Nonsilicon, Non-von Neumann Computing—Part II

By SANKAR BASU, *Life Fellow IEEE*
Guest Editor

RANDAL E. BRYANT, *Fellow IEEE*
Guest Editor

GIOVANNI DE MICHELI, *Fellow IEEE*
Guest Editor

THOMAS THEIS, *Fellow IEEE*
Guest Editor

LLOYD WHITMAN
Guest Editor

I. BACKGROUND AND TECHNICAL SCOPE

This is a special issue (Part II) expanding on a previous special issue [1] dedicated to future computing technologies. Part I appeared in the January 2019 issue of the PROCEEDINGS OF THE IEEE. Part II was inspired by our desire to highlight technological prospects for continued advances in computing despite diminishing returns from traditional investments aimed at increasing the device density and clock frequency in digital circuits. In our preface to Part I, we noted that the future of computing is at a crossroads—a statement still valid today.

The articles in this month's special issue provide insight into future computing technologies such as novel architectures, spintronic memories, and quantum computing.

Part I featured a selection of the many nascent alternative technologies and architectures being pursued in academic research labs and industry—the result of government and industry research and development (R&D) priorities at the highest level. The breadth and depth of this interdisciplinary R&D in "beyond-Moore's Law" computing was astonishing then and continues to grow today. This issue thus captures additional topics that were deferred from Part I due to space limitations as well as topics associated with new initiatives and expanded R&D programs that have emerged in the year and a half since the publication of Part I.

Computing technologies of the future remain one of the grand challenges not just for the engineering profession, but for the larger science and technology communities as well. We refer interested readers to Part I (see [2]) for details regarding the motivation. Here, we make some brief remarks on how the articles in Part II contrast with those in Part I and fill some gaps in the coverage of research.

The issue starts with two articles on the foundational issues of two very different types (Section A). The first one is on information theoretic techniques of error correction for resilient computing in the context of recognized computational modules in the architecture community for high performance computing (HPC) (the so-called "dwarfs"). The second article in this section is also foundational but from the viewpoint of device physics. The spectral shapes of the energy bands in semiconducting materials are of concern here and are found to be of critical importance in the design of tunnel field-effect transistors (tFETs) that have been researched for many years as a candidate for fast-switching yet low-voltage (thus low power) transistor. Although it does not discuss in detail the prior

Digital Object Identifier 10.1109/JPROC.2020.3001748

research on tFETs, the article does provide a useful bibliography for those interested.

In anticipation of the current interest in artificial intelligence (AI)-inspired hardware architectures, Part I provided a coverage of the issues intrinsic to cognitive, brain-inspired, and neuromorphic architectures as alternatives to traditional von Neumann architectures. In the second section of the present issue (Section B), however, we delve into deep neural networks (DNNs), and their implementation in different technologies (e.g., silicon CMOS as well as optical technologies) with a view toward applications. The focus here is not to mimic the behavior of the brain for energy efficiency as was the case in some articles in Part I (see [1, pp.144–164]), but in keeping with the current performance-driven trend in AI, to explore implementation of DNNs using architectural heterogeneity in silicon CMOS, or noncharge-based (optical) alternatives.

With the advent of our present digital age, analog computing receded to the background for the last decades. However, despite issues of scalability, variability, noise susceptibility, etc., analog circuits and systems are making a comeback due to several technical reasons—at least for specialized applications. The solo article on analog computing in Section C of Part II deals with current state of research on analog field programmable gate arrays (FPGAs) by a leading expert in the field. The article is also instructive for an overall outlook on analog computing in general.

Although several nonsilicon and noncharge-based technologies were addressed in Part I, spin-based systems did not receive coverage. While spin-based systems are already making their mark on memory technologies, new breakthroughs are in the offing and the limits of Flash and other memories are being approached. Considerable progress has been made recently in research centers (e.g., at Purdue University) toward spin-based circuits for computing architectures potentially with wider capabilities, thus altogether making the prospect of an all-spin computing machine closer to reality at least for some applications. These two articles form the contents in Section D on spintronics.

