8b to 136GOPS, 12.4TOPS/W@2b) mixed-precision DNN inference via 2/4b SIMD ISA extensions and a precision-scalable reconfigurable binary engine. The SoC leverages on-chip monitors and adaptive body biasing to boost energy efficiency by 30% with no performance reduction (6.38TOPS/W on ResNet-20, mixed-precision).

Paper 22.5 presents a complementary combination of a CNN and SNN accelerator for energy-efficient inference and training with high accuracy.

- The chip performs optimal workload allocation to the CNN core and SNN core in order to take advantage of the best properties of both cores. Workloads can be divided by layers or even within a layer based on factors such as spike sparsity and forward gradient magnitude. It also employs an attention module which can shrink the input distribution and increase the spike sparsity, allowing more workloads to be allocated to the SNN. The C-DNN processor is fabricated in 28nm CMOS technology and achieves 24.5TOPS/W inference, compared to 17.5TOPS/W in CNN-only mode and 20.8TOPS/W in SNN-only mode for an ImageNet classification task.

APPLICATIONS AND ECONOMIC IMPACT

- Emerging AI-enabled IoT SoCs for wearables or drones run highly heterogeneous workloads with a very low-power envelope. Applications include object detection and recognition, as well as audio/speech processing. Workloads may vary widely in computational complexity, as well as precision requirements for different applications and even within compute stages of the same application. Heterogeneous architectures handle such diverse workloads well in terms of end-to-end performance and power.

- Attention to SNNs has increased due to various benefits, such as on-device learning and low power consumption. In order to maximize energy efficiency at the expense of area efficiency, a SNN engine is combined with a CNN engine to be a complementary DNN processor for reducing both the inference and training power. Such a heterogenous approach leverages task switching between the ultra-low-power SNN core and the accurate CNN cores for highly efficient inference and accurate training.

Session 23 Overview: Analog Sensor Interfaces

Analog Subcommittee

Session Chair: Chinwuba Ezekwe, Robert Bosch, Sunnyvale, CA

Session Co-Chair: Minkyu Je, KAIST, Daejeon, Korea

Analog techniques continue to enable improvements in the resolution, accuracy, and energy efficiency of sensor interfaces. Noise cancellation and new CLS techniques reduce the energy of a capacitance-to-digital converter. Shunt- and hybrid-based current sensors achieve lower gain error by using temperature coefficient calibration and crossover frequency optimization. A stress sensor with a current-mirror readout architecture achieves high DR and sensitivity control. BJT-based high-accuracy temperature sensors are demonstrated by using capacitive biasing and continuous-time current-mode readout techniques, and a temperature sensor with built-in offset, gain and non-linearity calibration achieves low FoM and small area.

- In Paper 23.1, Peking University and the Advanced Institute of Information Technology of Peking University present a capacitance-to-digital converter with 37.12aF_{rms} resolution and FoM of 7.9fJ/conv.-step. It is achieved by cancelling kT/C noise and introducing incomplete-settling-based CLS with background calibration of its clock.
- In Paper 23.2, Delft University of Technology and MIRISE Technologies present a shunt-based current sensor with high energy efficiency and ±0.2% gain error over -40°C to 125°C. The temperature coefficient of the shunt resistance is corrected by a 2-step calibration using a stable on-chip reference current and temperature sensing.
- In Paper 23.3, Delft University of Technology and Infineon Technologies present a magnetic current sensor with 5MHz bandwidth and ±1.1% gain flatness based on a hybrid scheme. The sensor uses a Hall plate and a pick-up coil with optimized crossover frequency to widen bandwidth, and an area-efficient DC servo loop.
- In Paper 23.4, Hamburg University of Technology and the University of Freiburg reveal a closed-loop integrated CMOS stress sensor system. Using 32 integrated MOS stress sensors with both mechanical and non-mechanical offset suppressions, the system achieves a range of 730MPa and an SNR of 95.4dBFS, with tunable sensitivity between 10 and 173kPa/LSB.
- In Paper 23.5, Delft University of Technology and Tsinghua University present a BJT temperature sensor in 0.18µm CMOS, which achieves an inaccuracy of ±0.15°C (3σ) from -55°C to 125°C with 1-point trim by switching the base of a capacitively biased PNP. The sensor, which operates from a 0.95 to 1.4V supply, achieves a resolution FoM of 0.34pJ·K^2.
- In Paper 23.6, Eindhoven University of Technology presents a dynamic resistive temperature sensor with fully on-chip offset, gain and non-linearity calibration. With the correction, the temperature sensor achieves +0.7/-0.6°C inaccuracy with 2.98 pJ/conversion and occupies 0.0023mm^2 area.
- In Paper 23.7, Delft University of Technology and Tsinghua University present a BJT-based temperature sensor using a new current-mode readout- circuit architecture. This results in a highly accurate continuous-time BJT-based sensor with a 3σ inaccuracy of ±0.1°C over the military temperature range.

Session 23 Highlights: Analog Sensor Interfaces

[23.2] A 40A Shunt-Based Current Sensor with ±0.2% Gain Error From –40°C to 125°C and Self-Calibration

[23.4] A Closed-Loop 12bit CMOS-Integrated Stress Sensor System with 4bit Adjustable Sensitivity from 178 to 11kPa/LSB at up to 22.5kS/s and 5bit Dynamic Range Adjustment

[23.7] A BJT-Based Temperature Sensor with ±0.1°C (3σ) Inaccuracy from –55°C to 125°C and a 0.85pJ·K² Resolution FoM Using Continuous-Time Readout

Paper 23.2 Authors: Zhong Tang[1], Nandor G. Toth[1], Roger Zamparette[1], Tomohiro Nezuka[2], Yoshikazu Furuta[2], Kofi A.A. Mankinwa[1]

Paper 23.2 Affiliation: [1]Delft University of Technology, Delft, The Netherlands, [2]MIRISE Technologies, Aichi, Japan

Paper 23.4 Authors: Kim Allinger[1], Matthias Kuhl[2]

Paper 23.4 Affiliation: [1]Hamburg University of Technology, Hamburg, Germany, [2]University of Freiburg - IMTEK, Freiburg, Germany

Paper 23.7 Authors: Nandor G. Toth[1], Zhong Tang[1], Teruki Someya[1,3], Sining Pan[1,2], Kofi A. A. Makinwa[1]

Paper 23.7 Affiliation: [1]Delft University of Technology, Delft, The Netherlands, [2]Tsinghua University, Beijing, China
[3]now with SiTime, Tokyo, Japan

Subcommittee Chair: Maurits Ortmanns, University of Ulm, Germany, Analog

CONTEXT AND STATE OF THE ART

- Low-cost metal shunts have been used to realize current sensors with demonstrated accuracy up to 85°C at low current levels. Innovations in calibration help to extend their operating range to the high temperatures and current levels required in the field of automotive applications. Stress sensors with a high degree of tunability are needed to address the disparate requirements of a wide range of emerging applications from tactile sensing to predictive maintenance (PdM).
- BJT-based temperature sensors achieve better accuracy with 1-point trim than resistor-based sensors with 2-point trim, at the cost of worse energy efficiency. Continuous-time processing in BJT-based sensors offers a path towards breaking this tradeoff.

TECHNICAL HIGHLIGHTS

- **Delft University of Technology and MIRISE Technologies present a metal-shunt-based current sensor with >11dB improved energy efficiency and ±0.2% gain error.**
    - The temperature coefficient of the shunt resistor is corrected by a 2-step calibration using a stable on-chip reference current and temperature sensing.
    - It achieved a gain error of ±0.2% over a wider current (±40A) and temperature (-40 to 125°C) range after a 1-point gain trim.
    - The presented self-calibration scheme allows a gain error of 0.3% to be maintained even in the presence of shunt resistance drift.
- **Hamburg University of Technology and the University of Freiburg present a closed-loop 12bit readout architecture for a CMOS integrated stress sensor with 95.4dBFS SNR.**
    - The stress sensor system has a tunability of sensitivity between 10 and 173kPa/LSB with 4bit resolution, and 5bit analog offset compensation capability to reduce an induced mechanical stress during fabrication and assembly, increasing the measurement range to 730MPa.
    - High SNR of 95.4dBFS and small active sensor area for high spatial resolution are achieved.

- **Delft University of Technology and Tsinghua University present a BJT-based temperature sensor using a current-mode readout circuit, achieving a 3σ inaccuracy of ±0.1°C over the military temperature range.**
    - Component mismatch is reduced by a combination of dynamic element matching and a low-cost resistor-ratio calibration scheme.
    - The sensor achieves a 1-point trimmed inaccuracy of ±0.1°C (3σ) from –55°C to 125°C, and a resolution FoM of 0.85pJ·K^2, making it 4× more energy-efficient than state-of-the-art BJT-based sensors with similar accuracy.

APPLICATIONS AND ECONOMIC IMPACT

- An accurate gain of current sensors over a wide temperature and input current range enables safer battery management in automotive applications.
- Tunable CMOS integrated stress sensors offer application in IoT, robotics and predictive maintenance. Tunable specification of the stress sensor allows a wider range of application to be addressed with the same sensor.
- Temperature sensors are omnipresent in electronic devices. Improving efficiency and accuracy at the same time allows to save cost.

Session 24 Overview: THz Signal Generation

RF Sub-Committee

Session Chair: Ruonan Han, MIT, Cambridge, MA

Session Co-Chair: Yves Baeyens, Nokia Bell Labs, Murray Hill, NJ

This session presents innovations at circuit, electromagnetic, and system levels that push the boundaries of silicon-based THz sources. Advances are made in performance metrics including the output power, energy efficiency, phase noise, and bandwidth.

- In Paper 24.1, the City University of Hongkong presents a beam-steerable 0.64-0.69THz radiator array in 65nm CMOS with 9.1dBm of output power.
- In Paper 24.2, KAIST presents a low-phase-noise CMOS signal source that provides a 264 to 287GHz signal with 75fs jitter.
- In Paper 24.3, Tsinghua University and Sanechips Technology show a slotline-based SiGe frequency doubler achieving 1.1-to-4.7dBm output power between 200 and 350GHz.

Session 24 Highlights: THz Signal Generation

[24.2] A 264-to-287GHz, -2.5dBm Output-Power, and -92dBc/Hz 1MHz-Offset Phase-Noise CMOS Signal Source Adopting a 75fs$_{rms}$ Jitter D-Band Cascaded Subsampling PLL

Paper Authors: B-T. Moon, S-G. Lee, J. Choi

Paper Affiliation: KAIST, Daejeon, Korea

Subcommittee Chair: Jan Craninckx, imec, Leuven, Belgium

CONTEXT AND STATE OF THE ART

- The sub-THz frequency range is gaining attention for the next generation communications. LO signals with low phase noise are critical to achieve advanced modulation schemes.
- With the limitation of silicon technologies, it is challenging to reduce the phase noise of THz sources while maintaining high output power and low DC power consumption.

TECHNICAL HIGHLIGHTS

- **The work presents a 264-to-287GHz signal generation chain consisting of an H-band frequency doubler and a cascaded subsampling PLL.**
    - The signal source delivers -92dBc/Hz phase noise at 1MHz offset. The 1kHz-to-10MHz rms jitter is 75fs.
    - A peak output power of -2.5dBm is achieved at the DC power of 262mW.

APPLICATIONS AND ECONOMIC IMPACT

- Low-phase-noise sources are critical for the beyond-5G communications, sub-millimeter-wave radar imaging, spectroscopy, and other emerging high-frequency systems.

Session 25 Overview: RF Transceiver Building Blocks

RF Subcommittee

Session Chair: Jeffrey Walling, Virginia Tech, Blacksburg, VA

Session Co-Chair: Hongtao Xu, Fudan University, Shanghai, China

This session presents five papers whose focus is on emerging building blocks for RF transceiver systems. The first paper presents a CMOS digital PA achieving high output power and efficiency. The next paper presents a mm-wave PA that achieves linear operation across almost an octave of bandwidth, covering the FR2 bands in 5G wireless standards. The third paper presents a W-band LNA achieving a low noise figure using novel noise-canceling techniques. This is followed by a paper presenting an RF-DAC as an alternative to PLL-based FMCW transmitters. Finally, the session concludes with a paper describing a frequency translating N-path filter as an alternative to large passive filters in RF front-end circuits.

- In paper 25.1, Fudan University and South China University of Technology present a power-combined digital PA achieving >4W peak output power with 33.6% peak power-added efficiency.
- In paper 25.2, South China University of Technology and Fudan University introduce a millimeter-wave broadband linear PA that uses an asynchronously coupled resonator-based matching technique and a broadband linearization technique to achieve a 76% 3dB fractional bandwidth in the Ka-band.
- In paper 25.3, the University of Electronic Science and Technology of China presents a deep-noise-canceling mm-wave LNA to reduce the noise figure. The deep-noise-canceling LNA achieves a peak small-signal gain of 16.5dB and the minimum of NF of 4.8dB at >70GHz.
- In paper 25.4, Infineon Technologies and Johannes Kepler University introduce a 4b RF-DAC acting as a direct-digital frequency synthesizer and as an alternative to PLL-based chirp generation for FMCW radars, achieving chirps with a 4GHz bandwidth.
- In paper 25.5, Washington University in St. Louis and Columbia University present a tunable N-path filter using a charge-pump to boost the clock strength and embedded frequency translation to break traditional trade-offs in power consumption and operation frequency in such circuits and allow the operation at up to 5GHz.

Session 25 Highlights: RF Transceiver Building Blocks

[25.1] A 4.1W Quadrature Doherty Digital Power Amplifier with 33.6% Peak PAE in 28nm Bulk CMOS

Paper 25.1 Authors: Jiaxiang Li[1], Yun Yin[1], Hang Chen[2], Jie Lin[1], Yicheng Li[1], Xianglong Jia[1], Zhen Hu[1], Xiuyin Zhang[2], Hongtao Xu[1]

Paper 25.1 Affiliation: [1]Fudan University, Shanghai, China, [2]South China University of Technology, Guangzhou, China

Subcommittee Chair: Jan Craninckx, imec, Leuven, Belgium

CONTEXT AND STATE OF THE ART

- Power amplifiers are key building blocks in wireless transceivers for cellular and connectivity applications.
- Achieving high output power and efficiency simultaneously has been difficult to accomplish and is a limiting factor in full transceiver integration
- A compact 8-way power-combining transformer enabling voltage-mode Doherty operation enables high output power and high efficiency at high backoff, which is important for modern communications signals.

TECHNICAL HIGHLIGHTS

- Fudan University and South China University present a power combined SCPA achieving >4W peak output power with 33.6% peak power-added efficiency.
- An 8-way power-combined quadrature DPA with IQ-reuse and Doherty techniques is presented, which achieves 4.1W (36.1dBm) peak P_{out} and 33.6% peak PAE at 2.9GHz in 28nm bulk CMOS. It achieves 1dB RF bandwidth of 2.6 to 3.2GHz, and the peak PAE is more than 30% across 2.7 to 3.2GHz.

APPLICATIONS AND ECONOMIC IMPACT

- Compact CMOS PAs achieving high output power potentially enable the increased integration of transceivers
- RF-DACs, such as the SCPA, directly convert bits-to-RF, obviating the need for analog blocks that are challenging to realize in fine-line CMOS.

Session 26 Overview:
Display and User Interaction Technologies

IMMD Subcommittee

Session Chair: Joonsung Bae, Kangwon National University, Kangwon, Korea

Session Co-Chair: Johan Vanderhaegen, Google, Mountain View, CA

Display and user interaction technologies continue to evolve and specialize for next-generation applications. The first paper describes an 8K display driver IC with a high-speed display interface receiver, followed by a virtual rotating gesture recognizable touch readout IC with haptic feedback on a circular bezel area. The final paper shows a self-capacitance touchscreen controller IC with high common-mode noise immunity.

- In Paper 26.1, Samsung Electronics describes an 8K source-driver IC including a power-switching fast-slew-rate buffer and effective 3-tap DFE receiver. It achieves 4.9mV DVRMS, 17V/μs slew rate, and 8Gb/s data rate under 30.2dB interface channel insertion loss.
- In Paper 26.2, Kwangwoon University and Zinitix present a virtual rotating gesture-recognition system integrated with touch sensing and haptic driving circuits to provide coherent haptic feedback responses, achieving an SNR of 63dB with 8.1mW power consumption.
- In Paper 26.3, Samsung Electronics shows a touch-screen controller for a 2D self-capacitance sensor supporting an active stylus in a large capacitive load of 100pF. A 216-channel IC fabricated in a 45nm CMOS process achieves an SNR of 45.8 and 62dB with a 500fF signal and 1-mm-φ active stylus, respectively.

Session 26 Highlights:
Display and User Interaction Technologies

[26.1] A Source-Driver IC Including Power-Switching Fast-Slew-Rate Buffer and 8Gb/s Effective 3-Tap DFE Receiver Achieving 4.9mV DVRMS and 17V/μs Slew Rate for 8K Displays and Beyond

Paper Authors: Kyungho Ryu, Ji-Yong Jeong, Jung-Pil Lim, Kil-Hoon Lee, Kyongho Kim, Yongil Kwon, Seongjong Yoo, Siwoo Kim, Hyun-Wook Lim, Jae-Youl Lee

Paper Affiliation: Samsung Electronics, Hwaseong, Korea

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

CONTEXT AND STATE OF THE ART

- Next-generation TVs and gaming monitors are expected to accommodate 8K resolution with a refresh rate of 120Hz. These display systems require thousands of channels with fast settling times, and high-speed data transfer to the display drivers.

TECHNICAL HIGHLIGHTS

- A power-switching structure achieves 4.9mV deviation of root-mean-square voltage output, and a fast-slew-rate (FSR) buffer shows 17V/μs slew-rate under 5kΩ and 350pF panel load, which is an improvement of 50% compared to the FSR off case.
- An effective 3-tap DFE receiver in a 0.18μm CMOS process overcomes 30dB interface channel insertion loss. A receiver front-end comprises a passive equalizer and continuous-time linear equalizer utilizing the active inductor and negative-capacitance.

APPLICATIONS AND ECONOMIC IMPACT

- Fast slew rate and an efficient source driver suitable for higher resolution and faster refresh rate displays. An 8K 65-inch display driven by the proposed source driver-ICs with 8Gb/s effective 3-tap adaptive equalizer.
- High-resolution and large-format-display monitors offer better quality and realistic images, which drives the 8K market growth. The 8K market is projected to grow from USD 3 Billion in 2019 to USD 27 Billion by 2024.

Session 27 Overview:
Innovations from Outside the (ISSCC) Box

Technology Directions Subcommittee

Session Chair: Ali Hajimiri, California Institute of Technology, Pasadena, CA

Session Co-Chair: Firooz Aflatouni, University of Pennsylvania, Philadelphia, PA

There have been numerous new and important developments outside the circuit community that can have profound impact on the solid-state circuits society either through emergence of new applications or via offering of new platforms for processing, communications, and sensing. This invited session seeks to provide exposure to a few of such developments for the ISSCC audience who may not have day-to-day interactions with such subjects in the hope of inspiring new ways of thinking in circuit design and collaboration opportunities.

- In Paper 27.1, Northwestern University presents a review paper on recent bioelectronic systems for cardiac therapy which are composed of a wireless and implantable stimulator and energy harvester, wearable ECG recorder and power transmitter. Bioresorbable implantable devices are used in a closed feedback loop to provide autonomous and adaptive treatment. These results provide high-impact opportunities for future integrated circuits.
- In Paper 27.2, Caltech reviews DNA nanotechnology, which has grown from tens of nucleotides in the early 1980s to system designs specifying millions of them today. Research in DNA nanotechnology includes well-defined self-assembled structures that were produced in a multi-stage process, one-shot spontaneous self-assembly process, and creating molecular systems that perform all information processing steps autonomously. DNA nanotechnology can lead to new platforms for computation or new substrates for circuits.
- In Paper 27.3, imec presents the potential of 2-D materials to replace Si in scaled technologies as they retain high mobility at atomically thin (below 1nm) channel thickness. Transition metal dichalcogenides (TMD) technology, which consist of a layer of transition metal atoms sandwiched between two layers of chalcogen atoms, as a suitable candidate for scaled logic applications is discussed.
- In Paper 27.4, Stanford University provides a review of recent advances in inverse designed photonics, where a physics-guided search through the full parameter space is performed until an optimal device design is reached. Resulting devices typically outperform conventional designs by many orders of magnitude in footprint, efficiency, and stability. Co-design of electronics and photonics opens a plethora of new applications.

