Guest Editorial Introduction to the Special Issue on the 2022 IEEE International Solid-State Circuits Conference (ISSCC)

THE International Technical Program Committee (ITPC) of the IEEE International Solid-State Circuits Conference (ISSCC) selects outstanding articles from the papers presented at the conference and invites the authors to submit an extended manuscript to the special issue of IEEE JOURNAL OF SOLID-STATE CIRCUITS (JSSC). This January issue contains the selected papers from the Digital Architectures and Systems, Digital Circuits, Machine Learning, Memory, and Wireline subcommittees.

I. WIRELINE

The wireline section includes both traditional electrical signaling and optical signaling demonstrating 100+ Gb/s bandwidth with impressive energy efficiency. These papers describe advanced digital and low-power analog-mixed signaling and clocking techniques necessary to equalize electrical and optical signal impairments. This is complemented with high frequency clock generation techniques that can generate very low jitter at 56 GHz.

The first two papers describe traditional wireline signaling over electrical links covering short-reach to long-reach applications. In [A1], Khairi et al. achieve an impressive 224 Gb/s data rate using an ADC-DSP based solution. Implemented in a 5 nm FinFET the receiver can equalize up to 38 dBs of loss using 30 tap FFE with an optional DFE. In [A2], Ye et al. demonstrate that a combination of analog mixed-signal equalization is a viable alternative to ADC-DSP solutions for short to medium-reach channels. They employ a combination of 3-tap FFE on transmit side and a 4-tap FFE at the receiver side using a two-stage 16-way time-interleaved front end. Fabricated in 28-nm CMOS technology, the transceiver achieves a bit error rate $<$1e-11 at 112-Gb/s transmission with a channel loss of 20.8 dB and energy efficiency of 2.29 pJ/b.

The next two articles describe latest advances in silicon photonic solutions integration and circuits techniques. In [A3], Hashemi Talkhooncheh et al. present a 100 Gb/s PAM-4 optical transmitter integrating silicon photonic modulator with a CMOS driver. The metal-oxide-silicon capacitor (MOSCAP) based push-pull segmented Mach-Zehnder Modulator (MZM) is 3D integrated with 28 nm CMOS that delivers 1.2 Vppd swing to modulators using a 0.9 V supply. The optical eye achieves 4.3 dB extinction ratio. In [A4], Sheng et al. describe the analog mixed signal part of the coherent receiver. The authors introduce a four-way time-interleaved architecture that enables two-stage equalization functionality to address chromatic dispersion. The receiver implemented in 28 nm CMOS also includes a polarization demultiplexer and carrier recovery (CR) to eliminate the inter-symbol interferences induced by fiber impairments and the carrier frequency offset.

To complement the signaling techniques described in the above articles, the next article describes high-frequency clock generation techniques. In [A5], Zhao et al. report a rather impressive 110 fs in jitter integrated over a bandwidth of 0.01 to 40 MHz. The PLL employs a switched-current finite impulse response (FIR) filter to suppress the $\Delta\Sigma$ modulator noise. Implemented in 28 nm CMOS process, the PLL consumes 23 mW and occupies 0.1 mm^2 area while achieving -98 dBc/Hz phase noise at 1 MHz offset in fractional-N mode.

II. DIGITAL CIRCUIT TECHNIQUES (DCT)

The section on Digital Circuit Techniques brings four papers that provide insights into digital circuits and systems that span across individual circuit macros, design techniques and toolchains as well as full systems. In [A6], Huang et al. from the University of Washington present a power-management architecture to recycle the energy stored in the decoupling capacitor before a processor enters sleep mode, with run-time tracking to minimize sleep/resume energy losses. Measurements demonstrate a 56.7% reduction in total energy losses for a duty-cycled 65 nm ARM Cortex-M0 processor with 1000 active cycles.

In [A7], Park et al. from KAIST present an ultra-low-jitter ring-oscillator-based injection-locked clock multiplier using power-gating injection-locking and a background multi-functional digital calibrator in 65 nm CMOS. Illustrating breakthrough rms jitter and FoM of the output signal at 8.16 GHz are 97 fs and -248.7 dB, respectively.

The third article in this section is from researchers at Samsung. In [A8], Lee et al. describe a fully automated hardware-driven clock-gating architecture for their 5 nm mobile SoC. The clock components are linked by a handshake interface and the clocks are automatically switched off if there is no activity. This scheme achieves 10% to 40% power reduction, while only occupying 0.03% of the die area.

