Guest Editorial Introduction to the Special Section on the 2023 IEEE International Solid-State Circuits Conference (ISSCC)

THIS special section of the IEEE JOURNAL OF SOLID-STATE CIRCUITS (JSSC) highlights outstanding papers presented at the 2023 IEEE International Solid-State Circuits Conference (ISSCC), which was held from February 19 to 23, 2023 in San Francisco, USA, under the conference theme "Building on 70 years of Innovation in Solid-State Circuit Design." ISSCC is the foremost global forum for the presentation of advances in solid-state circuits and systems-on-a-chip (SoCs) and offers a unique opportunity for engineers working at the cutting edge of integrated circuit (IC) design and application. The conference includes several technical programs ranging from analog, digital, memory, wireline (WLN)/wireless, and power management circuits and systems with applications in various fields. This JSSC special section highlights selected papers from ISSCC, specifically on topics related to WLN circuits, digital circuit techniques (DCTs), digital architecture and systems (DASs), machine learning (ML) accelerators, and memory circuits.

In 2023, ISSCC had a ∼30% acceptance rate, and more than 2000 people attended the conference in person. This special section of JSSC features 25 selected papers from ISSCC 2023, of which two papers are from the WLN, three from DCTs, two from DASs, six from ML, and six from the memory subcommittees. These papers provide comprehensive materials extending their conference proceedings and were carefully selected from highly qualified accepted papers. The acceptance of these papers to JSSC was based on a thorough peer review process. The accepted papers cover a wide range of topics, as described below.

In the area of WLN circuits, the first paper by Zhang et al. from Broadcom presents a serial link transceiver designed for a wide range rate up to 112 Gb/s. The RX has a linear equalizer, three-tap FFE, and 18-tap DFE. The TX uses a 7-bit DAC-based driver with six-tap FFE. The TX shows an eye with an RLM of 0.999. The transceiver can compensate for up to 43.9-dB insertion loss channel at 112 Gb/s. It consumes 690 mW per RX/TX. It is fabricated in 7-nm FinFET technology and occupies 0.63 mm^2 per RX/TX. The second paper by Ye et al. from Peking University presents a 200-Gb/s five-tap delay-line-based receiver FFE in 28-nm 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 with distributed amplifiers to alleviate reflection. *RC* source degeneration in the variable transconductance cell provides low-frequency equalization. The RX FFE achieves a power efficiency of 0.43 pJ/b and compensates for a 17.2-dB loss channel.

The DCTs committee selected three outstanding articles from the papers presented at ISSCC 2023. The first paper by Seol et al. from the University of Michigan reports a fully integrated keyword spotting system implemented in 28 nm that consumes 1.5 μW and achieves 92.8% accuracy on a standardized dataset. The proposed architecture employs a skip recurrent neural network (skip-RNN) to control power dissipation in both the analog front end and the digital back end for overall system energy reduction. The RNN analyzes signal content to compute future audio frames to skip, during which time the analog front end, feature extractor, and the neural network classifier are all powered down. The second paper by Zhang et al. from The University of Southern California presents a fractional-N MDLL that employs injection-error scrambling and background DTC nonlinearity calibration to achieve 800-fs rms jitter and −67-dBc fractional spur at 1.5-GHz output frequency in 65-nm CMOS. The proposed MDLL architecture utilizes two cascaded DTCs to completely disrupt any periodicity in injection error. A third-order calibration technique for DTC offset, gain, and nonlinearity is proposed to minimize noise floor elevation attributable to the proposed injection-error scrambling while preserving low jitter. Finally, the third paper from the Georgia Institute of Technology and TSMC reports a heterogeneous RRAM CNN and SRAM SNN SoC, enabling 73.53 TOPS/W to support a 100-outputs/s event camera for hybrid frame and event-based target tracking. The hybrid architecture synergizes the advantages offered by both CNN and SNN architectures to realize fused frame and event vision. The application-level accuracy degradation due to RRAM variations in CIM applications is restricted by a triple-error-correction scheme.