Session 28 Overview:
High-Density Memories and High-Speed Interfaces

Memory Subcommittee

Session Chair: Seung-Jae Lee, SamsungHwaseong, Korea

Session Co-Chair: Dongkyun Kim, SK Hynix, Icheon, Korea

Innovations in 3D NAND Flash, very-high density 5b/cell, will be introduced. New challenges and solutions to improving reliability for DRAM will be presented: including probabilistic aggressor tracking against row-hammer attacks and core bias modulation to overcome process limitations. Meanwhile, the evolution of the memory high-speed interface continues: new technologies are introduced such as single-ended PAM4 signaling to achieve speeds exceeding 16Gb/s/pin, offset-calibration HBM3-interface technology to achieve a 1.15TB/s bandwidth, an input jitter filtering digital PLL technology, and an edge boosting equalizer using a t-coil.

- In Paper 28.1, Intel presents a 1.67-Tb 5b/cell Flash memory fabricated in a 192-layer floating-gate 3D-NAND technology, featuring a 23.3Gb/mm^2 bit density with a die capacity of 1.67Tb within an 73.3mm^2 area, and a t_R and t_{PROG} of 354µs and 5500µs.
- In Paper 28.2, SK Hynix presents a high-performance 1-Tb 3b/cell 3D-NAND Flash with 194MB/s write throughput for over 300 layers. Five new schemes are introduced and these design technologies enable a high-performance (a t_R of 34µs and a program throughput of 194MB/s) 1-Tb 3b/cell 3D-NAND Flash memory with a greater than 20Gb/mm^2 bit density, which uses the peripheral-circuit-under-cell-array architecture.
- In Paper 28.3, Samsung presents a 4-nm 16-Gb/s/pin single-ended PAM4 parallel transceiver with switching-jitter compensation and transmitter optimization. This paper achieves 0.764pJ/b within 0.0073mm^2. A relaxed TX termination of 20Ω and RX termination of 50Ω is adopted to maximize the eye opening.
- In Paper 28.4, Samsung presents a 4-nm 1.15-TB/s HBM3 interface with resistor-tuned offset calibration and in-situ margin detection. This work presents a read-valid-window-margin (VWM) improvement technique including an offset calibration, and an in-situ read VWM detection scheme. A compact slim-bit-slice architecture with a stacked I/O achieves a high-bandwidth of up to 1.15TB/s.
- In Paper 28.5, KAIST shows a 900-µW 1 – 4-GHz input-jitter-filtering digital-PLL-based 25%-duty-cycle quadrature-clock generator for ultra-low-power clock distribution for high-speed DRAM interfaces. This work presents a low-power clock-distribution scheme for DRAM, using a quadrature clock generator, which can generate accurate 25%-DC quadrature signals over a 1 - 4GHz range.
- In Paper 28.6, Seoul National University presents a 32-Gb/s/pin 0.51-pJ/b single-ended resistorless impedance-matched transmitter with a t-coil-based edge-boosting equalizer in 40nm CMOS. The 2-tap t-coil-based edge-boosting equalizer compensates for the high-frequency impedance drop and does not consume static current for non-transition sequences achieving a power efficiency of 0.51pJ/b.
- In Paper 28.7, Samsung presents a 1.1-V 6.4-Gb/s/pin 24-Gb DDR5 SDRAM with a highly-accurate duty-cycle corrector and an NBTI tolerant DLL. The 24-Gb density DDR5 occupies 71.8mm^2/channel, and is implemented in a 4th-generation 10-nm DRAM technology.
- In Paper 28.8, SK Hynix presents a 1.1-V 16-Gb DDR5 DRAM with probabilistic-aggressor tracking, a refresh-management function, per-row hammer tracking, a multi-step precharge, and core-bias-voltage modulation for security and reliability enhancement. This comprehensive scheme leads to a failure-probability reduction due to row hammer attacks by 93.1%, and an improvement to cell-retention time of 17%.

Session 28 Highlights:
High-Density Memories and High-Speed Interfaces

[28.1] A 1.67-Tb 5b/cell Flash Memory Fabricated in a 192-Layer Floating-Gate 3D-NAND Technology and Featuring a 23.3Gb/mm² Bit Density

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

[28.8] A 1.1-V 16-Gb DDR5 DRAM with Probabilistic-Aggressor Tracking, a Refresh Management Function, Per-Row Hammer Tracking, a Multi-Step Precharge, and Core-Bias-Voltage Modulation for Security and Reliability Enhancement

Paper 28.1 Authors: Ali Khakifirooz[1], Eduardo Anaya[2], Sriram Balasubrahmanyam[2], Geoff Bennett[1], Daniel Castro[2], John Egler[2], Kuangchan Fan[2], Rifat Ferdous[1], Kartik Ganapathi[1], Omar Guzman[2], Chang Wan Ha[1], Rezaul Haque[2], Vinaya Harish[2], Majid Jalalifar[2], Owen W. Jungroth[2], Sung-taeg Kang[1], Golnaz Karbasian[1], Jee-Yeon Kim[1], Siyue Li[2], Aliasgar S. Madraswala[2], Srivijay Maddukuri[2], Amr Mohammed[1], Shanmathi Mookiah[2], Shashi Nagabhushan[2], Binh Ngo[2], Deep Patel[2], Sai Kumar Poosarla[2], Naveen Prabhu V[2], Carlos Quiroga[2], Shantanu Rajwade[1], Ahsanur Rahman[2], Jalpa Shah[2], Rohit S. Shenoy[1], Ebenezer Tachie Menson[2], Archana Tankasala[1], Sandeep Krishna Thirumala[1], Sagar Upadhyay[2], Krishnasree Upadhyayula[2], Ashley Velasco[2], Nanda Kishore Babu Vemula[2], Bhaskar Venkataramaiah[2], Jiantao Zhou[1], Bharat M. Pathak[2], Pranav Kalavade[1]

Paper 28.1 Affiliation: [1]Intel, Santa Clara, CA, [2]Intel, Folsom, CA

Paper 28.3 Authors: Jahoon Jin, Soo-Min Lee, Kyunghwan Min, Sodam Ju, Jihoon Lim, Hyunsu Chae, Kwonwoo Kang, Yunji Hong, Yeongcheol Jeong, Sang-Ho Kim, Jongwoo Lee, Joonsuk Kim

Paper 28.3 Affiliation: Samsung Electronics, Hwaseong, Korea

Paper 28.8 Authors: Woongrae Kim, Chulmoon Jung, Seongnyuh Yoo, Duckhwa Hong, Jeongjin Hwang, Jungmin Yoon, Ohyong Jung, Joonwoo Choi, Sanga Hyun, Mankeun Kang, Sangho Lee, Dohong Kim, Sanghyun Ku, Donhyun Choi, Nogeun Joo, Sangwoo Yoon, Junseok Noh, Byeongyong Go, Cheolhoe Kim, Sunil Hwang, Mihyun Hwang, Seol-Min Yi, Hyungmin Kim, Sanghyuk Heo, Kyoungchul Jang, Shinho Chu, Yoonna Oh , Kwidong Kim, Junghyun Kim, Soohwan Kim, Jeongtae Hwang, Sangil Park, Junphyo Lee, Inchul Jeong, Joohwan Cho, Jonghwan Kim

Paper 28.8 Affiliation: SK hynix Semiconductor, Icheon, Korea

Subcommittee Chair: Meng-Fan (Marvin) Chang, National Tsing Hua University (NTHU), Taiwan

CONTEXT AND STATE OF THE ART

- Successful QLC development over many generations has paved the way for industry-wide adoption of QLC technology. The transition to a 5b/cell (PLC) will serve as another steppingstone to accelerate growth in bit-storage density.
- Meanwhile, the demand for high-speed interfaces continues. Recently, the four-level pulse-amplitude modulation (PAM4) signaling scheme has been adopted, as it throughput is doubled without an increase in the Nyquist frequency.
- DRAM manufacturers are facing technology scaling challenges beyond the 1a-nm DRAM process. The decrease in row-hammer tolerance and refresh-retention time, due to the reduced cell size, are regarded as critical challenges for device scaling.

TECHNICAL HIGHLIGHTS

- **Intel presents a 1.67-Tb 5b/cell flash memory fabricated in a 192-layer floating-gate 3D-NAND technology and featuring a 23.3Gb/mm^2 bit density.**
    - A seven-strobe fast/soft bit-read (FSBR) algorithm is implemented to improve error correction capabilities, and achieves a read time (t_R) of 354µs.
    - A program suspend/resume capability is implemented while supporting FSBR during suspend and without the need to additionally latch zthe page buffer.
- **Samsung presents a 4-nm 16-Gb/s/pin single-ended PAM4 parallel transceiver with switching-jitter compensation and transmitter optimization.**
    - This single-ended PAM4 transceiver uses a new switching-jitter compensation technique for the RX, which improves the timing margin by 0.06UI at 16Gb/s/pin, from 0.31 to 0.37UI, at the cost of a small additional circuit that occupies 75.2um^2.
    - A fractionally-spaced feedforward equalizer and a capacitive-peaked equalizer enhance the transmitted eye openings from 0.25 to 0.50UI.
- **SK Hynix presents a 1.1-V 16-Gb DDR5 DRAM with probabilistic-aggressor tracking, a refresh-management function, per-row hammer tracking, a multi-step precharge, and core-bias-voltage modulation for security and reliability enhancement.**
    - New schemes to improve row-hammer characteristics are introduced.
    - These schemes lead to a failure probability reduction due to row hammering by 93.1% and an improvement in cell retention time by 17%.

APPLICATIONS AND ECONOMIC IMPACT

- The transition to 5b/cell (PLC) expands Flash storage to wider markets, where a lower cost at a reasonable performance is paramount.
- The technology on PAM4 will be able to maximize the use of PAM4, which is in the spotlight as the future generation memory interface.
- A row-hammer protection and a refresh-management scheme are presented to guarantee DRAM security and reliability despite aggressive scaling from 1a-nm to sub-10-nm nodes; scaling-down can reduce manufacturing cost significantly.

Session 29 Overview:
Digital Accelerators and Circuit Techniques

Digital Circuits Subcommittee

Session Chair: Mingoo Seok, Columbia University, New York, NY

Session Co-Chair: Kazuki Fukuoka, Renesas, Tokyo, Japan

Subcommittee Chair: Keith Bowman, Qualcomm, Raleigh, NC

Digital accelerators and advanced circuit techniques remain as key drivers to support existing and emerging applications with high performance and energy efficiency. This session features seven new accelerators, namely two compute-in-memory-based k-SAT solvers, a neural engine leveraging sign-magnitude data representation, a heterogeneous RRAM CNN and SRAM SNN SoC, a microwatt keyword spotting system, a mixed-signal GPS accelerator and a universal soft-detection decoder. This session also features two papers proposing novel digital circuit techniques, namely a 3D-stacked capacitor-processor system and an 8T embedded non-volatile SRAM array.

- In Paper 29.1, Georgia Institute of Technology researchers present a k-SAT solver that features a mixed-signal stochastic recurrent neural network and unsupervised learning to search for an optimal solution. The 65nm test chip operates at 1.2V and 400MHz and consumes 32.5mW. It demonstrates 74.0% solvability and 11.25ms median runtime for 99% satisfiability in 30-variable and 126-clause hard 3-SAT problems.
- In Paper 29.2, the University of Texas at Austin describe a fast, reliable and scalable all-digital compute-in-memory SAT solver in 65nm CMOS. The solver allows one-shot parallel computation of up to 1024 clauses of 128 variables, to converge on satisfiability in only 710⌠s and 1.1⌠J on 60-variable 256-clause hard 3-SAT problems.
- In Paper 29.3, University of Michigan researchers present a 28nm neural engine that leverages sign-magnitude data representation and bit sparsity for increased energy efficiency over 2's complement implementations. The comparative implementation in a 2.18mm² chip achieves a 53% peak energy-efficiency improvement at 0.65V with 8.09TOPS/W.
- In Paper 29.4, Graphcore presents wafer-level stacking of high-density capacitors in a 130nm deep-trench-capacitor technology to enhance the performance and V_{MIN} of a 7nm 1472-tile multithreaded in-order processor. The stacked chips of 822mm² integrate 750⌠F capacitance to reduce peak-to-peak voltage swings from 250mV to 47mV, enabling a 450MHz higher (+32%) clock frequency at 0.75V or a 140mV (17%) V_{MIN} improvement at 1.575GHz.
- In Paper 29.5, a joint Georgia Institute of Technology and TSMC team reports a heterogeneous RRAM CNN and SRAM SNN SoC for hybrid frame and event-based target tracking. The system occupies 20.25mm^2 in a 40nm process, supports RRAM-CIM throughput of 14.74TOPS, and matches the bandwidth of the event camera for 100 outputs/s with efficiency of 73.53TOPS/W for the fused frame and event vision.
- In Paper 29.6, the University of Michigan describes a fully integrated keyword spotting system, which employs a skip recurrent neural network (RNN) to reduce the power consumption of the analog front-end and the digital back-end by adaptively sub-sampling audio input frames based on the signal content. The proposed system consumes 1.5⌠W in 28nm CMOS and achieves 92.8% accuracy on a standard dataset.
- In Paper 29.7, researchers from Shanghai Jiao Tong University and Columbia University present a GPS accelerator with charge-based correlation computing that employs an area-efficient, two-phase mixed-signal multiplier. The 28nm test chip achieves 114-394TOPS/W, demonstrating an 8.2· energy-efficiency improvement compared to the state-of-the-art design at the same throughput.
- In Paper 29.8, researchers from Boston University, Massachusetts Institute of Technology and Maynooth University present an integrated universal soft-detection decoder using ordered reliability bits guessing random additive noise decoding (ORBGRAND). The 40nm chip consumes 0.76pJ per decoded bit, while achieving a throughput of 6.5Gbps and an average latency of 40ns at a frame error rate of 10^{-7}.
- In Paper 29.9, the University of California presents a multi-time programmable 8T eNVSRAM macro, which can operate as a regular SRAM and an NVM in a 22nm FDSOI standard logic process. The proposed 8T bitcell of 0.86⌠m² implemented in a 1Kb array enables local non-volatile data storage within the SRAM bitcell with charge-trap transistors, achieving 0.55ns SRAM read and 2ns NVM recall.

Session 29 Highlights:
Digital Accelerators and Circuit Techniques

[29.4] Wafer-Level Stacking of High-Density Capacitors to Enhance the Performance of a Large Multicore Processor for Machine Learning Applications

Paper Authors: S. Felix[1], S. Morton[2], S. Stacey[1], J. Walsh[1]

Paper Affiliation: [1]Graphcore, Bristol, UK, [2]Graphcore, Adelaide, Australia

Subcommittee Chair: Keith Bowman, Qualcomm, Raleigh, NC, Digital Circuits Subcommittee

CONTEXT AND STATE OF THE ART

- Application-dependent load transients induce supply-voltage disturbances in high-performance machine learning (ML) processors. These disturbances require large supply-voltage margins to ensure correct operation, severely degrading processor energy efficiency.
- An approach to reduce the magnitude of these voltage disturbances is to integrate ultra-high decoupling capacitance within the packaged die.

TECHNICAL HIGHLIGHTS

- **Graphcore demonstrates the use of wafer-level stacking of high-density capacitors to enhance the performance and minimum operating voltage (V_{MIN}) of a 1472 multithreaded in-order processor for ML applications.**
    - The capacitors and processor are manufactured in 0.13μm and 7nm process technologies, respectively. The high-density capacitors enable a 450MHz higher (+32%) clock frequency at 0.75V or a 140mV (17%) V_{MIN} improvement at 1.575GHz.
    - This system employs TSMC's wafer-on-wafer technology to stack the 1472-tile processor chip with an approximately 750μF stacked capacitor die to suppress peak-to-peak supply-voltage swings from 250mV to 47mV, resulting from application-dependent load transients. The reduced voltage swing allows a reduction in the voltage margins, resulting in either higher performance or lower V_{MIN}.

APPLICATIONS AND ECONOMIC IMPACT

- Machine learning (ML) processors along with increasingly complex models consume large quantities of energy for training and inference. Circuits, architectures and algorithms for energy-efficient ML processors will have a significant and lasting impact on the carbon footprint of ML applications.

Session 29 Highlights:
Digital Accelerators and Circuit Techniques

[29.5] A 73.53TOPS/W 14.74TOPS Heterogeneous RRAM In-Memory and SRAM Near-Memory SoC for Hybrid Frame and Event-Based Target Tracking

[29.6] A 1.5μW End-to-End Keyword Spotting SoC with Content-Adaptive Frame Sub-Sampling and Fast Settling Analog Frontend

Paper 29.5 Authors: M. Chang[1], A. Lele[1], S. Spetalnick[1], B. Crafton[1], S. Konno[1], Z. Wan[1], A. Bhat[1], W-S. Khwa[2], Y-D. Chih[3], M-F. Chang[2], A. Raychowdhury[1]

Paper 29.5 Affiliation: [1]Georgia Institute of Technology, Atlanta, GA, [2]TSMC Corporate Research, Hsinchu, Taiwan, [3]TSMC Design Technology, Hsinchu, Taiwan

Paper 29.6 Authors: J-H. Seol, H. Yang, R. Rothe, Z. Fan, Q. Zhang, H-S. Kim, D. Blaauw, D. Sylvester

Paper 29.6 Affiliation: University of Michigan, Ann Arbor, MI

Subcommittee Chair: Keith Bowman, Qualcomm, Raleigh, NC, Digital Circuits Subcommittee

CONTEXT AND STATE OF THE ART

- Moving on from cloud applications, machine learning (ML) is now making progress all the way to edge and low-power embedded devices. Accelerators are starting to exploit various ML algorithm features as an opportunity to improve the processor energy efficiency for sensor applications at the edge, such as perception or speech recognition.
- High-density, non-volatile resistive RAM (RRAM) with compute-in-memory capability is a promising candidate for implementing highly efficient convolutional neural networks (CNNs), demonstrating high-memory-capacity integration to enable a broad range of complex real-world applications.
- In applications such as vision-based target identification and tracking, traditional frame cameras with CNN-based processing provide high accuracy at a cost of limited tracking capability due to inadequate throughput. Meanwhile, processing event streams from dynamic vision sensors (DVS) with spiking neural networks (SNNs) support high throughput, while suffering from low accuracy.

TECHNICAL HIGHLIGHTS

- **Georgia Institute of Technology and TSMC report a heterogeneous RRAM CNN and SRAM SNN SoC, enabling 73.53TOPS/W and 14.74TOPS to support a 100 outputs/s event camera to realize hybrid frame and event-based target tracking.**
    - The hybrid architecture effectively synergizes the advantages of: 1) a CNN frame pipeline enabled by RRAM-based compute cores, and 2) an SNN event pipeline enabled using SRAM-based cores to successfully realize fused frame and event vision.
    - A triple-error-correction (TEC) scheme provides tolerance for high variations in RRAM CIM operations to restrict the application-level accuracy degradation to less than 1%.
    - The architecture allows for SRAM-based continuous accesses with parallel-operating processing elements for processing up to 11.1 million events/s to handle real-time DVS data.

- **The University of Michigan reports a fully integrated keyword spotting system implemented in 28nm that consumes 1.5μW and achieves 92.8% accuracy on a standard dataset.**
    - The architecture combines energy-efficient design techniques with use of a skip recurrent neural network (RNN) to control power dissipation in both the analog front-end and the digital back-end for overall system power reduction.
    - The RNN adaptively sub-samples audio input frames based on the signal contents and computes an optimal number of 16ms audio frames to skip, while powering down the analog front-end, feature extractor and the neural network classifier that executes the RNN.
    - A two-stage amplifier combined with two-step frequency control achieves a reduction of the analog front-end settling time from 37ms to 1ms, allowing power down during frame skip.

APPLICATIONS AND ECONOMIC IMPACT

- These contributions demonstrate continued progress in the use of ML techniques to advance the state of the art in analog/mixed signal applications at the edge.
- High-capacity RRAM integrated with SRAMs in hybrid compute-in-memory/digital SoCs pave the way for an increasing number of real-world edge applications requiring intelligence on the edge.

Session 30 Overview: Power Management Techniques

Power Management Subcommittee

Session Chair: Xun Liu, Chinese University of Hong Kong, Shenzhen, China

Session Co-Chairs: Gael Pillonnet, CEA-Léti, Grenoble, France

Power-management techniques have been adopted to improve the system performance of emerging power-converter applications such as energy harvesting, high-voltage actuation, step-up-and-down power supply, and wireless power transfer (WPT). Efficient synchronized-switch harvesting-on-inductor (SSHI) extraction technique, maximum power point tracking, ultra-wide conversion ratio, scalable multi-chip stackable bias-flip technique, various DC-DC converters topologies (hybrid, buck-boost, switched capacitor), and advanced WPT systems are presented in this session to demonstrate the state-of-the-art performance of these power management systems.

