Phase Guided Profiling for Fast Cache Modeling

Andreas Sembrant, David Black-Schaffer and Erik Hagersten

UPMARC
Uppsala Programming for Multicore Architectures Research Center

UPPSALA UNIVERSITET

1. Cache Behavior Over Time

Why do we care?

- Performance profiling: optimize code
- Code optimization: optimize data
- Scheduling: optimize sharing

Average cache behavior is misleading:

- Different phases need different optimizations.

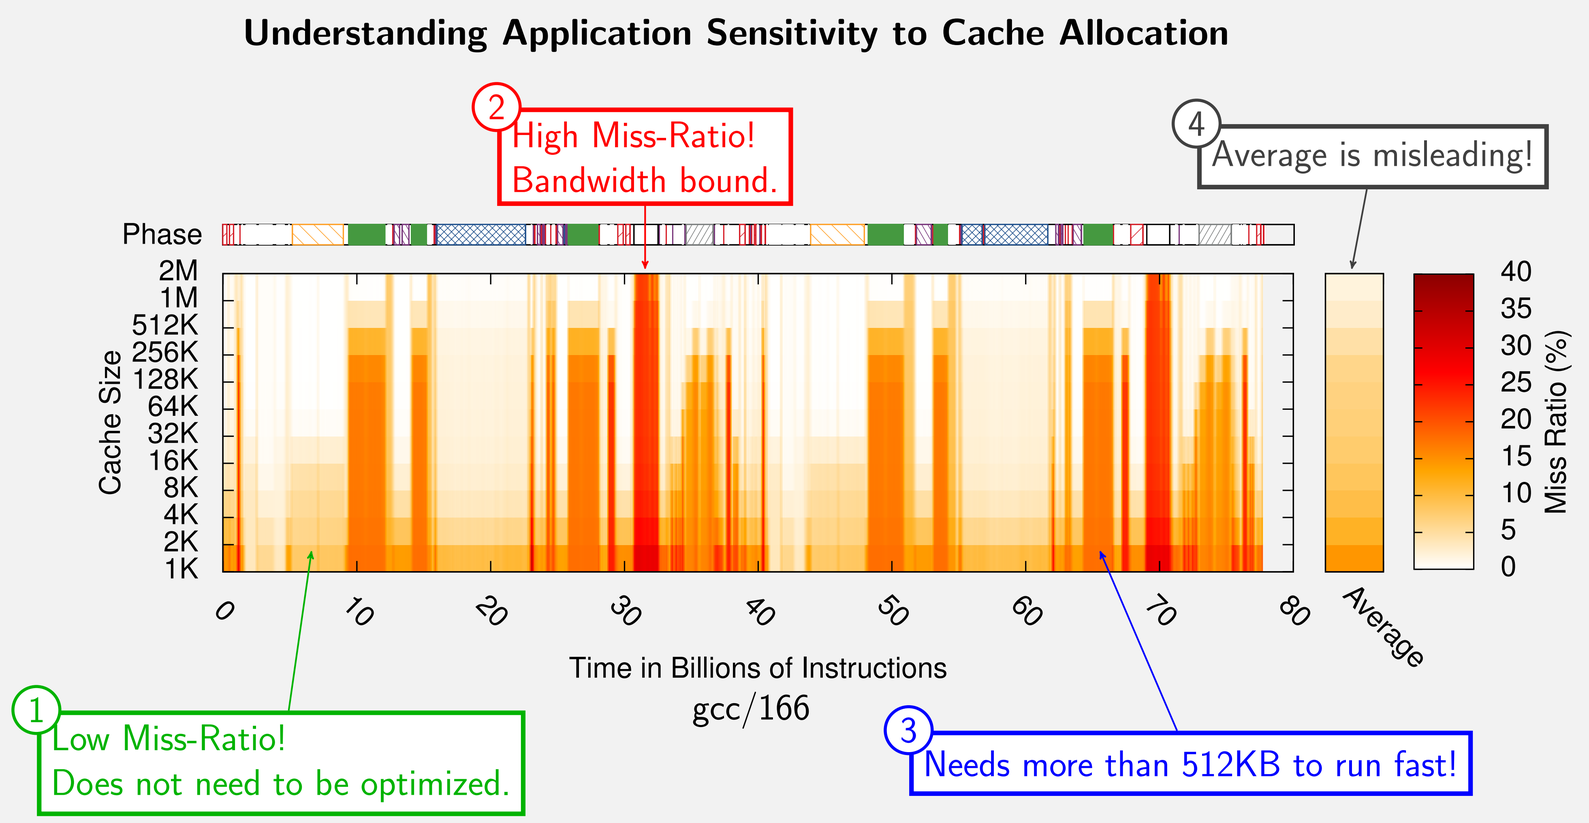

2. Method: StatCache + ScarPhase = Phase Guided Profiling for Fast Cache Modeling

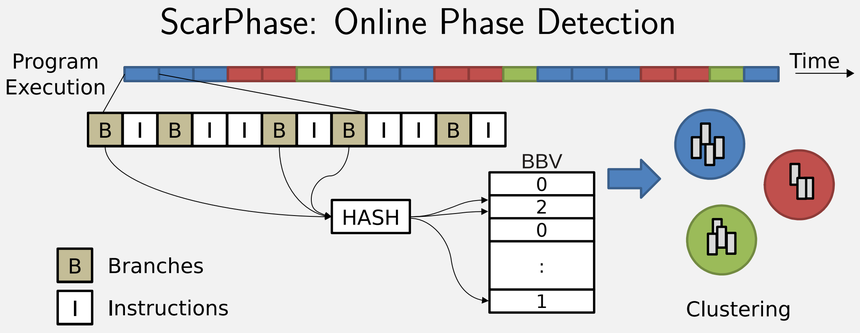

- StatCache: Statistical cache model based on sparsely sampled memory accesses.