Digital Object Identifier 10.1109/JSSC.2022.3225691

Finally, in the last article [A9] of this section, Vanderpool et al. from the IBM team present a 7 nm 16-core Power10 processor, featuring a series of architectural, design, and implementation improvements to ensure continued performance gains.

III. DAS

The digital architecture and systems section features general-purpose processors, security hardware accelerators, and domain-specific processors. In [A10], Nagi et al. from UCLA demonstrate a reconfigurable processor array that takes advantage of high-density die-to-die connectivity to deliver significant energy efficiency and throughput gains for data-driven computing applications. The article features a 2×2 multi-die reconfigurable processor fabricated in 16 nm with a multi-layer switch-box interconnect utilizes a silicon interposer with die-to-die IO bump pitch of 10 μm bump pitch with IO energy consumption of 0.38 pJ/b.

The next two articles demonstrate hardware accelerators for secure computing. In [A11], Chang et al. from Tsinghua University describe an agile crypto-processor supporting multiple 3rd-round finalists from NIST's post-quantum cryptographic (PQC) competition. The processor achieves 48 KOPS throughput at 3.4 μJ/Op energy efficiency and consumes an area of 3.6 mm^2 in 28 nm. As neural networks are increasingly deployed on the edge, confidentiality of proprietary models and inputs becomes increasingly important. In [A12], Hsieh et al. present a threshold-masking-based NN accelerator, fabricated in 28 nm which secures model parameters/inputs against power side-channel attacks using a threshold-based masking scheme along with a light-weight Trivium cipher with area and energy overhead limited to 64% and 6×, respectively. In [A13], [A14], and [A15], Im et al., Zhu et al., and Maji et al., respectively, showcase performance and efficiency advancements in domain-specific processors for a diverse range of applications, ranging from a 45 nm fully-integrated biometric smart card SoC enabling secure fingerprint authentication, a 40 nm neural signal processor for seizure prediction with a sensitivity of 92% and specificity of 99.1%, and finally a 28 nm depth signal processing SoC performing point-cloud based neural network for 3D object detection at 31.9 fps.

IV. Machine Learning

The ML subcommittee selected five outstanding articles from the 13 papers presented at the ISSCC 2022 ML sessions.

In [A16], Park et al. from Samsung present a multi-mode 8K-MAC neural processing unit for mobile SoC, manufactured in a leading edge 4 nm CMOS process. The NPU features a unified multi-precision MAC support from INT4/8/16 to FP16 in unified FP/INT datapath to boost area efficiency, and datapath reconfiguration to enhance utilization. It achieves 11.59 TOPS/W and 1.72 TOPS/mm^2 for MobileNetEdgeTPU with multi-mode support for low-latency mode or always-on mode.

In [A17], Houshmand et al., from KU Leuven and imec, present a digital and analog hybrid neural network SoC, featuring a RISC-V host coupled with a digital ML accelerator and an analog CIM core. The hybrid architecture enables concurrent execution of NNs across the digital and analog ML cores using a scheduler that considers required precision and efficiency. The 22 nm chip achieves top efficiency 790 TOPs/W (1.5b/8b, analog-in-memory), resp. 4T OPs/W (8b/8b, digital).

In [A18], Ju et al. from Northwestern University present a 65 nm Systolic Neural CPU Processor with 95% PE utilization for combined deep learning and general-purpose computing. It shares interesting architectural idea that reconfigures multiple processing elements (PE) to compose a RISC-V CPU. By reconfiguring the PEs to CPU or NN accelerator according to the given task, it increases area efficiency while exhibiting high utilization.

In [A19], Wang et al., from Tsinghua University and Tsing-Micro, present a 28 nm approximate-computing-based transformer processor with big-exact/small-approximate processing elements, bidirectional asymptotic sparsity speculation, and an out-of-order computing scheduler, toward energy-efficient computation of transformers with global attention.

In [A20], Tu et al., from Tsinghua University and the University of California at Santa Barbara, describe a reconfigurable digital CIM processor for large-scale deep learning engines. To accelerate both cloud inference and training, their chip supports both FP and INT computation in a unified pipeline, and features an in-memory Booth multiplication scheme, in-memory accumulation, and exponent pre-alignment.