- In Paper 30.1, Ulsan National Institute of Science and Technology presents a scalable N-step equal split SSHI that achieves an energy power extraction of 1170% with a low-Q 1μH inductor. This work operates with 91% flip efficiency and scalable 4-to-16 steps, and consumes only 22nA in its sleep mode.
- In Paper 30.2, National Yang Ming Chiao Tung University, Chip-GaN Power Semiconductor, and Realtek Semiconductor present a multi-input bipolar energy harvester with ultra-low-power comparator and zero current detection (ZCD) to achieve low input-voltage ripple, high accuracy, and 93.2% end-to-end efficiency.
- In Paper 30.3, Delft University of Technology presents a new duty-cycle-based MPPT algorithm, validated in a piezoelectric energy harvesting system with a bias-flip rectifier, that achieves 98% peak MPPT efficiency and 738% power-extraction enhancement compared to a passive rectifier.
- In Paper 30.4, Dartmouth College presents a modular/scalable switched capacitor (SC) converter with integrated auxiliary boost converter for high-voltage electrostatic and piezoelectric (PZT) actuators. With 3 chips stacked, the design provides up to $1kV_{pp}$ (VCR>270), delivering and recovering >1W reactive power at >96% efficiency.
- In Paper 30.5, Ulsan National Institute of Science and Technology and Samsung Electronics present a 3-level boost converter with a fully state-based phase selection technique and an adaptive slope generator to select the operation mode, balance the flying capacitor voltage and smoothen the mode transition, which achieves a high efficiency (~95.3%) and a wide output range (5 to 32V).
- In Paper 30.6, University of Science and Technology of China and Hefei CLT Microelectronics present a flying-capacitor-assisted buck-boost converter with only 4 switches (which are all immune from voltage stress issue) to reduce the average inductor current and to achieve 98.6% peak efficiency and $1.47A/mm^2$ on-die current density.
- In Paper 30.7, University of Macau, Zhuhai UM Science & Technology Research Institute, and Instituto Superior Tecnico/University of Lisboa reconfigure VCF steps at different VCRs in continuously scalable-conversion-ratio SC (CSC) converters to achieve an optimum efficiency under limited resources. It achieves a 90% peak efficiency with 33 VCF steps and the widest VCR range with a 75% efficiency.
- In Paper 30.8, National Yang Ming Chiao Tung University, Chip-GaN Power Semiconductor, and Realtek Semiconductor present a 3D wireless power transfer (WPT) system, including 3TXs and one central controller to arrange the power transfer according to the position of the RX. The design uses a noise cancellation technique to enhance the received signal by 52% with an auxiliary coil.
- In Paper 30.9, University of Texas at Richardson presents a 40.68MHz single-stage dual-output rectifier composed of only 3 active diodes to achieve the highest resonance frequency, the smallest RX coil size, and a competitive power efficiency of 90.1%. Synchronous PFM control is developed to provide fast load transient responses with unnoticeable output cross regulation.
- In Paper 30.10, Renesas Electronics presents a single-chip Qi-compliant 40W wireless-power-transmission controller, implementing dual RMS and DC current sensing, and adaptive ZVS method to increase the EMI figure by 4dB and to improve the efficiency by up to 1.7%.

Session 30 Highlights: Power Management Techniques

[30.1] A Scalable N-Step Equal Split SSHI Piezoelectric Energy Harvesting Circuit Achieving 1170% Power Extraction Improvement and 22nA Quiescent Current with a 1µH-to-10µH Low Q Inductor

[30.5] A 95.3% 5V-to-32V Wide Range 3-Level Current Mode Boost Converter with Fully State-based Phase Selection Achieving Simultaneous High-Speed VCF Balancing and Smooth Transition

[30.8] 3D Wireless Power Transfer with Noise Cancellation Technique for -62dB Noise Suppression and 90.1%-Efficiency

Paper 30.1 Authors: Yeon-Woo Jeong[1], Seung-Ju Lee[1], Jong-Hun Kim[1], Mun-Jung Cho[1], Hwa-Soo Kim[1], Se-Un Shin[1]

Paper 30.1 Affiliation: [1]Ulsan National Institute of Science and Technology, Ulsan, Korea

Paper 30.5 Authors: Seung-Ju Lee[1], Yeon-Woo Jeong[1], Mun-Jung Cho[1], Jong-Hun Kim[1], Hwa-Soo Kim[1], Jun-Suk Bang[2], Se-Un Shin[1]

Paper 30.5 Affiliation: [1]Ulsan National Institute of Science and Technology, Ulsan, Korea, [2]Samsung Electronics, Hwaseong, Korea

Paper 30.8 Authors: Fei Huang[1], Chi Yu Huang[1], Yu-Chun Luo[1], Ke-Horng Chen[1], Ying-Hsi Lin[2], Shian-Ru Lin[2], Tsung-Yen Tsai[2]

Paper 30.8 Affiliation: [1]National Yang Ming Chiao Tung University, Hsinchu, Taiwan, [2]Realtek Semiconductor, Hsinchu, Taiwan

Subcommittee Chair: Bernhard Wicht, University of Hannover, Hannover, Germany

CONTEXT AND STATE OF THE ART

- Circuits with piezoelectric transducers that are used for vibration energy harvesting typically use a full-bridge rectifier and to minimize the charge loss in the inhernet capacitor of the piezoelectric transduce, they require bulky inductors.
- Multi-level DC-DC boost converters are widely used for various applications that require a high supply voltage. However, in practice, such converters suffer from a flying capacitor voltage imbalance which causes not only large losses in the switches and the inductor, but also damage to the switches. To minimize the losses they require robust calibration techniques.
- The power efficiency of resonant-based wireless power transfer systems depends on the receiver coil positioning.

TECHNICAL HIGHLIGHTS

- **Ulsan National Institute of Science and Technology introduces a scalable N-step synchronized-switch harvesting-on-inductor (SSHI) showing the best extraction efficiency compared to previous state-of-the-art with low-Q flipping inductor.**
    - A scalable 4-to-16 steps equal split SSHI achieves 91% flip efficiency and 1170% power extraction improvement with a 1µH low-Q inductor and dedicated zero-current detector, consuming only 22nA in a sleep mode. The proposed approach enables reducing the need of bulky inductors.

- **Ulsan National Institute of Science and Technology presents robust capacitor voltage balancing technique for hybrid DC-DC converters without creating current sub-harmonics and allowing current-mode control.**
    - A 3-level boost converter with a fully state-based phase-selection technique and an adaptive slope generator to select the operation mode, balance the flying capacitor voltage and smoothen the mode transition introduces robust calibration and sensing techniques that help achieving a high efficiency (~95.3%) and a wide output range (5-to-32V).

- **National Yang Ming Chiao Tung University, presents a 3D wireless power transfer (WPT) system, including 3TXs and one central controller to arrange the power transfer according to the position of the RX.**
    - The noise cancellation technique uses the auxiliary coil to enhance the received signal by 52% while the noise can be reduced to –62.6dB. The demodulator and the vector detector at the central controller can get the received k factor of each TX to decide the power transfer of each TX. The peak efficiency can reach 90.1% and the maximum power is 69W from 3TXs.

APPLICATIONS AND ECONOMIC IMPACT

- Advanced circuit techniques for highly efficient energy extraction with low bill of material improve the power density of vibration energy harvesting system.
- Robust and high-performance capacitor voltage balancing paves the way for wide adoption of hybrid converters.
- The proposed 3D WPT can break through the spatial limitations to ensure RX can receive energy from TX in a 3D space.

Session 31 Overview:
Energy-Efficient Radios for UWB, BMI, and IoT Systems

Wireless Subcommittee

Session Chair: Negar Reiskarimian, Massachusetts Institute of Technology, Cambridge, MA

Session Co-Chair: Jan Prummel, Renesas Electronics Corporation, 's-Hertogenbosch, The Netherlands

This session presents energy-efficient wireless systems for communications and ranging applications. The first two papers focus on ultra-wideband systems for positioning. The next two papers focus on data links for neural implants, and the final four papers cover various techniques to achieve low-power, high TX efficiency and high RX sensitivity in NB-IoT and BLE radios.

- In Paper 31.1, Tsinghua University presents a quadrature uncertain-IF IR-UWB transceiver with Twin-OOK modulation, achieving -71dBm sensitivity at 10Mb/s and 0.96cm ranging resolution.
- In Paper 31.2, Samsung Electronics & Pusan National University present a fully-integrated IEEE 802.15.4/4z-Compliant 28nm CMOS UWB SoC that demonstrates a highly integrated solution with 14.25dBm TX peak power, and -102.5dBm RX sensitivity.
- In Paper 31.3, Tsinghua University & Beijing Ningju Technology Co. show a 1.8Gb/s, 2.3pJ/bit, crystal-less IR-UWB transmitter for neural implant applications, featuring a hybrid modulation scheme combining Differential 16-Pulse-Position-Modulation (D16PPM), PWM and DBPSK.
- In Paper 31.4, Zhejiang University demonstrates a 128-channel 2mm×2mm battery-free neural dielet, which merges the channels through multi-carrier orthogonal backscatter, achieving a data-rate of 20.16Mb/s and an energy efficiency of 0.8pJ/bit.
- In Paper 31.5, Zhejiang University and Microaiot demonstrate an integrated passive BLE tag realizing battery-free communication in tag-to-tablet/smartphone, tablet/smartphone-to-tag, and tag-to-tag modes, achieving a data-rate of 1Mb/s and a 160m backscatter range towards a BLE commodity tablet.
- In Paper 31.6, the University of Macau & Instituto Superior Tecnico/University of Lisboa present an ultra-low-power (ULP) long-range active-RF Tag in 65nm CMOS achieving a 20.5% TX efficiency at -22dBm EIRP, and a -60.4dBm RX sensitivity at 17.8nW.
- In Paper 31.7, Fudan University presents a 0.7-to-2.5GHz sliding digital-IF quadrature digital transmitter in 28nm CMOS for multimode/multiband NB-IoT/BLE applications, achieving 40.9% average system efficiency with 25.4dBm average output power and -15.9dB EVM at 0.9GHz NB-IoT.
- In Paper 31.8, Oregon State University presents a distributed N-path mixer-based noise-cancelling architecture with integrated passives combined with spread-spectrum two-tone modulation in a ULP 0.4-to-0.95GHz RX with modulated interferer and LO jitter tolerance. The 22nm RX achieves 100kb/s with -85dBm sensitivity, SIR of -14 and -41dBm at 2MHz and 25MHz offset, respectively for 2MS/s 16-QAM blockers, and consumes 100 to 149uW.

Session 31 Highlights:
Energy-Efficient Radios for UWB, BMI, and IoT Systems

[31.2] A Fully Integrated IEEE 802.15.4/4z-Compliant 6.5-to-8GHz UWB System-on-Chip RF Transceiver Supporting Precision Positioning in a CMOS 28nm Process

[31.3] A 1.8Gb/s, 2.3pJ/bit, Crystal-Less IR-UWB Transmitter for Neural Implants

Paper 31.2 Authors: Wan Kim[1], Hyun-Gi Seok[1], Geunhaeng Lee[1], Jaekeun Lee[1], Sinyoung Kim[1], Chanho Kim[1], Wonkang Kim[1], Youngsea Cho[1], Seungyong Bae[1], Jongpil Cho[1], Wonjun Jung[1], Hyeokju Na[1], Byoungjoong Kang[1], Honggul Han[1], Chiyoung Ahn[1], Hoon Kang[1], Sukjin Jung[1], Hyukjun Sung[1], Seunghyun Oh[1], Ji-Seon Paek[2], Jongwoo Lee[1], Joonsuk Kim[1]

Paper 31.2 Affiliation: [1]Samsung Electronics, Hwaseong, Korea, [2]Pusan National University, Pusan, Korea

Paper 31.3 Authors: Jiaxin Lei[1], Xiliang Liu[2], Wei Song[1], Heng Huang[1], Xiaoyan Ma[2], Junliang Wei[2], Milin Zhang[1]

Paper 31.3 Affiliation: [1]Tsinghua University, Beijing, China, [2]Beijing Ningju Technology,Beijing, China

Subcommittee Chair: Chih-Ming Hung, MediaTek, Taipei, Taiwan

CONTEXT AND STATE OF THE ART

- To support the growing UWB market for emerging applications, low-cost SoCs are needed that can integrate all required functions within a single chip.
- Implantable brain-machine interfaces (BMI) require low power operation with multi-Gb/s throughput that can be achieved with complex modulation schemes in IR-UWB systems.

TECHNICAL HIGHLIGHTS

- **Samsung Electronics & Pusan National University introduce a fully integrated IEEE 802.15.4/4z-Compliant 2T/3R UWB SoC achieving 14.25dBm TX peak power and -102.5dBm RX sensitivity**
    - A non-uniform quantization DAC is used to enhance pulse quality. DC offset cancellation with an averaging algorithm is implemented to improve SNR and a low-cost RX I/Q mismatch calibration technique is embedded for precise ranging measurements.
- **Tsinghua University and Beijing Ningju Technology present a crystal-less IR-UWB transmitter with 1.8Gb/s data-rate and 2.3pJ/bit TX energy efficiency for neural implants**
    - A hybrid modulation scheme combining Differential 16-Pulse-Position-Modulation, PWM and DBPSK is proposed to balance the data-rate and transmission range. The proposed hybrid modulation is realized by a ring-oscillator-based all-digital transmitter. A transcutaneous transmission range of 20cm is measured with an 18mm pork tissue applied.

APPLICATIONS AND ECONOMIC IMPACT

- UWB promises high data-rate and precise positioning over a small distance, which can be useful in many consumer electronic and emerging brain-machine interface applications.
- IEEE802.15.4a and 4z standardize impulse radio for data communication and indoor local positioning.
- Smartphones have already started to adopt UWB technology since 2021 to enhance spatial awareness. New chipsets for sensor nodes will release the full potential of UWB, including low-power, secure, Gb/s data-rate, short-range communications.

Session 32 Overview:
Intelligent Biomedical Circuits and Systems

IMMD Subcommittee

Session Chair: Jun-Chau Chien, National Taiwan University, Taipei, Taiwan

Session Co-Chair: Mahsa Shoaran, EPFL, Lausanne, Switzerland

This session covers recent advances in intelligent biomedical interface designs for various applications. The papers discuss innovative technologies to improve the energy efficiency and form factor of wearable or implantable systems, enable more reliable bio-potential recording and effective stimulation-based therapies, and realize smart machine-learning-enabled neural interface SoCs. The first paper presents a highly integrated behind-the-ear wearable platform for multimodal mental-health monitoring. The second paper describes a sub-retinal prosthesis with reduced stimulus scattering, followed by wide input-range energy-efficient interfaces for current measurement and bio-potential recording. A 384-channel online spike-sorting IC comes next, followed by a patient-independent Convolutional Neural Network-SoC for seizure detection in epilepsy. The last paper describes a wireless bidirectional interface for machine-learning-enabled fascicle-selective peripheral nerve stimulation.

- In Paper 32.1, Ulsan National Institute of Science and Technology presents a behind-the-ear patch-type multimodal integrated interface with 275-fold input impedance boosting for continuous mental health monitoring.
- In Paper 32.2, Korea University describes a 505-channel pixel-sharing sub-retinal prosthesis SoC with 35.8dB dynamic range via time-based photodiode sensing. The SoC adopts a modular active/return electrode scheme to reduce stimulus scattering by 95.4%, and per-pixel dynamic voltage scaling to save stimulator power by up to 64%.
- In Paper 32.3, Daegu Gyeongbuk Institute of Science and Technology presents a 1V 136.6dB-DR 4kHz-BW ΔΣ current-to-digital converter that employs a truncation-noise-shaped baseline-servo-loop for precise current measurement in emerging applications such as PPG monitoring.
- In Paper 32.4, Daegu Gyeongbuk Institute of Science and Technology introduces a 2^{nd}-order noise-shaping SAR-ADC with enhanced input impedance, 1kHz BW, and 181.9dB FOM_{DR}, suitable for ExG recording in wearable applications.
- In Paper 32.5, imec and Fudan University present a 384-channel online spike-sorting IC with unsupervised clustering, achieving software-comparable accuracy, an area of 0.0013mm^2/channel, and power consumption of 1.78μW/channel. The chip is suitable for real-time high-density spike-sorting applications.
- In Paper 32.6, National University of Singapore presents a 0-shot-retraining machine-learning SoC for patient-independent seizure detection in epilepsy. The chip integrates a seizure-cluster-inception Convolutional Neural Network (SciCNN) and 22-channel double-data-rate Nyquist AFE, and is verified on human EEG and intracranial EEG recordings.
- In Paper 32.7, University of Toronto describes a 64-channel bidirectional peripheral nerve interface for minimally invasive fascicle-selective stimulation. The wireless chip features 173dB FOM noise-shaping SAR ADCs and 1.38pJ/b frequency-multiplying current-ripple transmitter, and is verified *in vivo* in a closed-loop machine-learning-enabled nerve activation scenario.

Session 32 Highlights:
Intelligent Biomedical Circuits and Systems

[32.1] A Behind-The-Ear Patch-Type Mental Healthcare Integrated Interface with 275-Fold Input Impedance Boosting and Adaptive Multimodal Compensation Capabilities

[32.2] A Stimulus-Scattering-Free Pixel-Sharing Sub-Retinal Prosthesis SoC with 35.8dB Dynamic Range Time-Based Photodiode Sensing and Per-Pixel Dynamic Voltage Scaling

Paper 32.1 Authors: Hyunjoong Kim[1], Myeongwoo Kim[1], Kwangmuk Lee[2], Sanghyeon Cho[1], Chan Sam Park[1], Solwoong Song[3], Dae Sik Keum[4], Dong Pyo Jang[3], Jae Joon Kim[1]

Paper 32.1 Affiliation: [1]Ulsan National Institute of Science and Technology, Ulsan, Korea, [2]Samsung Electronics, Hwaseong, Korea, [3]Hanyang University, Seoul, Korea, [4]SOSO H&C Daegu, Korea

Paper 32.2 Authors: Kyeongho Eom[1], Minju Park[1], Han-Sol Lee[1], Seung-Beom Ku[1], Namju Kim[2], Seongkwang Cha[3], Yong Sook Goo[3], Sohee Kim[2], Seong-Woo Kim[4], Hyung-Min Lee[1]

Paper 32.2 Affiliation: [1]Korea University, Seoul, Korea, [2]Daegu Gyeongbuk Institute of Science and Technology, Daegu, Korea, [3]Chungbuk National University, Cheongju, Korea, [4]Korea University Guro Hospital, Seoul, Korea

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

CONTEXT AND STATE OF THE ART

- Conventional brain activity monitoring devices, including head bands and scalp caps, are not appropriate for daily-life monitoring applications because of their discomfort and electrode-related artifacts. Therefore, a wearable platform suitable for daily monitoring is needed to enable mental healthcare applications.
- Retinal prostheses treat age-related macular degeneration (AMD) and retinal pigment degeneration to restore visual function. However, they generally require an external camera to capture the visual information before applying current stimuli to the retinal cells. A sub-retinal implant integrating light detectors and simulators will enhance the functionality and safety of retinal prostheses.

TECHNICAL HIGHLIGHTS

- **Ulsan National Institute of Science and Technology introduces a behind-the-ear wearable patch SoC with multimodal functionality for mental health applications.**
    - The SoC showcases the highest level of integration in a behind-the-ear wearable platform measuring EEG, ECG, PPG, GSR, and BioZ, and can perform transcutaneous vagus-nerve stimulation (VNS). The system employs an auxiliary compensation path in the ExG readout circuit and achieves an input impedance of 2.06GΩ at an input-referred noise of 0.22μVrms under 54μW of power consumption.

- **Korea University presents a sub-retinal prosthesis SoC that integrates photodiode arrays and power-efficient stimulators with reduced stimulus scattering through active/return electrode configurations.**
    - The SoC includes 505 pixel-sharing channels integrating stimulators with per-pixel dynamic voltage scaling technique and 35.8dB dynamic-range time-based photodiode sensors for 19klux light detection, enabling a sub-retinal prosthesis for visual restoration without the need for a bulky camera.