The issue ends with Section E dedicated to the subject of quantum computing. Interest in quantum information sciences in general, and in quantum computing, in particular, is increasing worldwide, with the United States recently legislating a National Quantum Initiative [10]. Research in this area has accelerated in both academic and industrial laboratories, driven by increasing government funding and corporate investments in both industrial R&D and start-up companies. Benchmarks for progress have been established, and a number of companies have made their nascent computing platforms available to the academic research community. Despite this broadening of interest, presentations of the research that are accessible to nonspecialist engineers, physicists, and computer scientists are still uncommon and valuable. We are, therefore, pleased to have two such articles in this issue; one from a leading industrial research group and another from an outstanding university research group. A Point of View article on quantum computing that appeared in Part I [3] provides a "quantum" connection between Part I with Part II of this special issue.

Lastly, one may wonder how speculative and feasible the different technologies discussed in both Parts I and II are at the present time. Indeed, there was a request from the larger editorial board of the PROCEEDINGS OF THE IEEE that it may be beneficial for the nonspecialists if the guest editors provide comments on ratings of all the technologies covered in both Parts I and II in terms of their status, potential, feasibility, etc. While this is a much desired and a lofty goal, it is fair to say that the situation is in flux, and while many industry practitioners are attempting to gauge the situation (e.g., benchmarking is being attempted in industry for specific problem domains, DNNs being one), it will be a while before a clearer picture emerges (it took about half a century for silicon CMOS technology with von Neumann architectures to get where we are at now!). The two-part issue presented a sampling from a broad range of emerging technologies with possible implementations in diverse types of architectures. Not every possible idea could be incorporated due to reasons beyond our control (e.g., nano-electro-mechanical circuits and systems). As of this writing, still more ideas keep emerging (dubbed sensible computing, thermodynamic computing, so on and so forth). The proper match between technology and architecture for optimal performance (e.g., speed, power, latency, accuracy, etc.), the latter of which itself remains an elusive metric depending on the application, is the goal of future research. However, it is perhaps not unfair to remark that the landscape of computing in general is changing, and newer machines will be adapted and optimized more and more for special purpose applications making use of a diverse set of technologies, some of which we hope to have presented in this two-part special issue.

As in Part I, we categorize the articles included in Part II of the special issue into several broad categories, and present Editors' perspectives on specific articles next. Since many of these articles survey large ongoing work they may not provide full coverage of the extent of such projects. Most projects have open and extensive websites of their own, available to the readership of the journal for further reference.

A. Foundational Issues

In the first section of this issue, we group together two articles with somewhat foundational appeal, but from two extreme ends of the computing stack. The first article addresses reliability and fault tolerance of high performance computing (HPC)

frameworks from an information theory/error correction standpoint, whereas the second article explores the fundamental physical bottlenecks for the energy efficient design of tunnel FETs as the basic computing elements at the lowest end of the computing stack.

The section opens with the article titled "Addressing unreliability in emerging devices and non-von Neumann architectures using coded computing," by Dutta *et al*. The article reviews the status of current research on error correction and fault tolerance—an issue that is becoming more important with the advent of miniaturization of the basic building blocks to the extent of nanoscale. This is done primarily in a distributed computing scenario enabled by cloud computing—a scenario important in HPC applications. The article focuses on a number of basic blocks of scientific computing, called "dwarfs," from linear algebra, spectral theory, and Monte Carlo (MapReduce) methods etc. Discussions center around the use of "coded computation" for error correction and fault tolerance as methods that are far superior to naïve use of traditional hardware fault tolerance/error correction methods, for example, the triple modular redundancy (TMR). A performance metric for the coding schemes, namely the "recovery threshold," is introduced and four specific "dwarfs" are evaluated via this metric. As more "dwarfs" have been added to the first seven enunciated some time ago, the vision is to lay out a formalism on which information theorists, systems experts, and circuit designers can work together to build resilience in the next-generation computing systems. To this end, several problems for future research, including how the techniques developed may potentially apply to analog and emerging technology-based computations are pointed out as well. The article has a flavor similar to the one that appeared in Part I of the issue [4], and in some sense may be classified in the same family of information theory-inspired design.