V. Memory

The memory subcommittee selected six outstanding articles from the 19 papers presented at the ISSCC 2022 memory sessions. The first group of this special issue includes three papers on advancements of HBM, LPDDR5 and GDDR6 DRAMs for high-performance graphics, server, and machine learning. The second group includes two papers on advancement of in-memory computation aiming to improve energy efficiency of machine learning applications in data center and AI edge devices. The third group of paper is devoted to advancement in 3D NAND Flash memory.

In [A21], Park et al. from SK Hynix present a 12-high stack, 192 Gb HBM3 with a 7 Gb/s/pin I/O speed to reach an 896 GB/s bandwidth. New design features include an in-DRAM ECC, internal NN-DFE, TSV auto-calibration, and machine-learning-driven layout optimization. These features enable efficient data transfer among stacked dies while maintaining low power consumption.

In [A22], Jung et al. from KAIST present a supply-noise-induced jitter (SIJ) cancellation technique, aiming to address the impact of SIJ created by clock distribution network and data transmitting and receiving paths on a LPDDR5 mobile DRAM. The proposed technique is based on an adaptive filter using a least mean square (LMS) algorithm to cancel the jitter along the CDN and TX/RX paths. The RMS of jitter for RDQS clock is reduced by more than 80%, which increases data eye-opening margin by four times at 6.4 Gb/s.

In [A23], Lee et al. from Samsung Electronics present a 16 Gb GDDR6 achieving 27 Gb/s/pin in NRZ signaling. A T-coil using a single metal layer is designed for the first time in a DRAM process to increase maximum I/O speed for NRZ signaling. A merged-MUX TX increases the operating frequency and decreases power and area consumption. Quad skew training adjusts the quadrature clock skew to maximize the sampling margin, covering a wide frequency range.

In the second section, two articles are devoted to introduce advancements in compute-in-memory designs using GDDR6 DRAM to accelerate deep learning application and non-volatile ReRAM for edge-AI devices.

In [A24], Kwon et al. from SK Hynix describe an 1 ynm, GDDR6-based accelerator-in-memory solution with a command set for deep-learning operation. The 8 Gb design achieves 1 GHz MAC operations with a peak throughput of 1 TFLOPS and supports major activation functions to improve computation accuracy. Compared to commercial offerings based on HBM DRAM, the GDDR6-based design enables a relatively low-cost solution with an I/O bandwidth of 16 Gb/s. This solution enables up to a 10× speedup relative to emulated GPU systems using HBM2 bandwidth.

In [A25], Hung et al. from National Tsing Hua University present a 22 nm 8 Mb ReRAM compute-in-memory (CIM) macro targeting low-power edge applications with a DC-current-free time-domain MAC operation. The nonvolatile CIM design is the first to adopt a time-space readout scheme to achieve 0.76 ns/b for TAC-OUT and 61.8 TOPS/W for an 8 b-input, 8 b-weight, and 19 b-output MAC computation.

The third section is devoted to introduce the advancement in 3D NAND Flash memory which performance and density continue to improve for storage applications. In [A26], Yuh et al., from Western Digital and KIOXIA, describe a 162-word-line-layer 1 Tb 4b/cell 3D Flash memory with a 15 Gb/mm^2 bit density. A program throughput of 60 MB/s and a tR of 65 μs are achieved by employing an 8 kB WL central stair structure and a contact-through-WL architecture. The design supports a 2.4 Gb/s IO speed with a low-tapped termination (LTT) and center-tapped termination (CTT) combo driver. A time-division peak-power management feature is employed to improve system parallelism and high-speed wafer-level testing.

MASUM HOSSAIN, *Guest Editor*
University of Alberta
Edmonton, AB T6G 2R3, Canada

ARIJIT RAYCHOWDHURY, *Guest Editor*
ECE Department
Georgia Institute of Technology
Atlanta, GA 30332 USA

SANU K. MATHEW, *Guest Editor*
Intel Corporation
Hillsboro, OR 97124 USA

YAKUN SOPHIA SHAO, *Guest Editor*
EECS Department
University of California
Berkeley, CA 94720 USA

YIH WANG, *Guest Editor*
Design Technology Platform
Taiwan Semiconductor Manufacturing Corporation
Hsinchu 300-096, Taiwan
e-mail: eric_wang@tsmc.com

APPENDIX: RELATED ARTICLES

[A1] A. Khairi et al., "A 1.41-pJ/b 224-Gb/s PAM 4 6-bit ADC-based SerDes receiver with hybrid AFE capable of supporting long reach channels," *IEEE J. Solid-State Circuits*, early access, Oct. 26, 2022, doi: 10.1109/JSSC.2022.3211475.