APPLICATIONS AND ECONOMIC IMPACT

- A breakthrough behind-the-ear wearable system with multimodal recording and stimulation enables mental healthcare applications from daily monitoring.
- A low-power-consumption and high-input-impedance recording front-end enables robust measurements of reduced biopotential signals.
- A sub-retinal implant integrating both high dynamic-range light detectors and retinal simulators with low stimulus scattering enhances the patient's visual acuity without the need for an external bulky camera.

Session 33 Overview:
Non-Volatile Memory and Compute-In-Memory

Memory Subcommittee

Session Chair: Hidehiro Shiga, KIOXIA, Tokyo, Japan

Session Co-Chair: Takashi Ito, Renesas Electronics, Tokyo, Japan

Non-volatile memory continues to be applied to more advanced process node and continues to expand into a wider range of applications. The first paper introduces 16nm STT-MRAM macro, which can operate at automotive temperatures. Next, a 22nm near-memory-computing-macro using STT-MRAM is presented, followed by 9Mb FeRAM macro with 10^{12} endurance. The last paper presents compute-in-memory macro using STT-MRAM.

- In Paper x.1, TSMC presents 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.
- In Paper x.2, National Tsing Hua University shows a 22nm 8Mb STT-MRAM near-memory-computing macro with 8b-precision and 46.4-160.1 TOPS/W.
- In Paper x.3, Institute of Microelectronics of the Chinese Academy of Sciences presents a 9Mb HZO-based embedded FeRAM macro with 10^{12} cycle endurance and 5ns-read 7ns-write.
- In Paper x.4, Southeast University introduces a 28nm 2Mb STT-MRAM computing-in-memory macro with refined bit-cell and 22.4-41.5 TOPS/W.

Session 33 Highlights:
Non-Volatile Memory and Compute-In-Memory

[33.1] A 16-nm 32-Mb Embedded STT-MRAM with a 6ns Read-Access Time, 1M-Cycle Write Endurance, and a 20-Year Retention at 150ºC and an MTJ-OTP Solution for Magnetic Immunity

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

Paper 33.1 Authors: Po-Hao Lee, Chia-Fu Lee, Yi-Chun Shih, Hon-Jarn Lin, Yen-An Chang, Cheng-Han Lu, Yu-Lin Chen, Chieh-Pu Lo, Chung-Chieh Chen, Cheng-Hsiung Kuo, Tan-Li Chou, Chia-Yu Wang, J.J. Wu, Roger Wang, Harry Chuang, Yih Wang, Yu-Der Chih, Tsung-Yung Jonathan Chang

Paper 33.1 Affiliation: TSMC, Hsinchu, Taiwan

Paper 33.2 Authors: Yen-Cheng Chiu[1], Win-San Khwa[2], Chung-Yuan Li[1], Fang-Ling Hsieh[1], Yu-An Chien[1], Guan-Yi Lin[1], Po-Jung Chen[1], Tsen-Hsiang Pan[1], De-Qi You[1], Fang-Yi Chen[1], Andrew Lee[1], Chung-Chuan Lo[1], Ren-Shuo Liu[1], Chih-Cheng Hsieh[1], Kea-Tiong Tang[1], Yu-Der Chih[3], Tsung-Yung Chang[3], Meng-Fan Chang[1,2]

Paper 33.2 Affiliation: [1]National Tsing Hua University, Hsinchu, Taiwan, [2]TSMC Corporate Research, Hsinchu, Taiwan, [3]TSMC, Hsinchu, Taiwan

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

CONTEXT AND STATE OF THE ART

- A high-density embedded MRAM with high endurance and long data retention, using an MTJ-based one-time-programmable solution for magnetic immunity.
- A STT-MRAM-based near-memory-computing non-volatile CIM macro using a system/software/circuitry co-design approach, showing the highest (thus far) read bandwidth and energy efficiency.

TECHNICAL HIGHLIGHTS

- **TSMC introduces a 32-Mb embedded STT-MRAM with 20-year retention at 150ºC.**
    - Using 16-nm FinFET technology, the 32-Mb embedded STT-MRAM with a 6ns read access time and a 1M-cycle write endurance has been fabricated, with a 0.033μm^2 cell size.
- **National Tsing Hua University and TSMC present an STT-MRAM near-memory-computing macro.**
    - A 22-nm 8-Mb STT-MRAM CIM achieves a 436GB/s read bandwidth and 46.4 - 160.1TOPS/W energy efficiency for 576 accumulations for 8b-inputs and 8b-weights.

APPLICATIONS AND ECONOMIC IMPACT

- High-endurance and long-retention MRAM is essential for microcontrollers that are used in automotive applications.
- Non-volatile CIM is necessary for edge-AI devices due to its low-power characteristics.

Session 34 Overview:
Cryo-CMOS for Quantum Computing

Technology Directions Subcommittee

Session Chair: Fabio Sebastiano, Delft University of Technology, Delft, The Netherlands

Session Co-Chair: Giorgio Ferrari, Politecnico di Milano, Milano, Italy

This session focuses on cryo-CMOS ICs for quantum computing. The first three papers describe cryo-CMOS quantum control ICs for superconducting quantum processors. Next, a cryo-CMOS VCO achieving a state-of-the-art FoM is described. The final paper of the session describes a low-power 260GHz cryo-CMOS backscatter transceiver for communication between 300K to 4K.

- In Paper 34.1, MIT presents a 22nm cryo-CMOS IC enabling wireless communication at 260GHz between 300K and 4K. Using OOK modulation, data rates for a link from 300K to 4K and vice-versa of 4.4Gbps and 4Gbps are realized while achieving energy efficiencies of 34fJ/bit and 176fJ/bit, respectively.
- In Paper 34.2, Google presents a 28nm Cryo-CMOS IC for full control of a two-qubit superconducting quantum processor unit cell. The IC dissipates 4mW/qubit under active control at 3K and is benchmarked for both single- and two-qubit gates.
- In Paper 34.3, Tsinghua University, Nanyang Technological University, and Beijing Academy of Quantum Information Sciences present a cryo-CMOS IC for XY and Z control of superconducting qubits, with the XY and Z controllers dissipating 13.4mW and 0-to-7.5mW per qubit under active control, respectively.
- In Paper 34.4, Pohang University of Science and Technology presents a 40nm cryo-CMOS XY quantum control IC using a superposition of sinusoids to realize DRAG-compensated sinusoidally shaped RF pulses without the need for a separate envelope memory. The IC dissipates 8.2-to-19.1mW/qubit under active control at 3.5K.
- In Paper 34.5, UESTC and Chengdu Data Automation System Technologies Co., Ltd., show a calibration-free cryo-CMOS Class-$F_{2,3}$ VCO implemented in a 65nm bulk-CMOS process achieving a FoM as high as 202dBc/Hz at 4.2K.

Session 34 Highlights:
Cryo-CMOS for Quantum Computing

[34.1] THz Cryo-CMOS Backscatter Transceiver: A Contactless 4 Kelvin-300 Kelvin Data Interface

[34.2] A 28nm bulk-CMOS IC for full control of a superconducting quantum processor unit-cell

Paper 34.1 Authors: Jinchen Wang, Mohamed I. Ibrahim, Isaac B. Harris, Nathan M. Monroe, Muhammad Ibrahim Wasiq Khan, Xiang Yi, Dirk R. Englund, Ruonan Han

Paper 34.1 Affiliation: Massachusetts Institute of Technology, Cambridge, MA

Paper 34.2 Authors: Juhwan Yoo[1], Zijun Chen[1], Frank Arute[1], Shirin Montazeri[1], Marco Szalay[1], Catherine Erickson[1], Evan Jeffrey[1], Reza Fatemi[1], Marissa Giustina[1], Markus Ansmann[1], Erik Lucero[1], Julian Kelly[1], Joseph C Bardin[1,2]

Paper 34.2 Affiliation: [1]Google Quantum AI, Goleta, CA, [2]University of Massachusetts, Amherst, MA

Subcommittee Chair: Fabio Sebastiano, Delft University of Technology, Delft, The Netherlands

CONTEXT AND STATE OF THE ART

- Large-scale quantum computers have the theoretical capability to solve problems untractable using classical processors.
- Solid-state quantum processors require a cryogenic operation to limit perturbations on the quantum operations.
- Innovations in providing and controlling the signals necessary to the quantum processor inside the cryostat are required to scale up their size

TECHNICAL HIGHLIGHTS

- **The Massachusetts Institute of Technology presents a THz transceiver for a contactless interface of cryogenic quantum processors, reducing the heat load of the cryostat.**
    - A passive backscatter technique is implemented in 22nm FinFET process, reaching an energy efficiency of 34fJ/b (downlink) and 176fJ/b (uplink) at 260GHz when operated at 4-5 K.
- **Google introduces a 28nm CMOS quantum controller enabling the control of superconducting two-qubit gates.**
    - The cryo-CMOS chip integrates two microwave pulse generators for qubit rotation and the DACs for controlling two-qubit gates, maintaining a total power consumption of less than 4mW per qubit. The chip was extensively validated using superconducting quantum gates.

APPLICATIONS AND ECONOMIC IMPACT

- The computing power of quantum processors is quickly approaching the level required to solve practical problems, generating high expectations and economic excitement.
- Quantum computers are expected to revolutionize the development and discovery of drugs and materials, the optimization of complex systems, and cybersecurity

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:

- 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.
- 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.

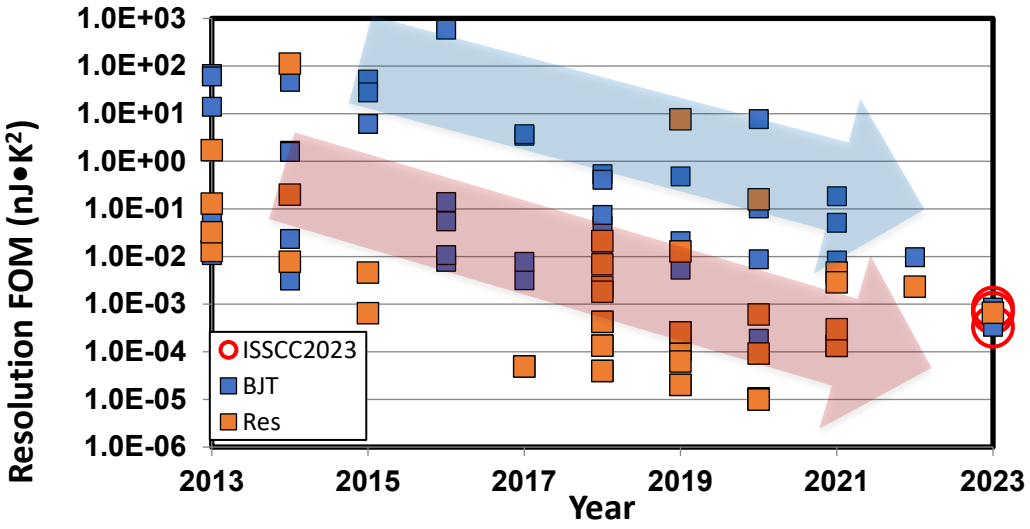

Figure 1: Trends in power efficiency of temperature sensors.

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.

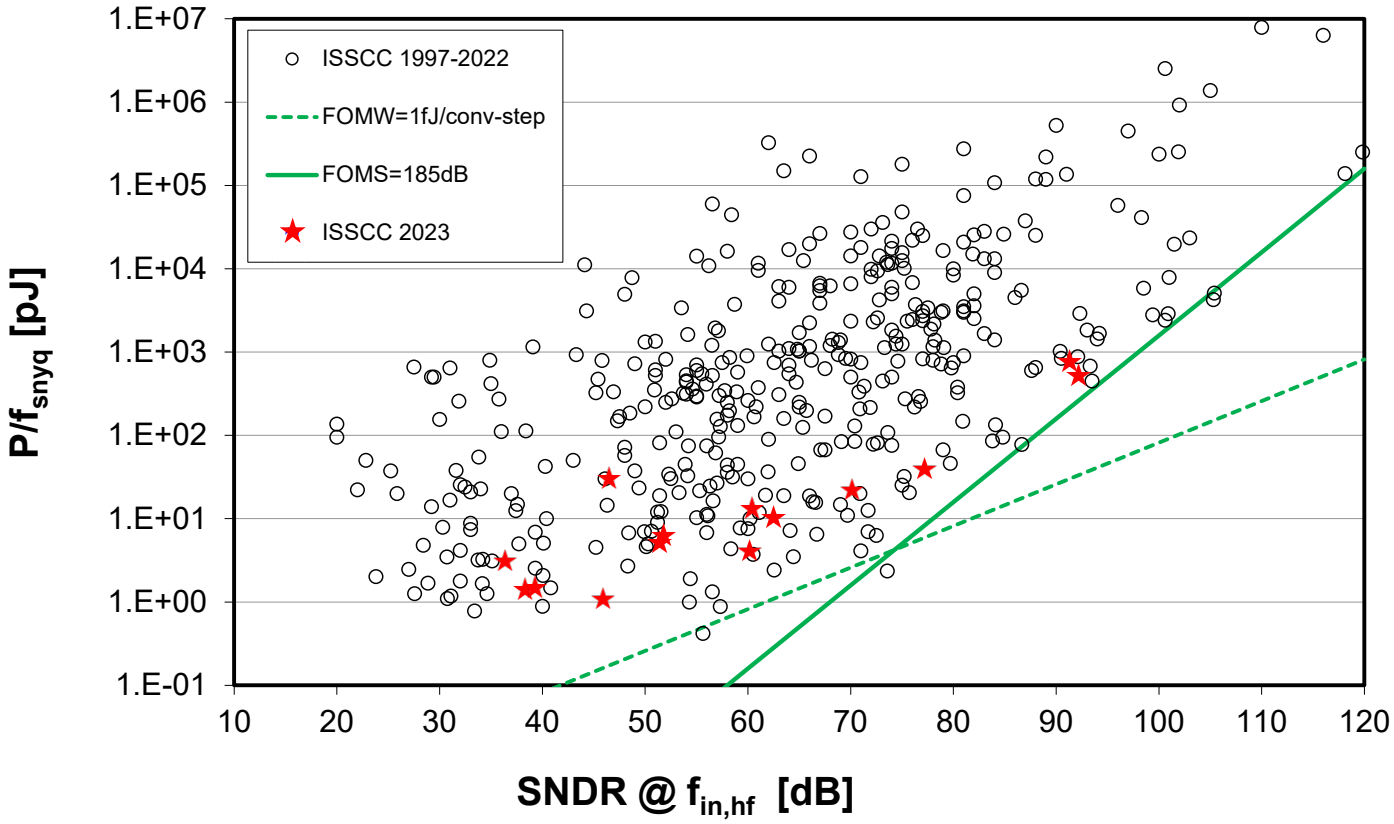

Figure 1: ADC power efficiency (P/f_{snyq}) as a function of SNDR.

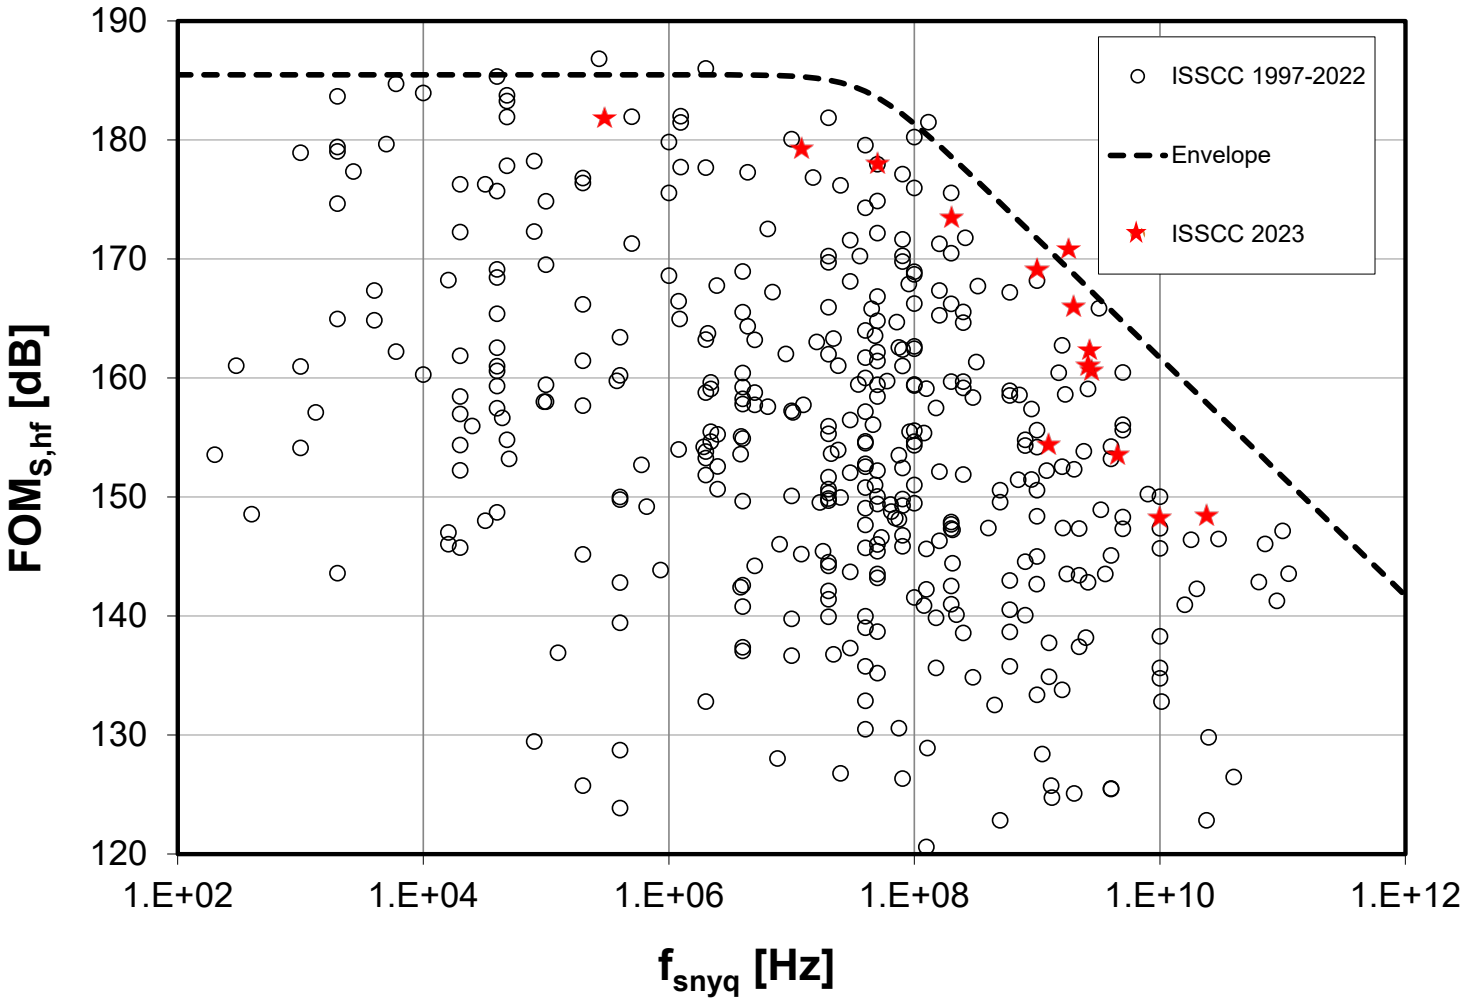

Figure 2: Power normalized noise and distortion vs. the Nyquist sampling rate.

