Guest Editorial: Current Trends in Low-Power Design

NAEHYUCK CHANG

Seoul National University

and

JÖRG HENKEL

Karlsruhe Institute of Technology

ACM Reference Format:

Chang, N. and Henkel, J. 2010. Guest editorial: Current trends in low-power design. *ACM Trans. Des. Autom. Electron. Syst.* 16, 1, Article 1 (November 2010), 8 pages.

DOI = 10.1145/1870109.1870110. <http://doi.acm.org/10.1145/1870109.1870110>.

Low-power consumption of semiconductor devices, circuits, or systems is not only a constraint but a goal of the design process. Initially, low-power design mainly focused on dynamic power consumption. Later, leakage and standby power consumption became important as semiconductor scaled. Recently, low-power design has expanded its focus to thermal management and green computing. The range of low-power systems now includes power management of large-scale data centers, Grid-scale energy generation, and storage systems as well.

The International Symposium on Low-Power Electronics and Design (ISLPED) focuses on recent advances in all aspects of low-power electronics and design, ranging from process and circuit technologies, simulation and synthesis tools, to system-level design and optimization. This special section invited extended versions of distinguished papers at ISLPED 2009. Besides submissions from ISLPED 2009 regular papers, any other high-quality submissions that fit the topics of the special section were invited.

Authors' addresses: N. Chang, Seoul National University; J. Henkel, Karlsruhe Institute of Technology; email: henkel@informatik.uni-karlsruhe.de.

Permission to make digital or hard copies part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from the Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.
© 2010 ACM 1084-4309/2010/11-ART1 \$10.00 DOI: 10.1145/1870109.1870110.
<http://doi.acm.org/10.1145/1870109.1870110>.

ACM Transactions on Design Automation of Electronic Systems, Vol. 16, No. 1, Article 1, Pub. date: November 2010.

This special section consists of six articles covering a range of important issues such as nanometer MOSFET minimum-energy point, power gating considering negative bias temperature instability, low-power behavioral synthesis, memory architecture for 3D integration, a virtual machine for green computing, and firmware management of networked embedded systems.

Nanometer MOSFET Effects on Minimum-Energy Point

Lowering the supply voltage results in a quadratic reduction in dynamic power. However, as CMOS technology scaling increases leakage current, short-channel effects, and process variation, the longer gate delay caused by lowering supply voltage again results in extended operation time, which eventually increases leakage energy over the execution time. Therefore, there exists an optimum supply voltage that achieves the minimum energy per operation [Calhoun et al. 2005]. This minimum-energy point operation concept is widely applied to leakage-aware Dynamic Voltage Scaling (DVS) [Calhoun and Chandrakasan 2006; Zhai et al. 2005].

Along with high-performance applications, there are strong needs for ultra-low-power applications such as Radio-Frequency IDentification (RFID) tags, biomedical devices and sensor networks, fabricated in nanometer CMOS technologies [Kwong et al. 2009; Kaul et al. 2010; Pu et al. 2009]. Subthreshold logic fulfills ultra-low-power operation compromising speed performance [Soeleman and Roy 1999]. The minimum-energy operating point is extremely important due to the longer gate delay as well as large leakage power proportion.

It is widely known that variability cannot be neglected in nanometer technologies, especially for the threshold voltage variations due to Random Dopant Fluctuations (RDF). Consequently, variability strongly affects the minimum-energy operating point scaling trend [Bol et al. 2009]. The article entitled “Nanometer MOSFET Effects on the Minimum-Energy Point of Sub-45nm Subthreshold Logic - Mitigation at Technology and Circuit Levels” analyzes the effects that increase the minimum energy per operation in nanometer technologies. This article also proposes two directions for mitigating these new effects at 45nm node to keep the minimum-energy operating point under control. This enables the optimum MOSFET selection for circuit designers and fully depleted Silicon-On-Insulator (SOI) for technology developers. The results show that the use of low threshold voltage medium gate length MOSFETs in a 45nm technology yields 35% reduction in minimum energy per operation when compared with baseline MOSFETs.