[A2] B. Ye et al., "A 2.29-pJ/b 112-Gb/s wireline transceiver with RX four-tap FFE for medium-reach applications in 28-nm CMOS," *IEEE J. Solid-State Circuits*, early access, Nov. 29, 2022, doi: 10.1109/JSSC.2022.3223052.

[A3] A. H. Talkhooncheh et al., "A 100-Gb/s PAM 4 optical transmitter in a 3-D-integrated SiPh-CMOS platform using segmented MOSCAP modulators," *IEEE J. Solid-State Circuits*, early access, Nov. 3, 2022, doi: 10.1109/JSSC.2022.3210906.

[A4] K. Sheng et al., "A 4.6 pJ/b 200 Gb/s analog DP-QPSK coherent optical receiver in 28 nm CMOS," *IEEE J. Solid-State Circuits*, early access, Feb. 20, 2022, doi: 10.1109/JSSC.2022.3211347.

[A5] Y. Zhao, O. Memioglu, L. Kong, and B. Razavi, "A 56-GHz fractional-*N* PLL with 110-fs jitter," *IEEE J. Solid-State Circuits*, early access, Nov. 21, 2022, doi: 10.1109/JSSC.2022.3220547.

[A6] C. H. Huang, A. Mandal, D. Peña-Colaiocco, E. P. Da Silva, and V. S. Sathe, "Regenerative breaking: Optimal Energy recycling for energy minimization in duty-cycled domains," *IEEE J. Solid-State Circuits*, early access, Nov. 23, 2022, doi: 10.1109/JSSC.2022.3221143.

[A7] S. Park, S. Yoo, Y. Shin, J. Lee, and J. Choi, "A Sub-100 fs-jitter 8.16-GHz ring-oscillator-based power-gating injection-locked clock multiplier with the multiplication factor of 68," *IEEE J. Solid-State Circuits*, early access, Oct. 10, 2022, doi: 10.1109/JSSC.2022.3210212.

[A8] J. G. Lee, H. Jeon, Y. Choi, and A. Kim, "Fully automated hardware-driven clock-gating architecture with complete clock coverage for 5 nm Exynos mobile SoC," *IEEE J. Solid-State Circuits*, early access, Feb. 20, 2022, doi: 10.1109/JSSC.2022.3219410.

[A9] B. Vanderpool et al., "Deterministic frequency and voltage enhancements on the power10 processor," *IEEE J. Solid-State Circuits*, early access, Dec. 6, 2022, doi: 10.1109/JSSC.2022.3225378.

[A10] S. S. Nagi, U. Rathore, K. Sahoo, T. Ling, S. S. Iyer, and D. Markovic, "A 16-nm 784-core digital signal processor array, assembled as a 2 × 2 dielet with 10-μm pitch interdielet I/O for runtime multiprogram reconfiguration," *IEEE J. Solid-State Circuits*, early access, Oct. 27, 2022, doi: 10.1109/JSSC.2022.3212685.

[A11] J. S. Chang et al., "A 1.05-A/m minimum magnetic field strength single-chip, fully integrated biometric smart card SoC achieving 792.5-ms transaction time with anti-spoofing fingerprint authentication," *IEEE J. Solid-State Circuits*, early access, Nov. 15, 2022, doi: 10.1109/JSSC.2022.3220081.

[A12] Y. Y. Hsieh, Y. C. Lin, and C. H. Yang, "A 96.2 nJ/class neural signal processor with adaptable intelligence for seizure prediction," *IEEE J. Solid-State Circuits*, early access, Feb. 20, 2022, doi: 10.1109/JSSC.2022.3218240.

[A13] D. Im et al., "DSPU: An efficient deep learning-based dense RGB-D data acquisition with sensor fusion and 3-D perception SoC," *IEEE J. Solid-State Circuits*, early access, Nov. 9, 2022, doi: 10.1109/JSSC.2022.3218278.

[A14] Y. Zhu et al., "RePQC: A 3.4-uJ/Op 48-kOPS post-quantum crypto-processor for multiple-mathematical problems," *IEEE J. Solid-State Circuits*, early access, Dec. 9, 2022, doi: 10.1109/JSSC.2022.3216758.