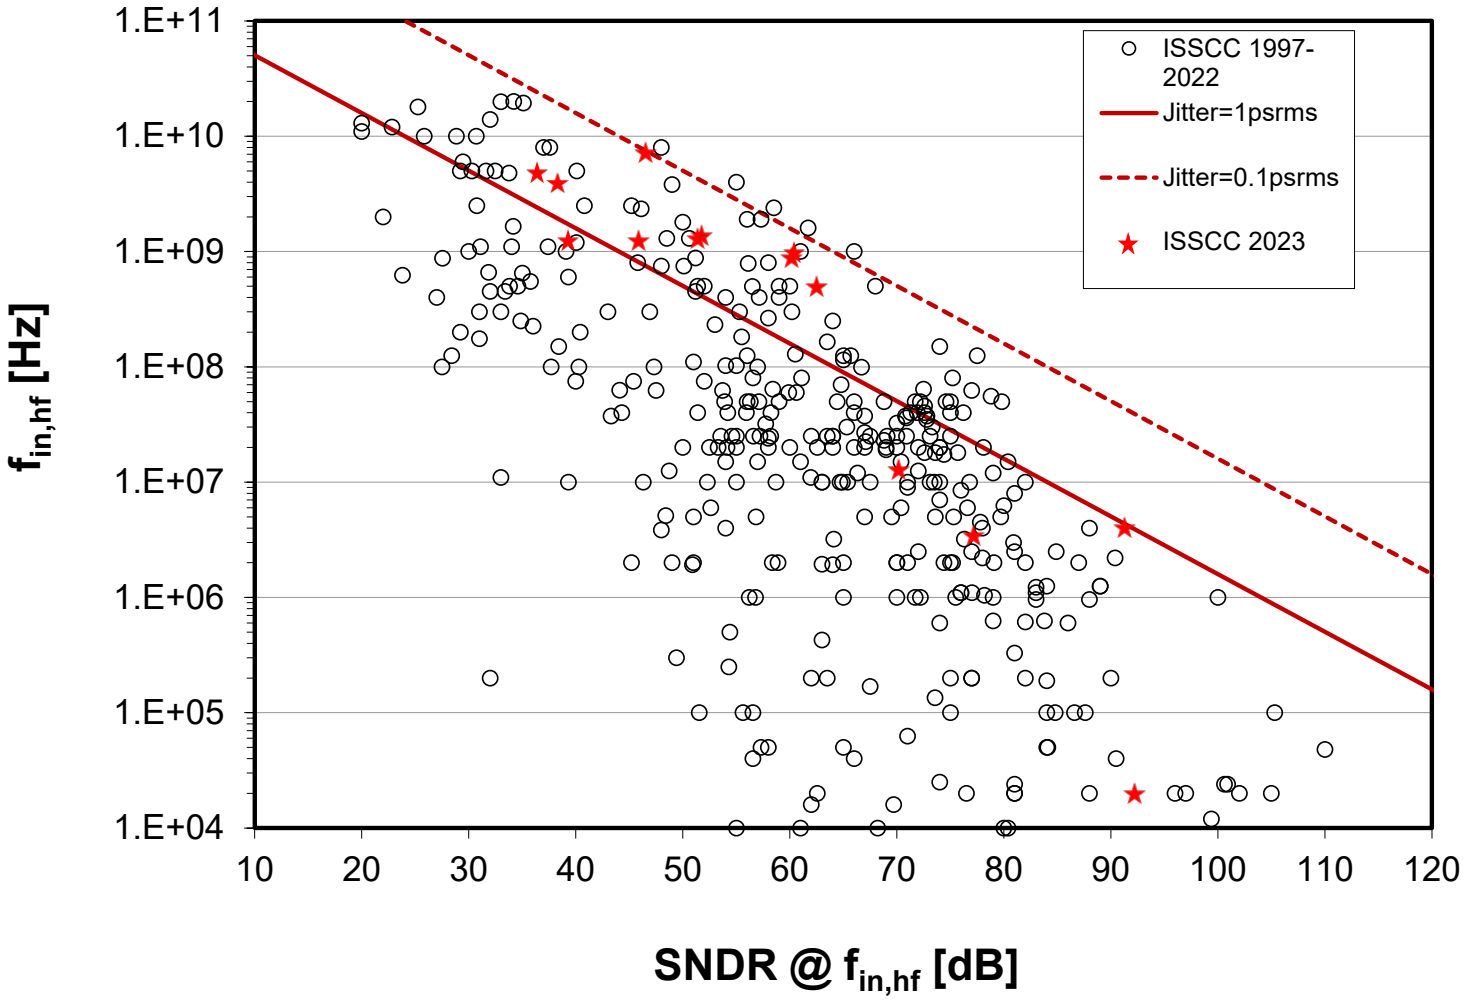

Figure 3: Bandwidth vs. SNDR.

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.

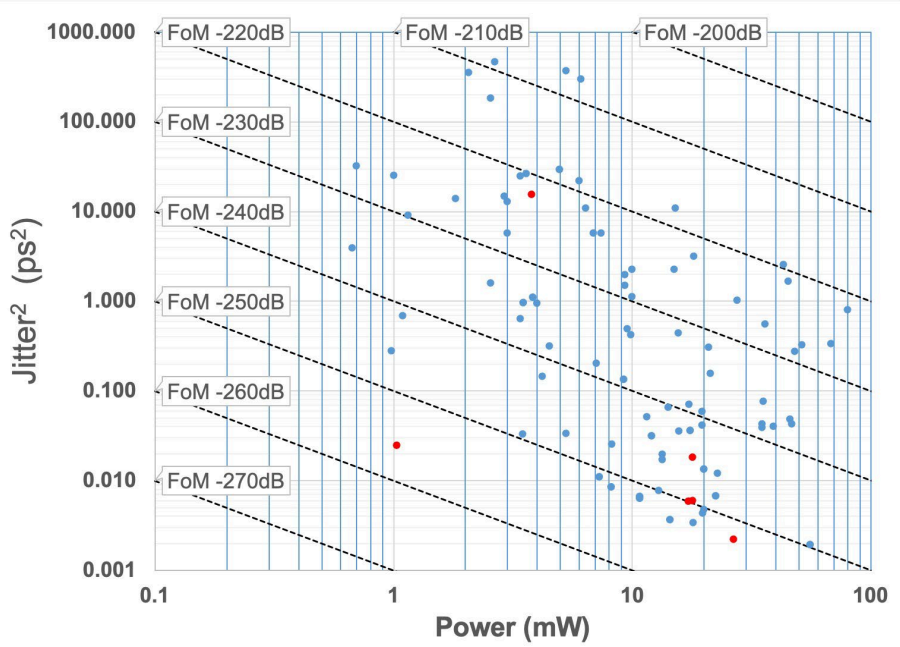

Fig. 1. PLL trends.

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.

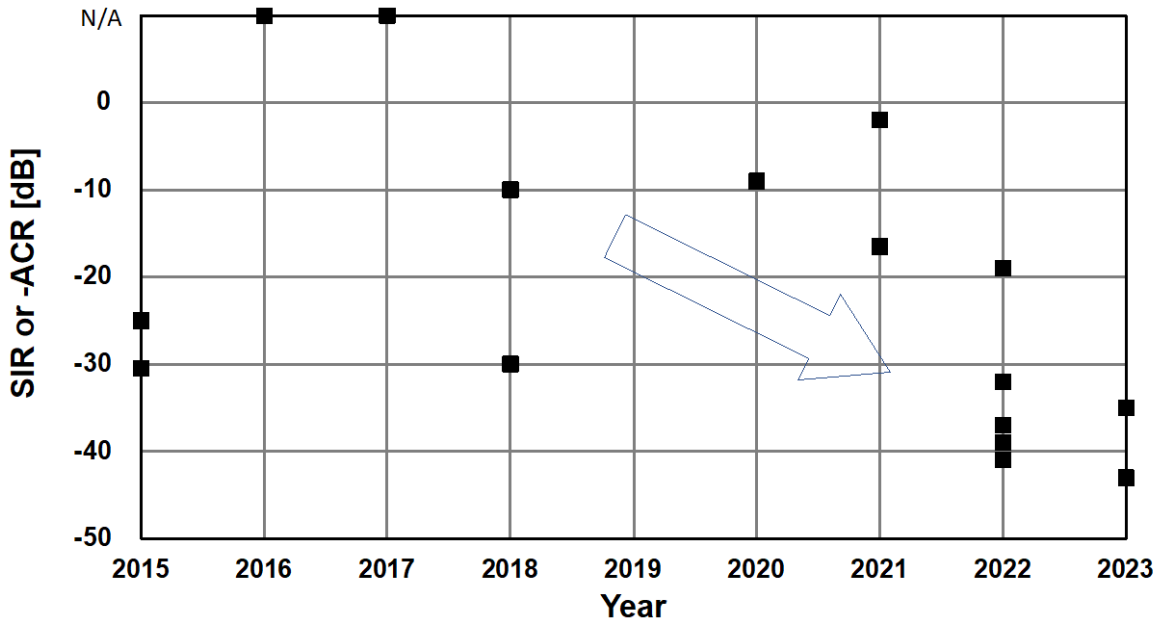

Figure 1: Plot of signal-to-interference ratio (SIR) or adjacent channel rejection (-ACR) vs. year as reported by ULP receivers published at ISSCC.

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.

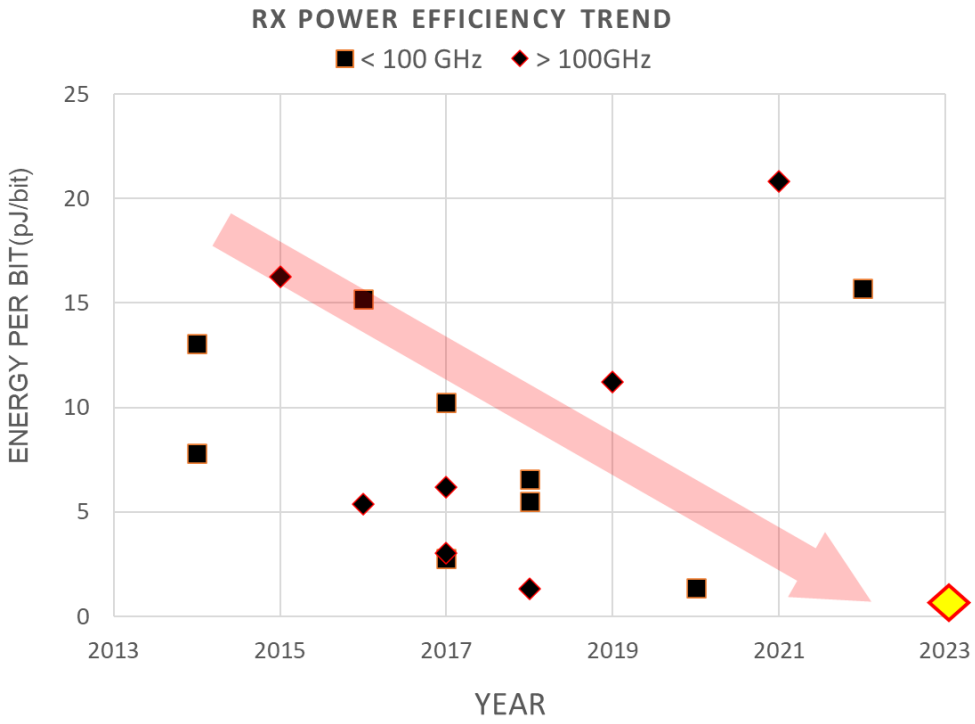

Figure 2: Receiver power efficiency trend of mm-wave and sub-THz high-speed wireless communications

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.

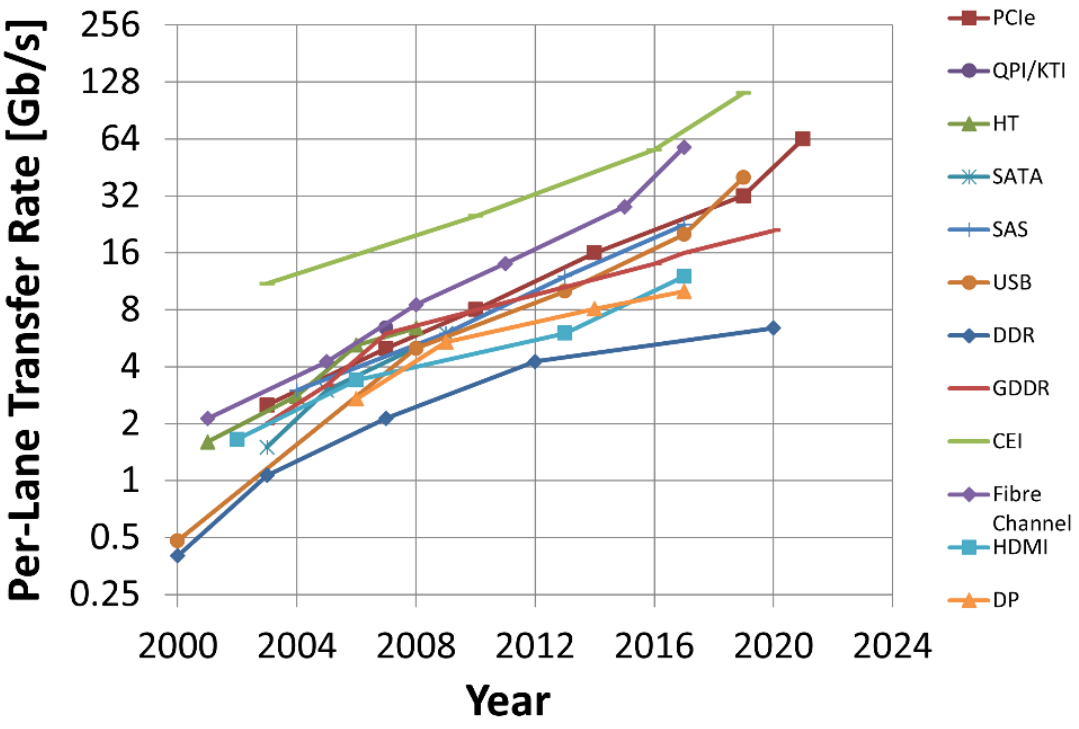

Figure 1: Per-lane data rate vs. Year for a variety of common I/O standards.

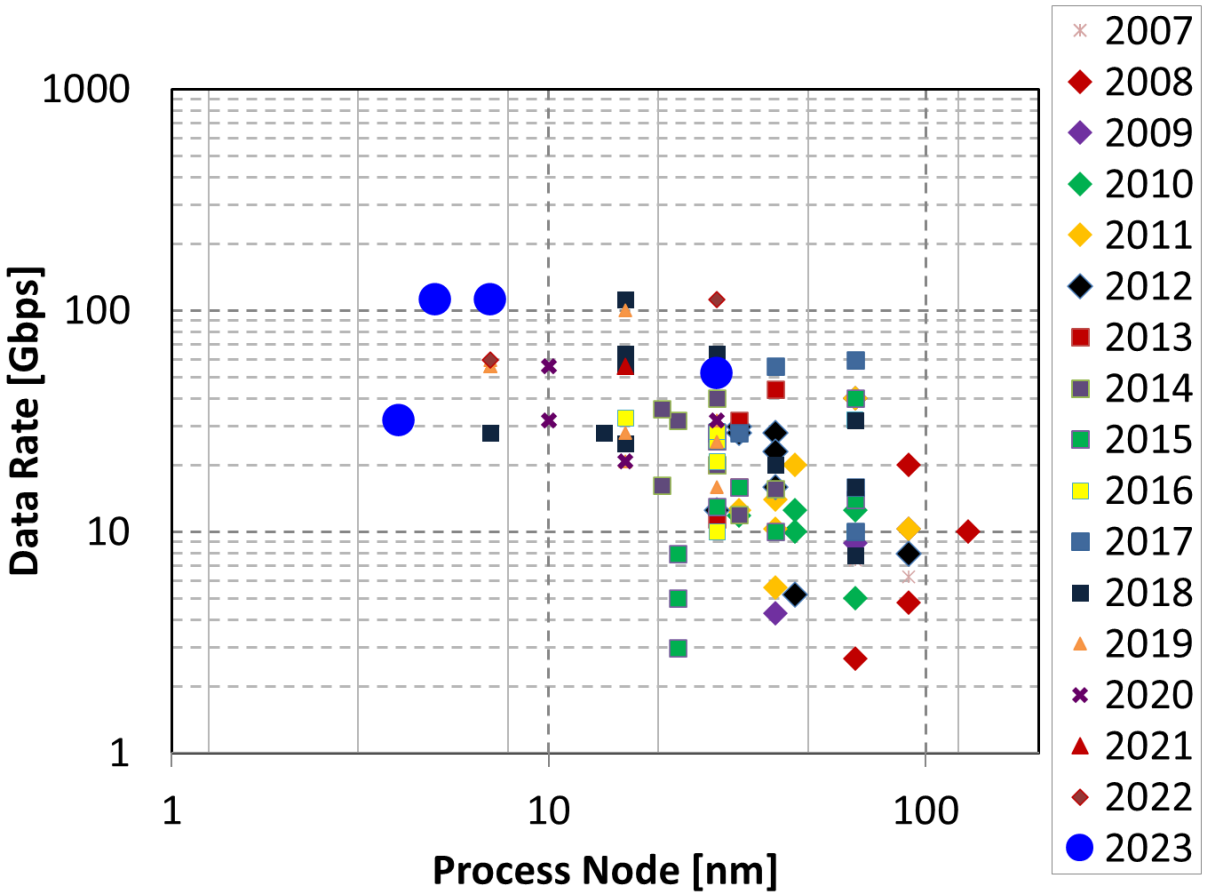

Figure 2: Data-rate vs. process node and year.

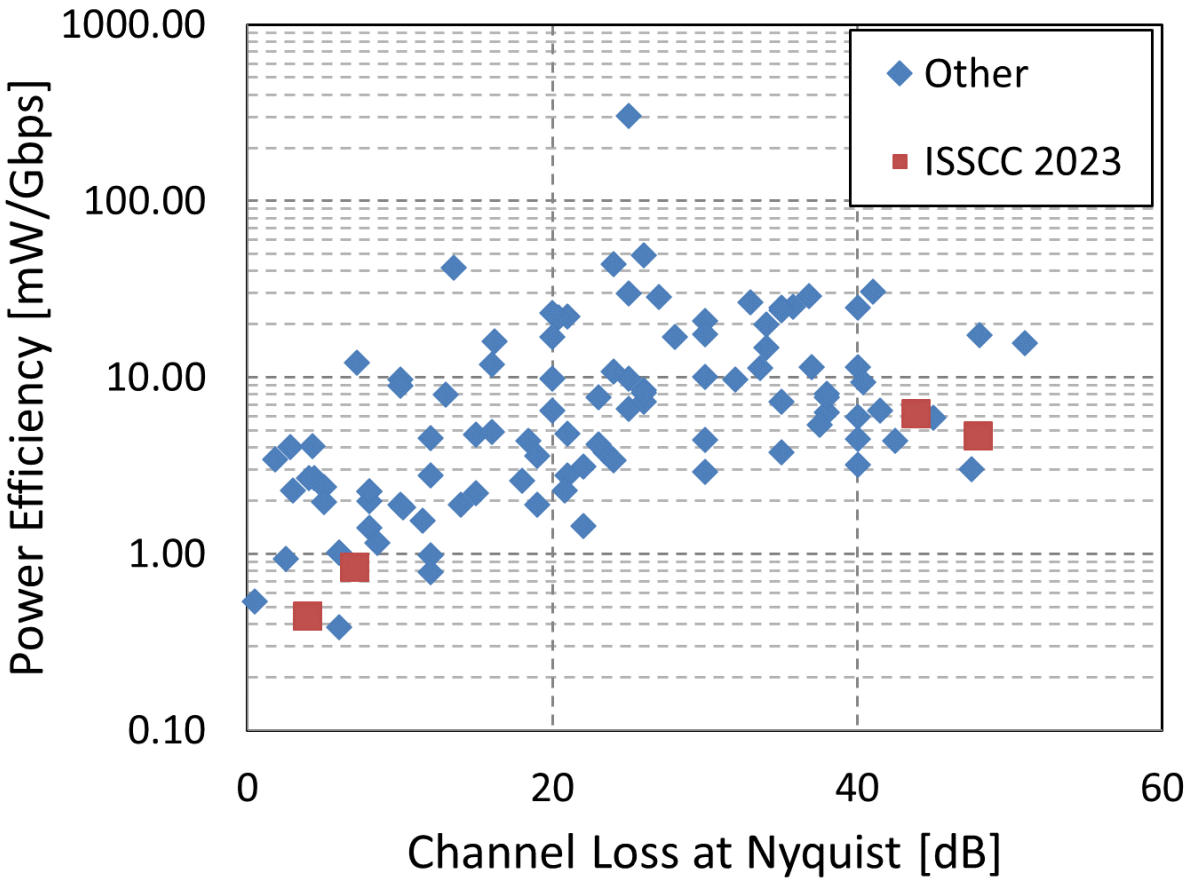

Figure 3: Power Efficiency vs. channel loss and year.

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.

