

Speculative Precomputation

What's the Problem?

What's the Opportunity?

Opportunity?




Figure 2. Cumulative L1 data cache misses contributed by the worst behaving static loads.




Figure 3. Speedup when 10 worst behaving static loads are assumed to always hit in cache.

What do you have to get right to make this work (base case)?

What do you have to get right to make this work (base case)?

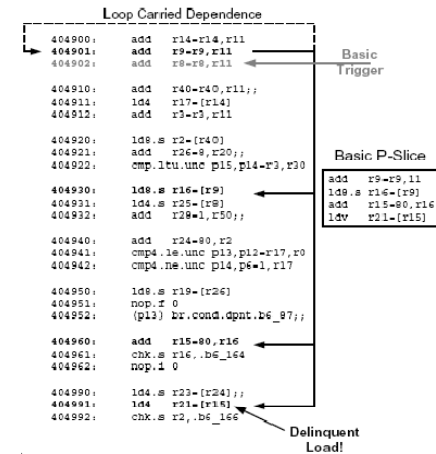


Figure 5. The p-slice for an example delinquent load.

Big weakness of the base case?

Chaining Triggers?

- How do they work?
- Easier way?
- Why not used here?



Looping Chaining Helper Prefetching Threads

- What's the big win?
- What new problems?



Speculative Precomputation vs. Inline Software Prefetching?



Speculative Precomputation vs. Hardware Prefetching?



Balanced Multithreading



What's the conflict?



What are major costs of SMT?



Why combine SMT with coarse-grain multithreading?



Why is CG mtg cheaper?

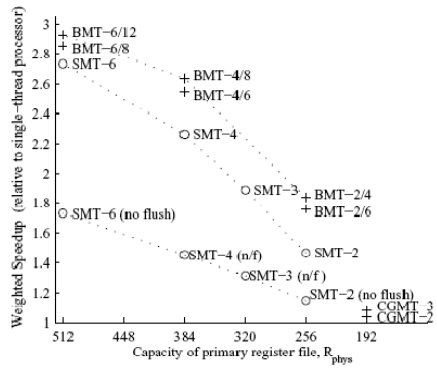


Figure 1. Speedup vs register file size.




Figure 8. Speedup vs instruction window size. Weighted speedup is relative to single-thread execution with 128 renaming registers.