The next article turns to tFET as a promised candidate for low power electronics. It is natural to think that one of the main ingredients of an energy-efficient computing machine is the design of most energy efficient switching devices (transistors) for building energy-efficient logic circuitry. The tFETs have enjoyed being a major potential candidate for such a device. Although tFETs generated a profusion of research in recent times, it has not yet fully materialized in terms of its promise. Indeed, an entire NSF research center—the Center for Energy Efficient Electronic Science (https://e3s-center.berkeley.edu/) at the University of California at Berkeley—with a lifespan of a decade or more had been dedicated for this purpose. With this in mind, the Center Director and a leading expert in the field, Prof. Eli Yablonovitch, was invited to write a survey article dedicated to the challenges and opportunities of the tFET research. In the course of formulating that article, the research team came up with somewhat more in-depth analysis of the requirements for tFET to be successful switching device with promised subthreshold ON–OFF characteristics. The article "Tunnel-FET switching is governed by non-Lorentzian spectral line shape" by Vadlamani *et al.* from the University of California Berkeley breaks new grounds by arguing that the fundamental difficulty appeals to more basic physics, has to do with spectral shapes of the energy bands involved, and exponentially decaying tails of spectral shapes would be necessary among other requirements. The authors expect that devices satisfying these requirements will be realized experimentally during the next years. The article points to the relevance of similar difficulties to other problems of device physics, and thus the considerations involved should have a broader appeal than just design of tFETs. The article also includes a bibliography of important publications on tFETs, right from the late 1970s to the state-of-the-art experimental demonstrations in the last couple of years, and thus it could also serve as a bibliographical reference to the tFET literature.

B. Deep Learning Neural Network Architectures

As the first issue was being edited and went into press, the topic of AI and hardware associated with it began to emerge as the next topic of interest in the research community. Since then AI hardware has dominated the interest of the hardware community more than ever before, machine learning (ML) as the theoretical/analytical underpinning of the field led the way, and virtually every industry and academic research group became involved with research inspired by AI in one way or another. While AI as a field is almost as old as modern computing, the push to build low level hardware for AI applications is relatively recent and has been enabled by advances in hardware technologies of the last decades. On the other hand, limits to power of computing encumbered by the end of Moore's law for traditional von Neumann machines built using silicon CMOS technologies diverted attention toward different types of architectures, algorithms, and hardware that enable learning, prediction, and inference tasks. Indeed, this was the genesis of the Grand Challenge in Computing announced by the White-House OSTP in 2015 (see [2]). Attention was drawn to building computing machines mimicking the power-efficiency of the mammalian brain, on the one hand, and statistical ML-based architectures, on the other.

These two somewhat dichotomous aspects in the development of hardware continue to abate. Both brain research and AI/ML research have their own independent histories. Attempts to understand the human brain from a biological perspective has been pursued in the neuroscience community, from which recent hardware research has derived cues, for example, by adopting the spiking

neurons as an important object of study (there exist several other examples, some of which appeared in Part I of the present two-part special issue). Machine learning, on the other hand, has a different history originating in early computer science (e.g., aspects of AI, pattern recognition), statistics, cybernetics, human language technologies, etc.

Admittedly, the recent surge in interest in AI/ML hardware was spurred by the unprecedented success of DNNs. Originally conceived by early attempts to emulate the brain, current neural networks, including the DNNs in particular, have very little to do with the biological neurons, and at the same time reasons for the success of DNN is poorly understood on analytical methods cultivated in ML research. In this respect, DNNs occupy an interesting middle ground because they are neither brain-inspired nor is it fully amenable to ML methodology, and yet enjoys huge success. While Part I dealt with brain and cognitive science-inspired paradigms, Part II of the present issue deals with two articles on different implementations of DNNs, which lie exactly at the middle ground just mentioned. Recent articles [7], [8] in the PROCEEDINGS OF THE IEEE outside this two-part special issue have addressed DNN implementations in GPUs, as well as neuromorphic networks in emerging resistive RAM technologies [8]. The reader willing to have a more detail alternative perspectives on this topic may consult the articles just mentioned.

By contrast, in the present issue, the article titled "The heterogeneous deep neural network processor with a non-von Neumann architecture" by Shin and Yoo from KAIST, South Korea, suggests a non-von Neumann architecture with multiple learning functionalities but specialized for DNN applications only. While up until now most DNNs have been implemented in CPUs and GPUs, the article makes a detailed case that a non-von Neumann architecture using ML-specific integrated circuits (MSICs) can be used for different applications (e.g., classification, recognition of speech, image, and video), and has benefits of operating at lower power than GPU/CPU-based DNNs. A heterogeneous architecture involving computation-limited convolutional neural networks and memory bandwidth-limited recurrent neural networks is used in their application-specific implementation, and implementation results of the fabricated chip are presented in great detail. It must be reemphasized that the DNN hardware discussed in this article is very different from architectures directly inspired by the brain in terms of structure and operation mechanism.

