

Introduction to Special Issue on Circuit Technology for ULP

By **ROBERT H. REUSS**, *Senior Member IEEE*
Guest Editor

MICHAEL FRITZE
Guest Editor

I. BACKGROUND

It is well known that device scaling as popularized by Moore's Law and systematized by the ITRS roadmap has led to a steady record of progress over the past decades. This trend has seen device performance steadily increase together with the number of transistors per chip, all the while keeping cost per chip essentially constant.

This relentless pursuit of higher performance and capability, has allowed the electronics industry to provide more functionality per chip at equal or even reduced cost. Only recently has the downside of this astounding progress begun to receive the attention it deserves. Specifically, the power consumed by integrated circuits has become an increasingly important issue over the last several years.

Initially driven by the need to increase "talk time" for cell phones and the "time off the plug" for portable laptops in the early 1990s, great efforts were made in power management of the electronics. These efforts paid tremendous dividends and adequate battery operation for all but the most over-worked road warrior became manageable. However, this success proved short-lived because it spawned applications and equipment with ever more demand for additional features and higher performance levels. The power per chip of some components has now reached the point that inability to dissipate the waste heat has created significant system design issues, with design for power management now becoming a dominant requirement.

In recent years, power requirements for even our many non-mobile electronic applications are receiving growing attention as part of the global warming and "green" concerns [1], [2] Current attention is primarily focused on power requirements for high performance computing centers and server "farms." Significant reduction in power dissipation for these applications is important, not only for the environmental factors associated with power generation, but also because of the difficulties in disposal of the waste heat. Finally, rising cost factors

associated with these issues are helping drive electronics towards a larger focus on more power efficient operation. While improvements in this area must be developed, the operational constraints of these "fixed" systems are in some ways not as demanding as those of "portable" electronics. The highest possible performance is essential for these "fixed" applications; however system mobility is not a requirement, so the quantity of, not access to, power is the major question. In contrast, applications such as mobile communications and unattended sensor networks must keep power consumption to an absolute minimum because there is frequently no readily available source of power.

The category of applications we have focused on for this Special Issue are those involving mobile/portable communications and sensor systems. While performance cannot be ignored in such applications (no one wants a cell phone that does not operate in real time!), the electronics requirements for these applications are frequently not as demanding as those for high performance computing. What is demanding is extending operational life as long as possible while operating on as few (or no) batteries as possible. While military applications are a major beneficiary of these requirements, other applications such as security systems and emergency responder's equipment have similar low power electronics needs. An example of how low power design methods to be described in this Special Issue might be applied to a military radio application is

This Special Issue focuses on ultralow power (ULP) solutions in chips and emphasizes power management design primarily in commercial CMOS applications for mobile/portable communications and sensor systems.

in [3], where it was found that dramatic power savings in the baseband signal processing of as much as 150X (depending on the waveform) might be possible. Extended operational life is also an important consideration for commercial applications, but not nearly so much because wall plugs are usually readily available (except it seems in airports where the demand is the highest!).

Driving requirements is not the same as driving technology. Commercial needs provide impetus for solutions, but rarely to the degree required for applications where a convenient source of power is not readily available. As others have described, the military/high performance market is now only a small percentage of the total semiconductor market. Thus, to keep cost within reason, commercial off the shelf (COTS) electronics is the preferred, if not mandated, solution for achieving ultralow power. Given the dynamics of the semiconductor industry, now and for the foreseeable future, COTS means CMOS. Therefore, in this Special Issue we emphasize ultralow power (ULP) solutions based on commercial CMOS.

II. EXTREME VOLTAGE SCALING FOR ULTRALOW POWER

Numerous efforts have been undertaken to manage the growing power problem; however, here the discussion will be limited to techniques involving significant voltage downscaling. This we define as ultralow power (ULP) operation. The emphasis of such approaches is based on a drastic reduction of chip power requirements, even if performance is therefore somewhat degraded. The challenge here is to scale the voltage as far as possible while adversely impacting performance for the intended application as little as possible.

From semiconductor physics it is known that power is proportional to the square of the operating voltage (V). This suggests that reducing voltage could have a significant impact on power consumption (e.g., scaling V from 1.2 V to 0.3 V might provide a > 10X power savings). Indeed, it does and can, up to a

point. But, as will be described in these Special Issue papers, doing so creates a host of problems. Beyond a certain point, power actually starts to increase because of leakage currents [4]. The other critical drawback in scaling voltage is the accompanying degradation of performance [4], which generally degrades linearly with voltage above the device threshold V_t (i.e., nominally on), and then exponentially as the device is operated in subthreshold (i.e., nominally off). So, the biggest power savings (subthreshold, approaching 100 mV bias) results in the biggest performance degradation. While some medical, human I/O and (well known) watch ULP chips are valuable in spite of slow operation in MHz (even KHz) regime, most applications cannot tolerate such low performance.

