

## Work In Progress: *Conquest*—Combining Battery-Backed RAM and Threshold-Based Storage Scheme to Conserve Power

RuGang Xu, Computer Science Department, UCLA, rxu@cs.ucla.edu

An-I Andy Wang, Department of Computer Science, FSU, awang@cs.fsu.edu

Geoffrey Kuenning, Computer Science Department, Harvey Mudd College, geoff@cs.hmc.edu

Peter Reiher, Computer Science Department, UCLA, reiher@cs.ucla.edu

Gerald J. Popek, Computer Science Department, UCLA, popek@cs.ucla.edu

Energy conservation has become a pressing research topic in recent years as users of mobile devices demand longer battery life while looking for better performance and more features. Due to rising energy costs, the server market is also becoming sensitive to power consumption. We are developing a system that satisfies both needs. Our focus is on reducing disk accesses, which account for 8% of the power usage for laptops, 20% for web servers, and up to 70% for disk-intensive servers.

Our approach minimizes disk accesses by efficiently using an abundance of persistent RAM. In particular, we are applying the *Conquest*<sup>1</sup> disk/RAM hybrid filesystem, which has much better performance than disk caching and or simple RAM filesystems. We seek to combine these benefits with power reduction.

In *Conquest*, RAM is the permanent store for all metadata (file attributes), executables, and small files. Disk stores only the content of large files. Because large-file accesses are infrequent, *Conquest* can often put the disk into a lower power state.

We examined the power benefits of *Conquest* using a 150-day trace-based study of our main lab server. As we increased the number of small files stored in RAM, larger time gaps were found between disk accesses. Given a small-file threshold of 32KB and a spin-down threshold of 1 second, we found 300,000 opportunities for spin-down. Accounting for the energy break-even times of spinning the disk down and back up, the effective spin-down time totaled 110 days. In other words, we could have saved about 70% of the disk power consumption by using *Conquest*.

However, *Conquest* requires an abundance of persistent memory. For example, our server requires 2.5GB of persistent RAM to store all files that are below 32KB. The gains for performance and power savings versus the RAM requirements quickly dropped when we raised the bar to 64KB; the sweet spot for our trace seems to be 32KB. At a threshold of 64KB, we required 1 GB more RAM and saved only 5%

more disk power. The increased RAM requirements have low dollar costs but unclear power costs.

To evaluate the power trade-offs between disk and RAM, we need to understand RAM power management as it applies to RAM-based filesystems. The penalties for idle-active transitions are relatively low compared to disk, allowing for aggressive power-down algorithms. A low-power RAM filesystem will still be much faster than conventional filesystems.

In *Conquest*, the penalties incurred for disk access will only apply to large files. Fortunately, large-file access is infrequent and usually sequential, allowing aggressive prefetching and prediction. Prediction is simplified because we only need to know when to spin up the disk, which is a matter of knowing which applications require large files. For example, when we start a multimedia player, we can simultaneously spin the disk up. We can also take advantage of human expectations: users are more willing to wait for loading a long movie than for starting an editor.

The *Conquest* approach is novel because previous disk-based power savings research focused on delaying disk updates, aggressive prefetching, custom caching, file access prediction, and various threshold policies designed to complement the underlying filesystem, while *Conquest* is a filesystem built for performance and energy reduction. Unlike these other approaches, *Conquest* allows the disk to stay spun down even when the system is in heavy use. Many previous approaches yield no power savings and have significantly lower performance during heavy use.

To complement our trace studies, we plan to run benchmarks on our working *Conquest* prototype to gather power and performance measurements. The next step will be to investigate power management for RAM-based filesystems and to identify common applications that require large file use. This part of the work will highlight the trade-offs of RAM performance versus power consumption and their impact on users.

---

<sup>1</sup> A. Wang, G.H. Kuenning, P. Reiher, G.J. Popek. [Conquest: Better Performance Through a Disk/Persistent-RAM Hybrid File System](#). *Proceedings of the 2002 USENIX Annual Technical Conference*, Monterey, June 2002