NSF Workshop on Energy-Efficient Data Management, May 2011.

Exploiting Dark Silicon for Energy Efficiency

Nikos Hardavellas
Northwestern University, EECS, nikos@northwestern.edu

EXECUTIVE SUMMARY

Application datasets grow faster than Moore's Law [7,8], both in personal and desktop computing, as well as in the scientific domain. For example, during March 2011, over 1.6PB of data were generated and transferred for processing among the Tier-1 sites of CERN's Large Hadron Collider computing grid [3]. Processing all these data requires unprecedented amounts of computation, with commensurate energy demands. Computing's insatiable appetite for energy has created an IT industry with approximately the same carbon footprint as the airline industry, accounting for 2% of the global greenhouse gas emissions. Computers consumed 408 TWh globally in 2010, and with a forecasted 10% annual growth on installed computing systems, this figure will rise, negatively impacting the environment and the economics of computing [1]. Unfortunately, a large fraction of this energy is wasted. Modern processors account for 38% of the power consumption of a computing system [5], but the general-purpose computing substrate they provide is hugely power inefficient. For example, compared to an optimized ASIC design, conventional multicores consume 157–707 times more energy when executing various phases of an H.264 encoder [6].

We postulate that a promising path towards energy-efficient data processing is the development of multicore processors with specialized cores. To implement these cores, we propose to utilize silicon that would otherwise remain powered-off, or "dark", due to the physical constraints of processor chips. Specialized multicores can leverage the underutilized die area to overcome the power barrier, delivering significantly higher performance for the same power envelope, thereby significantly reducing energy consumption. In the remaining of the paper we describe these constraints and how they give rise to dark silicon, explain why utilizing dark silicon is a viable solution, estimate the possible energy benefits that could be achieved by specialized computing on commercial workloads (e.g., online transactional processing, decision support, and web server workloads), and conclude with a set of research directions for the software community.

THE RISE AND FALL OF DARK SILICON

As transistors shrink following Moore's Law, exponentially more transistors fit on a single chip with each process generation. However, transistor supply voltages do not scale fast enough: a 10-fold increase in transistor counts over the last decade was followed by only a 30% drop in supply voltage [2,9]. The power consumption of the additional transistors can no longer be mitigated through circuit-level techniques [4]. Voltage-frequency scaling (VFS) may decrease power consumption by lowering the operating frequency which allows a drop in supply voltage. However, the shrinking range of operational voltages dampen the impact of VFS, as the difference between V_{th} and V_{dd} narrowed by 70% during the last decade [9]. At the same time, the transistors' switching speed cannot be reduced sufficiently to keep the power consumption at bay, and simultaneously deliver reasonable performance [7,8]. Whereas the power requirements grow, chip power delivery and cooling limitations remain largely unchanged [2]. With the power envelope remaining constant, and voltages scaling much slower than the exponential growth in transistor density, we'll soon be unable to power all transistors simultaneously. Short of a technological miracle, we head towards an era of "*dark silicon",* able to build dense devices we cannot afford to power [7,8].

Instead of fighting dark silicon, we propose to embrace it and harness it. By analyzing the impact of physical constraints (area, power, thermal, bandwidth) on multicore processor designs across technologies, we postulate that server chips will not scale efficiently beyond a few tens to low hundreds of cores, while upwards of 1,000 cores can fit in a single chip [7,8]. Thus, an increasing fraction of the chip in future technologies will be left unutilized. We observe that the abundant, power-constrained and underutilized die real-estate can be effectively used to implement customized heterogeneous cores, where each core is able to perform a narrow set of tasks at significantly higher speed and lower energy.

We envision an architecture that provides a sea of specialized cores, with the executing workload powering up only the most application-specific hardware, while the rest of the chip is switched off to conserve energy. While these cores are highly specialized, and thus of limited functionality, the expanse of the unused area allows the implementation of a large number of them, increasing the probability that a suitable core will be available [7,8]. Certain area could also be devoted to reconfigurable hardware, allowing greater flexibility at a modest overhead. Industry leaders have already shown signs of moving towards this direction, with AMD's Fusion and Intel's Knights Ferry and Stellarton chips combining general-purpose cores together with vector units, SIMD engines, and even FPGAs. Our preliminary estimates indicate that specialized multicores deliver up to a 12-fold decrease in energy consumption when running commercial workloads [8].

RESEARCH DIRECTIONS

To realize this vision, we need to revisit the software infrastructure to utilize efficiently the diverse hardware resources. As a first step, we need to characterize the behavior of the execution phases of commercial server workloads and applications commonly run on large computing installations, and identify the dominant performance and energy hogs; these will be candidates for off-loading to specialized hardware. Then, we should analyze the execution requirements of each candidate to determine the hardware architecture that best maps to these requirements. Identifying commonalities across candidates that make them amenable to execution on similar hardware will avoid core over-specialization. Finally, we need to develop novel language constructs and runtime techniques to facilitate and drive the execution across heterogeneous cores, while exposing the data access and sharing patterns to the execution environment so the scheduler can minimize data movements.

REFERENCES

[1] Financial Times. Computing: powerful argument for cutting IT energy consumption. *http://www.ft.com/cms/s/0/4e926678-bf90-11df-b9de-00144feab49a.html#axzz18EH%YGujk*, September 2010.

[2] The international technology roadmap for semiconductors (ITRS). *http://www.itrs.net/*, 2008.

[3] Worldwide LHC Computing Grid, Realtime VO-wise Data Transfer. *http://lcg.web.cern.ch/lcg/*, 2011.

[4] S. Borkar. Microarchitecture and design challenges for gigascale integration. In Proceedings of the 37th International Symposium on Microarchitecture, December 2004.

[5] X. Fan, W.-D. Weber, and L.A. Barroso. Power provisioning for a warehouse-sized computer. In Proceedings of the 34th International Symposium on Computer Architecture, June 2007.

[6] R. Hameed, W. Qadeer, M. Wachs, O. Azizi, A. Solomatnikov, B.C. Lee, S. Richardson, C. Kozyrakis, and M. Horowitz. Understanding sources of inefficiency in general-purpose chips. In Proceedings of the 37th International Symposium on Computer Architecture, June 2010.

[7] N. Hardavellas. Chip Multiprocessors for Server Workloads. PhD Thesis, Carnegie Mellon, Aug. 2009.

[8] N. Hardavellas, M. Ferdman, A. Ailamaki, and B. Falsafi. Power scaling: the ultimate obstacle to 1K-core chips. Technical Report NWU-EECS-10-05, Northwestern University, October 2010.

[9] Y. Watanabe, J.D. Davis, and D.A. Wood. WiDGET: Wisconsin decoupled grid execution tiles. In Proceedings of the 37th International Symposium on Computer Architecture, June 2010.