Large-Reach Memory Management Unit Caches

Abhishek Bhattacharjee
Department of Computer Science, Rutgers University

1. Address Translation Overhead

- Nominally-sized applications: 10% - 15%[1,2]
- Virtualization: 89%[1,2]
- Big-data workloads: 5-85%[3]

2. Address Translation Hardware

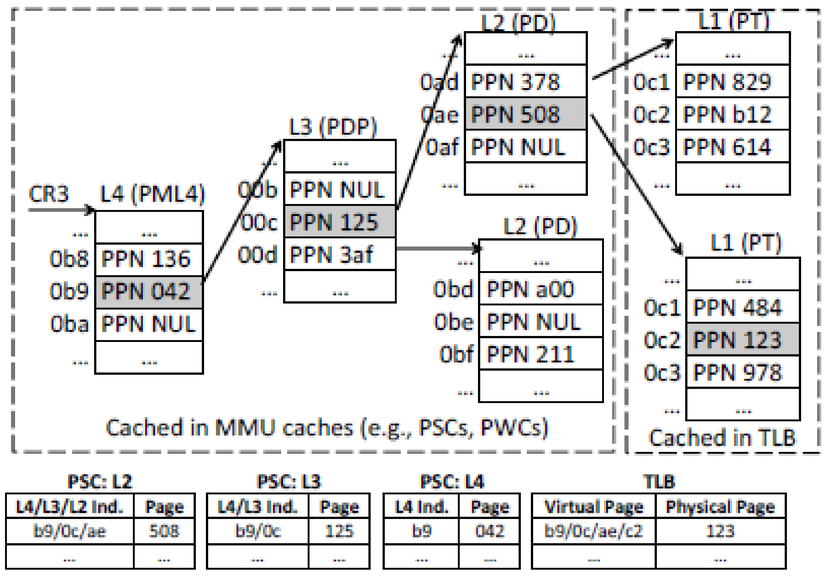

3. Real-System Measurements

- Intel i7, 8 cores, 8GB memory, 512-entry TLBs
- 32-entry L2 PSC, 4-entry L3 PSC, 2-entry L4 PSC

