




# Leveraging Cache Coherence to Detect and Repair False Sharing On-the-fly

Vipin Patel, Swarnendu Biswas, Mainak Chaudhuri (Indian Institute of Technology Kanpur)



Scan for artifacts

57th IEEE/ACM International Symposium on Microarchitecture 2024

Poster Flow: Top to Bottom, Left to Right

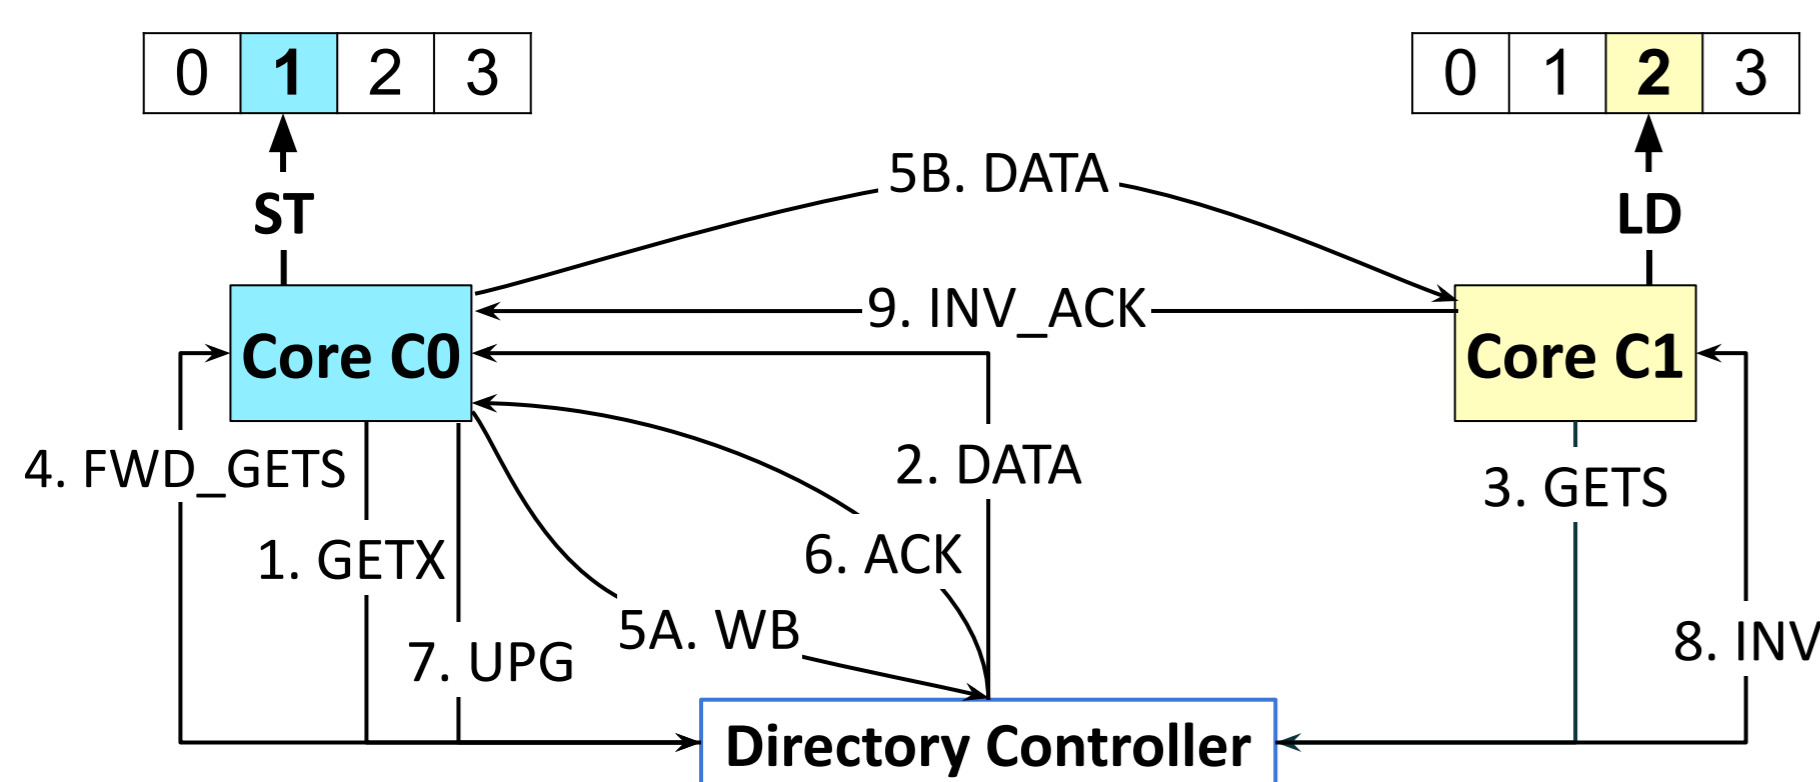
## False Sharing in Real World Applications

