

Amdahl's Law in the Multicore Era

Mark D. Hill

Computer Sciences Department
University of Wisconsin—Madison
<http://www.cs.wisc.edu/~markhill/>

Over the last several decades computer architects have been phenomenally successful turning the transistor bound provided by Moore's Law into chips with ever increasing single-threaded performance. During many of these successful years, however, many researchers paid scant attention to multiprocessor work [1].

Now as vendors turn to multicore chips, researchers are reacting with more papers on multi-threaded ideas. While this is good, we are concerned that further work on single-thread performance will be squashed.

In this talk, based in part on an upcoming paper with Michael Marty [2], we apply Amdahl's Law to several multicore chips variants: symmetric cores, asymmetric cores, and dynamic techniques that allow cores to work together on sequential execution. Starting with Amdahl's simple software model, we add a simple hardware model based on fixed chip resources.

Our simple results encourage multicore designers to view performance of the entire chip rather than focusing only on core efficiencies. Moreover, we observe that obtaining optimal multicore chips performance requires further research in *both extracting more parallelism and making sequential cores faster*.

This talk seeks to stimulate discussion and future work, as well as temper the current pendulum swing from the past's under-emphasis on parallel research to a future with too little sequential research.

References

[1] Mark D. Hill and Ravi Rajwar, *The Rise and Fall of Multiprocessor Papers in the International Symposium on Computer Architecture (ISCA)*, <http://www.cs.wisc.edu/~markhill/mp2001.html>, March 2001.

[2] Mark D. Hill and Michael R. Marty, *Amdahl's Law in the Multicore Era*, to appear in IEEE Computer, 2008.

Biography

Mark D. Hill (<http://www.cs.wisc.edu/~markhill>) is professor in both the computer sciences department and the electrical and computer engineering department at the University of Wisconsin-Madison, where he also co-leads the Wisconsin Multifacet project with David Wood. He earned a PhD from University of California, Berkeley. He is an ACM Fellow and a Fellow of the IEEE. His past work ranges from refining multiprocessor memory consistency models to developing the *3C model* of cache behavior (*compulsory*, *capacity*, and *conflict* misses).