

Technical Perspective

FAWN: A Fast Array of Wimpy Nodes

By Luiz André Barroso

INNOVATION IN COMPUTING systems thrives at the beginning and the end of technology cycles. When facing the limits of an existing technology or contemplating the applications of a brand new one, system designers are at their creative best. The past decade has been rich on both fronts, particularly for computer architects. CMOS technology scaling is no longer yielding the energy savings it used to provide across generations, resulting in severe thermal constraints leading to increased attention to so called “wimpy processors.” These processors achieve high performance and energy efficiency by using a larger number of low-to-modest-speed CPU cores. Also in the past decade, the consumer electronics industry’s investment in non-volatile storage technologies has resulted in NAND FLASH devices that are becoming competitive for general-purpose computing usage as they fit nicely within the huge cost/performance gap between DRAM and magnetic disks. FLASH-based storage devices are over 100 times faster than disks, although at over 10 times the cost per byte stored.

The emergence of wimpy processors and FLASH met a promising deployment scenario in the field of large-scale data centers for Internet services. These warehouse-scale computing (WSC) systems tend to run workloads that are rich in request-level parallelism—a match for the increased parallelism of wimpy CPUs—and are very data intensive—a match for the high input-output rates that are possible with FLASH technology. The energy efficiency potential of both these tech-

nologies could help lower the substantial energy-related costs of WSCs.

Given all this potential, how can we explain the rather slow pace of adoption of these technologies in commercial WSCs? At first glance, wimpy processors and FLASH seem compelling enough to fit within existing data center hardware and software architectures without the need for substantial redesign of major infrastructure components, thus facilitating rapid adoption. In reality, there are obstacles to extracting the maximum value from them. Hölzle¹ summarized some of the challenges facing wimpy cores in commercial deployments, including parallelization overheads (Amdahl’s Law) and programmer productivity concerns. FLASH adoption has also suffered due to software related issues. FLASH will not fully replace disks for most workloads due to its higher costs, therefore storage system software must be adapt-

FAWN combines wimpy cores and FLASH to create an efficient, high-throughput, key-value storage system.

ed to use both FLASH and disk drives effectively.

The lesson here is that to extract the most value from compelling new technology one often needs to consider the system more broadly, and rethink how applications and infrastructure components might be changed in light of new hardware component characteristics. This is precisely what the authors of the following article on FAWN have done.

FAWN presents a new storage hardware architecture that takes advantage of wimpy cores and FLASH devices, but does so alongside a new datastore software system infrastructure (FAWN-DS) that is specifically targeted to the new hardware component characteristics. The system is not a generic distributed storage system, but one that is specialized for workloads that require high rates of key-value lookup queries. By co-designing the hardware and software, and by targeting the system for a particular (but compelling) use case, the authors present a solution that has greater potential to realize the full value of new energy-efficient components. Their approach, which includes building and experimenting with actual software and hardware artifacts, is a model worthy of being followed by future systems research projects. 

Reference

1. Hölzle, U. Brawny cores still beat wimpy cores, most of the time. *IEEE Micro* (Aug/Sept. 2010).

Luiz André Barroso (luiz@google.com) is a Distinguished Engineer at Google.

© 2011 ACM 0001-0782/11/07 \$10.00