[A15] S. Maji, U. Banerjee, S. H. Fuller, and A. P. Chandrakasan, "A threshold implementation-based neural network accelerator with power and electromagnetic side-channel countermeasures," *IEEE J. Solid-State Circuits*, early access, Nov. 1, 2022, doi: 10.1109/JSSC.2022.3215670.

[A16] J. S. Park et al., "A multi-mode 8K-MAC HW-utilization-aware neural processing unit with a unified multi-precision datapath in 4 nm flagship mobile SoC," *IEEE J. Solid-State Circuits*, early access, Feb. 20, 2022, doi: 10.1109/JSSC.2022.3205713.

[A17] P. Houshmand et al., "DIANA: An end-to-end hybrid digital and analog neural network SoC for the edge," *IEEE J. Solid-State Circuits*, early access, Oct. 31, 2022, doi: 10.1109/JSSC.2022.3214064.

[A18] Y. Ju and J. Gu, "A systolic neural CPU processor combining deep learning and general-purpose computing with enhanced data locality and end-to-end performance," *IEEE J. Solid-State Circuits*, early access, Oct. 27, 2022, doi: 10.1109/JSSC.2022.3214170.

[A19] Y. Wang et al., "An energy-efficient transformer processor exploiting dynamic weak relevances in global attention," *IEEE J. Solid-State Circuits*, early access, Oct. 25, 2022, doi: 10.1109/JSSC.2022.3213521.

[A20] F. Tu, "ReDCIM: Reconfigurable digital computing-in-memory processor with unified FP/INT pipeline for cloud AI acceleration," *IEEE J. Solid-State Circuits*, early access, Dec. 2, 2022, doi: 10.1109/JSSC.2022.3222059.

[A21] M. J. Park et al., "A 192-Gb 12-high 896-GB/s HBM 3 DRAM with a TSV auto-calibration scheme and machine-learning-based layout optimization," *IEEE J. Solid-State Circuits*, early access, Feb. 20, 2022, doi: 10.1109/JSSC.2022.3193354.

[A22] Y. Jung, S. Lee, H. Kim, and S. Cho, "A supply-noise-induced jitter canceling adaptive filter for LPDDR5 mobile DRAM," *IEEE J. Solid-State Circuits*, early access, Sep. 27, 2022, doi: 10.1109/JSSC.2022.3203221.

[A23] D. Lee et al., "A 16Gb 27Gb/s/pin T-coil based GDDR6 DRAM with merged-MUX TX, optimized WCK operation, and alternative-data-bus," *IEEE J. Solid-State Circuits*, early access, Dec. 1, 2022, doi: 10.1109/JSSC.2022.3222203.

[A24] S. Lee et al., "A 1 ynm 1.25 V 8 Gb, 16 Gb/s/pin GDDR 6-based accelerator-in-memory supporting 1TFLOPS MAC Operation and various activation functions for deep-learning applications," *IEEE J. Solid-State Circuits*, early access, Feb. 20, 2022, doi: 10.1109/JSSC.2022.3200718.

[A25] J. M. Hung et al., "8-b Precision 8-Mb ReRAM compute-in-memory macro using direct-current-free time-domain readout scheme for AI edge devices," *IEEE J. Solid-State Circuits*, early access, Sep. 2, 2022, doi: 10.1109/JSSC.2022.3200515.

[A26] J. H. Yuh et al., "A 1-Tb 4-b/cell 4-plane 162-layer 3-D flash memory with 2.4-Gb/s IO interface," *IEEE J. Solid-State Circuits*, early access, Aug. 11, 2022, doi: 10.1109/JSSC.2022.3193326.

Masum Hossain (Member, IEEE) received the B.Sc. degree from the Bangladesh University of Engineering and Technology, Dhaka, Bangladesh, the M.Sc. degree from Queen's University, Kingston, ON, Canada, and the Ph.D. degree from the University of Toronto, Toronto, ON, Canada, in 2002, 2005, and 2010, respectively.

From 2008 to 2010, he was with the Analog and Mixed Signal Division, Gennum Corporation, Burlington, ON, Canada. From 2010 to 2012, he was with Rambus Laboratory, Sunnyvale, CA, USA, as a Senior Member of Technical Staff, In 2013, he joined the Department of Electrical and Computer Engineering, University of Alberta, Edmonton, AB, Canada.