NBTI and NBTI-Aware Design

Nowadays time-dependent degradation of a semiconductor is much more difficult to handle than process variation because the deviation cannot be detected at the manufacturing stage, but it develops later. Bias Temperature Instability (BTI) results in the degradation of the oxide, thus causing a drift of the threshold voltage over time. The threshold voltage drift again results in performance change and leakage power variation. Negative BTI (NBTI) affects

PMOS transistors while Positive BTI (PBTI) affects NMOS devices. Logic “1” results in PMOS transistor degradation due to NBTI while logic “0” partly recovers the degradation process. Unfortunately, NBTI eventually incurs irreversible effects that cause reliability issues because the recovery process is limited. As current technologies mostly keep PMOS transistors negatively biased, thus NBTI gives more direct impact on reliability than PBTI. NBTI results in increase of the threshold voltage over time and consequent decrease in drain current and transconductance. The degradation exhibits logarithmic dependence on time, and experimental data reports variation of the threshold voltage of about 10–15% per year.

Overall, it is very common to observe highly biased inputs for some PMOS transistors, which will degrade them very quickly [Abella et al. 2007]. Therefore, it would be possible to limit NBTI-induced aging by keeping most of the circuit signals to the logic “1” state, by proper implementation of the local functions and/or by isolation of nodes with high zero probability during technology mapping [Kumar et al. 2007]. Minimum leakage vectors lead to minimum circuit performance degradation and maintain maximum leakage reduction rate during the standby mode [Wang et al. 2007].

NBTI also affects reliability of sequential circuits. NBTI tightens the setup and hold timing constraints imposed on the flip-flops. Different types of flip-flops exhibit different levels of susceptibility to NBTI-induced change in their setup/hold time values. An elaborated NBTI-aware transistor sizing technique can minimize the NBTI effect on timing characteristics of the flip-flops [Abrishami et al. 2008]. Furthermore, memory-like structures have a special characteristic in NBTI degradation. As bit cells consist of two inverters arranged in a ring manner, there is always one of the inverters with negative voltage (logic input “0”). Thus, PMOS transistor NBTI degradation is unavoidable. Instead of elimination of the chances of the PMOS degradation in memory structures, the second best case degradation happens when the value at the output of each inverter is “0” 50% of the time [Abrishami et al. 2008].

The article entitled “NBTI-Aware Clustered Power Gating” proposes a new solution for concurrently optimizing power consumption (static power, in particular) and aging. This article introduces power gating to mitigate NBTI. The authors observe that the conventional way of implementation of power gating for leakage reduction provides a natural way of reducing the NBTI effects. When a circuit is in a standby state, it is intrinsically immune to NBTI-induced aging. However, although power gating is always advantageous from the aging perspective, its actual effectiveness in terms of performance depends on its implementation.

The authors elaborate a clustered power gating strategy which allows a better control of the initial delay degradation. The authors implement trade-off analysis with an automated tool that leverages standard EDA tools in an industry-strength design flow, using an industrial 45nm technology. Results on a set of standard benchmarks show average lifetime extensions of about 2.2X while saving, on average, 87% of leakage power with respect to the nonpower-gated circuit.

Low-Power Behavioral Synthesis

Many low-power techniques in RTL synthesis rely on Observability Don't-Care (ODC) conditions. ODC conditions effectively identify unnecessary operations in a Boolean network [De Micheli 1994; Devadas et al. 1994; Hassoun and Sasao 2002]. ODC conditions are widely used to eliminate unnecessary operation of digital logics. ODC conditions are useful for power-friendly RTL descriptions in behavioral synthesis. Operation gating explicitly adds predicate to an operation based on its ODC condition so that the operation can be avoided using RTL power optimization techniques such as clock gating.

The problem of computing ODC conditions in a sequential RTL model has been approached in a number of ways. The article entitled “Behavior-Level Observability Analysis for Operation Gating in Low-Power Behavioral Synthesis” introduces the concept of behavior-level observability and its approximations in the context of behavioral synthesis. This article also describes an efficient procedure to compute an approximated behavior-level observability of every operation in a dataflow graph. Unlike previous techniques which work at the bit level in Boolean networks, the proposed method is able to perform analysis at the word level, and thus greatly reduces the computation overhead with a reasonable approximation.