Two papers were selected in the area of DASs. The first paper (Han et al.) presents a mobile neural 3-D rendering processor for metaverse applications, mitigating the complexity

Digital Object Identifier 10.1109/JSSC.2023.3328667

of the compute-intensive neural radiance fields rendering. The brain-inspired MetaVRain architecture accelerates spatial attention, temporal familiarity, and top-down attention. Removal of inessential computations and acceleration through 1-D and 2-D neural engines supports sparsity exploitation and data reuse. Positional encoding is carried out via periodic polynomial sinusoidal function approximation for power and area reduction. A 28-nm test chip validated with multiple datasets achieves a peak 118-frames/s frame rate at 99.95% lower power than recent GPUs. The second paper by Kumar et al. from Intel introduces a source-agnostic fault injection attack-resistant AES accelerator based on run-time fault detection through arithmetic checkers in the nonlinear portion of AES and parity-based checkers in the linear one. Composite-field inverse checker, redundant affine parity circuits, and byte-interleaved register placement enable 99.1% error coverage against raster and box-scan laser attacks. Test-chip measurements show 111× and 10 000× improvement in the minimum traces to disclosure against laser and undervoltage attacks.

In the area of ML, six papers were selected. The paper by Tu et al. from Tsinghua University presents digital compute-in-memory (CIM)-based multimodal transformer accelerator, which consumes 2.24 μJ/token at 0.7 V and 160-MHz frequency for the ViLBERT-based model with INT8/INT16 mixed precision. The paper by Kim et al. from the Korea Advanced Institute of Science and Technology (KAIST) reports DynaPlasia, an eDRAM CIM-based ML accelerator with a novel triple-mode 3T2C cell, and a dynamic reconfigurable core architecture. DynaPlasia achieves an energy efficiency of 37.2 TOPS/W and a compute density of 2.03 TOPS/mm^2 at 1.0 V and 250 MHz for INT4/INT5 activation/weight precision. The paper by Hung et al. from National Tsing Hua University reports a 22-nm nonvolatile AI processor that features a hybrid computing architecture (in-/near-memory computing) and ReRAM devices of mixed precision (SLC/MLC memory states) to balance NVM capacity, energy, and accuracy. For MobileNet-V2 with INT8 precision, 51.4/25.1 TOPS/W was achieved for CIFAR-10/ImageNet inference task. The paper by Conti et al. from the University of Bologna and ETH Zürich presents a 22-nm heterogeneous SoC with 16 RISC-V cores and precision-scalable DNN accelerators, achieving an energy efficiency of 12.4/1.64 TOPS/W at INT2/INT8 precision values. The paper by Moon et al. from the Pohang University of Science and Technology presents 1-to-8-b scalable-precision ML accelerator in 28-nm CMOS, which features flexible quantization compression support to handle different precisions and data formats. The prototype chip achieves up to 95.9/127.8 TOPS/W with 30%/90% sparsity. The paper by Kim et al. from KAIST presents a complementary deep neural network (C-DNN) processor, a hybrid CNN and SNN accelerator, which allocates training/inference tasks to CNN or SNN core in inter-/intralayer manner. C-DNN processor achieves 83.4/17.2-TOPS/W energy efficiency for CIFAR-10/ImageNet training task.

In the memory (MEM) area, six papers were selected. The paper by Bae et al. from the University of California presents the first continuous-time latch-based Ising Machine with 1440 spins. The proposed Ising machine enables fully parallel continuous-time operations and achieves a 1051× faster time to solution (<20 ns) than prior discrete-time solutions, while consuming only 0.2–3 nJ. The paper by Jin et al. from Samsung Electronics introduces a 4-nm 16-Gb/s/pin 0.764-pJ/b single-ended PAM4 parallel transceiver with switching-jitter compensation and transmitter optimization. A relaxed TX termination of 20 Ω and an RX termination of 50 Ω are adopted to maximize the eye opening. The paper by Wu et al. from National Tsing Hua University describes the first true floating-point hybrid-domain (analog and digital) SRAM CIM macro. A 22-nm 832-kb macro achieves 72.14 TFLOPS/W at BF16-IN and BF16-W with FP32-OUT and up to 128 accumulations. The paper by Wu et al. from Fudan University presents a 9-Mb HZO-based embedded FeRAM macro with a 10^{12}-cycle endurance and 5-ns read and 7-ns write access time. The 700-nm-diameter TiN/HZO/TIN ferroelectric capacitors are integrated in 130-nm CMOS process and showed ten-year data retention at 85 °C. The paper by You et al. from National Tsing Hua University introduces a 22-nm 8-Mb STT-MRAM near-memory-computing macro. It achieves a 436-GB/s read bandwidth and 53.6–190.2-TOPS/W energy efficiency when performing 8-b input and 8-b weight operation with 26-b output and 576 accumulations. The paper by Chae et al. from Samsung Electronics describes a 4-nm FinFET third-generation high-bandwidth memory (HBM3) interface with 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 1.15 TB/s.

