# Guest Editorial for the Special Issue on Ultra-Low-Voltage VLSI Circuits and Systems for Green Computing

THE WIDE and increasing interest in green computing is not fully justified by the slowly growing sensibility of global citizens to energy-related and environmental issues. The interest in green computing actually reflects the recent societal demand for computing systems that fit both individuals' lifestyle and their interaction in professional/social/economic/interest-based networks. To fulfill the increasing demand for mobility and continuous connection, and thanks to major technical advances, computing systems have become more pervasive both in space and time. This major change in the evolution of the semiconductor industry is noticeable, for example, in the shrinking market share of personal computers in favor of smaller and always-connected devices. Also, green very large scale integration (VLSI) circuits and systems are key enablers of exciting, prospectively impactful and profitable applications such as implantable bio-medical networks, sensors for ambient intelligence, and so on.

To sustain this demand for mobility, pervasiveness, and ceaseless connection, form factor (space) and battery autonomy (time) will likely be the technology drivers of the semiconductor industry in the next decade, in lieu of more traditional figures like performance and functionality. In this context, the energy efficiency will play an even more crucial role, as it directly defines both the form factor and autonomy of computing systems through their energy sources.

After more than 20 years of intense research on low-power circuits and systems, voltage scaling is still a very effective approach to enhance the energy efficiency [1]–[5]. However, it is well known that aggressive voltage scaling poses many challenges that need to be overcome to push the energy envelope of mass-produced VLSI systems, as well as to enable robust operation close to the minimum energy point (MEP) [1]–[3]. Some of the most crucial challenges are related to: 1) the degradation of performance at ultra-low voltages; 2) the difficulty of ensuring best energy efficiency under a wide performance/voltage range; 3) the degradation of yield due to high sensitivity to process/voltage/temperature (PVT) variations and transient/soft errors, or equivalently the increased energy cost related to the additional design margin needed to ensure adequate yield; 4) the increased contribution of leakage to the overall energy budget; 5) the large design effort needed to meet performance/energy constraints under increased uncertainty and complexity, including preliminary design exploration, chip design, and verification [1]–[5].

All these challenges related to voltage/energy scalability need to be coherently addressed at all levels of abstraction, from the application and software level down to the circuit and possibly device level (although to a limited extent, due to economic considerations). The research trends clearly indicate that across-level thinking is key to successfully face those challenges. However, in the end, we still have a long way to go before design flows really embrace co-design at nonadjacent levels of abstraction and enable global energy-centric optimization of VLSI systems. At the same time, considerable energy efficiency improvements are expected to be achieved by leveraging run-time reconfiguration/adaptation and heterogeneity with specialized blocks for energy-hungry frequent tasks. Further energy efficiency improvements are certainly achievable through explicit exploitation of approximate computation under algorithmic noise tolerance, data locality-aware processing, and co-design of processing and power delivery.

In perspective, the main goal in green VLSI systems is the truly global energy minimization under widely varying time-dependent operating conditions and user expectations, which requires the availability of components whose energy can be widely scaled by trading it off for other assets (e.g., performance, signal quality) [4], [5]. In my vision of green VLSI Systems-on-Chip, predesigned and highly heterogeneous components (e.g., IPs) will be seamlessly integrated by letting them exchange information on the instantaneous requirements that they are expected to meet, as well as their energy state. This will require an additional specialized communication channel that we will call the "energy management channel" (EMC), which is asynchronous and has negligible throughput (and hence energy overhead). Through the EMC channel shown in Fig. 1, each module receives the instantaneous requirements that it must meet (e.g., performance, effective arithmetic precision—including occasional errors) and adjusts its internal knobs to minimize their energy under those constraints. At the same time, each component sends out information on its internal energy-related parameters (e.g., timing margin, bit error rate) that reflect internal PVT variations and support global decisions for energy minimization. Depending on the system complexity, dynamic requirements of individual components can be assigned by either a centralized energy manager or a hierarchical management scheme. I envision this approach as an efficient compromise between the need for global energy optimization under a wide range of conditions/uncertainty/applications and the ability to reuse individual components that scale and adapt