The challenge then is to achieve the advantage obtained by aggressive voltage scaling (operation at 120 mV could enable a 100X power reduction over standard 1.2 V CMOS), but with minimal compromise to the speed and robustness associated with commercial CMOS operating in the standard, super-threshold regime.

Due to the exponential change in current with voltage in the subthreshold regime, process and temperature caused variations result in significant performance degradation [5]. Significant variation even in standard, above-threshold devices is a major reason voltage scaling has slowed recently and is only expected to reduce by 2X (a maximum 4X in power savings) over the next decade.

If subthreshold operation holds the most promise for ultralow power, how can the serious issues with this approach be addressed such that performance degradation and other adverse aspects are minimized? No standard procedures for subthreshold design have yet emerged, but there are a variety of techniques currently being explored. Parallelism [6] is one approach to recovering lost performance. This is essentially trading area for power savings wherein more circuits are operating at much reduced power to recover some of the lost throughput.

This is essentially the approach used to reduce power in multi-core designs. We believe the authors of the papers herein have done an excellent job of describing the other challenges and additional potential solutions for enabling sub V_t , ULP operation.

The sequence of papers in this Special Issue is intended to flow beginning with best practices of leading edge commercial technology to an outlook of many of the approaches that might be applied in the future to achieve extreme voltage scaling for ULP operation. This is followed by a description of the application of some of these ULP methods to specific cases. There follows a series of papers that present methods to significantly scale voltage for power savings while maintaining high performance capabilities. Finally, several papers address specific solutions to overcome issues with subthreshold and ULP operation.

The first paper of the issue is from **Texas Instruments**. It provides the reader with an excellent background of the current status of technology scaling as applied to wireless device applications with an emphasis on the need to achieve process and circuit co-optimization to minimize power requirements. Techniques for reducing leakage currents under static conditions and dynamic power during operation are reviewed. Much more sophisticated power management techniques were developed as scaling moved from > 90 nm down to 45 nm. Many of the techniques described and issues resolved relate to lower voltage operation in either operational or standby mode. Such methods may not be generic, but must be carefully tailored relevant to specific functional requirements. Finally, a critical feature of ULP is also covered—how to do designs in a standard manner to minimize time and cost. Industry standardization of design methods will not only facilitate ULP designs, but also will enable new tools and methods to be integrated into the design flow as they are matured.

In contrast to what is available commercially, authors from **Purdue**

offer their view of opportunities in deeply voltage scaled operation. Their premise is that while sub Vt operation for ULP has many advantages, it also presents many challenges. Amongst the challenges is that device operation in the sub Vt regime is quite different than in super-threshold. These differences must be understood and device-circuit-architecture optimization strategies tailored differently than would be the case for operation at standard voltages. In addition, increased susceptibility of subthreshold circuits to process variations makes it imperative to consider design techniques to enhance their robustness at ultralow voltages. A further challenge for sub Vt system operation is not just logic, but also robust SRAM design and operation. Advances in this arena are also discussed.

The challenge for any technical approach is to apply it to practical applications. This is the thrust of the third paper from a team of **ULP researchers from MIT**. The challenge they tackled involved more than “simply” operating in the subthreshold regime. While ultralow-voltage operation has considerable challenges, even more difficult is dynamic voltage scaling during operation where voltage scales to provide just the required performance. To demonstrate the voltage-scalable capabilities of logic cells, SRAMs, ADCs, and DC-DC converters that they have developed, they report on the application of these circuits to two energy-constrained chip designs. One is a ULP video decoder that can support a 100X performance difference under various operational scenarios. The other is 0.3 V microcontroller with on-chip SRAM and DC-DC converter. This device can be used for sensor applications where the workload changes by as much as a factor of 80X.

While ultradynamic voltage scaling offers attractive features, as described by the MIT group, performance in the low power mode may also be critical. The next paper, from **IBM** focuses on approaches to improve power efficiency in high-performance parallel systems. Specif-

ically, rather than scale voltage into the sub-Vt regime and encounter the many problems this causes, a more modest scaling to the near-threshold region is proposed. About an 8X improvement in power efficiency can be attained with no loss in system-level performance by addressing appropriate device, circuit, and architecture considerations. The critical elements (logic, memory, I/O and power delivery) of a low-power computing system are explained in the context of appropriate design optimizations. This leads to a consideration of how such solutions might be applied to complex systems such as Blue Gene.