Discussion then turns to optical architectures (again in the specific context of DNNs)—a technology that did not find its legitimate place in Part I. Efforts to build computing machines using optical technology is, in some sense, decades old. Despite initial promises, early attempts at optical computing did not pan out for a number of reasons, and interest died out. While the natural parallelism offered by light seemed attractive, among other difficulties, signal conversion from electronics to optics and *vice versa* proved cumbersome and bulky in implementation. However, recent improvements in technology and emphasis on architectures of different nature have brought optical implementations within reach. Use of optics, for example, in the form of silicon photonics as low power interconnects is one way that optics has reemerged. We have also witnessed recent attempts by various groups toward optical implementations of Ising machines with promise for hardware approaches to solving combinatorially hard problems of the type that usually challenge quantum-like computing machines (more on it later).

While we do not delve into optical architectures of this latter type in the present issue, the next article "Silicon photonics codesign for deep learning" by Cheng *et al*. from Columbia University deals with implementation of DNNs in optical technology. The article takes advantage of the fact that vector–matrix multiplication is a key element of the arithmetic operations in large deep (spiking) neural networks. The authors propose the codesign of an integrated silicon photonics architecture for this purpose. A comprehensive system-level codesign is presented by taking full advantage of both the optical speed-up and electronic manipulation of the parallelism and memory, and by paying careful attention to interactions between the two technologies. Specifically, the article explores the mapping of abstract to device-level mathematics, the asymptotic runtime of the optical vector matrix multipliers, the abstract device complexity, scaling, and the control of the device. In addition, the latency, power penalty, scaling, and overall footprint of the physical layer are studied with respect to matrix multiplication performance in the context of DNNs.

C. Analog Computing Architecture

Analog computing is perhaps at least as old as modern digital computing. However, over the years, analog and digital VLSI designs have come to constitute two significantly different areas of technology in objectives, knowledge, skills, and deliverables, whereas mixed-signal circuits leverage advantages of both (the RF circuits and systems are also viewed as analog endeavor in terms of tools and techniques employed). The differences in knowledge bases arising from the facts that an analog circuit designer needs to know more of semiconductor technology/physics, electrical circuit theory, control and feedback, whereas a digital circuit designer's tool-bag consists of Boolean algebra, synchronous and asynchronous circuits, etc., further contributed to this divergence. The dearth of expertise in analog design, accompanied by progressively less emphasis on it in academic curriculum, has led some to call out analog design as more of an art form. Large-scale analog circuits were not just hard to design, the inaccuracies resulting from the large signal-to-noise ratio,

lack of scalability, etc., were also of major concerns. More recently, however, due to new ways of thinking, in some instances the dichotomy between analog and digital has begun to be blurred, and primarily due to its potential in energy efficiency, analog computing is beginning to make a comeback.

While classic applications of analog computing constituted attempts to solve ordinary differential equations by simulating them in analog circuitry, potentially newer applications have also emerged (e.g., sensors and biomedical applications; there are others at the very cutting of technology, which time and space will not allow us to discuss in this special issue). Meanwhile, the analog design principles have also gained technical inspiration from the advances in digital technologies in its tools and techniques, as well as in asking questions of more foundational nature. As an example of the former development, considerations of analog field-programmable analog arrays (FPAAs) are inspired by the FPGAs of the digital world, and are, in fact, the main subject of an article in the present special issue. As an example of the latter type of development, while early advances in digital computing led researchers to raise issues of decidability, computability, and classification of problems according to complexities of computation, similar questions have now been asked for analog computing as well. An early attempt in this direction is the discussion of complexities of computations over the field of real numbers (see [9]). We are aware of ongoing attempts to solve computationally hard combinatorial problems, say, by mapping them to the study of systems of ordinary differential equations, which are in turn implemented in analog hardware of various types. Such investigations not only open new ways of approaching otherwise difficult problems by analog means, but also shed new light on our understanding of analog dynamical systems as well (e.g., computationally hard problems like the NP complete problems often get mapped to nonlinear analog dynamical systems that exhibit chaotic behavior with extreme sensitivity to parameter variations, thus in some sense, conceptually bridging the analog-discrete divide).