Color versions of one or more of the figures in this brief are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TCSII.2012.2231011

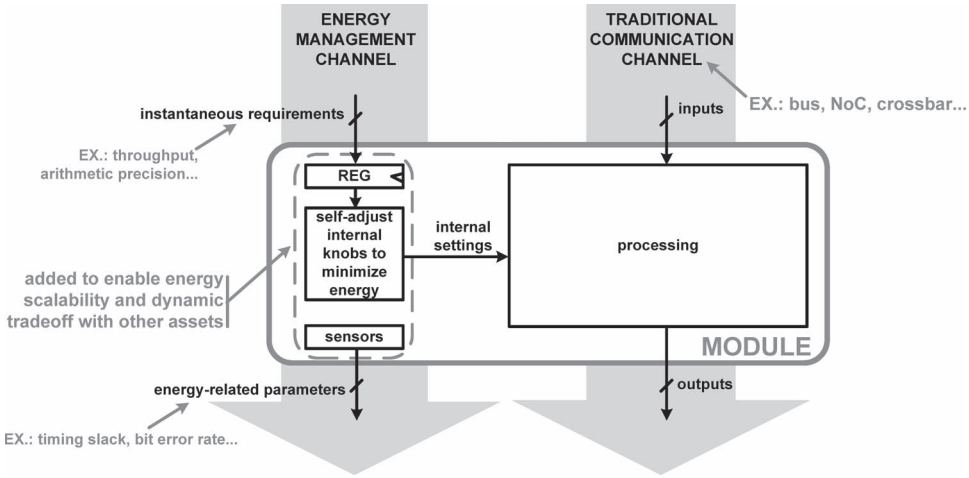


Fig. 1. Herein the envisioned scheme for energy-scalable reusable VLSI modules with the additional EMC to exchange information on the energy state and assigned requirements with power management.

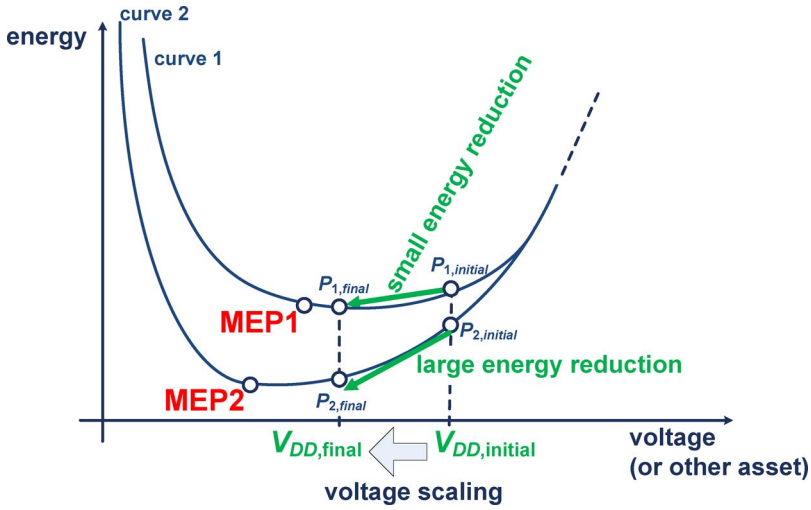


Fig. 2. Energy per operation of a generic VLSI module/system versus voltage (or other knob). The picture shows energy savings after voltage scaling in two cases: curve 1 (curve 2) has MEP close to (farther from) the operating points of interest. When scaling voltage from $V_{DD,\mathrm{initial}}$ to $V_{DD,\mathrm{final}}$, energy savings in curve 2 (from operating point $P_{2,\mathrm{initial}}$ to $P_{2,\mathrm{final}}$) are much larger than curve 1 (from $P_{1,\mathrm{initial}}$ to $P_{1,\mathrm{final}}$) thanks to the higher energy sensitivity to voltage. This explains why MEP needs to be controlled and pushed to bottom-left to keep practical operating points far from the flat minimum.

to the externally assigned requirements. In this way, overdesign and energy overhead are avoided both at the component level (thanks to the ability to self-tune for externally assigned requirements) and the system level (thanks to global optimization through information exchange on the EMC). As a further benefit of the above approach, the design exploration/design/verification effort is considerably reduced since each module can be reused in multiple designs/platforms.