The next paper follows similar thinking; near-threshold operation provides power savings while mitigating performance losses. The insight the **UC Berkeley and UCLA** team brings to the problem is the need for a modeling framework that covers the entire transistor operation region from weak to strong inversion, enabling the understanding of the entire energy-delay space. Of particular relevance is the near-threshold region which has not been well modeled in conventional circuit applications. With accurate models, design optimization tools become possible. The authors then move to describe the use of such tools to identify and explore a logic family uniquely suited for near-threshold operation. Furthermore, the tools can also be used to explore various architectural optimizations in a more systematic and quantitative manner.

While Moore’s law continues to provide additional transistors, power budgets have begun to prohibit those devices from actually being used. The paper from ULP researchers at the **University of Michigan** defines the barriers to the widespread adoption of near-threshold operation and describes current work aimed at overcoming these obstacles. They focus on key challenges with respect to low voltage operation which include: 1) > 10X loss in performance, 2) 5X increase in performance variation, and 3) a drastic increase in functional failure rate of memory and

logic elements. Approaches to overcoming these barriers via a synergistic approach combining methods at the algorithm and architecture levels with circuit and technology levels are reviewed. In addition, solutions based on adaptive techniques in which processors and mixed-signal circuits dynamically adjust to meet the constraints imposed by process variation, changing environments, and aging are considered. The overall goal is the instantiation of various methods to provide robust, error-tolerant, self-correcting ULP circuits.

While the first set of papers describe the advantages and issues with voltage scaling and how far into the subthreshold regime it is practical to go, the next series of papers explores ULP concepts that address other aspects of ULP operation to provide more capabilities and functionality. Researchers from **University of Virginia** assert that a large range of ULP applications continue to face performance constraints at certain times that exceed the capabilities of subthreshold operation. In their paper, they describe two different examples of designing flexibility and adaptivity into ULP systems across the architecture and circuit levels to meet both ULP requirements as well as performance demands. Their approach expands on ultradynamic voltage scaling (UDVS) methods to combine multiple supply voltages with component level power switches to provide much more power efficient operation. Such an approach can enable practical circuit operation across the wide phase space from maximum performance to minimum energy. They then describe how such reconfigurable subthreshold circuits might be made applicable to ULP embedded systems such as a field programmable gate array (FPGA) to offer cost effective custom solutions across a wide range of applications.

An essential element of most electronic systems is some form of memory component. Several papers described earlier consider the issues associated with SRAMs operated in ULP mode. In the next paper the

Adesto Technologies team describes how non-volatile memory technologies might be applied to the ULP challenge. Discrete and embedded non-volatile memory (NVM) technologies have become an integral part of numerous personal media devices. Despite the expanded use of non-volatile memory, little has changed with respect to the core technology that hold the data when power has been turned off. Today, all of the non-volatile memory devices in production require high voltage, hardly consistent with the push towards ultralow power operation. Future ULP applications will likely require operation at 500 mV or less. Several emerging technologies are currently competing to become the next generation of non-volatile memory solution. This paper reviews the fundamental characteristics of current non-volatile memory technologies as well as several promising emerging technologies from a unique perspective of power requirements and specifically discusses the suitability of each one for use in ultralow power and subthreshold CMOS applications with particular attention given to conductive bridging memory (CBRAM) as a means to achieve low voltage operation consistent with ULP CMOS.

The focus of the preceding papers is how subthreshold operation can be exploited to stretch low power circuit designs beyond the normal modes of operation. However, because of significant device variation in this region (current is exponential with voltage), the uncertainty in timing closure generated by operating in subthreshold represents a major challenge that must be dealt with. A potential solution to overcome issues related to clocked logic is to simply eliminate the clock. Asynchronous logic has great potential in this regard, but as clocked design has been the standard for many years, considerable effort will be needed to bring clockless logic into the mainstream. In this paper **from Camgian and Boeing**, initial exploration of the benefits and challenges of asynchronous design for enabling sub-Vt operation is presented.

First, some background on clockless logic design approaches giving a brief overview of some of the characteristics of the different design styles and focusing on NULL Convention Logic is provided. Then, application of a clockless logic approach from a system perspective vs. the corresponding clocked approach is presented and the results of the comparisons summarized. An overview of additional research and development that is needed to make asynchronous design techniques available to subthreshold designers is also presented.