**Seeing through hardware counters: a journey to threefold performance increase**

Created: 2017-05-16  
Updated: 2024-09-16  
Resolved: 2024-04-16


secondary\_super\_cache does not scale well

## Cache Contention due to False Sharing



- Degrades performance
- Inflates interconnect traffic
- Increases energy consumption

## Speedup by Manually Fixing False Sharing



Avg Speedup **1.34x** and Max speedup of **3.06x**

Savings: **84%** interconnect traffic and **25%** in energy consumption

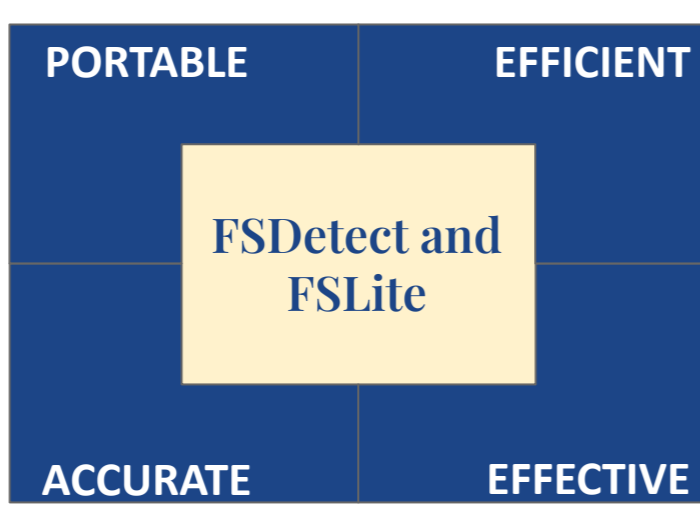
## Challenges in False Sharing Elimination

- Heap Organization<sup>1</sup>
- Cache Organization
- Compiler Optimization<sup>2</sup>
- Application Language
- Runtime Environment

```
struct {
  pthread_t tid; POINT_T*
  points;
  int num_elems; long long SX;
  long long SY; long long SXX;
  long long SXY; long long SYY;
  } reg_args;
```