The proposed algorithm exploits the observability-masking nature of some Boolean operations, as well as the select operation, and allows certain forms of other knowledge, once uncovered, to be considered for stronger observability conditions. The approximation is proved optimal for (acyclic) dataflow graphs when non-Boolean operations other than select are treated as black boxes. The behavior-level observability condition obtained by the algorithm can be used to guide the operation scheduler to optimize the efficiency of operation gating, demonstrating an average of 33.9% reduction in total power.

3D Architectures

Migrating to new silicon technology nodes provided in the past advantages in many design criteria. However, silicon scaling will eventually come to an end as feature sizes reach physical limits. A promising means to sustain scaling for some while is integration in 3D, that is, utilizing the third dimension in form of, for example, stacked layers of silicon dies with direct vertical tunneling [Xie et al. 2006; Loh et al. 2007]. While 3D architectures promise to provide advantages in some key architectural concerns like decreased latency in on-chip communication (since layer-to-layer distance is very short), a very high integration density, and an additional dimension providing new opportunities for routing signals on a chip, etc., disadvantages also exist. The most important disadvantage among others is the dissipation of heat since the ratio of number of transistors per surface area (i.e., higher thermal densities) is far smaller than in conventional 2D architectures. At the current state, first 3D architectures have been successfully built and CAD design tool chains are increasingly being developed in order to address the specific problems associated with 3D and exploit their advantages. The article entitled “Low-Power Hypercube Divided Memory FFT Engine Using 3D Integration” aims at

exploiting the architectural advantages of 3D integration. The article deals with an FFT (Fast Fourier Transformation) processor with a hypercube memory architecture to reduce power consumption and to decrease latency. The deployed architecture is a memory-on-logic design comprising TSVs (Through Silicon Vias). The basic idea of the approach is to divide the whole memory into smaller memories in order to avoid the long wires typically associated with large memories. Each small memory block can thereby be accessed simultaneously. As a result, the authors report benefits in almost all relevant design aspects compared to a 2D implementation. In particular, it is reported that memory access energy is significantly reduced, the total wire length is less than half, and the whole architecture can be implemented onto a significantly smaller die area.

Green Computing and Green Virtualization

Low-power consumption is no longer a critical design consideration of battery-operated systems only. The Environmental Protection Agency (EPA) report states that the energy consumption of servers and data centers has doubled in the past five years, and is expected to quadruple in the next five years to more than 100 billion kWh at a cost of about \$7.4 billion annually [Pakbaznia and Pedram 2009]. It goes without saying that power reduction of data center is crucial even though it is powered by the Grid. There are a number of different techniques currently employed to reduce the energy cost and power density in data centers.

Due to the high power density of the data centers, about 30% of the total energy cost of a data center is used for cooling [Pelley et al. 2009; Rasmussen 2007]. The hot-aisle/cold-aisle structure, which has become common practice these days, is one of the attempts to improve the cooling efficiency of data centers [Patel et al. 2002]. About 10–15% power loss is due to power distribution and conversion losses in the data center [Pelley et al. 2009]. DC power delivery systems are viable, can be 20% or more efficient than current AC delivery systems, can be more reliable, and potentially cost less in the long run [Ton et al. 2008].

Recently, it turns out that scheduling and consolidation are efficient data center power management schemes without infrastructure investment and/or modification. There are numerous prior works on temperature-aware task scheduling [Moore et al 2005; Tang et al. 2008], which can reduce the server temperature and thus save cooling cost. Server consolidation, which refers to assigning incoming tasks to the minimum number of active servers in the data center and shutting down unused servers, is another approach for power reduction of data centers [Intel 2006]. Consolidation of VMs originally distributed in many different servers into a minimum number of servers enables effective power management.