As another direction that I believe needs to be pursued toward greener electronics, an explicit and better control of the MEP of VLSI systems is needed. For example, in the last decade, a great deal of attention has been devoted to MEP in the trend of energy per operation versus voltage, as the preferred operating point to maximize energy efficiency. From my perspective, we should recognize that the position of MEP actually affects the energy efficiency of a much wider range of operating points around the MEP. Indeed, energy tends to be pretty constant for operating points that are not too far from MEP (MEP is typically a flat minimum), so all operating points next to the MEP have basically the same energy as the MEP. This means that no or very little energy can be saved when scaling the voltage around those operating points, as shown in Fig. 2 (see curve 1 when voltage is scaled from $P_{1,\mathrm{initial}}$ to $P_{1,\mathrm{final}}$). On the other hand, better energy efficiency requires a large sensitivity of energy on voltage, so that voltage scaling can bring significant energy advantages. To keep such sensitivity high on a broad range of voltages at the right of MEP, it is necessary to push MEP toward lower optimum voltages, and not just toward lower energies (see curve 2 in Fig. 2 when voltage is scaled from $P_{2,\mathrm{initial}}$ to $P_{2,\mathrm{final}}$). This requires a better understanding of the MEP and how to control the knobs that affect its position at all levels of abstraction. Clearly, these considerations are immediately extended to the tradeoff between energy and any other parameter of interest.

This special issue aims to address many of the aforementioned challenges, with the ultimate goal of contributing to push the boundaries of energy consumption in VLSI circuits and systems. The 21 papers that form this special issue

were selected to meet the targeted acceptance rate of 32%, which matches the recent historical minimum of International Solid-State Circuits Conference (ISSCC). Each paper received between three and five reviews, with 3.9 reviews per paper on average. The papers are focused on ultra-low-voltage circuit and architectural level techniques, as well as design exploration and design flows to enable energy reduction in the near-threshold and sub-threshold domains. Eleven testchips are presented throughout the papers to demonstrate new ideas in ultra-low-voltage memories, clock generation/distribution, processing modules, interfaces, body sensors, and field-programmable gate arrays (FPGAs). Although this special issue focuses on digital or mostly-digital circuits, a paper on analog circuits is also included to discuss the feasibility of operation below the thermal voltage, to give a broader perspective on the ultimate limits to aggressive voltage scaling.

This special issue is organized by grouping papers by topic. The first group of papers addresses challenges related to the design of robust and energy efficient ultra-low-voltage SRAM memories and is, unsurprisingly, the most numerous. In particular, this group of papers is opened by the work of UCBerkeley/CEA-LETI on yield analysis of 28-nm arrays under assist techniques [6], as well as the work of MIT/Rochester on error correction [7]. Papers [8] (National Chiao Tung University), [9] (NTU-Singapore), and [10] (BGU) discuss of new techniques to improve the performance or the resiliency of SRAM arrays through novel bitline organization and bitcell topology. Paper [11] (National Chiao Tung University) analyzes the robustness of sensing schemes in FinFET technology and [12] (UBuffalo) proposes a novel error-aware array organization for voltage overscaling.