- ScarPhase: Uses Intel PEBS to sample branches to identify phases (2% overhead).

3. Modeling Cache Behavior Over Time

Option 1: StatCache on all execution

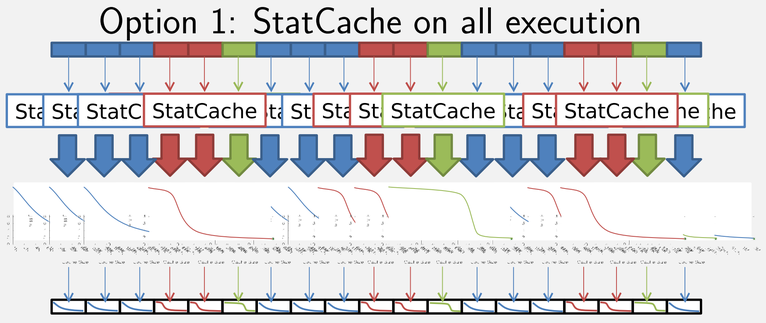

Very high overhead!

Option 2: Periodic Profiling

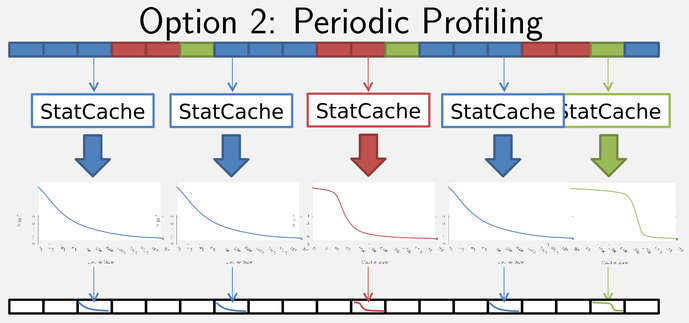

Lower overhead, but lower resolution.

Option 3: Phase Guided Profiling

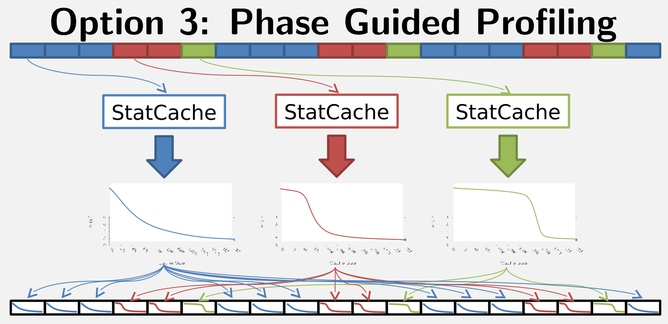

Overhead proportional to number of phases.

4. Accuracy and Overhead

Accuracy:

- Strictly better accuracy for phase guided profiling.
- Much better for some applications because of smarter profiling.

39% better accuracy!

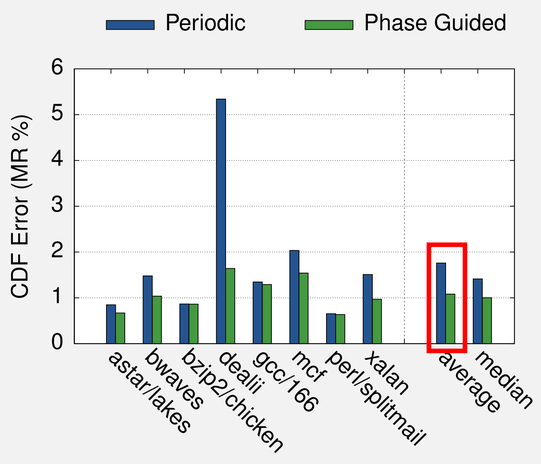

Overhead:

- No redundant samples from the same phase.
- 21% overhead on average.