We highly appreciate the time and efforts of all the authors and reviewers, which ensured the production of high-quality manuscripts in this special issue. We would also like to express our gratitude to the members of the WLNs, DCTs, DASs, ML, and memory technical program committees, whose contributions played a crucial role in the conference's success. In addition, we wish to extend our deepest appreciation to Dennis Sylvester, the Editor-in-Chief of JSSC, for his valuable guidance. We also sincerely thank Danielle Marinese and the dedicated JSSC administration for their indispensable assistance in publishing this Special Section.

MIKE SHUO-WEI CHEN
Department of Electrical and Computer Engineering
University of Southern California
Los Angeles, CA 90089 USA
e-mail: swchen@usc.edu

VISVESH S. SATHE
School of Electrical and Computer Engineering
Georgia Institute of Technology
Atlanta, GA 30332 USA
e-mail: sathe@gatech.edu

MASSIMO ALIOTO
Department of Electrical and Computer Engineering
National University of Singapore
Singapore 117583
e-mail: massimo.alioto@gmail.com

JAE-SUN SEO
School of Electrical and Computer Engineering
Cornell Tech
New York, NY 10044 USA
e-mail: jseo28@asu.edu

HIDEHIRO SHIGA
KIOXIA Corporation
Yokohama 247-8585, Japan
e-mail: hidehiro.shiga@kioxia.com

Mike Shuo-Wei Chen (Senior Member, IEEE) received the B.S. degree from National Taiwan University, Taipei, Taiwan, in 1998, and the M.S. and Ph.D. degrees from the University of California at Berkeley, Berkeley, CA, USA, in 2002 and 2006, respectively, all in electrical engineering.

As a graduate student, he proposed and demonstrated the asynchronous SAR ADC architecture, which has been adopted in industry today for low-power high-speed analog-to-digital conversion products. He is currently a Professor with the Electrical Engineering Department, University of Southern California (USC), Los Angeles, CA, USA, where he leads the Analog Mixed-Signal Circuit Group, focusing on high-speed low-power data converters, frequency synthesizers, RF/mm-wave transceiver designs, analog circuit design automation, bioinspired computing, and nonuniformly sampled circuits and systems.

Visvesh S. Sathe (Member, IEEE) received the B.Tech. degree from the Indian Institute of Technology Bombay, Mumbai, India, and the M.S. and Ph.D. degrees from the University of Michigan, Ann Arbor, MI, USA.

He has held positions at the University of Washington, Seattle, WA, USA, and the Low-Power Advanced Development Group, AMD. His work at AMD involved inventing and translating energy-efficient circuit technologies, such as resonant clocking and adaptive clocking for supply droop mitigation into production microprocessors. He is currently an Associate Professor with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA, USA, where his group, the Processing Systems Laboratory (PSyLab), focuses on energy-efficient computing, implantable electronics, and run-time hardware optimization for digital and mixed-signal systems.

Dr. Sathe was a recipient of the NSF Career Award in 2019 and the Intel Outstanding Researcher Award in 2021. He serves on the Technical Program Committee of the IEEE Custom Integrated Circuits Conference (CICC) and the International Solid-State Circuits Conference (ISSCC). He also serves in the SSCS Webinar Committee to promote professional and student development.

