



A Study of the Effect of Prefetching in Shared Memory Resource Contention

Tanima Dey, Wei Wang, Jack W. Davidson, Mary L. Soffa

University of Virginia

{td8h, ww6r, jwd, soffla}@virginia.edu


Motivation

- **Challenge:** Managing shared resource contention in CMPs
- **Problem:** Contention for the memory hierarchy resources, especially the shared caches
- **Focus:** Hardware prefetcher – can have potential impact on memory resource contention and application's performance
- **Goal:** Study how prefetching affects an application's performance for shared-memory resource contention

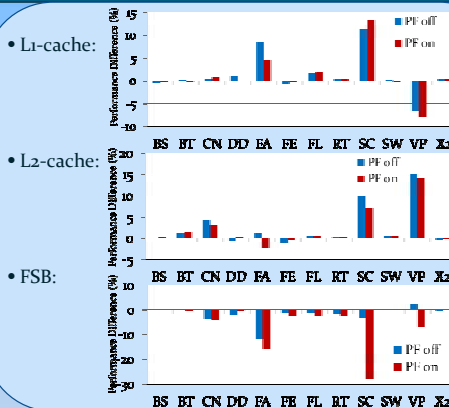
Methodology

- Compare the application's performance
 - Prefetcher: enabled/disabled
 - There is contention for the targeted resource
- Use two types of contention for multi-threaded applications
 - Intra-application
 - Inter-application
- Measure contention for a targeted resource:
 - Baseline: threads do not share resource
 - Contention : threads share resource


Intra-application Contention



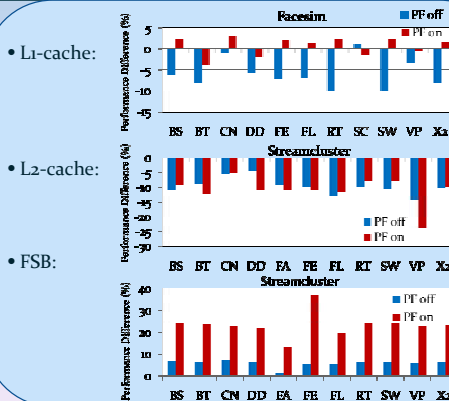
Characterization Results



Inter-application Contention



Characterization Results



Summary and Future Work

- Prefetching degrades (25%) and improves (4%) application's performance depending on data access pattern and co-runners
- Dynamic prefetching control for maximum performance and throughput

References

[1] T. Dey, W. Wang, J. Davidson, M. Lou Soffa, Characterizing Multi-threaded Applications based on Shared-Resource Contention, ISPASS 2011.
 [2] S. Zhuravlev, S. Blagodurov, and A. Fedorova. Addressing shared resource contention in multicore processors via scheduling. ASPLOS 2010.