

# BIG CHIPS


..... For many years, classical silicon scaling (that is, Dennard scaling) provided three benefits: doubling of integration capacity to deliver twice the number of transistors per unit area every two years at constant price, increased transistor performance, and reduced energy per transistor. With the steady erosion of the latter two benefits, the low-hanging fruit for performance increase has been to exploit integration capacity via multithreading, multicore, and now many-core. Indeed, a basic premise of multicore has been that, by scaling supply voltage and using more transistors, one can deliver improved throughput with lower power. On the other hand, the erosion of energy per transistor raises the specter of *dark silicon*: with constant power budgets but increasing power densities, an ever-smaller portion of the chip can be on at any given moment. This suggests that architects can shrink die sizes, rather than integrate unusable transistors, as transistor budgets meet hard power limits across both mobile and server product spaces.

Why, then, does the industry still produce big chips? In a personal communication, Shekhar Borkar, an Intel Fellow, notes the steady stream of creative ways in which integration capacity has been exploited: aggressive supply-voltage scaling with more transistors, new paradigms in computation such as special-purpose hardware and accelerators for higher power efficiency, and implementation of more functions on a chip than can be afforded power-wise, along with fine-grained power management to activate selectively what can be afforded at any given

point in time. Borkar concludes, "All of this will require tremendous transistor integration capacity, and there will be no dark silicon. Therefore, big chips will continue to have a future."

Indeed, recent chip sizes from IBM and Intel have been in the 500 to 600 mm<sup>2</sup> range, exceeding the values posited in recent editions of the *International Technology Roadmap for Semiconductors*. At the 45-nm technology node, with eight cores per socket, IBM's Power7 chip is 567 mm<sup>2</sup>, and Intel's Nehalem-EX chip is 684 mm<sup>2</sup>. Tileria has integrated 64 cores in a single chip, and Intel's research vehicle, the Single Chip Cloud, has 48 cores. As the semiconductor industry strives to keep Moore's law on track with respect to utility in future processor products, futures that avoid "big" (transistor counts) and "many" (cores) are difficult to envision. Thus, a fundamental challenge for the future will be the intelligent use of silicon area to maximize performance while keeping power and packaging costs within budget.

As guest editors, we are privileged to introduce this *IEEE Micro* special issue on Big Chips. This special issue aims to share some of the trade-offs inherent in big chips with respect to performance capability, power density challenges, scalability of communication, and packaging and cooling costs. This issue also explores new technologies that can achieve the effect of big chips, such as 3D integration, along with associated trade-offs. Submissions were solicited across a broad range of issues involved in building and sustaining big chips in the future,



**Andrew B. Kahng**  
University of California  
at San Diego

**Vijayalakshmi Srinivasan**  
IBM T.J. Watson  
Research Center

including how to partition the chip between caches and CPU, network scalability for inter-core communication, impact on reliability and sparing of cores and cache, impact of yield of different configurations, and system integration with heterogeneous special-purpose cores, accelerators, and field-programmable gate arrays (FPGAs). The articles in this issue cover only a small subset of challenges related to building effective big chips. Nevertheless, we hope that this issue will lead to further dialogue and contemplation that will help define a roadmap for future big chips.

### Models and fundamental challenges for big chips

In “Toward Dark Silicon in Servers,” Nikos Hardavellas et al. begin with the premise of dark silicon—that server chips will be inherently unscalable beyond several hundreds of cores, with mounting fractions of dark silicon that can't be powered up. The authors propose a combination of heterogeneous, specialized compute engines in combination with bandwidth-mitigating approaches to break the near-term performance energy-efficiency barriers.

In “Scaling with Design Constraints: Predicting the Future of Big Chips,” Wei Huang et al. use an analytical model to explore the scalability of current multicore designs based on design constraints including voltage/frequency scaling, thermal design power (TDP), core count growth, and increase in cache size. They argue for architectural innovations to bridge the gap between promised throughput due to technology scaling, and the TDP-constrained throughput realized in practice. They predict adoption of new packaging solutions, including 3D integration to realize throughput growth.

### Scalable big chips

In “Rigel: A 1,024-Core Single-Chip Accelerator Architecture,” Daniel Johnson et al. target a broad class of data- and task-parallel applications with Rigel, a single-chip 1,024-core accelerator. The authors walk us through the evolution of Rigel and highlight the challenges of scaling the memory systems and maintaining coherent caches. The article

closes with a discussion of future opportunities and design challenges for such large-scale designs.

“MOPED: Accelerating Data Communication on Future CMPs” by Junli Gu et al. presents the challenges of data movement within many cores in a big chip. MOPED is an explicit hardware communication mechanism that enables off-loading synchronization and communication from CPUs. Detailed simulation results are presented to study MOPED's benefits in reducing effective memory latency and network latency and in effectively overlapping computation with communication to improve overall performance.

### Design methodologies for big chips

In “Physical Synthesis with Clock-Network Optimization for Large Systems on Chips,” David Papa et al. describe the deficiencies of traditional physical-synthesis methodologies and present a new methodology to improve timing closure through clock-network synthesis and careful placement of flip-flops and latches. The authors demonstrate the proposed methodology's effectiveness on IBM's large CPU designs and document significant improvements in timing, wire length, and area.

### Enabling technologies for big chips

“Attaining Single-Chip, High-Performance Computing through 3D Systems with Active Cooling” by Ayse K. Coskun et al. makes a case for efficient 3D liquid-cooled systems based on design-time thermal analysis with respect to flow rate, dynamic voltage and frequency scaling (DVFS), and job scheduling decisions. The authors use an integrated controller that monitors temperature and core utilization and tunes the control knobs to dynamically set the DVFS setting of each core and the coolant flow rate of the chip, based on a set of predefined rules.

This special issue provides a snapshot and a sampling of the tremendous activity going on in the area of design of big chips. It is useful for microarchitects to understand the limitations and challenges of present technologies as they explore new ideas to realize scalable big chips in the future.

We hope that you will enjoy the selected articles, and we encourage you to provide feedback on this issue.

MICRO

## Acknowledgments

We are grateful to be able to present such a strong collection of ongoing work related to the challenges of designing big chips. We thank all the authors who submitted articles, and we thank the reviewers for providing detailed and constructive comments during the review process.

**Andrew B. Kahng** is a professor of CSE and ECE at the University of California at San Diego. His research interests include integrated-circuit physical design and performance analysis, the integrated-circuit design-manufacturing interface, combinatorial algorithms and optimization, and the road-mapping of systems and technology. Kahng holds a PhD in computer science from the

University of California at San Diego. He is an IEEE Fellow.

**Vijayalakshmi Srinivasan** is a research staff member at the IBM T.J. Watson Research Center. Her research interests are in computer architecture, especially processor microarchitecture, and multicore and multiprocessor memory systems. Srinivasan holds a PhD in electrical engineering and computer science from the University of Michigan. She is a senior member of IEEE.

Direct questions and comments about this article to Andrew B. Kahng at [abk@cs.ucsd.edu](mailto:abk@cs.ucsd.edu) or to Vijayalakshmi Srinivasan at [viji@us.ibm.com](mailto:viji@us.ibm.com).



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.



**ONLINEPLUS™**  
publishing evolved

A new publication model that will provide subscribers with features and benefits that cannot be found in traditional print such as:

- More Rapid Publication of Research
- Online Access to the CSDL
- Interactive Disk and a Book of Abstracts
- Lower Price

Available Transactions Titles by 2012:

- TDSC
- TMC
- TPAMI
- TPDS
- TVCG

For more information about OnlinePlus™, please visit <http://www.computer.org/onlineplus>.

IEEE computer society