Dr. Hossain was a recipient of the Best Student Paper Award at the 2008 IEEE Custom Integrated Circuits Conference and the Analog Device's Outstanding Student Designer Award in 2010. In 2021, he received the EPS Society Nominated Best Paper Award in IEEE TRANSACTIONS IN COMPONENTS, PACKAGING AND MANUFACTURING.

Arijit Raychowdhury (Fellow, IEEE) is the Steve W. Chaddick Chair of the School of Electrical and Computer Engineering, Georgia Institute of Technology, where he was previously the Motorola Foundation Professor. From 2013 to July 2019, he was an Associate Professor and held the ON Semiconductor Junior Professorship. His industry experience includes five years as a Staff Scientist with the Circuits Research Lab, Intel Corporation, and two years as an Analog Circuit Researcher with Texas Instruments Inc. His research interests include low-power digital and mixed-signal circuit design and exploring interactions of circuits with device technologies. He has authored over 200 articles in journals and refereed conferences and holds more than 26 U.S. and international patents.

Dr. Raychowdhury's students have won several prestigious fellowships and 16 best paper awards over the years. He currently serves on the technical program committees of ISSCC, CICC, VLSI Symp, and DAC. He is an IEEE SSCS Distinguished Lecturer for 2021–2022.

Sanu K. Mathew (Fellow, IEEE) received the B.Tech. degree in electronics and communications engineering from the College of Engineering, Trivandrum, India, in 1993, and the M.S and Ph.D. degrees in electrical and computer engineering from the State University of New York at Buffalo, Buffalo, NY, USA, in 1996 and 1999, respectively.

He is a Senior Principal Engineer with the Circuits Research Lab, Intel Corporation, Hillsboro, OR, USA, where he leads research in computer arithmetic data-path circuits and hardware accelerators for cryptography and security. He has been with Intel Corporation since 1999. He holds 109 issued/pending patents and has published 95 conference/journal articles and authored two book chapters. He also mentors Intel and SRC-funded research projects in leading universities.

Dr. Mathew has served on the program committees of ISSCC, ARITH, ISLPED, DAC, and SOCC conferences.

Yakun Sophia Shao (Member, IEEE) received the B.S. degree in electrical engineering from Zhejiang University, Hangzhou, China, and the M.S. and Ph.D. degrees in computer science from Harvard University, Cambridge, MA, USA, in 2014 and 2016, respectively.

She is currently an Assistant Professor with the Electrical Engineering and Computer Sciences Department, University of California at Berkeley, Berkeley, CA, USA. Before Berkeley, she was a Senior Research Scientist with NVIDIA, Inc., Santa Clara, CA, USA. Her research interest is in the area of computer architecture, with a special focus on specialized accelerators, heterogeneous architecture, and agile VLSI design methodology.

Dr. Shao is a recipient of an Intel Rising Stars Faculty Award, a Google Faculty Rising Stars Award in Systems Research, a Facebook Research Award, the inaugural Dr. Sudhakar Yalamanchili Award, and the 2022 IEEE TCCA Young Computer Architect Award.

Yih Wang (Senior Member, IEEE) received the B.S. degree in electrical engineering from National Tsing-Hua University, Hsinchu, Taiwan, in 1993, and the M.S. and Ph.D. degrees in electrical and computer engineering from University of Florida, Gainesville, FL, USA, in 1998 and 2000, respectively.

He was the Pittman-Eminent Scholar Predoctoral and Postdoctoral Fellow of Florida-Solid-State-Electronics Laboratory, University of Florida, from 1999 to 2001, guided by Prof. Chih-Tang Sah. From 2001 to 2018, he was with Intel Logic Technology Development, Hillsboro, OR, USA, where he was a Senior Principal Engineer, leading the development of embedded memory technologies including the high-speed and low-power embedded SRAM, 22nm embedded DRAM, and embedded non-volatile memory (eNVM) such as STT-MRAM and ReRAM. He joined Taiwan Semiconductor Manufacturing Corporation (TSMC) in 2018 and is currently an Engineering Director in the Design and Technology Platform at TSMC. He has authored and coauthored more than 50 journal and conference publications and has more than 130 issued and pending U.S. patents on semiconductor process, memory devices, and memory circuit designs.

Dr. Wang is on the technical program committee of the ISSCC. He served on the technical program committee of the VLSI Symposium from 2016 to 2018. He received three Intel Achievement Awards for his works on developing embedded memory technologies for high-performance microprocessors.