The article entitled “vGreen: A System for Energy-Efficient Management of Virtual Machines” introduces vGreen, a multitiered software system to manage VM scheduling across different physical machines with the objective of managing the overall energy efficiency and performance. vGreen

exploits the relationship between the architectural characteristics of a VM such as instructions per cycle, memory accesses, etc., and its performance and power consumption. This method is elaborated to avoid misleading the VM management policies into making decisions that can create hot spots of activity, and degrade overall performance and energy efficiency.

This article shows that vGreen is able to dynamically characterize VMs, and accurately models their resource utilization levels. Based on these characteristics, it can intelligently manage the VMs across the physical cluster by issuing VM migration or DVFS commands. This improves the overall performance up to 100% and energy efficiency up to 55% compared with state-of-the-art VM scheduling and power management policies.

Networked Embedded Systems

Networked embedded systems have gained significant relevance within the last decade. Application areas are manifold and span from automotive embedded systems for safety-critical functions to wireless sensor networks that, for instance, monitor environmental data from a network of thousands of nodes distributed throughout many square miles. In many networked embedded systems, design for low power is a prime constraint. In the automotive field, for example, a low-power design of the ECUs (Electronic Control Units) allows for a small dimensioning of the battery, which in turn leads to a reduced carbon dioxide footprint through a reduced weight that needs to be moved by the vehicle. In wireless sensor networks on the other side, low-power consumption is often vital for the entire lifetime of the system since a limited amount of energy might have been provided in form of a primary battery. Hence, when the battery state of charge is lost, the system might not be used any further since recharging is too expensive and energy harvesting [Raghunathan and Chou 2006] is often not possible. The article entitled “Energy-Efficient Progressive Remote Update for Flash-Based Firmware of Networked Embedded Systems” deals with a novel and efficient technique to update the firmware of networked embedded systems. The background is that the process of updating the firmware is expensive, both in terms of (limited) bandwidth provided by the system and the energy that is associated with updating the firmware. Besides the energy consumed by the radio part, it also refers to the energy associated with the flash memory that typically requires to erase/overwrite a whole page before writing to it. This work therefore introduces a technique, applied at link time of the firmware code, that reduces the number of caller-to-callee dependencies crossing memory page limits and hence it reduces the before-mentioned erasing cycles of entire pages of the flash memory. The approach encompasses two techniques: (i) grouping functions into one page that have a high caller-to-callee dependency and (ii) reordering the function calls within a certain page. In addition, the authors apply a sophisticated patching method of the firmware code that reduces the otherwise necessary erasing/overwriting of entire pages. As a result, the authors report significant energy savings compared to current state-of-the-art in flash-based firmware updates.