Work at **MIT** provides insight into how power requirements in the analog domain might be addressed. Traditionally, operational amplifiers (op-amps) have been widely used in many high-performance analog circuits. However, the requirements for effective performance, especially in scaled CMOS, make op-amp based circuits power inefficient. As an alternative to op-amp based circuits, comparator and zero-crossing based circuits have been developed. This paper discusses the technical issues unique to these A/D converters as well as solutions that have been developed to improve their performance and practicality. The ultimate low power potential of these A/D converters is compared with various different types of ADC's from fundamental standpoint.

By now, we hope to have made the point that Si CMOS technology based on ultralow-power transistors is an enabling technology for many important ULP applications. With an operating voltage of 0.3 V, and an on-current of less than $1 \mu\text{A}/\mu\text{m}$, subthreshold transistors use orders of magnitude less power than transistors operated in strong inversion. Further, since it is possible to have equal NMOS and PMOS drive currents per micrometer of transistor width, equal sizing of NMOS and PMOS transistors becomes possible. However, as discussed in several preceding papers, simply lowering the operating voltage of a conventional high-performance transistor will not produce optimal subthreshold operation. By designing a fabrication

process optimized for subthreshold performance from the starting substrate through the final metallization, it is possible to realize an optimum device solution with minimum switching energy and leakage, without significant impact to the energy-delay product. Unfortunately, the tremendous cost and lead-time associated with custom technology, from a practical perspective precludes anything other than an approach compatible with CMOS. Silicon on insulator (SOI) CMOS, especially fully-depleted (FDSOI) is in that category. Currently, SOI is in production for (ironically) high performance ICs such as custom processors. However, the characteristics of SOI are very attractive for also achieving excellent performance under sub-Vt conditions providing these devices are properly engineered for ULP operation. **Workers from MIT Lincoln Laboratory** provide an excellent discussion of the opportunities provided by FDSOI for ULP while highlighting the challenges to be overcome in developing the process technology for mainstream production.

The series of papers in this Special Issue is by no means complete. There are other important aspects of power reduction that are not addressed in this special issue. A major factor that must be considered in future ULP technologies, but that we were unable to cover in this Special Issue, is on-chip interconnect. Metal interconnect has become a significant issue for conventional technologies with on-chip optical interconnect receiving serious consideration for the future [7]. At this time it is not clear if sub Vt operation will require an alternative interconnect solution, but this may be the case.

We would be remiss if we did not mention the pioneering work in this field by Vittoz [8] and recent work that includes ISSCC 2009 papers such as Itoh's on Adaptive Circuits for 0.5 V Nanoscale CMOS [9], Borkar *et al.* on 300 mV Reconfigurable Dual-Supply, 4-Way SIMD Accelerator [10], Pu *et al.* on Ultra-low Energy JPEG Co-Processor with Sub/Near-Vt Power Supply [11], and workshops on ultralow

voltage circuit design. We believe this is a strong indication that although sub Vt is certainly not new, as power dissipation becomes an ever larger problem for demanding applications, new techniques for low voltage operation with at least modest performance will be an important part of a comprehensive ULP solution. ■

REFERENCES

- [1] M. LaPedus, "Green-memory movement takes root," *Electron. Eng. Times*, pp. 20–24, May 18, 2009.
- [2] R. Merritt, "Researchers eye road to energy-efficient electronics," *Electron. Eng. Times*, pp. 8–9, Jun. 22, 2009.
- [3] B. Fette, "Analysis of energy starved electronics—Power savings opportunities in software defined radios," in *GOMACTech—07*, Mar. 2007, pp. 87–90.
- [4] A. Wang, B. H. Calhoun, and A. P. Chandrakasan, *Sub-Threshold Design for Ultra Low-Power Systems*. New York: Springer, 2006.
- [5] J. Kwong, Y. Ramadass, N. Verma, and A. Chandrakasan, "A 65 nm sub-Vt

Acknowledgment

We would like to thank DARPA for its continued support of microelectronics research that identifies and pushes the boundary of challenging areas such as subthreshold operation for ULP. We also appreciate the interest by and assistance of the IEEE PROCEEDINGS staff in making this issue

possible. The authors obviously are most responsible for this work. It is not only their creativity but dedication that has made this activity so rewarding. Finally, we must offer our sincere thanks and appreciation to the many reviewers for their diligence, expertise, and insights that completed the work on this Special Issue.