The article titled "Large-scale field-programmable analog arrays" by Hasler (Georgia Tech) undertakes analog design. Tools enabling design of large-scale FPAAs as ubiquitous analog-mixed-signal, low-power sensors similar to the digital FPGAs are considered. The article starts with a leisurely discussion of the FPAA concepts and its history in contrast to similar developments for the FPGA in the digital domain. It illustrates the potential capabilities of FPAAs leading up to large-scale Systems on Chip (SoCs) FPAA devices offering potential opportunities, for example, by enabling numerics, studying architecture complexity, and applications. Issues of scalability, programmability, reconfigurability, reuse, and memory techniques are discussed for analog systems, and future projections in terms of analog abstraction hierarchy, hardware infrastructure, performance metrics, etc., are also made in the context of the specific problems discussed.

D. New Spin on Spintronics

Using electron spin as an alternative to charge transfer in electronics was proposed in the 1990s and has been working on since then. In fact, the PROCEEDINGS had published two special issues on the entire field of spintronics—first [5] in 2003, and then 13 years later again in 2016 [6]. The 2016 issue predicted (cf. Scanning the Issue article [6]) enough advancement in another 13 years to warrant yet another special issue. Indeed, spintronic memories have meanwhile advanced to the stage of commercialization, but the desire to make logic from spintronic devices have not yet materialized for a number of technical reasons. However, in the interim, new advances have also been made in spintronic memories, and new types of computations, namely the probabilistic computations have appeared in the scene to which spintronics may indeed have much more to contribute.

The first article of this group of two articles on spintronics in the special issue, "Magnetic racetrack memory: From physics to the cusp of applications within a decade," authored by Bläsing *et al.*, a leading group of researchers from the Max Planck Institute (MPI) and the Center for Advanced Electronics (cfaed) Dresden, shows how race track memories (RTMs) that were originally proposed decades ago can benefit from recent results on spintronic research. The project is a team effort involving device physicists from MPI and computer architects from cfaed. Although RTMs were proposed more than a decade ago, recent findings in the field of spin-orbitronics have successfully overcome major impediments to make them practicable. Thus, as the two main memory technologies used today to store data—Flash and magnetic hard disk drives—approach their fundamental physical limits, the article argues that RTMs that store digital data in the form of magnetic domain walls in magnetic nanowires could be the next solution. The article discusses recent developments enabling this progress and compares RTM to other memory and storage technologies.

The last article of this section involves p-bit computing and can be viewed as a nice Segway to the next section on quantum computing. Probabilistic bits (p-bits) conceptually lie somewhere in between classical deterministic bits and the q-bits of quantum computing, and has recently made its way into the exploration of a novel computing paradigm. The p-bits are classical entities fluctuating between 0 and 1, first implemented as spins in magnetic tunnel junctions (MTJs). While different from quantum computing, the technology promises to have some similarities, and may not suffer from some disadvantages of quantum computing. The p-bit computing paradigm makes use of the fact

that a bit may be in a 0 or 1 state with a certain probability somewhat like a q-bit, but unlike quantum computing it does not have the properties of superposition or entanglement. Although the first p-bit implementation was in terms of spintronic technology using MTJs, other implementations are also possible including, for example, using the existing silicon CMOS technology. One aspect that sets apart such circuits and systems based on probabilistic spin logics (PSLs) from quantum computing is that they do not require cryogenic temperatures for successful operation, and yet may address integer factorization and annealing-type problems, which have been targeted by emerging quantum computers. The p-bit technology can be a vehicle for efficient implementation of stochastic ML algorithms as well, thus providing a conceptual bridge between two active but disjoint fields of research (e.g., ML and QC). For the interested reader, more details on PSLs are available from the website of the group (https://www.purdue.edu/p-bit/).

The article titled "From charge to spin and spin to charge: Stochastic magnets for probabilistic switching," authored by Camsari *et al.*, the PSL group from Purdue University, undertakes a comprehensive review of the topic and makes future predictions regarding where the technology might lead to. In particular, it discusses in detail how various types of p-circuits can be realized from p-bits using MTJs, how p-circuits can be implemented in hardware, and points out the potential applications of such hardware for integer factorization, inference on Bayesian networks, etc.

E. Quantum Computing

A quantum computer is a machine in which information is represented by quantum states, so that computation proceeds according to the laws of quantum mechanics. A quantum computer can, therefore, exploit quantum resources such as superposition and entanglement in its computations, allowing it to solve many problems with fewer operations than a conventional computer which operates according to the laws of classical physics. For some important problems, a quantum computer can be exponentially more efficient than its conventional counterpart.