REFERENCES

- ABELLA, J., VERA, X., AND GONZALEZ, A. 2007. Penelope: The nbti-aware processor. In *Proceedings of the 40th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO 40)*. IEEE Computer Society, 85–96.
- ABRISHAMI, H., HATAMI, S., AMELIFARD, B., AND PEDRAM, M. 2008. Nbti-Aware flip-flop characterization and design. In *Proceedings of the 18th ACM Great Lakes Symposium on VLSI (GLSVLSI'08)*. ACM, New York, 29–34.
- BOL, D., AMBROISE, R., FLANDRE, D., AND LEGAT, J.-D. 2009. Interests and limitations of technology scaling for subthreshold logic. *IEEE Trans. VLSI Syst.* 17, 10, 1508–1519.
- CALHOUN, B. AND CHANDRAKASAN, A. 2006. Ultra-dynamic voltage scaling (udvs) using subthreshold operation and local voltage dithering. *IEEE J. Solid-State Circ.* 41, 1, 238–245.
- CALHOUN, B., WANG, A., AND CHANDRAKASAN, A. 2005. Modeling and sizing for minimum energy operation in subthreshold circuits. *IEEE J. Solid-State Circ.* 40, 9, 1778–1786.
- INTEL. 2006. Server consolidation using quad-core processors. <http://www.intel.com/it/pdf/consolidate-using-quadcore.pdf>.
- KAUL, H., ANDERS, M., MATHEW, S., HSU, S., AGARWAL, A., KRISHNAMURTHY, R., AND BORKAR, S. 2010. A 300 mv 494gops/w reconfigurable dual-supply 4-way simd vector processing accelerator in 45 nm cmos. *IEEE J. Solid-State Circ.* 45, 1, 95–102.
- KUMAR, S. V., KIM, C. H., AND SAPATNEKAR, S. S. 2007. Nbti-Aware synthesis of digital circuits. In *Proceedings of the 44th Annual Design Automation Conference (DAC'07)*. ACM, New York, 370–375.
- KWONG, J., RAMADASS, Y., VERMA, N., AND CHANDRAKASAN, A. 2009. A 65 nm sub-vt microcontroller with integrated sram and switched capacitor dc-dc converter. *IEEE J. Solid-State Circ.* 44, 1, 115–126.
- LOH, G. H., XIE, Y., AND BLACK, B. 2007. Processor design in 3d die-stacking technologies. *IEEE Micro* 27, 3, 31–48.
- MOORE, J., CHASE, J., RANGANATHAN, P., AND SHARMA, R. 2005. Making scheduling “cool”: temperature-aware workload placement in data centers. In *Proceedings of the Annual Conference on USENIX Annual Technical Conference (ATEC'05)*. USENIX Association, 5–5.
- PAKBAZNIA, E. AND PEDRAM, M. 2009. Minimizing data center cooling and server power costs. In *Proceedings of the 14th ACM/IEEE International Symposium on Low-Power Electronics and Design (ISLPED'09)*. ACM, New York, 145–150.
- PATEL, C., SHARMA, R., BASH, C., AND BEITELMAL, A. 2002. Thermal considerations in cooling large scale high compute density data centers. In *Proceedings of the 8th Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM'02)*. 767–776.
- PELLEY, S., MEISNER, D., WENISCH, T. F., AND VANGILDER, J. W. 2009. Understanding and abstracting total data center power. In *Proceedings of the Workshop on Energy Efficient Design (WEED'09)*.
- PU, Y., DE GYVEZ, J., CORPORAAL, H., AND HA, Y. 2009. An ultra-low-energy/frame multi-standard jpeg co-processor in 65nm cmos with sub/near-threshold power supply. In *IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC'09)*. 146–147, 147a.
- RAGHUNATHAN, V. AND CHOU, P. H. 2006. Design and power management of energy harvesting embedded systems. In *Proceedings of the International Symposium on Low-Power Electronics and Design (ISLPED'06)*. ACM, New York, 369–374.
- RASMUSSEN, N. 2007. Calculating total cooling requirements for data centers. American Power Conversion, White Paper #25.
- SOELEMEN, H. AND ROY, K. 1999. Ultra-Low power digital subthreshold logic circuits. In *Proceedings of the International Symposium on Low-Power Electronics and Design (ISLPED'99)*. 94–96.
- TANG, Q., GUPTA, S. K. S., AND VARSAMOPOULOS, G. 2008. Energy-Efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach. *IEEE Trans. Parall. Distrib. Syst.* 19, 11, 1458–1472.

- TON, M., FORTENBERY, B., AND TSCHUDI, W. 2008. *DC Power for Improved Data Center Efficiency*. Lawrence Berkeley National Laboratory.
- WANG, Y., LUO, H., HE, K., LUO, R., YANG, H., AND XIE, Y. 2007. Temperature-aware nbti modeling and the impact of input vector control on performance degradation. In *Proceedings of the Conference on Design, Automation and Test in Europe (DATE'07)*. 546–551.
- XIE, Y., LOH, G. H., BLACK, B., AND BERNSTEIN, K. 2006. Design space exploration for 3d architectures. *J. Emerg. Technol. Comput. Syst.* 2, 2, 65–103.
- ZHAI, B., BLAAUW, D., SYLVESTER, D., AND FLAUTNER, K. 2005. The limit of dynamic voltage scaling and insomnia dynamic voltage scaling. *IEEE Trans. VLSI Syst.* 13, 11, 1239–1252.