The second group of papers is about the design of resilient ultra-low-voltage building blocks including All-Digital Phase-Locked Loop (AD-PLL) [13] (National Central University/NTU/ITRI), adaptive clocking for enhanced resiliency to supply voltage droops [14] (Georgia Tech), logic styles for resilient operation in the sub-threshold [15] (KULeuven), current-mode logic for widely tunable power/performance tradeoff [16] (EPFL/EM), ultra-low-voltage register file [17] (CCU), error-correcting codes for resilient operation in the sub-threshold [18] (USU/Lund), the impact of temperature on random variations in sub-threshold logic [19] (UTokyo/STARC/VDEC), and level shifters for resilient wide voltage conversion ratio [20] (University of Calabria).

Then, [21] proposes a VLSI architecture based on L1 cache clusters for near-threshold computing (UniBO) and [22] (UFSC) discusses the potential of analog blocks such as oscillators and rectifiers for operation below the thermal voltage. The next two papers are oriented to applications: [23] (UVirginia/UW) is about reconfigurable ultra-low-voltage digital filtering for body sensor nodes and [24] (MIT/Northeastern) demonstrates latch-based sub-threshold FPGA.

The final two papers are about design methodologies for ultra-low voltages. In particular, [25] (KULeuven) proposes an efficient design flow for operation at the MEP. The special issue is closed by paper [26] (Umich), which introduces error-aware design methodologies for VLSI systems under voltage overscaling.

The contribution of this special issue can be summarized with a chromatic metaphor based on "dark silicon," i.e., the fraction of the die area that cannot be used to ensure that power does not exceed the maximum power rating [5], [27]. Dark silicon is commonly seen as a limitation to the utilization of available devices in the spatial dimension, i.e., only a fraction of the space within a power-limited high-performance chip is actually usable at a time. However, we should also observe that energy-limited mobile and ubiquitous systems also suffer from dark silicon, although in the temporal dimension: indeed, such systems cannot be used from time to time due to the limited available energy. Hence, dark silicon actually corresponds to the inability to take full advantage of the available devices (i.e., Moore's law) in both space and time, due to limited energy efficiency.

Chromatically speaking, this special issue aims to contribute to replace some "dark silicon" with "green silicon," thanks to the innovative ideas herein proposed by all authors. Our sincere hope is that green silicon will not be just hype, but an upcoming change of paradigm in mainstream technology and a definite change of mindset for users.

This special issue would not have been possible without the outstanding contribution of all authors, despite the very tight schedule imposed by the editorial process. I am deeply indebted with all reviewers, who constantly exceeded my expectations in performing a fair, deep, and accurate selection of papers within a short cycle. I am also grateful to TCAS-II Editor-in-Chief Prof. Yong Lian (National University of Singapore) for his highly responsive, wise, and constant support. Last but not least, I would like to thank Alberto Grosso for his tireless technical support.

MASSIMO ALIOTO, *Guest Editor*
EECS-University of Michigan
Ann Arbor, MI 48109-2122 USA
alioto@umich.edu

DII-University of Siena
53100 Siena, Italy
malioto@dii.unisi.it

## REFERENCES

[1] M. Alioto, "Ultra-low power VLSI circuit design demystified and explained: A tutorial," *IEEE Trans. Circuits Syst. I, Reg. Papers*, vol. 59, no. 1, pp. 3–29, Jan. 2012.

[2] R. G. Dreslinski, M. Wieckowski, D. Blaauw, D. Sylvester, and T. Mudge, "Near-threshold computing: Reclaiming Moore's law through energy efficient integrated circuits," *Proc. IEEE*, vol. 98, no. 2, pp. 253–266, Feb. 2010.

[3] A. P. Chandrakasan, D. C. Daly, D. F. Finchelstein, J. Kwong, Y. K. Ramadass, M. E. Sinangil, V. Sze, and N. Verma, "Technologies for ultradynamic voltage scaling," *Proc. IEEE*, vol. 98, no. 2, pp. 191–214, Feb. 2010.

[4] J. Rabaey, "The swarm at the edge of the cloud—A new perspective on wireless," in *VLSI Symp. Tech. Dig.*, Honolulu, HI, Jun. 2011, pp. 6–8.