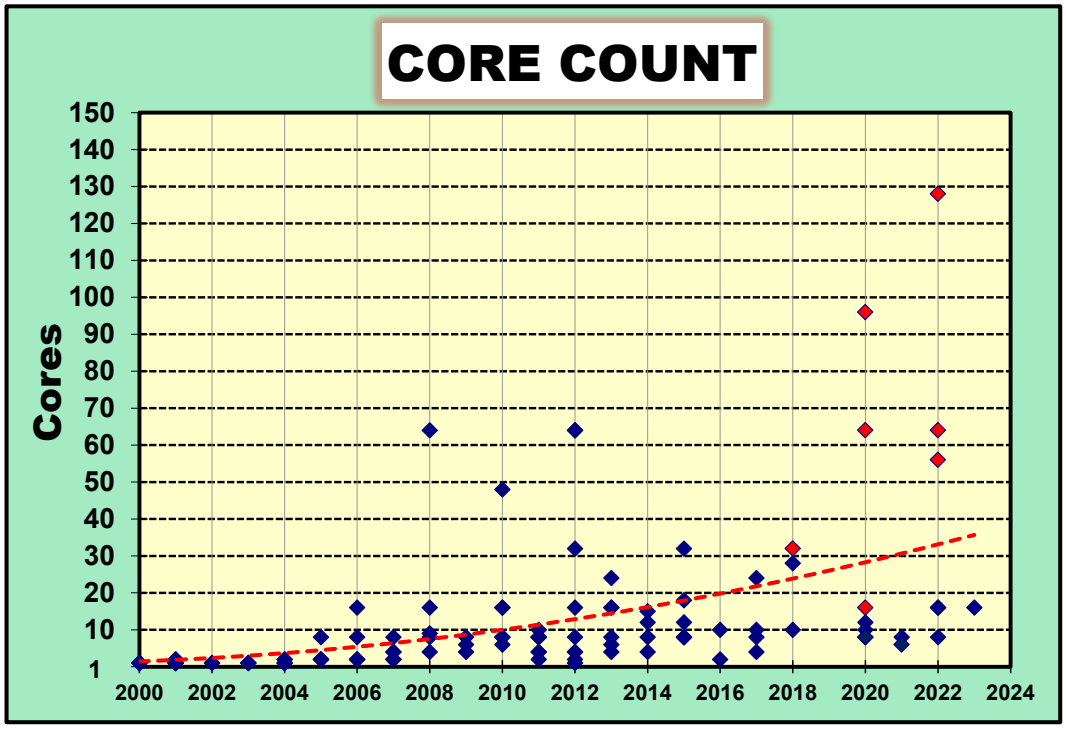

Figure 1: Core-count trends (red diamond designates multi-chip module).

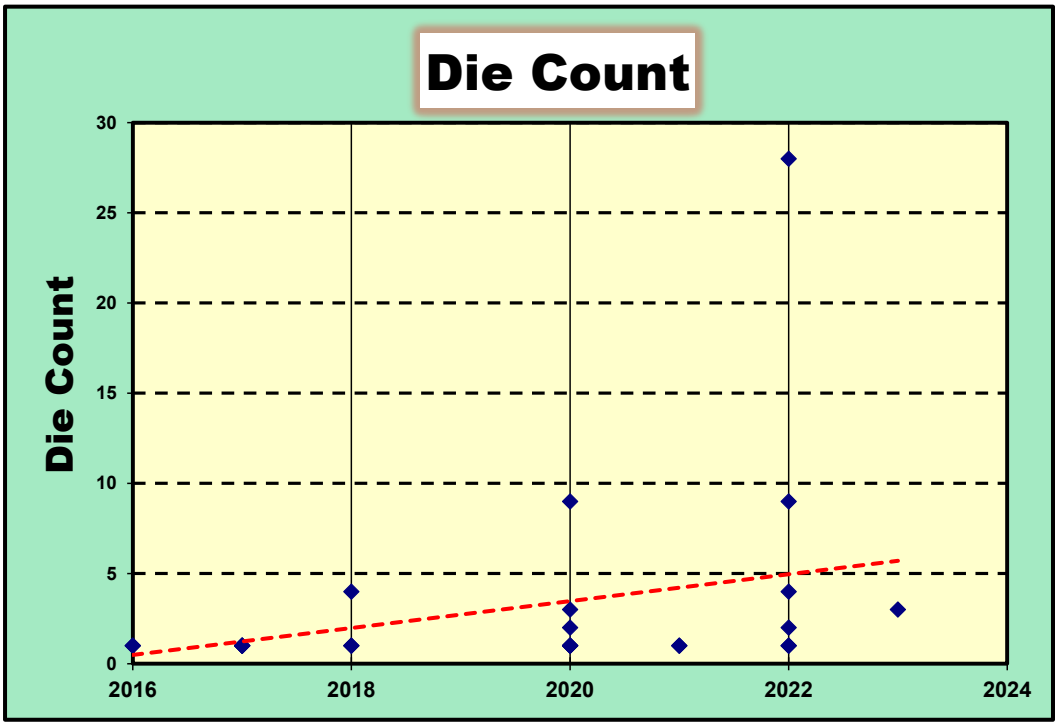

Figure 2: Die counts in a system trends.

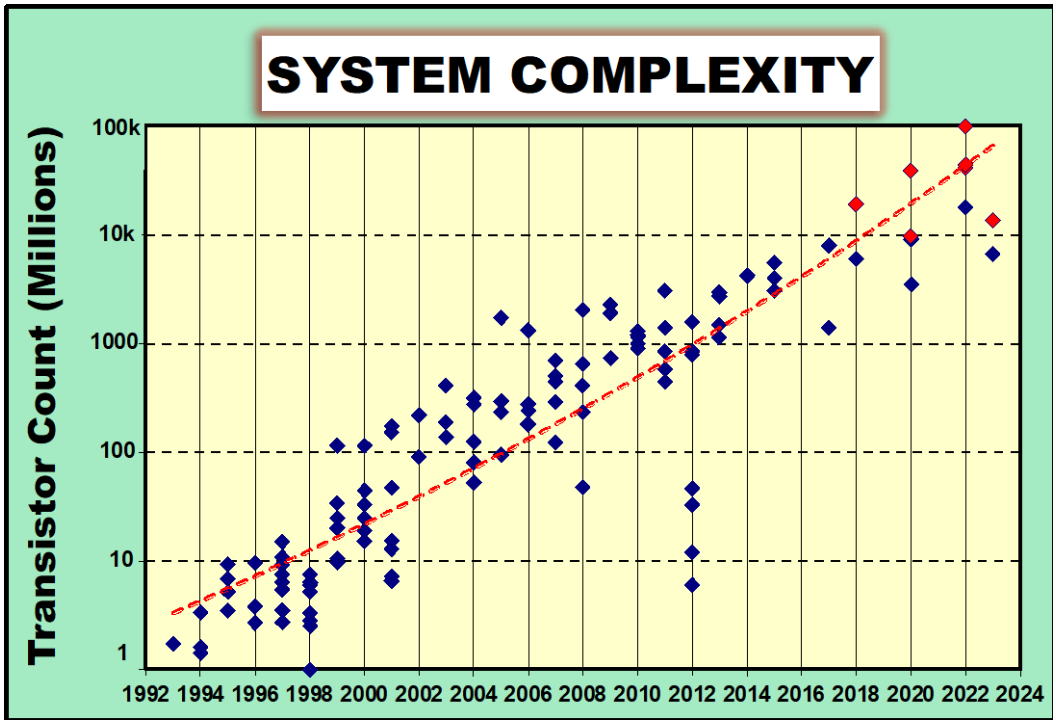

Figure 3: Chip-complexity scaling trends (red diamond designates multi-chip module).

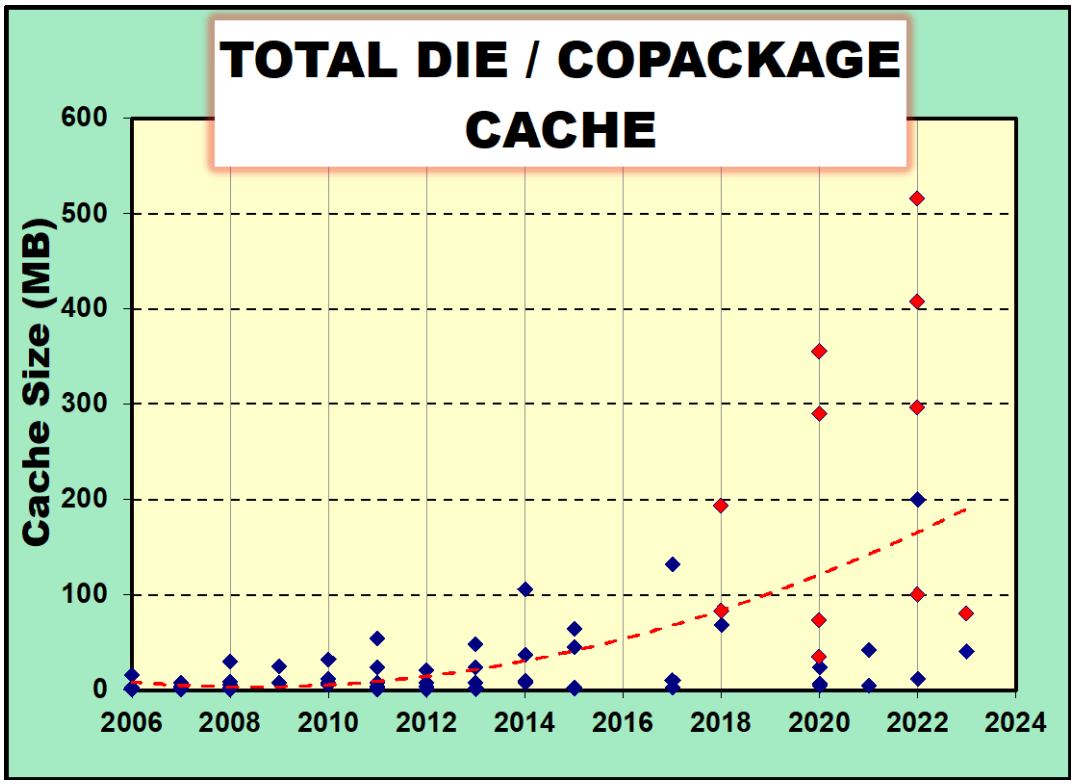

Figure 4: On-die cache-size trends (red diamond designates multi-chip module).

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.

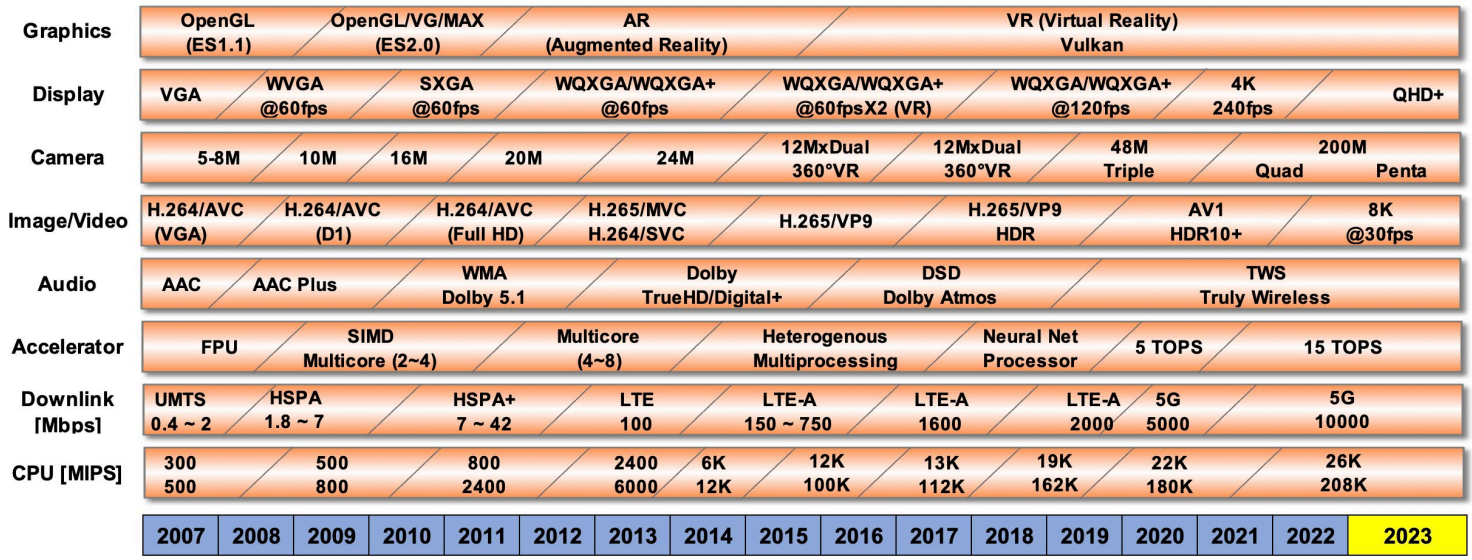

Figure 5: Application-processor trends in smartphones.

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.

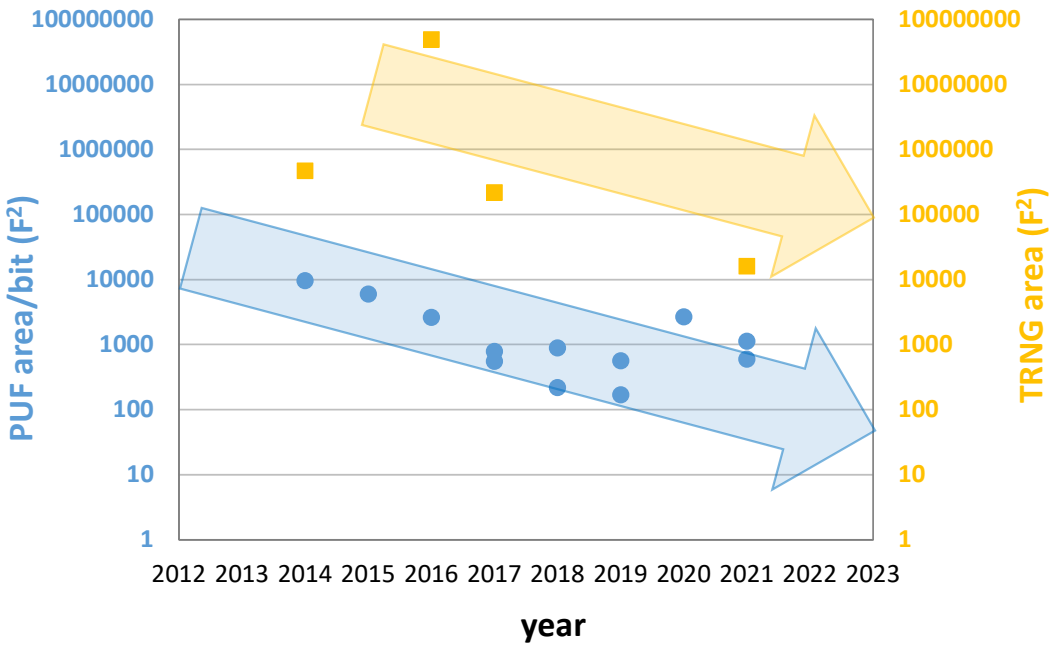

Figure 6. Area/bit trends for physically unclonable functions (PUFs) and area trends for true random number generators (TRNGs) published recently at ISSCC.

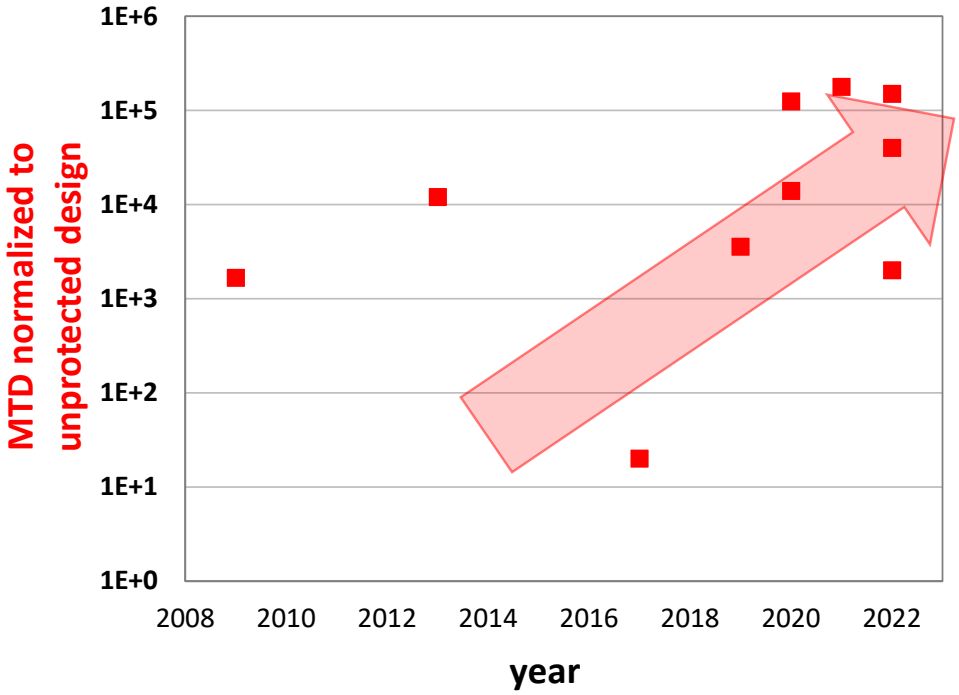

Figure 7: Improvement in measurements-to-disclosure (MTD) of cryptographic keys of side-channel counteraction techniques (normalized to unprotected design).

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.

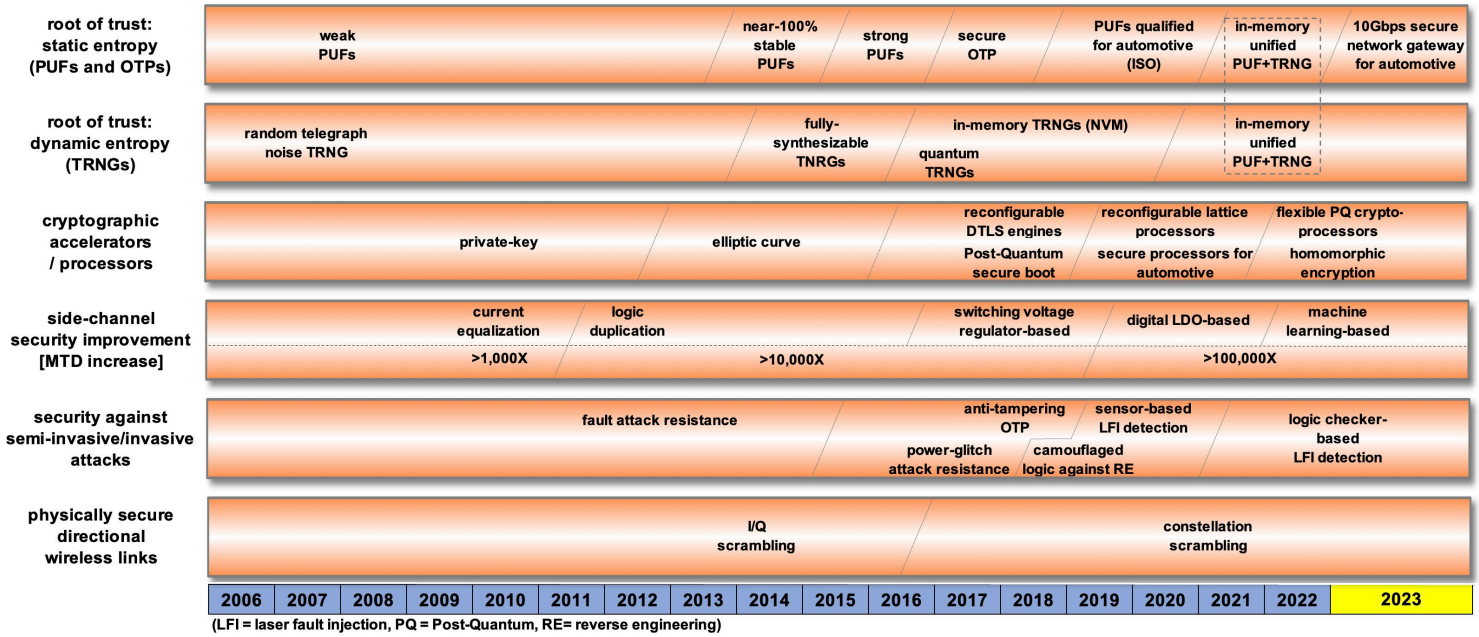

Figure 8: Security trends.

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.