6× faster!

- **Both faster and more accurate!**

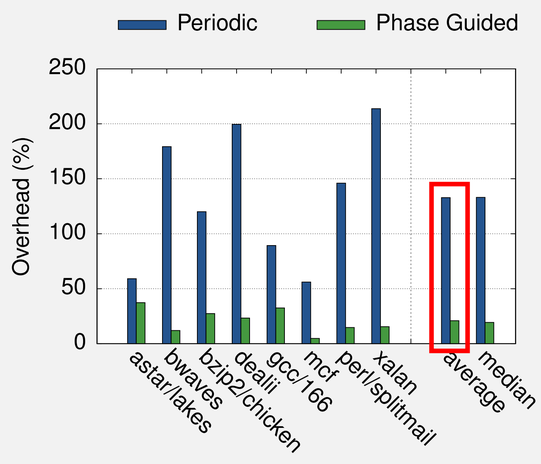

5. Phase Behavior

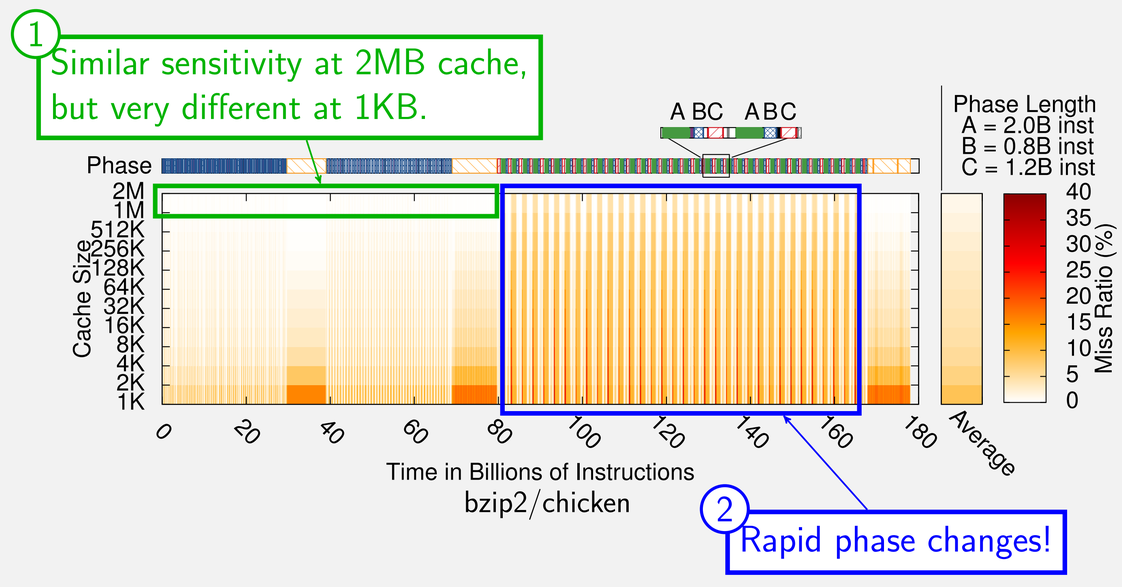

Intra-Phase Variations

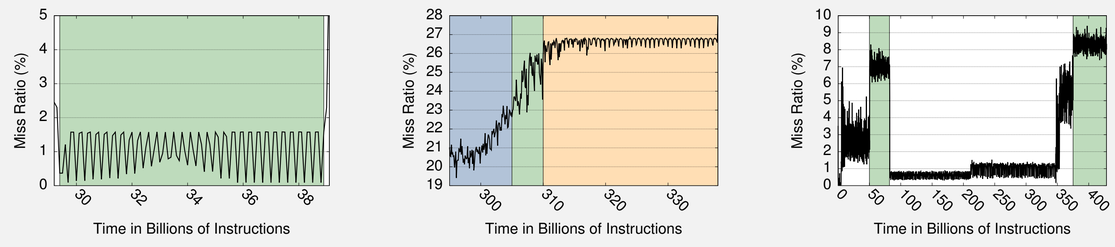

Periodic (bzip2/chicken) | Transition (mcf) | Instance (astar/lakes)

Solution: Spread out the memory samples over the whole phase.