[5] S. Borkar, "Truths and myths of embedded computing," in *Proc. 48th DAC*, San Diego, CA, Jun. 2011.

[6] B. Zimmer, S. O. Toh, H. Vo, Y. Lee, O. Thomas, K. Asanovic, and B. Nikolic, "SRAM assist techniques for operation in a wide voltage range in 28 nm CMOS," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 853–857, Dec. 2012.

[7] M. Zhang, V. Stojanovic, and P. Ampadu, "Reliable ultra-low voltage cache design for many-core systems," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 858–862, Dec. 2012.
[8] C.-Y. Lu, M.-H. Tu, H.-I. Yang, Y.-P. Wu, H.-S. Huang, Y.-J. Lin, K.-D. Lee, Y.-S. Kao, C.-T. Chuang, S.-J. Jou, and W. Hwang, "A 0.33 V, 500 kHz, 3.94 $\mu$W 40 nm 72 Kb 9 T subthreshold SRAM with ripple bit-line structure and negative bit-line write-assist," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, Dec. 2012.
[9] A. T. Do, Q. N. Truc, and T. T. Kim, "Sensing margin enhancement techniques for ultra-low voltage SRAMs utilizing bitline boosting current and equalized bitline leakage," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 868–872, Dec. 2012.
[10] A. Teman, A. Mordakhay, J. Mezhibovsky, and A. Fish, "A 40 nm sub-threshold 5 T SRAM bit cell with improved read and write stability," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 873–877, Dec. 2012.
[11] M.-L. Fan, V. P.-H. Hu, Y.-N. Chen, P. Su, and C.-T. Chuang, "Variability analysis of sense amplifier for FinFET subthreshold SRAM applications," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 878–882, Dec. 2012.
[12] N. Gong, S. Jiang, A. Challapalli, S. Fernandes, and R. Sridhar, "Ultra-low voltage split-data aware embedded SRAM for mobile video applications," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 883–887, Dec. 2012.
[13] K.-H. Cheng, J.-C. Liu, and H.-Y. Huang, "A 0.6-V 800-MHz all-digital phase-locked loop with a digital supply regulator," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 888–892, Dec. 2012.
[14] K. Chae and S. Mukhopadhyay, "All-digital adaptive clocking to tolerate transient supply noise in low voltage operation," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 893–897, Dec. 2012.
[15] N. Reynders and W. Dehaene, "Variation-resilient building blocks for ultra-low-energy sub-threshold design," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 898–902, Dec. 2012.
[16] A. Tajalli and Y. Leblebici, "Wide-range dynamic power management in low-voltage low-power subthreshold SCL," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 903–907, Dec. 2012.
[17] P.-Y. Chang, T.-J. Lin, J.-S. Wang, and Y.-H. Yu, "A 4 R/2 W register file design for UDVS microprocessors in 65 nm CMOS," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 908–912, Dec. 2012.
[18] C. Winstead and J. Neves Rodrigues, "Ultra-low-power error correction circuits—Technology scaling and sub-$V_{\mathrm{T}}$ operation," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 913–917, Dec. 2012.
[19] R. Takahashi, H. Takata, T. Yasufuku, H. Fuketa, M. Takamiya, M. Nomura, H. Shinohara, and T. Sakurai, "Large within-die gate delay variations in sub-threshold logic circuits at low temperature," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 918–921, Dec. 2012.
[20] M. Lanuzza, P. Corsonello, and S. Perri, "Low-power level shifter for multi supply voltage designs," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 922–926, Dec. 2012.
[21] M. Reza Kakoee, I. Loi, and L. Benini, "Variation-tolerant architecture for ultra low power shared-L1 processor clusters," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 927–931, Dec. 2012.
[22] C. Galup-Montoro, M. C. Schneider, and M. B. Machado, "Ultra-low-voltage operation of CMOS analog circuits: Amplifiers, oscillators, and rectifiers," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 932–936, Dec. 2012.
[23] A. Klinefelter, Y. Zhang, B. Otis, and B. H. Calhoun, "A programmable 34 nW/channel sub-threshold signal band power extractor on a body sensor node SoC," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 937–941, Dec. 2012.
[24] P. J. Grossmann, M. E. Leeser, and M. Onabajo, "Minimum energy analysis and experimental verification of a latch-based subthreshold FPGA," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 942–946, Dec. 2012.
[25] D. Bol, C. Hocquet, and F. Regazzoni, "A fast ULV logic synthesis flow in many-$V_t$ CMOS processes for minimum energy under timing constraints," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 947–951, Dec. 2012.
[26] D. Jeon, M. Seok, Z. Zhang, D. Blaauw, and D. Sylvester, "A design methodology for voltage overscaled ultra-low power systems," *IEEE Trans. Circuits Syst. II*, vol. 59, no. 12, pp. 952–956, Dec. 2012.
[27] H. Esmaeilzadeh, E. Blemz, R. S. Amant, K. Sankaralingamz, and D. Burger, "Dark silicon and the end of multicore scaling," in *Proc. ISCA*, San José, CA, Jun. 2011, pp. 365–376.