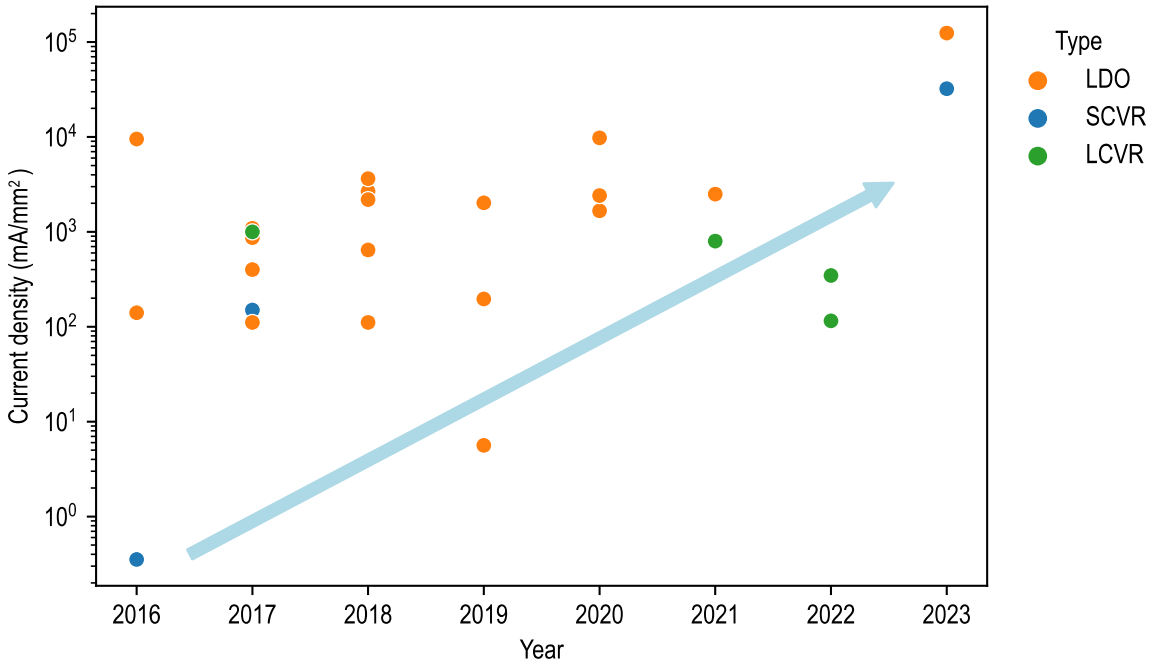

Figure 1. Integrated voltage regulator trends in current density with respect to active area.

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.

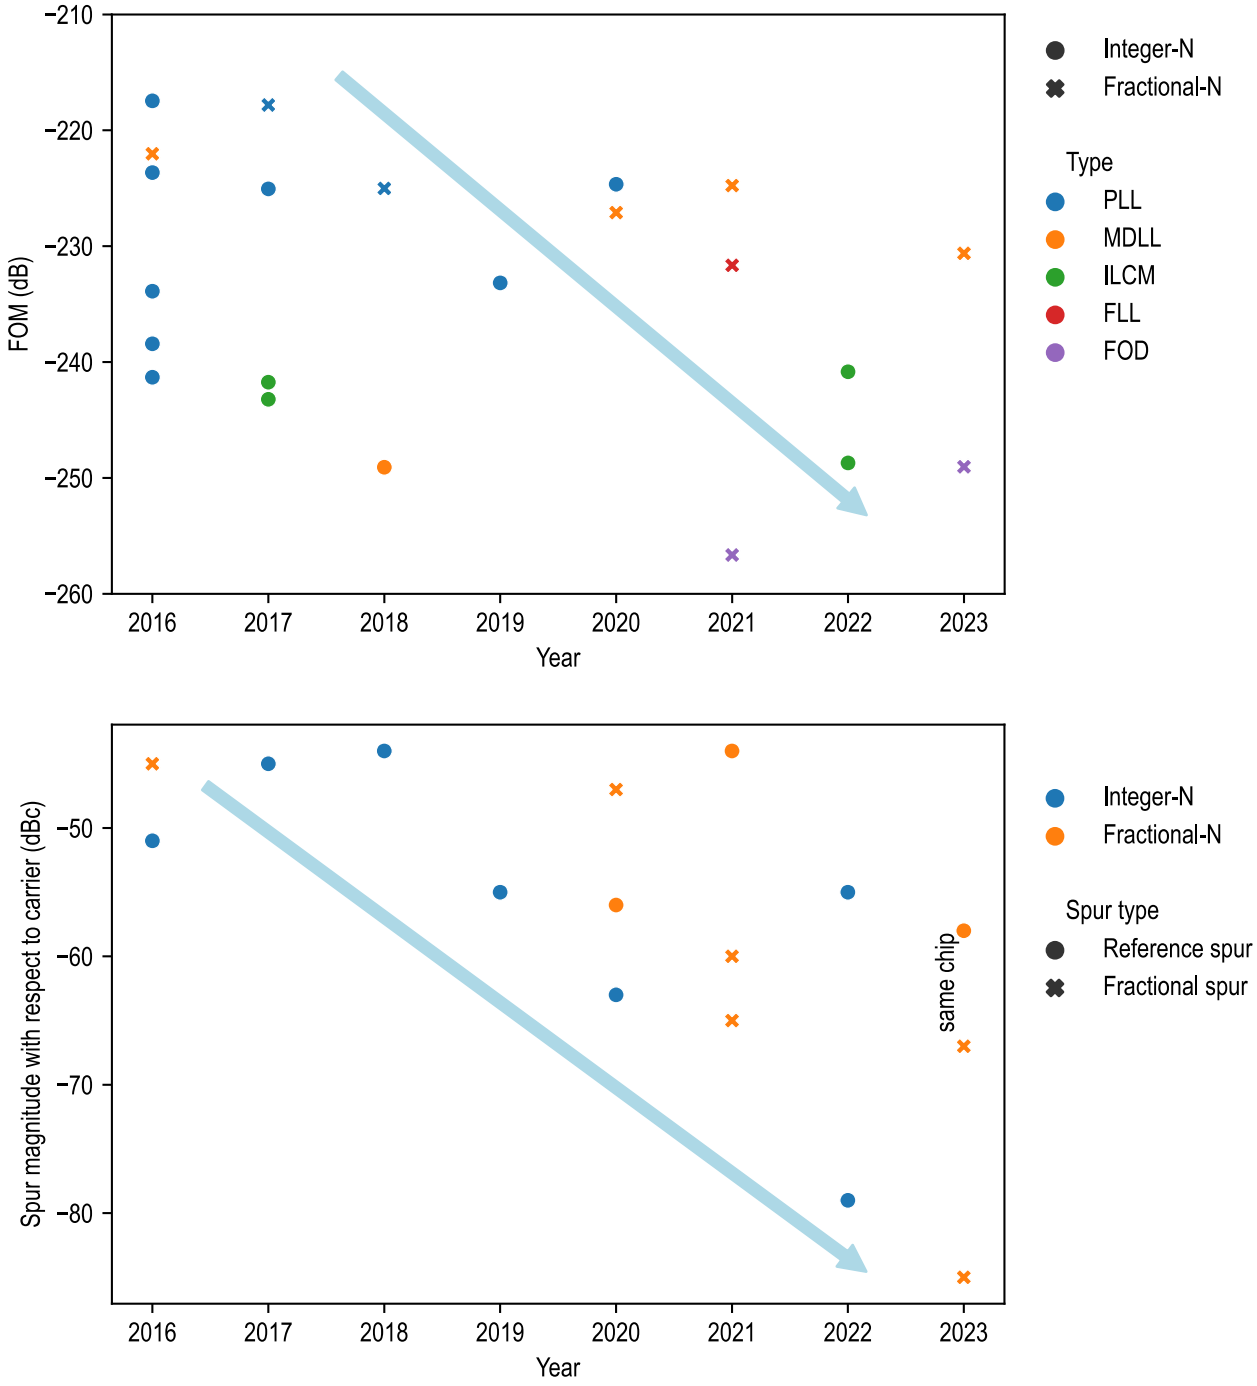

Figure 2. Top: digital clock generators key figure of merit (FoM) across recent years, defined as: $\text{FoM} = 10 \times \log_{10}\{ (\text{Jitter}_{\text{RMS}}/1\text{s})^2 \times (\text{Power}/1\text{mW}) \}$; bottom: spurious tone reduction trends.

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.

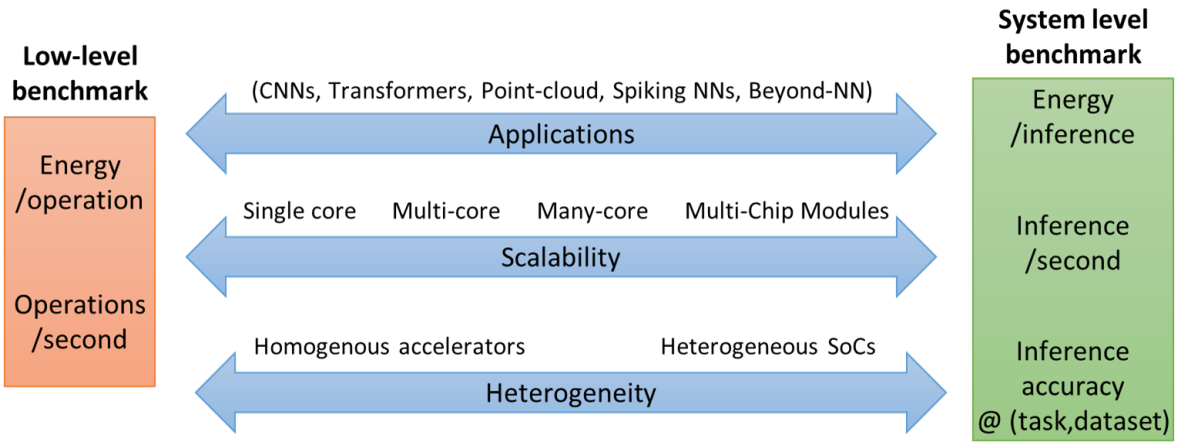

Figure 1: Various parameters impacting low-level and system-level benchmarking metrics.

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).

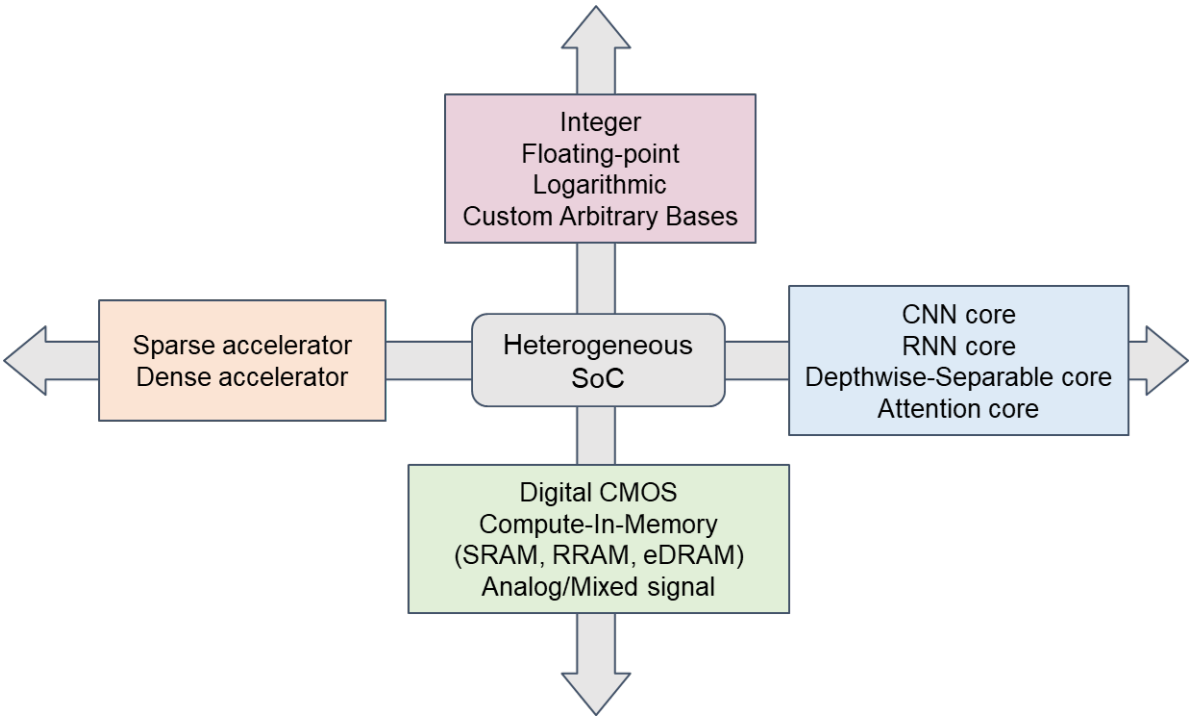

Figure 2: Heterogeneity in deep-learning accelerators.

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.

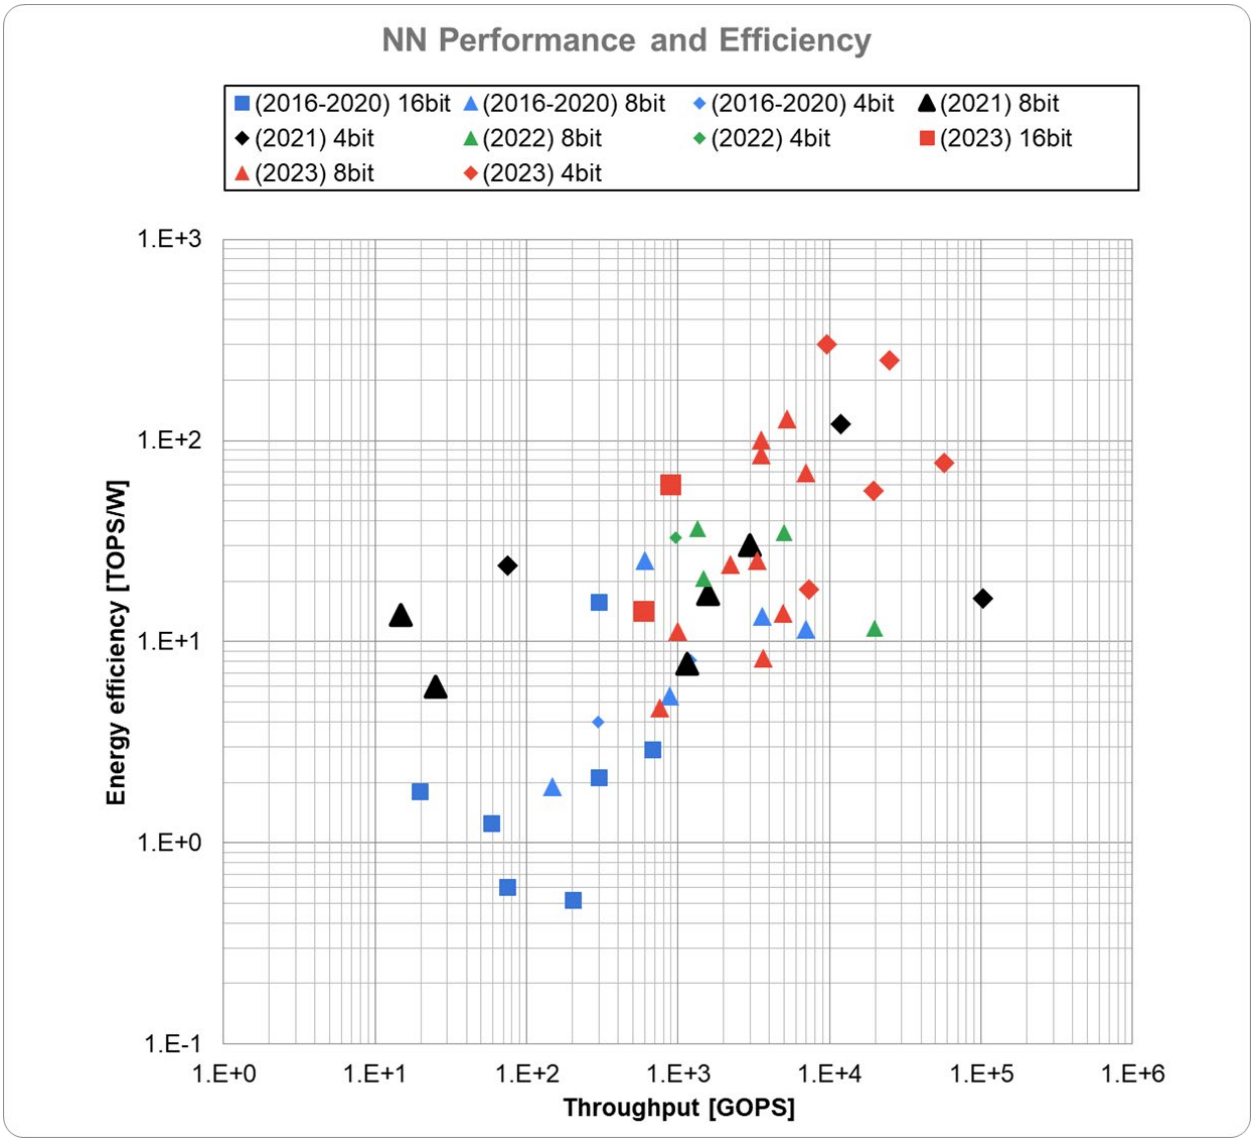

Figure 3: Deep-learning processor energy-efficiency (TOPS/W) and throughput (GOPS).

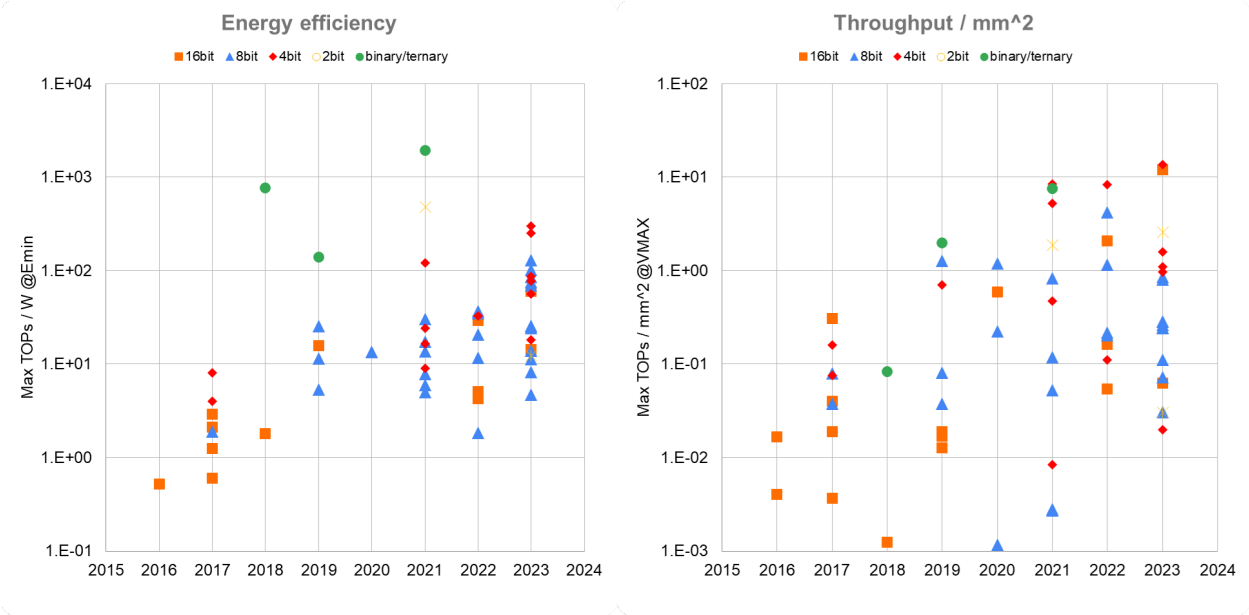

Figure 4: Evolution in energy efficiency (TOPS/W) and throughput per unit area (TOPS/mm^2) of ML inferencing processors.

[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.