But quantum states are famously fragile. Although the time evolution of a perfectly isolated quantum system is deterministic, it becomes uncertain when interactions with the surrounding environment are considered. The state is said to decohere on a time scale that depends on the strength of its interactions with the noisy environment. In the last three decades, researchers have made strides in choosing comparatively robust quantum systems engineering them to be highly isolated from noise. Coherence times sufficient for the execution of many successive computational operations (with a useful chance of success) are now routinely achieved. This progress still leaves a vast gap between the capabilities today's quantum computers and the low error rates and high reliability that we expect from any conventional computer.

Fortunately, there are methods for correcting errors in quantum systems that are roughly analogous to error correction methods in conventional computers. Experimental demonstrations are limited in scope, but many in the research community believe that further development of error correction methods and circuits, combined with continued progress in combating decoherence, will eventually allow the construction of highly reliable general-purpose quantum computers. But as they pursue this long-term goal, they are also exploring the capabilities of the quantum computers already running in their labs—characterized by physicist John Preskill as "noisy intermediate scalable quantum systems." The results are very promising, as exemplified by the two articles in this section.

In "Challenges and opportunities of near-term quantum computing systems," Córcoles *et al.* with the IBM T. J. Watson Research Center review their team's structured approach to system architecture and design. Provisions are made for access to the hardware at various levels to meet the widely varying needs of users. A complete software developer's toolkit facilitates exploration of new applications by those who prefer not to be acquainted with the hardware. In a broad sense, the resulting hardware and software abstractions will be familiar to experts in information technology. In contrast, in "Resource-efficient quantum computing by breaking abstractions," Shi *et al.* explore the value of breaking some of the basic abstractions traditionally used in the design of computational hardware and software. (The authors' affiliations include University of Chicago, Princeton University, Duke University, IBM T. J. Watson Research Center, Amazon Web Services Center for Quantum Computing, and California Institute of Technology.) For a variety of cases, the ability to solve more difficult problems is demonstrated, but at some cost in user friendliness. Together, those two articles demonstrate that today's still-primitive quantum computers are on the verge of solving interesting problems that challenge the capabilities of conventional computers. The types of problems that can be usefully solved are currently limited, but progress quickening. ■

REFERENCES

[1] "Nonsilicon, non-von Neumann computing—Part I," *Proc. IEEE*, vol. 107, no. 1, Jan. 2019.

[2] S. Basu, R. E. Bryant, G. De Micheli, T. Theis, and L. Whitman, "Nonsilicon, non-von Neumann computing—Part I," *Proc. IEEE*, vol. 107, no. 1, pp. 11–18, Jan. 2019.

[3] D. Maslov, Y. Nam, and J. Kim, "An outlook for quantum computing [point of view]," *Proc. IEEE*, vol. 107, no. 1, pp. 5–10, Jan. 2019.

[4] N. R. Shanbhag, N. Verma, Y. Kim, A. D. Patil, and L. R. Varshney, "Shannon-inspired statistical computing for the nanoscale era," *Proc. IEEE*, vol. 107, no. 1, pp. 90–107, Jan. 2019.

[5] S. A. Wolf and D. M. Treger, "Scanning the issue—Special issue on spintronics," *Proc. IEEE*, vol. 91, no. 5, pp. 647–651, May 2003.

[6] H. Ohno, M. D. Stiles, and B. Dieny, "Spintronics," *Proc. IEEE*, vol. 104, no. 10, p. 1786, Oct. 2016.

[7] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proc. IEEE*, vol. 105, no. 12, pp. 2295–2329, Dec. 2017.

[8] S. Yu, "Neuro-inspired computing with emerging nonvolatile memories," *Proc. IEEE*, vol. 106, no. 2, pp. 260–285, Feb. 2018.

[9] L. Blum, F. Cucker, M. Shub, and S. Smale, *Complexity and Real Computation*. New York, NY, USA: Springer-Verlag, 1998.