## Shortcomings of the Existing Approaches

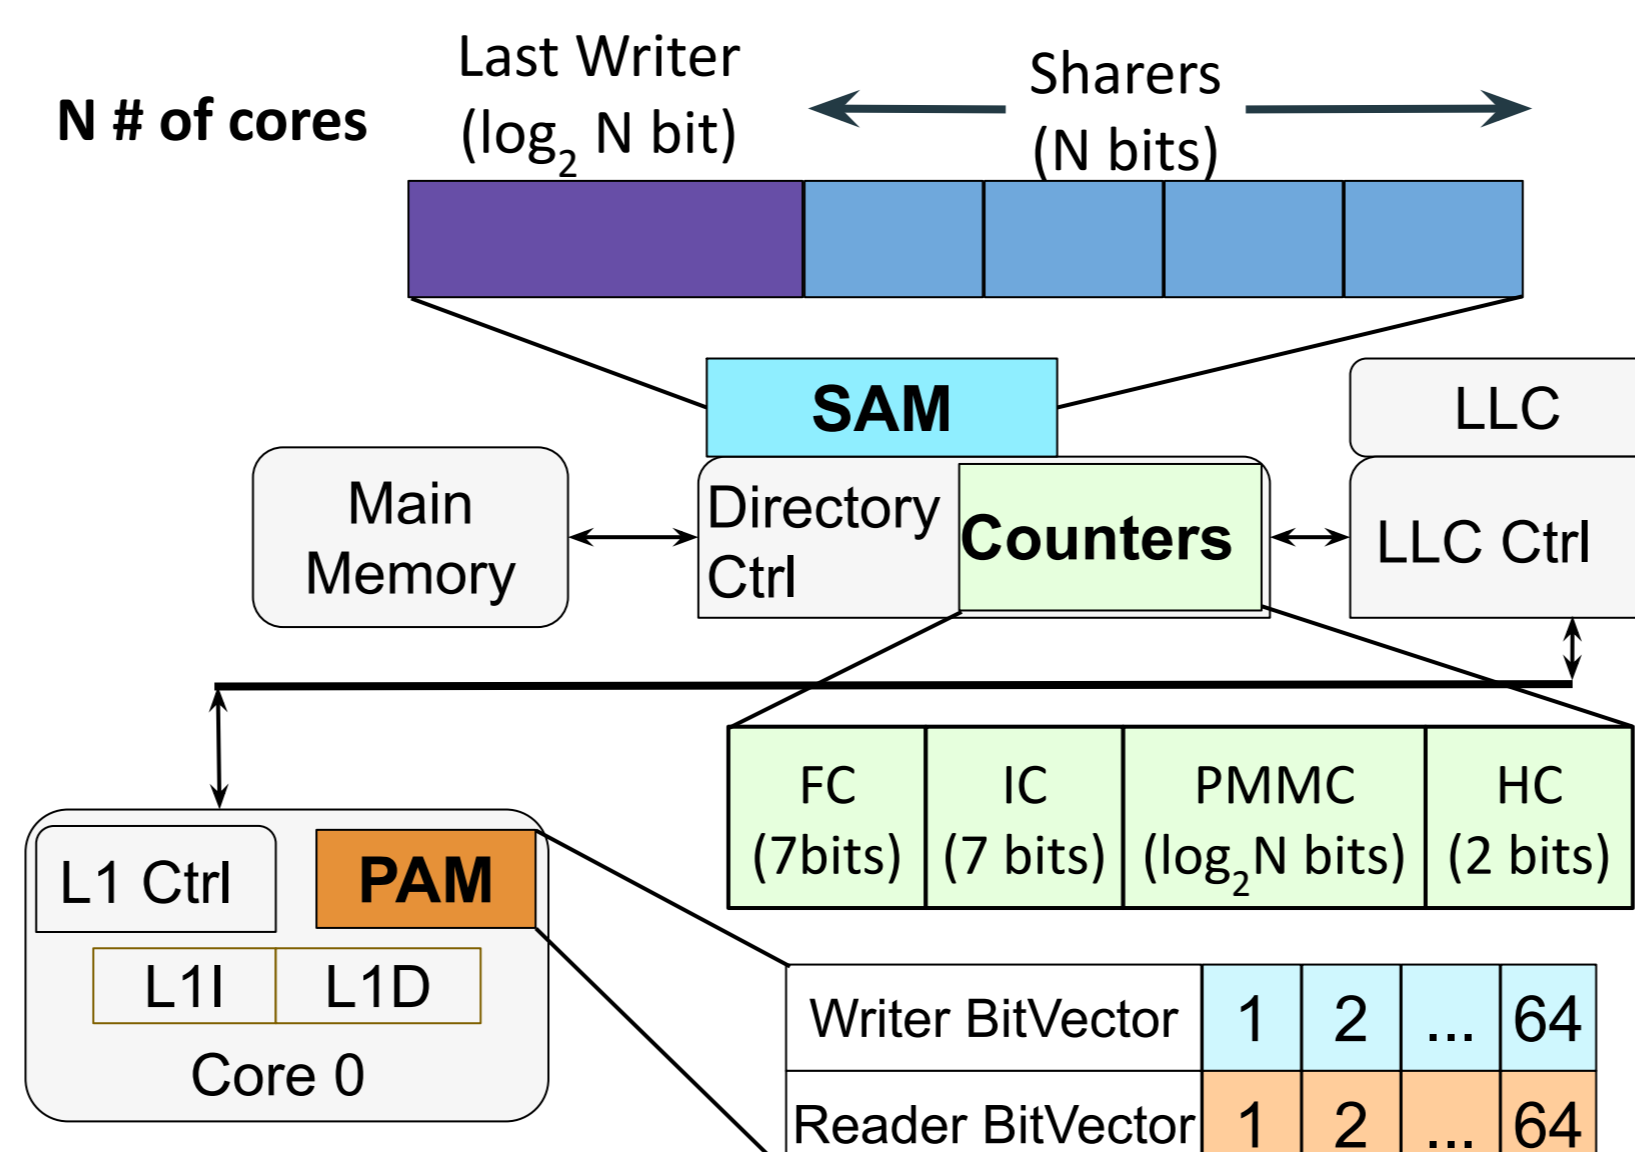
- Extending support to every language
- Inflation of memory footprint
- OS and Kernel modification<sup>3</sup>
- Cache organization modification
- Access to application source code
- Introduction of additional instructions



