Large-Reach Memory Management Unit Caches

Abhishek Bhattacharjee

Rutgers University

International Symposium on Microarchitecture-46

December 2013

Address Translation and MMU Caches

On a typical TLB miss:

- x86: 1-4 memory references
- ARM: 1-2 memory references
- Sparc: 1-2 memory references

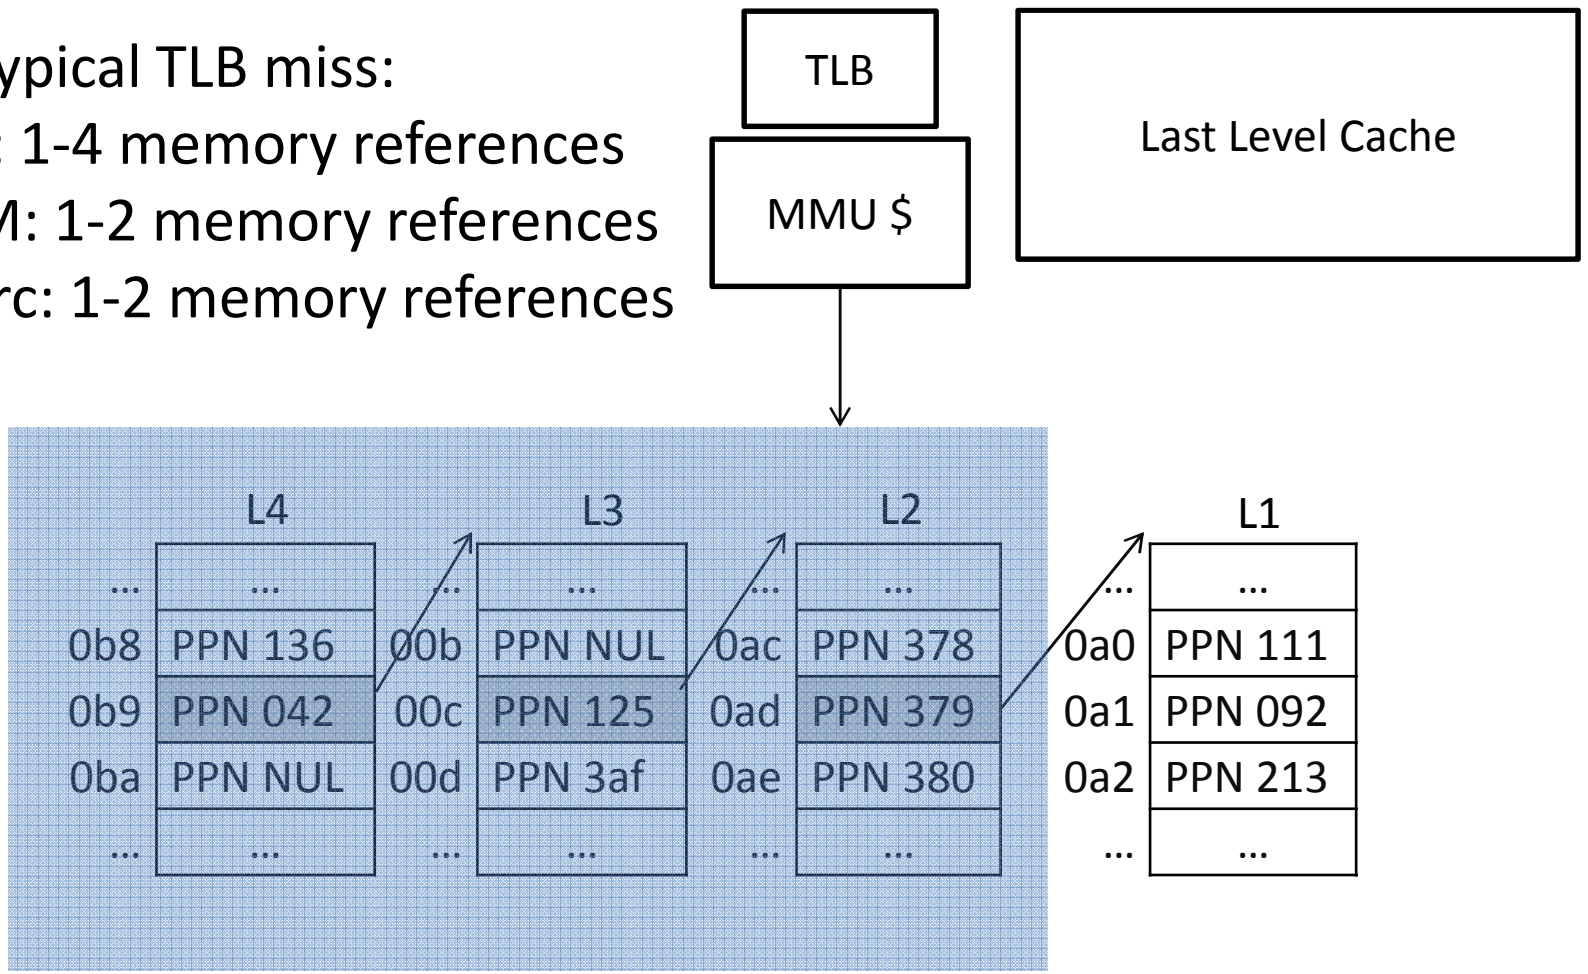

Approaching an Ideal MMU Cache

- Intel i7: 8 cores, 8GB memory, 512-entry L2 TLB, and 8MB LLC