[10] (2018). *National Quantum Initiative Act, Public Law No: 115-368, H.R.6227—115th Congress (2017-2018)*. [Online]. Available: https://www.congress.gov/bill/115th-congress/house-bill/6227

ABOUT THE GUEST EDITORS

Sankar Basu (Life Fellow, IEEE) is a Program Director of the National Science Foundation (NSF), Alexandria, VA, USA, in the Computer and Information Sciences and Engineering Directorate, where he is involved in the broad areas of computing including design automation for microand nano-electronic systems. Prior to NSF, he was with the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA. He has held visiting appointments at Ruhr University, Bochum, Germany, as an Alexander von Humboldt Fellow, and the Laboratory for Information and Decision Systems (LIDS), Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, for extended periods. During 2012, he served as an Embassy Science Fellow at the U.S. Embassy, Berlin, Germany, on detail to the U.S. Department of State. During his tenure at NSF, he has been involved with several interagency joint programs with, for example, the Semiconductor Research Corporation (SRC), DARPA, including those in response to the OSTP Grand Challenge in computing. He authored over one hundred refereed publications, a dozen patents, edited several volumes, and has given keynote talks at a number of international conferences. His scientific interests include nano-computing circuits, systems, signal processing, and machine learning.

Dr. Basu is a Fellow and a Council Member of the American Association of the Advancement of Science (AAAS). In 2011, the SRC recognized him with an award for enhancing its mission through collaboration. He co-chaired the OSTP Committee on Nanotechnology-Inspired Grand Challenge for Future Computing. He served on the editorial boards for about ten journals, including being the Editor-in-Chief of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS, IEEE Press, and served on the editorial board of the PROCEEDINGS OF THE IEEE from 2012 to 2018. He has organized and chaired, several major conferences, including the NATO Advanced Study Institute on Statistical Learning and Applications, resulting in a coedited book, *Advances in Learning Theory: Methods, Models, and Applications*.

Randal E. Bryant (Fellow, IEEE) received the B.S. degree in applied mathematics from the University of Michigan, Ann Arbor, MI, USA, in 1973, and the Ph.D. degree from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 1981.

He was an Assistant Professor at Caltech, Pasadena, CA, USA, from 1981 to 1984. He is currently a University Professor with the Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA. He has been on the faculty at Carnegie Mellon University since 1984, starting as an Assistant Professor and progressing to his current rank of the University Professor of Computer Science. He also holds a courtesy appointment with the Electrical and Computer Engineering Department. He served as the Dean of the School of Computer Science from 2004 to 2014. He is well known for the development of the ordered binary decision diagram (OBDD) data structure, used not just for formal hardware and software verification but also digital circuit testing and synthesis, AI planning, and combinatorial optimization. In addition, he has developed several techniques to verify circuits by symbolic simulation, with levels of abstraction ranging from transistors to models based on first-order logic. He teaches courses in computer systems. Along with David R. O'Hallaron, he developed a novel approach to teaching about the hardware, networking, and system software that comprise a system from the perspective of an advanced programmer, rather than from those of the system designers. He spent the 1989–1990 academic year as a Visiting Research Fellow at Fujitsu Laboratories Ltd., Kawasaki, Japan, and the 2014–2015 academic year as an Assistant Director of the Information Technology Research and Development at the White House Office of Science and Technology Policy, where together with Lloyd Whitman, he architected the OSTP Computing Grand Challenge. He has published the book *Computer Systems: A Programmer's Perspective* (Third Edition), is in use at over 320 universities worldwide, with translations into Chinese, Korean, Macedonian, and Russian. His research interests include the methods for formally verifying digital hardware and more recently some forms of software.

Dr. Bryant has received widespread recognition for his work. He is a Fellow of ACM as well as a member of the National Academy of Engineering and the American Academy of Arts and Science. His awards include the 1997 ACM Kanellakis Theory and Practice Award (shared with Edmund M. Clarke, Ken McMillan, and Allen Emerson) for contributing to the development of symbolic model checking, as well as the 1989 IEEE W. R. G. Baker Prize for the best paper appearing in any IEEE publication during the preceding year. He received the 2010 ACM/IEEE A. Richard Newton Technical Award in Electronic Design Automation, the 2009 EDAC/IEEE Phil Kaufman Award, and the 2007 IEEE Emmanuel R. Piore Award, reflecting the impact his work has had on tools used by the semiconductor industry for verifying their designs prior to manufacture.

Giovanni De Micheli (Fellow, IEEE) is currently a Professor and the Director of the Integrated Systems Laboratory, EPFL Lausanne, Lausanne, Switzerland. His current research interests include several aspects of design technologies for integrated circuits and systems, such as synthesis for emerging technologies.