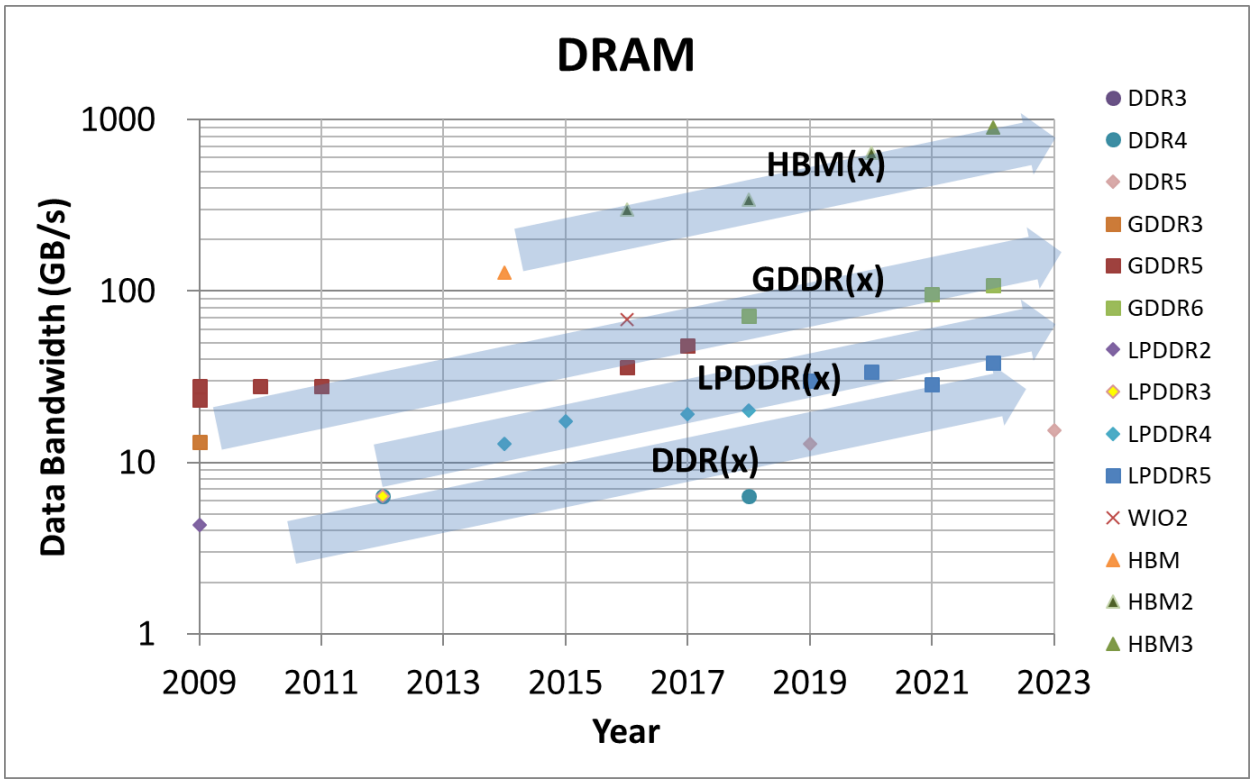

Figure 1 - DRAM data bandwidth growth

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.

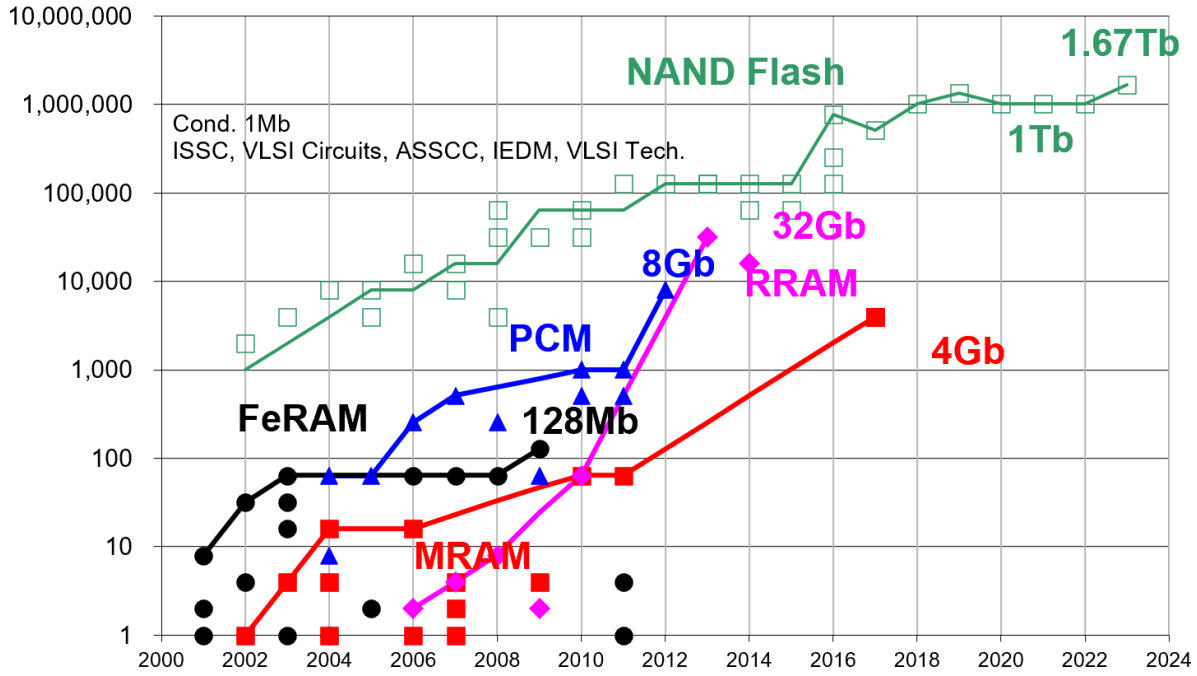

Figure 2 - Non-volatile memory capacity trend.

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).

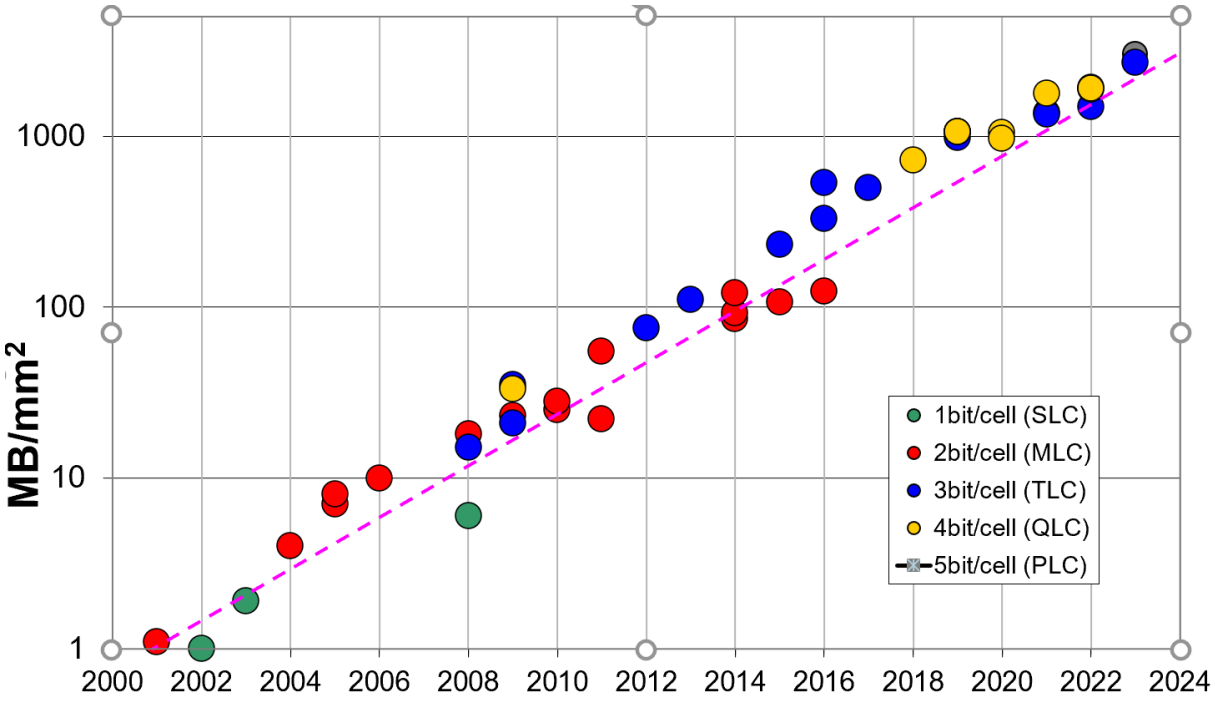

Figure 3 - NAND flash memory density trend.

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.

INDEX

INDEX

A paper number mentioned in this section follows the convention S.P, where S is the session number and P is the paper number. For example, 23.2 will be the second paper in the twenty-third session. You can refer back to the TECHNICAL SESSION OVERVIEWS in this Press Kit for additional details on any given paper. Some of the papers will also be available in the "not-so-technical" SESSION HIGHLIGHTS part of this Press Kit. All sessions and papers are in ascending order in both the Session Overviews and the Session Highlights sections of the Press Kit.

Technical Topics Mapped to Papers

Technical Topic	All papers in the following Sessions
Communication Systems includes Wireless, RF, and Wireline Subcommittees	**4, 6, 8, 12, 18, 19, 24, 25, 31**
Analog Systems includes Analog, Power Management and Data Converter Subcommittees	**3, 10, 11, 17, 20, 23, 30**
Digital Systems includes Memory, Digital Circuits, Machine Leaning and AI, Digital Architectures and Systems Subcommittees	**2, 7, 9, 14, 15, 16, 22, 28, 29, 33**
Innovative Topics includes Imagers/MEMS/Medical Devices/Displays and Technology Directions Subcommittees	**5, 13, 21, 26, 27, 32, 34**

Selected Presenting Companies/Institution Mapped to Papers

Chart 4.1

Affiliation	Paper#'s
Advanced Institute of Information Technology of Peking University	3.8, 7.8, 23.1
AMD	2.1, 12.1
Apple	21.3, 32.6
ARM	22.9
Asahi Kasei Microdevices	10.4
Autosilicon	21.6
Axelspace	19.3, 19.4
Beijing Academy of Quantum Information Sciences	34.3
Beijing Ningju Technology	31.3
BIRENTECH	16.2
Boston University	29.8
Broadcom	6.1
California Institute of Technology	13.5, 21.1, 27.2
Chengdu Data Automation System Technologies	34.5
Chinese Academy of Sciences	4.7, 5.5, 7.3, 22.2
Chip-GaN Power Semiconductor	11.4, 11.7, 20.1, 20.3, 20.4, 30.2, 30.8
Chungbuk National University	32.2
CISCO Systems	12.2

City University of Hong Kong	24.1
Columbia University	22.9, 25.5, 29.7
Cornell University	22.9
Daegu Gyeongbuk Institute of Science and Technology	32.2, 32.3, 32.4
Dartmouth College	30.4
Delft University of Technology	3.1, 3.2, 3.4, 19.1, 23.2, 23.3, 23.5, 23.7, 30.3
Delta Electronics	22.4
Dolphin Design	22.1
Eindhoven University of Technology	23.6
ETH Zürich	17.2, 21.3, 22.1
Everactive	15.2
Fudan University	16.2, 25.1, 25.2, 31.7, 32.5, 33.3
Fukuoka University	13.1
GeneASIC Technologies	2.4
Georgia Institute of Technology	21.3, 29.1, 29.5
Goodix Technology	3.1
Google Quantum AI	34.2
Graphcore	29.4
Gyeongsang National University	22.3
Hamburg University of Technology	23.4
Hanyang University	6.8, 32.1
Harvard University	22.9
Hefei CLT Microelectronics	30.6
Hefei University	3.7
Helmholtz-Zentrum Berlin für Materialien und Energie	21.4
Hokkaido University	2.3
Huawei Technologies	32.6
IBM Zurich Research Laboratory	17.2
imec	4.1, 21.2, 27.3, 32.5
imec - Holst Centre	21.2
Indian Institute of Technology Madras	3.3, 3.9
Industrial Technology Research Institute	7.1
Infineon Technologies	4.3, 4.5, 23.3, 25.4
Institute for Microelectronics Stuttgart	21.4
Institute of Microelectronics of the Chinese Academy of Sciences	16.3, 33.3
Instituto Superior Tecnico/University of Lisboa	3.6, 8.2, 8.4, 10.2, 10.3, 10.5, 10.7, 11.10, 11.2, 11.6, 11.9, 13.2, 17.1, 30.7, 31.6
Intel	9.4, 11.3, 14.3, 14.4, 15.5, 18.2, 28.1
Johannes Kepler University	25.4
Kobe University	13.3
Korea Advanced Institute of Science and Technology	2.7, 4.2, 4.6, 16.5, 21.6, 22.5, 24.2, 28.5
Korea Institute of Science and Technology	5.6
Korea University	6.6, 32.2
Korea University Guro Hospital	32.2
KU Leuven	32.5
KU Leuven ESAT-MICAS	18.4
Kun Gao Xinxin Technologies	17.8
Kwangwoon University	26.2
Kyoto University	2.3
Massachusetts Institute of Technology	19.2, 29.8, 34.1
Maynooth University	29.8
MediaTek	2.2, 6.2, 6.5, 7.6, 10.1, 17.3
Microaiot	31.5
MIRISE Technologies	23.2
MIT Lincoln Laboratory	15.1

Murata	11.1
Nanjing Prochip Electronic Technology	33.4
Nanjing University of Posts and Telecommunications	3.7
Nano Core Chip Electronic Technology	7.8
Nanyang Technological University	17.8, 34.3
National Cheng Kung University	17.6, 17.7
National Taiwan University	2.4, 2.5, 22.4
National Tsing Hua University	2.6, 7.1, 16.6, 22.6, 33.2
National University of Singapore	5.7, 32.6
National Yang Ming Chiao Tung University	2.4, 11.4, 11.7, 20.1, 20.3, 20.4, 30.2, 30.8
Northwestern University	27.1
now at Ain Shams University	12.3
now with Amazon	18.2
now with Apple	18.2
Now with MPS Spain	21.2
Nvidia	9.3
OmniVision Technologies	5.1
ONEsemiconductor	28.6
Oregon State University	10.4, 31.8
Osaka University	13.3, 13.4
Peking University	3.8, 6.3, 6.7, 7.3, 7.8, 10.6, 22.9, 23.1
Peng Cheng Laboratory	16.2
Pi2star Technology	22.2
Pohang University of Science and Technology	22.3, 34.4
Polar Bear Tech	15.4
Politecnico di Milano	4.3, 4.5
Pusan National University	31.2
Qualcomm	21.3
Rayence	5.6
Realtek Semiconductor	11.4, 11.7, 20.1, 20.3, 20.4, 30.2, 30.8
Renesas Electronics	15.3, 30.10
Samsung Electronics	5.4, 6.4, 26.1, 26.3, 28.3, 28.4, 28.7, 29.6, 30.5, 31.2, 32.1
Sanechips Technology	24.3
Semiconductor Energy Laboratory	13.1
Seoul National University	22.8, 28.6
Shanghai Jiao Tong University	29.7
Silicon Integrated	3.4
SiTime	23.7
SK hynix Semiconductor	28.2, 28.8
Socionext	17.4
Socionext Europe	17.4
Sony Advanced Visual Sensing	5.2, 5.3
Sony Semiconductor Solutions	5.2, 5.3
SOSO H&C	32.1
South China University of Technology	25.1, 25.2
Southeast University	7.2, 7.5, 33.4
Stanford University	18.1, 27.4
STMicroelectronics	16.7
SynSense	22.6
Syntiant	9.2
Taiwan Semiconductor Research Institute	2.4, 2.5, 22.4
Tesla	9.1
Texas A&M University	21.1
Texas Instruments	3.3, 3.9, 20.2

The N.1 Institute for Health	32.6
Tianjin University	5.5
Tokyo Institute of Technology	2.3, 4.4, 19.3, 19.4
Tsinghua Shenzhen International Graduate School	22.2
Tsinghua University	3.4, 7.3, 7.7, 8.1, 10.6, 14.2, 15.4, 16.1, 16.3, 16.4, 17.5, 22.2, 22.6, 23.5, 23.7, 24.3, 31.1, 31.3, 34.3
TSMC	7.4, 16.6, 33.1, 33.2
TSMC Corporate Research	16.6, 29.5, 33.2
TSMC Design Technology	29.5
Ulsan National Institute of Science and Technology	30.1, 30.5, 32.1
University of Texas	20.2, 30.9
University of Bologna	22.1
University of California	7.9, 11.1, 11.5, 21.5, 29.9
University of Delaware	13.6
University of Electronic Science and Technology of China	8.3, 18.3, 22.7, 25.3, 34.5
University of Freiburg - IMTEK	23.4
University of Illinois	3.5, 12.2, 12.3
University of Macau	3.6, 4.7, 5.5, 8.2, 8.4, 10.2, 10.3, 10.5, 10.7, 11.10, 11.2, 11.6, 11.9, 13.2, 17.1, 30.7, 31.6
University of Massachusetts	34.2
University of Michigan	29.3, 29.6
University of Science and Technology of China	30.6
University of Southern California	14.1
University of Stuttgart	21.4
University of Texas	29.2
University of Toronto	5.8, 32.7
University of Virginia	15.1
University of Waterloo	14.1
Vrije Universiteit Brussel	4.1
Washington University in St. Louis	25.5
Xi'an JiaoTong University	15.4
Xidian University	10.2, 10.7
Yonsei University	5.6
Zhejiang Lab	33.3
Zhejiang University	5.5, 11.8, 31.4, 31.5
Zhuhai UM Science & Technology Research Institute	11.10, 30.7
Zinitix	26.2

CONTACT INFORMATION

ANALOG

Subcommittee Chair: Maurits Ortmanns
University of Ulm
Work Phone: +49-731-50-26200
Email: maurits.ortmanns@uni-ulm.de

Press Designates: Ippei Akita
Chinwuba Ezekwe

DATA CONVERTERS

Subcommittee Chair: Jan Westra
Broadcom
Work Phone: +31-30-6084517
Email: jan.westra@broadcom.com

Press Designates: Yan Zhu
Nima Maghari

DIGITAL ARCHITECTURES & SYSTEMS

Subcommittee Chair: Thomas Burd
AMD
Work Phone: 408-749-2805
Email: tom.burd@amd.com

Press Designates: Hugh Mair
Sugako Otani

DIGITAL CIRCUITS

Subcommittee Chair: Keith Bowman
Qualcomm
Work Phone: 919-237-4750
Email: kbowman@qti.qualcomm.com

Press Designates: Eric Fang
Yvain Thonnart

IMAGERS, MEMS, MEDICAL & DISPLAYS

Subcommittee Chair: Rikky Muller
University of California, Berkeley
Work Phone: +1 617-519-8508
Email: rikky@berkeley.edu

Press Designates: Jerald Yoo
Johan Vanderhaegen

MACHINE LEARNING & AI

Subcommittee Chair: SukHwan Lim
Samsung Electronics
Work Phone: +1 408-630-9740
Email: sh19.lim@samsung

Press Designates: Yongpan Liu
Rangharajan Venkatesan

MEMORY

Subcommittee Chair: Meng-Fan Chang
National Tsing Hua University
Work Phone: +886-3-516-2181
Email: mfchang@ee.nthu.edu.tw

Press Designates: Violante Moschiano
Hye-Ran Kim

POWER MANAGEMENT

Subcommittee Chair: Bernhard Wicht
University of Hannover
Work Phone: +49-511-762-19690
Email: bernhard.wicht@ims.uni-hannover.de

Press Designates: Li Geng
Jason Stauth

RF

Subcommittee Chair: Jan Craninckx
imec
Work Phone: +32-16-28-87-56
Email: jan.craninckx@imec.be

Press Designates: Jeremy Dunworth
Wei Deng

TECHNOLOGY DIRECTIONS

Subcommittee Chair: Ali Hajimiri
Caltech
Work Phone: +1-626-487-8463
Email: hajimiri@caltech.edu

Press Designates: Denis Daly
Shawn Hsu

WIRELESS

Subcommittee Chair: Chih-Ming Hung
MediaTek
Work Phone: +1 512-887-2080
Email: hungcm@hotmail.com

Press Designates: David Wentzloff
Byung-Wook Min

WIRELINE

Subcommittee Chair: Yohan Frans
AMD
Work Phone: 1 408-879-2797
Email: yohanf@amd.com

Press Designates: Mike Shuo-Wei Chen
Kenny Shieh

Program Chair, ISSCC 2023
Piet Wambacq
imec
Work Phone: +32-16-281-218
Email: wambacq@imec.be

Program Vice-Chair, ISSCC 2023
Frank O'Mahony
Phone: +1 503-784-0549
Email: frank.omahony@intel.com

Press Coordinator
Shahriar Mirabbasi
University of British Columbia
Email: shahriar@ece.ubc.ca

Press-Relations Liaison
Kenneth C. Smith
University of Toronto
Work Phone: 416-418-3034
Email: lcfujino@aol.com

Editor-in-Chief: Shahriar Mirabbasi

Editor-at-Large: Kenneth C (KC) Smith

Publisher: Laura Chizuko Fujino