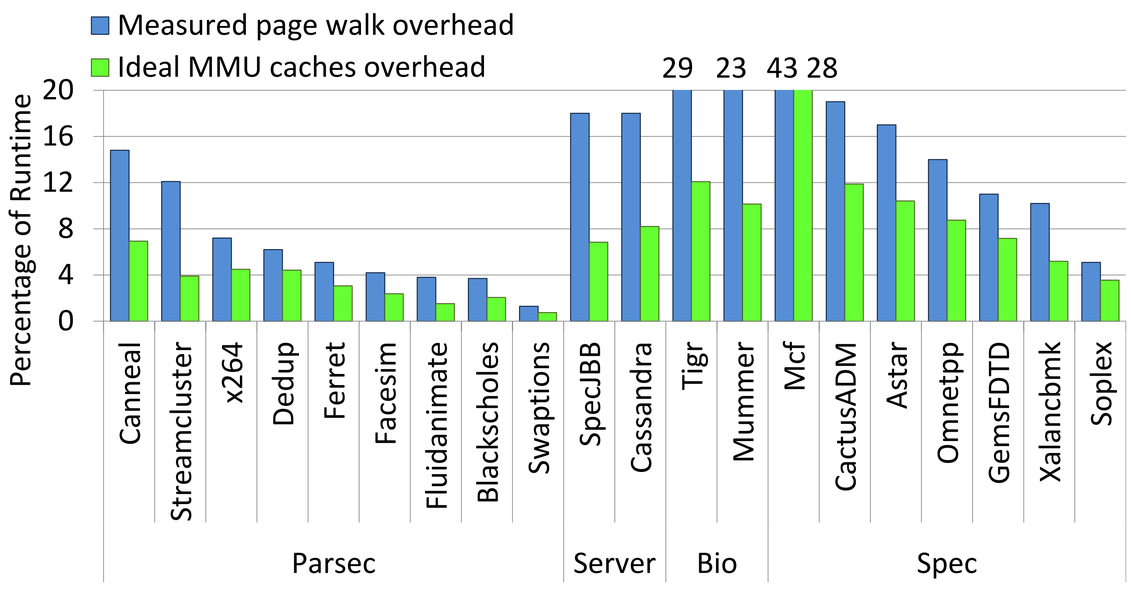

4. Coalesced MMU Caches

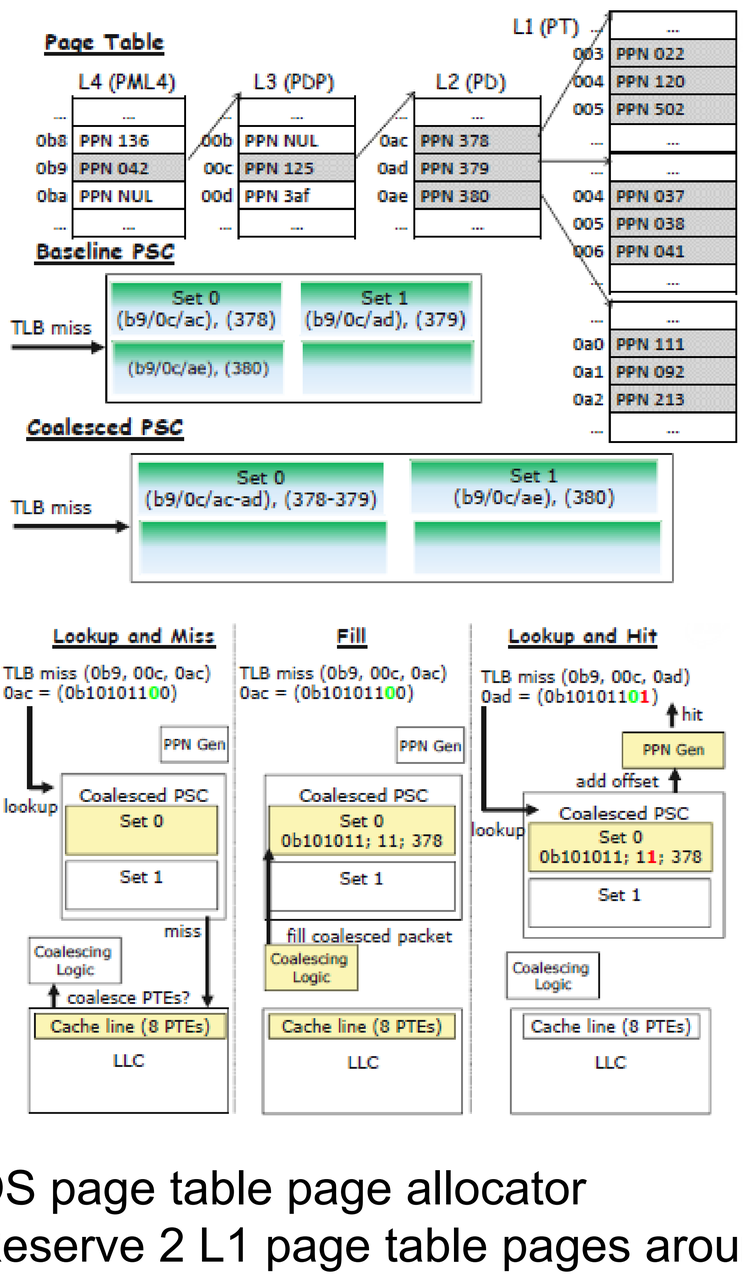

- OS page table page allocator
- Reserve 2 L1 page table pages around faulting L2 entry
- If reserved pages are used, coalescing possible
- Every page fault, clear reserved pages unused since last 4 faults