## Key Insight of Our Protocol Extensions

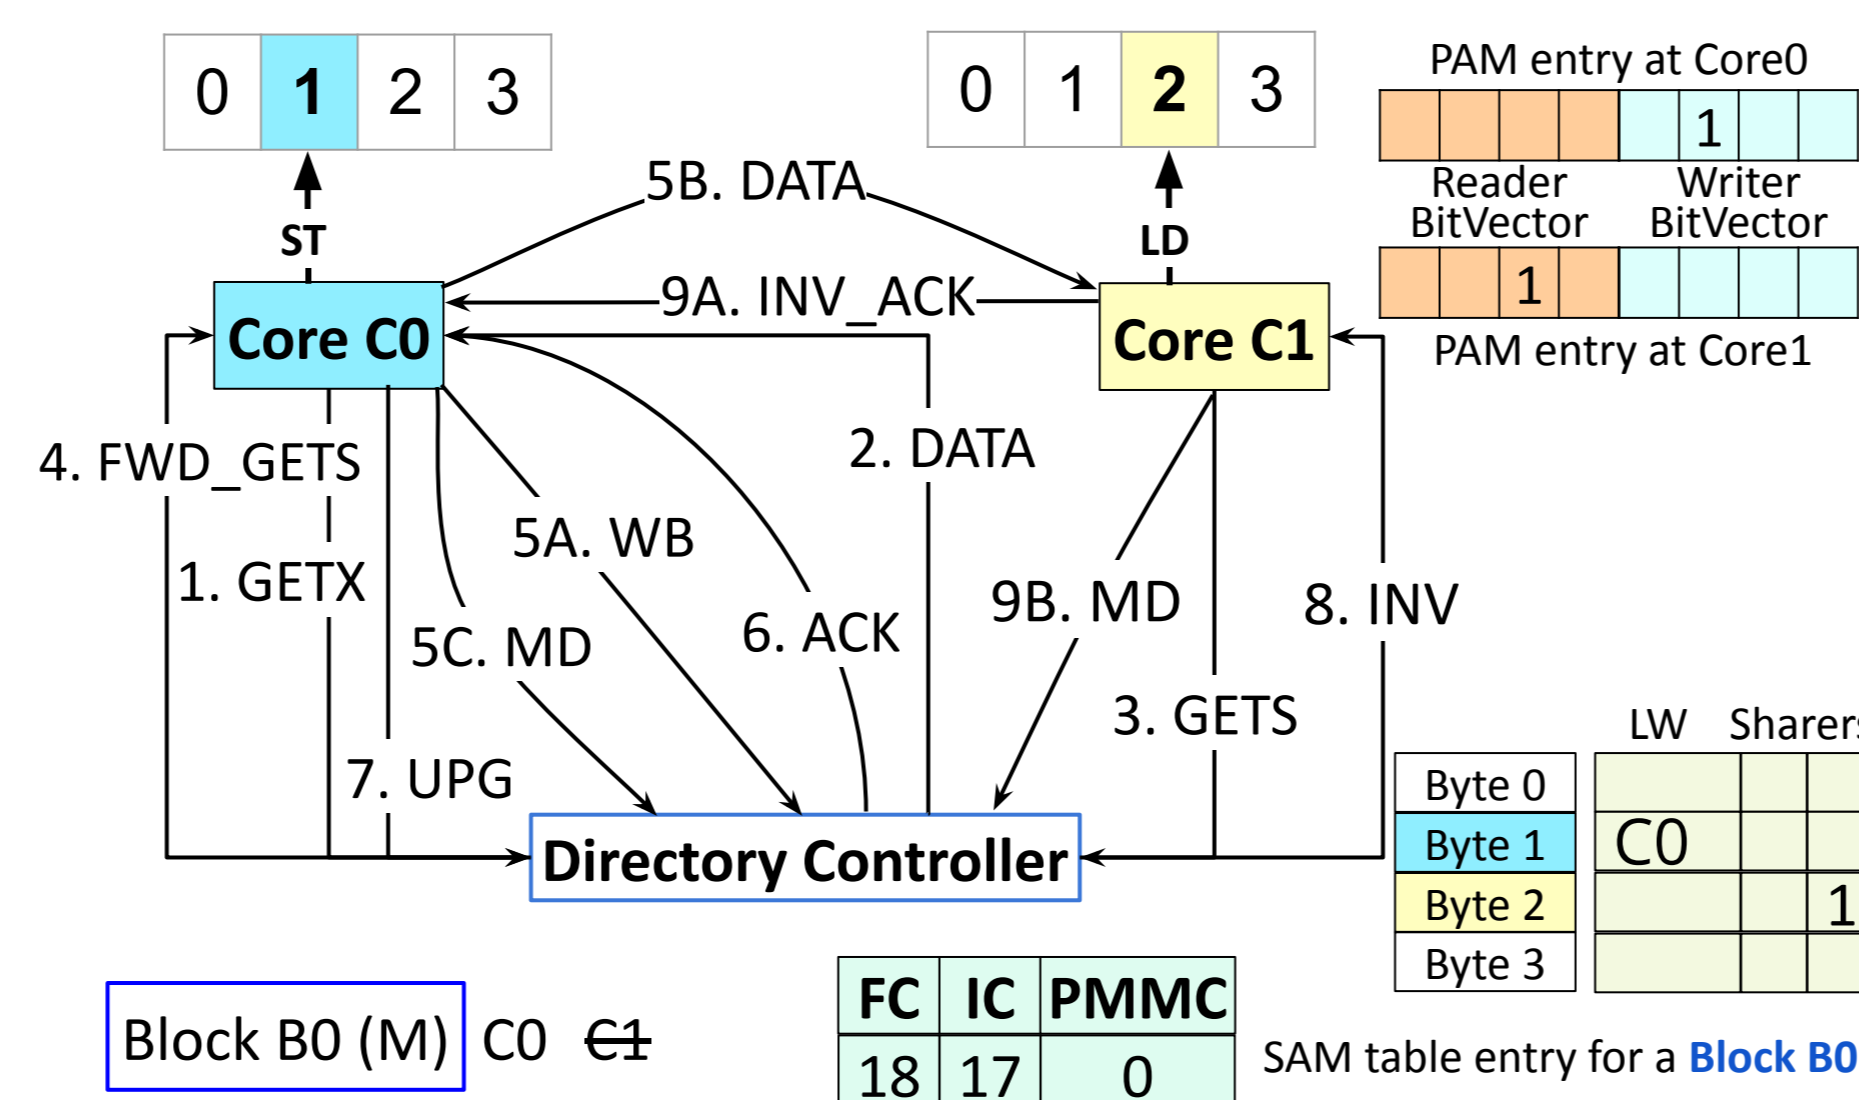
- FSDetect**
  - Tracks the sharing patterns of block
  - Identifies impactful instance of false sharing
- FSLite**
  - Eliminates false sharing by **privatizing** the blocks
  - Allows multiple writers for the disjoint bytes

