

MIT OpenCourseWare  
<http://ocw.mit.edu>

6.189 Multicore Programming Primer, January (IAP) 2007

Please use the following citation format:

Saman Amarasinghe and Rodric Rabbah, *6.189 Multicore Programming Primer, January (IAP) 2007*. (Massachusetts Institute of Technology: MIT OpenCourseWare). <http://ocw.mit.edu> (accessed MM DD, YYYY).  
License: Creative Commons Attribution-Noncommercial-Share Alike.

Note: Please use the actual date you accessed this material in your citation.

For more information about citing these materials or our Terms of Use, visit:  
<http://ocw.mit.edu/terms>

One trend that is helping the shift from superscalars to multicores is the “Memory Wall”, i.e. the increasing gap between processor and DRAM performance. How can multicores circumvent the memory wall issue better than superscalars?

Multicores architectures decentralize resources, including memory, and thus reduce contention on a shared global memory. The availability of local storage on each core serves to reduce contention on a globally shared main memory, and if data is distributed adequately among the cores, then there is an effective increase in overall concurrency.

Superscalars try to hide memory access latency with increased speculation and prediction (e.g., value and address prediction) which can be wasteful with respect to power. If a superscalar processor cannot completely hide the long access to memory, then it cannot make any progress. In multicores, a single core may idle until its data is fetched but that will not prevent other cores from continuing to compute.