**Massimo Alioto** (M'01–SM'07) was born in Brescia, Italy, in 1972. He received the Laurea degree in electronics engineering and the Ph.D. degree in electrical engineering from the University of Catania, Catania, Italy, in 1997 and 2001, respectively.

In 2002, he joined the Department of Information Engineering of the University of Siena as a Research Associate and in the same year as an Assistant Professor. In 2005, he was appointed Associate Professor of Electronics. In the summer of 2007, he was a Visiting Professor at *EPFL*—Lausanne (Switzerland). In 2009–2011, he held a Visiting Professor position at BWRC—*UCBerkeley*, investigating next-generation ultra-low-power circuits, and wireless nodes. In 2011–2012, he also was a Visiting Professor at the University of Michigan, Ann Arbor, investigating active techniques for resiliency in near-threshold processors, error-aware VLSI design for wide energy scalability, ultra-low-power circuits, and energy scavenging. In 2013, he holds a visiting faculty position at the Intel Circuit Research Lab (CRL), Oregon.

He has authored or coauthored more than 170 publications in journals (more than 60, mostly IEEE, Transactions) and conference proceedings. Two of them are among the most downloaded *TVLSI* papers in 2007 (respectively, 10th and 13th). He is a coauthor of the book *Model and Design of Bipolar and MOS Current-Mode Logic: CML, ECL and SCL Digital Circuits* (Springer, 2005). His primary research interests include ultra-low-power VLSI circuits and wireless nodes, near-threshold circuits for green computing, error-aware and widely energy-scalable VLSI circuits, and circuit techniques for emerging technologies. He is the Director of the Electronics Laboratory at the University of Siena (site of Arezzo).

Prof. Alioto is a member of the *HiPEAC* network of excellence and was a member of the *MuSyC* FCRP center. In 2010–2012, he was the Chair of the "VLSI Systems and Applications" Technical Committee of the IEEE Circuits and Systems Society, for which he was also Distinguished Lecturer in 2009–2010 and a member of the DLP Coordinating Committee in 2011–2012. He serves or has served as a Track Chair for a number of conferences (ISCAS, ICCD, ICECS, APCCAD, ICM). He was the Technical Program Chair of the ICM 2010 and NEWCAS 2012 conferences. He serves as an Associate Editor of the IEEE TRANSACTIONS ON VLSI SYSTEMS, the *ACM Transactions on Design Automation of Electronic Systems*, IEEE TRANSACTIONS ON CAS—PART I, the *Microelectronics Journal*, *Integration—The VLSI Journal*, the *Journal of Circuits, Systems, and Computers*, the *Journal of Low Power Electronics*, as well as the *Journal of Low Power Electronics and Applications*.