- microcontroller with integrated SRAM and switched capacitor DC-DC converter," *IEEE J. Solid-State Circuits*, vol. 44, no. 1, pp. 115–126, Jan. 2009.
- [6] M. Horowitz and W. Dally, "How scaling will change processor architecture," in *IEEE Int. Solid-State Circuits Conf.*, Feb. 2004, pp. 102–106.
- [7] D. A. B. Miller, "Device requirements for optical interconnects to silicon chips," *Proc. IEEE*, vol. 97, no. 7, pp. 1166–1184, Jul. 2009.
- [8] E. A. Vittoz, "Weak inversion for ultimate low-power logic," in *Low-Power Electronics Design*, C. Piguet, Ed. Boca Raton, FL: CRC, 2005.
- [9] K. Itoh, "Adaptive circuits for the 0.5 V nanoscale CMOS era," in *IEEE Int. Solid-State Circuits Conf.*, 2009, pp. 142–143.
- [10] H. Kaul, M. A. Anders, S. K. Mathew, S. K. Hsu, A. Agarwal, R. K. Krishnamurthy, and S. Borkar, "A 300 mV 494GOPS/W reconfigurable, dual-supply, 4-way SIMD vector processing accelerator in 45 nm CMOS," in *IEEE Int. Solid-State Circuits Conf.*, 2009, pp. 260–261.
- [11] Y. Pu, J. P. de Gyvez, H. Corporal, and Y. Ha, "An ultra-low-energy/frame multi-standard JPEG co-processor in 65 nm CMOS with sub/near-threshold power supply," in *IEEE Int. Solid-State Circuits Conf.*, 2009, pp. 146–147.

ABOUT THE GUEST EDITORS

Robert H. Reuss (Senior Member, IEEE) was a DARPA Program Manager in the Microsystems Technology Office from 2001 to 2006. He was responsible for several research thrusts into fabrication of flexible, large area electronics including high mobility TFTs for digital and RF applications and organic photovoltaics, as well as conventional microelectronics efforts that included reconfigurable, multi-core processor design, asynchronous logic design methodology, and subthreshold, ultralow-power operation. Since September 2007, he has been an independent consultant. Prior to joining DARPA, Dr. Reuss spent twenty years in various technology and research management positions with Motorola. Earlier, he worked for the U.S. government as a research and development manager for seven years and was a Research Faculty member at the University of Colorado for three years. Dr. Reuss received a Ph.D. in Chemistry from Drexel University in 1971. He has published over 50 papers and has been awarded 13 U.S. patents. His technology interests lie in the area of application of materials and electrochemistry technologies for advanced microelectronic applications and microsystems integration as well as large area electronics.



Michael Fritze joined the DARPA Microsystems Technology Office (MTO) as a Program Manager in August 2006. His interests are in the areas of ultralow-power electronics, graphene-based RF-electronics, 3D integrated circuits and low volume nanofabrication. Currently, Dr. Fritze is responsible for the 3D Integrated Circuits (3-D IC), Steep-Subthreshold-slope Transistors for Electronics with Extremely-Low Power (STEEP), Radiation Hardening by Design (RHBD), Carbon Electronics for RF Applications (CERA), and Leading Edge CMOS access Program (LEAP).



Dr. Fritze has published over 60 papers and articles in professional journals and holds several U.S. patents.

Prior to joining DARPA, Dr. Fritze was a staff member from 1995 to 2006 at MIT Lincoln Laboratory in Lexington, Massachusetts, a Federally-Funded Research and Development Center (FFRDC). At Lincoln Laboratory, he worked on fully-depleted silicon on insulator (FDSOI) technology development with an emphasis on novel devices. Particular interests included highly scaled, tunneling-based, and ultralow-power devices. Dr. Fritze also worked in the area of silicon-based integrated optics. Another research interest at Lincoln Laboratory was in the area of resolution-enhanced optical lithography and nanofabrication with particular emphasis on low volume technological solutions.

Prior to Lincoln Laboratory, Dr. Fritze held a post-doctoral appointment at AT&T Bell Laboratories in Holmdel New Jersey, where he worked in the areas of femtosecond optical spectroscopy and compound semiconductor optical properties. From 1984 to 1986, Dr. Fritze worked as a technical assistant at Bell Communications Research in Murray Hill, New Jersey, in the area of compound semiconductor quantum well electrical properties.

Dr. Fritze received a Ph.D. in Physics from Brown University in 1994, working in the area of compound semiconductor quantum well physics. He received a B.S. in Physics in 1984 from Lehigh University. Dr. Fritze is an elected member of Tau Beta Pi and Sigma Xi. He is a member of IEEE and SPIE and is active on the program committees of the EIPBN (3Beams) & GOMAC conferences. Dr. Fritze has published over 60 papers and articles in professional journals and holds several U.S. patents.