5. Shared MMU Caches

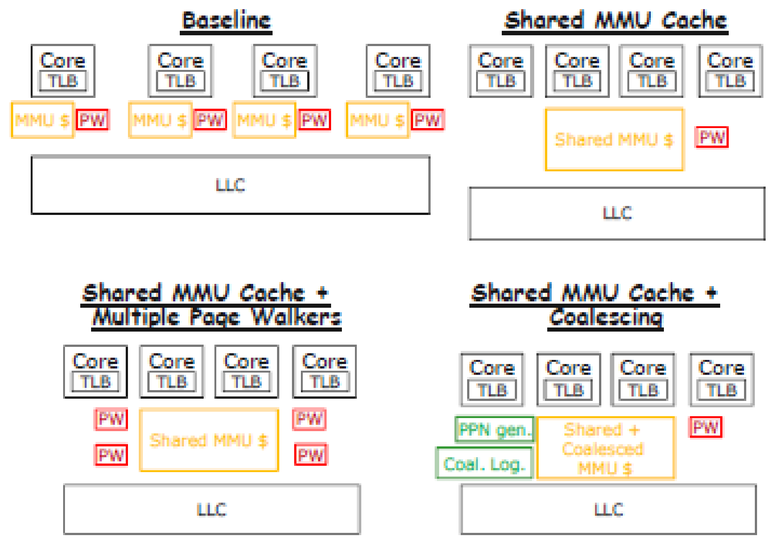

6. Results

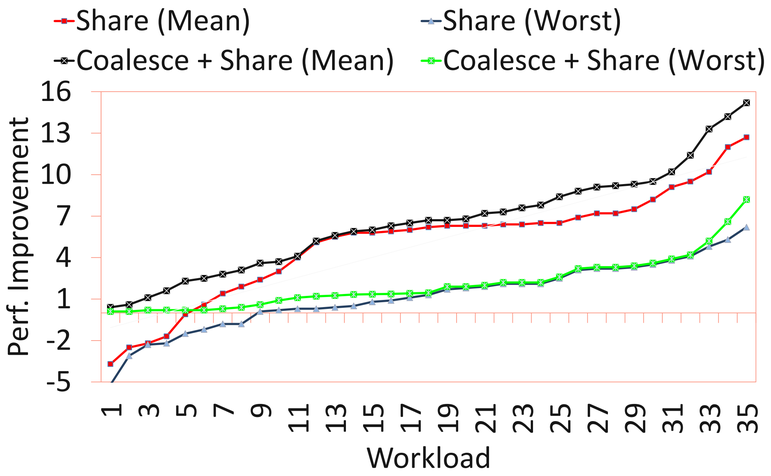

- Multiple ports and PTWs give 1-2% perf on 8-core CMP
- Combined approach within 3-5% of ideal case (3x larger MMU cache)

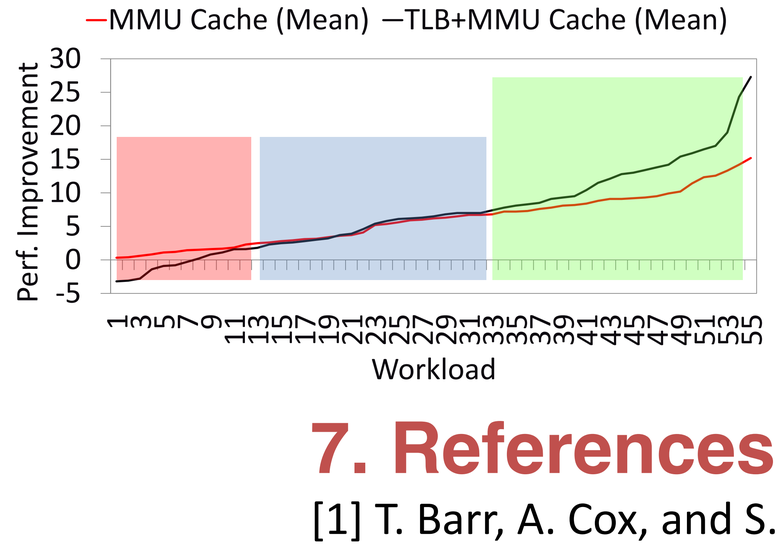

- Coalesced MMU caches provide more predictable performance improvements (10 – 12% on average)

7. References

[1] T. Barr, A. Cox, and S. Rixner, *SpecTLB: A Mechanism for Speculative Address Translation*, ISCA, 2011.
[2] R. Bhargava et al., *Accelerating Two-Dimensional PageWalks for VirtualizedSystems*, ASPLOS, 2008.
[3] A. Basu et al., *Efficient Virtual Memory for Big Memory Servers*, ISCA 2013.