Massimo Alioto (Fellow, IEEE) held positions at the University of Siena, Siena, Italy; Intel Labs—CRL; the University of Michigan, Ann Arbor, MI, USA; the University of California at Berkeley, Berkeley, CA, USA; and the École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland. He is currently a Professor with the Department of Electrical and Computer Engineering, National University of Singapore, Singapore, where he leads the Green IC Group and is the Director of the Integrated Circuits and Embedded Systems Area. He has authored or coauthored more than 350 publications on journals, conference proceedings, and five books. His primary research interests include widely energy-scalable systems, sustainable and self-powered silicon systems, data-driven systems for machine intelligence, and hardware security, among others.

Prof. Alioto served as an International Solid-State Circuits Conference (ISSCC) and ASSCC Subcommittee Member. His research has received various best paper awards, such as ISSCC 2023, and recognitions, such as top academic contributor at the VLSI Symposium on Circuits 2023. He served as the Technical Program Chair for ISCAS, SOCC, ICECS, and NEWCAS. He is the Chair of the Distinguished Lecturer Program for the IEEE Circuits and Systems Society. He was a Distinguished Lecturer of the CAS and the SSC Society. He served as a Guest Editor for several IEEE journal special issues, such as the IEEE JOURNAL OF SOLID-STATE CIRCUITS (JSSC), IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—I (TCAS-I), IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—II (TCAS-II), and IEEE JETCAS. He was an Editor-in-Chief of IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (TVLSI) SYSTEMS and an Deputy Editor-in-Chief of the IEEE JOURNAL ON EMERGING AND SELECTED TOPICS IN CIRCUITS AND SYSTEMS (JETCAS).

Jae-Sun Seo (Senior Member, IEEE) received the B.S. degree from Seoul National University, Seoul, South Korea, in 2001, and the M.S. and Ph.D. degrees from the University of Michigan, Ann Arbor, MI, USA, in 2006 and 2010, respectively, all in electrical engineering.

From 2010 to 2013, he worked at the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA. From 2014 to 2023, he was an Assistant/Associate Professor with the School of ECEE, Arizona State University, Tempe, AZ, USA, where he was a Visiting Researcher at the Intel Circuits Research Lab in 2015 and the Meta Reality Lab from 2022 to 2023. In July 2023, he joined the School of Electrical and Computer Engineering, Cornell Tech, New York, NY, USA, as an Associate Professor. His research interests include efficient hardware design of machine learning/neuromorphic algorithms and integrated power management.

Dr. Seo was a recipient of the IBM Outstanding Technical Achievement Award in 2012, the NSF CAREER Award in 2017, the Intel Outstanding Researcher Award in 2021, and the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (TVLSI) SYSTEMS Best Paper Award in 2022. He has served on the technical program committees for International Solid-State Circuits Conference (ISSCC), MLSys, DAC, and DATE. He served as an Associate Editor for IEEE TRANSACTIONS ON BIOMEDICAL CIRCUITS AND SYSTEMS (TBioCAS) and IEEE OPEN JOURNAL OF THE SOLID-STATE CIRCUITS SOCIETY (OJ-SSCS).

Hidehiro Shiga (Member, IEEE) received the B.S., M.S., and Ph.D. degrees from the University of Tokyo, Tokyo, Japan, in 1997, 1999, and 2002, respectively, all in electronic engineering.

In 2002, he joined the System-on-a-Chip Center, Toshiba Corporation, Yokohama, Japan, where he was engaged in the development of high density emerging memories. He was involved in the embedded FeRAM Design Team and the High-Speed FeRAM Development Team, where he successfully developed the world's largest 128-Mb chain FeRAM chip and presented the results in the International Solid-State Circuits Conference (ISSCC) 2009. In 2009, he joined the 3-D Flash Memory (BiCS) Development Team, where he has been a Leading Engineer in the Memory Core Circuits Design Team of the multiple generations of BiCS chips, including the first generation prototype chip. He is currently working as a Senior Specialist at the Device Technology Research and Development Center, KIOXIA Corporation, Yokohama.

Dr. Shiga has been a member of the Technical Program Committee of the IEEE ISSCC since 2021.