Prof. De Micheli is a Fellow of ACM, a member of the Academia Europaea, and an International Honorary Member of the American Academy of Arts and Sciences. He was a recipient of the 2019 ACM/SIGDA Pioneering Achievement Award, the 2016 IEEE/CS Harry Goode Award for seminal contributions to design and design tools of networks on chips, the 2016 EDAA Lifetime Achievement Award, the 2012 IEEE/CAS Mac Van Valkenburg Award for contributions to theory, practice, and experimentation in design methods and tools, and the 2003 IEEE Emanuel Piore Award for contributions to computer-aided synthesis of digital systems. He also received the Golden Jubilee Medal for outstanding contributions to the IEEE CAS Society in 2000, the D. Pederson Award for the Best Paper on the IEEE TRANSACTIONS ON COMPUTER-AIDED DESIGN/INTEGRATED CIRCUITS AND SYSTEMS (CAD/ICAS) in 1987 and 2018, and several best article awards, including DAC in 1983 and 1993, DATE in 2005, Nanoarch in 2010 and 2012, and Mobihealth in 2016. He has been serving IEEE in several capacities, namely: Division 1 Director, from 2008 to 2009, a Co-Founder and the President Elect of the IEEE Council on EDA from 2005 to 2007, the President of the IEEE CAS Society in 2003, and the Editor-in-Chief of the IEEE TRANSACTIONS ON CAD/ICAS from 1997 to 2001. He has been the Chair of several conferences, including Memocode in 2014, DATE in 2010, pHealth in 2006, VLSI SOC in 2006, DAC in 2000, and ICCD in 1989.

Thomas Theis (Fellow, IEEE) is currently an IBM Distinguished Research Staff Member Emeritus with the IBM T. J. Watson Research Center, Ossining, NY, USA, and the Chief Strategist for Utopus Insights, Inc. He joined the IBM T. J. Watson Research Center in 1978 to study electronic properties of materials and held various senior management and executive positions from 1984 to 2015. In the late 1990s, as a Senior Manager, Silicon Science and Technology, he coordinated the transfer of copper interconnection technology from IBM Research to the IBM Microelectronics Division. The replacement of aluminum chip wiring by copper was a microelectronics industry first, and remains the worldwide standard. As IBM's Strategist for exploratory research worldwide from 1998 to 2012 and as the Director, Physical Sciences, from 1998 to 2010, he conceived and initiated successful research programs in silicon nanophotonics and Josephson junction-based quantum computing, and championed research in nanoelectronics, exploratory memory devices, and applications of information technology to address societal needs in energy, infrastructure, and the environment. From 2010 to 2012, he organized research projects aimed at greatly improved energy-efficiency in computing. On assignment from IBM from 2012 to 2016, he led the Semiconductor Research Corporation's Nanoelectronics Research Initiative, a private-public partnership funding university research. He joined Columbia University, New York City, NY, USA, in 2016, as a Professor of practice in electrical engineering and the Executive Director of the Columbia Nano Initiative. He retired from Columbia University in 2018 and joined Utopus Insights Inc., in 2019 under a long-term consulting contract.

Prof. Theis is a Fellow of the American Physical Society (APS). He was the 2015 recipient of the George E. Pake Prize of the APS.

Lloyd Whitman received the B.S. degree in physics from Brown University, Providence, RI, USA, and the M.S. and Ph.D. degrees in physics from Cornell University, Ithaca, NY, USA.

He is currently serving as an Assistant to the Director for Science Policy and Planning at the National Science Foundation (NSF), Alexandria, VA, USA, on detail from his position as the Chief Scientist with the National Institute of Standards and Technology (NIST), Gaithersburg, MD, USA. At NSF, he works to identify new opportunities for cooperation and collaboration between NSF and other organizations, coordinates the development and implementation of NSF-wide science policies, and advises the Director on strategic planning, assessment, and policy analysis. He spent most of his research career at the United States Naval Research Laboratory, Washington, DC, USA, where he led a diverse portfolio of research studying nanostructures and their integration into advanced sensor systems. He was subsequently the Founding Deputy Director of the Center for Nanoscale Science and Technology, NIST. He joined NSF in 2019 after serving for over five years at the White House Office of Science and Technology Policy (OSTP) in both the Obama and Trump administrations, most recently as the Principal Assistant Director for Physical Sciences and Engineering and a co-leader of the science portfolio. At OSTP, he oversaw a broad range of S&T areas and national research and development programs and strategies, including advanced materials, nanotechnology, and semiconductors. He has over 160 publications and multiple patents, and has been recognized with numerous media citations and awards.