## Overview of Architectural Modification



- Shared Access Metadata (SAM) to track the access history
- Private Access Metadata (PAM) to track the local access
- Fetch (FC) and Invalidation (IC) counters to filter the impactful instance
- Pending Metadata Message (PMMC) to track the inflight metadata
- Hysteresis Counter (HC) to prevent frequent trigger and termination of privatization

## FSDetect: Detection of False Sharing

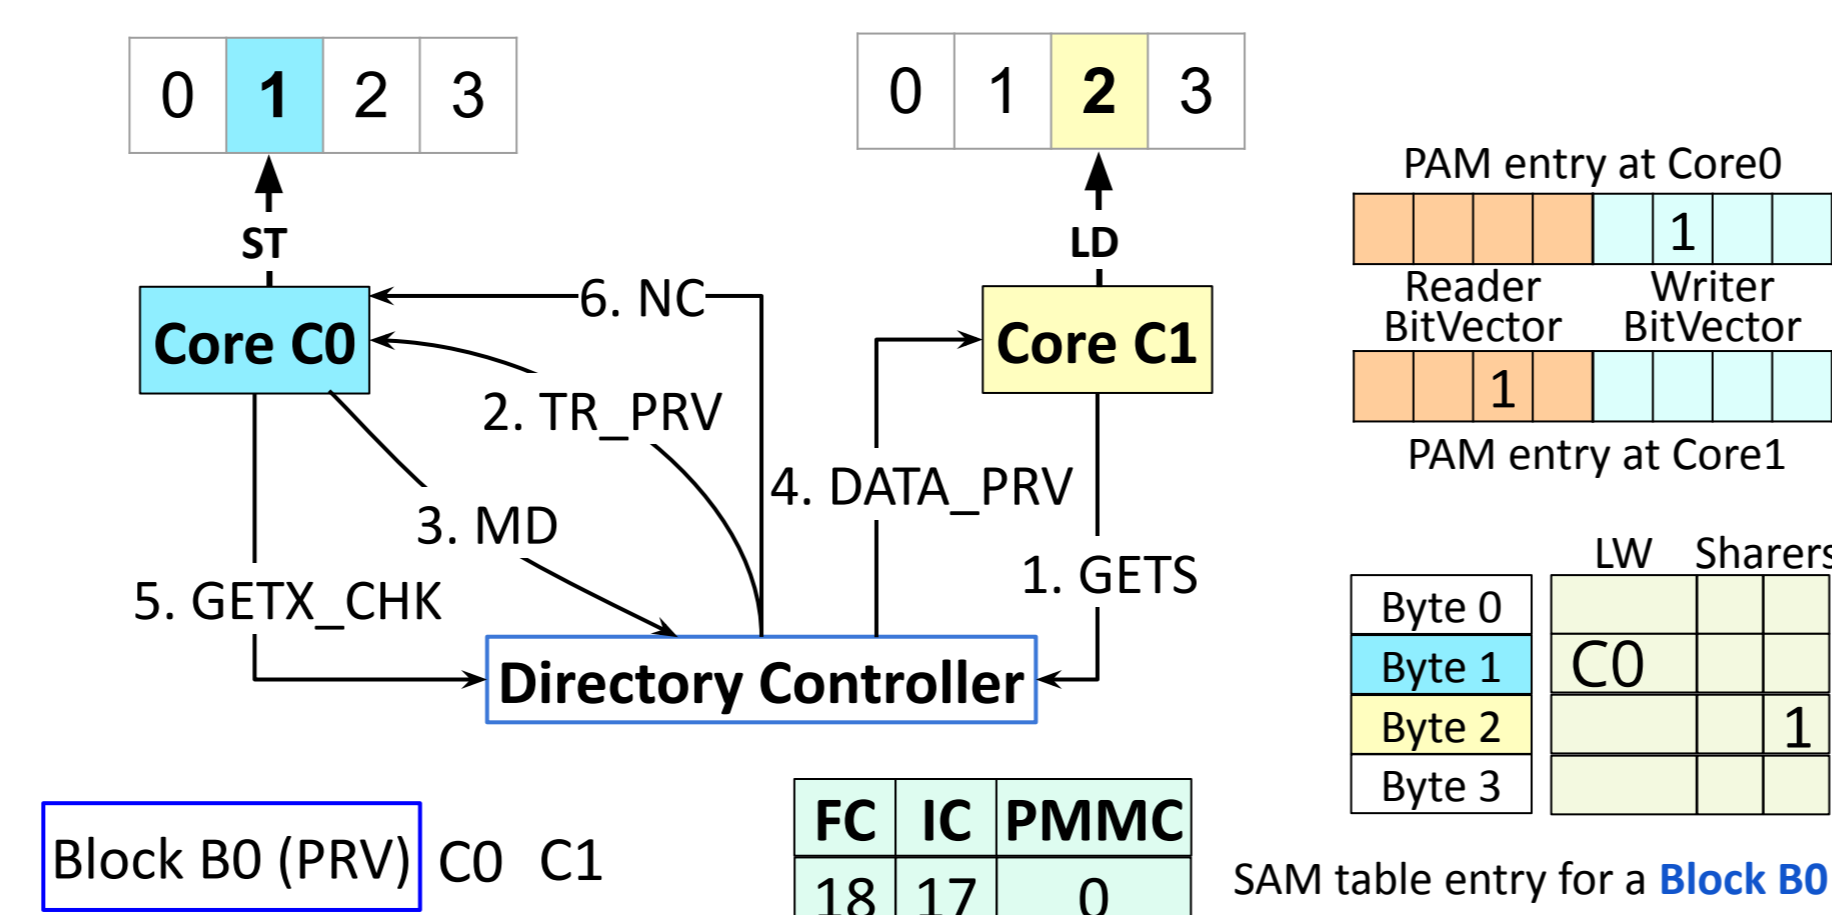


Impactful instance of false sharing

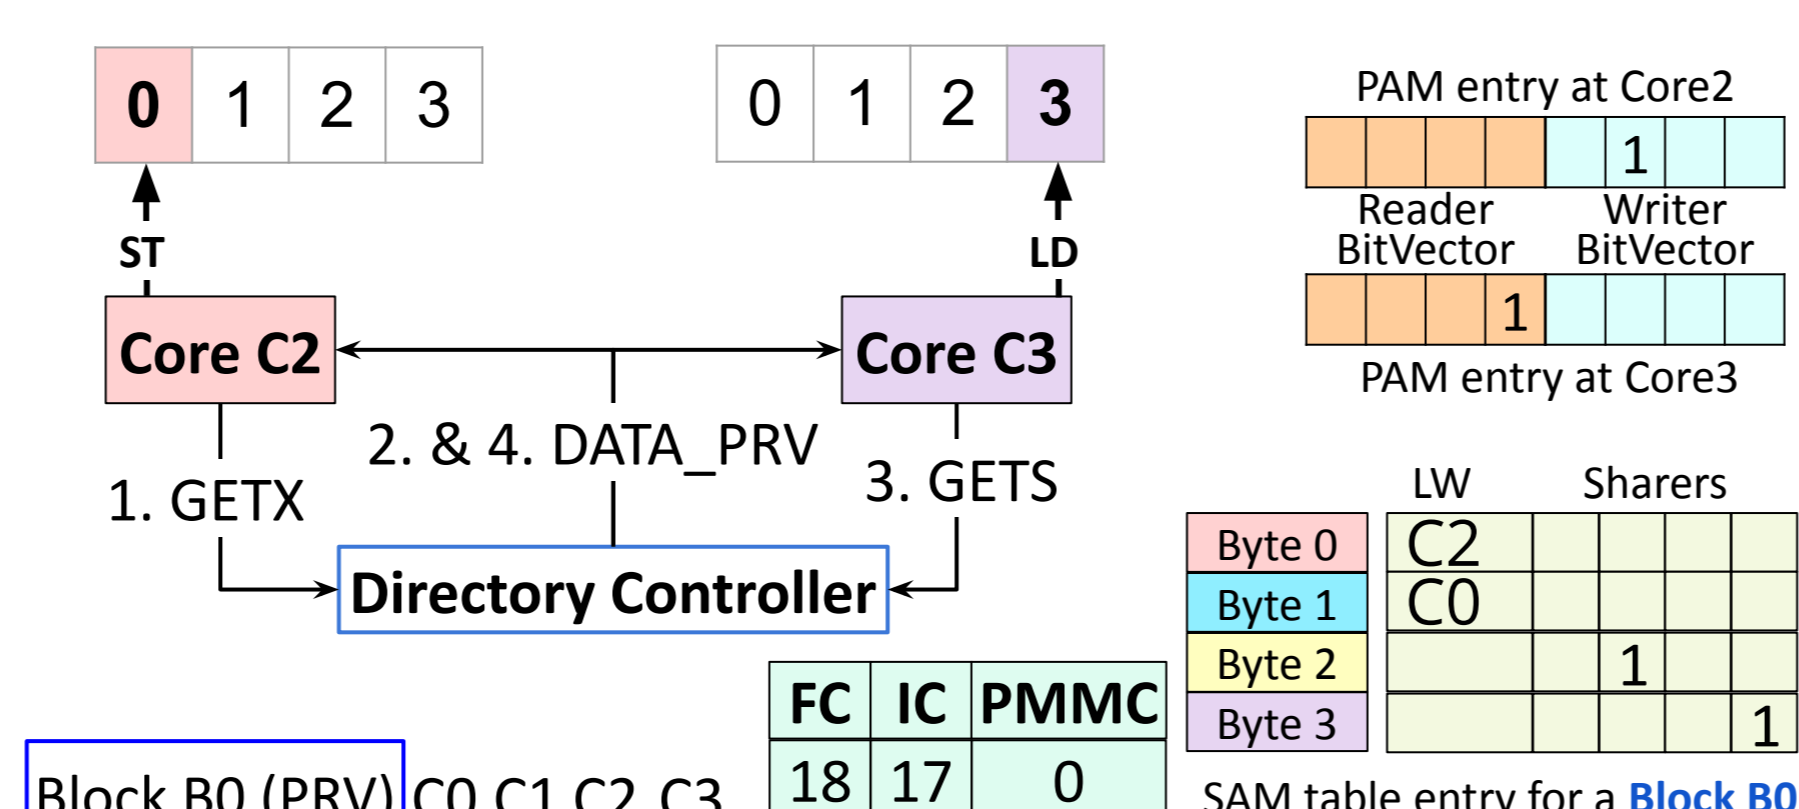
## Necessary conditions to trigger Privatization

- Fetch Counter >  $\tau_p$
- Invalidation Counter >  $\tau_p$
- Pending Metadata Messages == 0
- Directory performs a conflict detection check
- Access history in SAM has no overlapping accesses

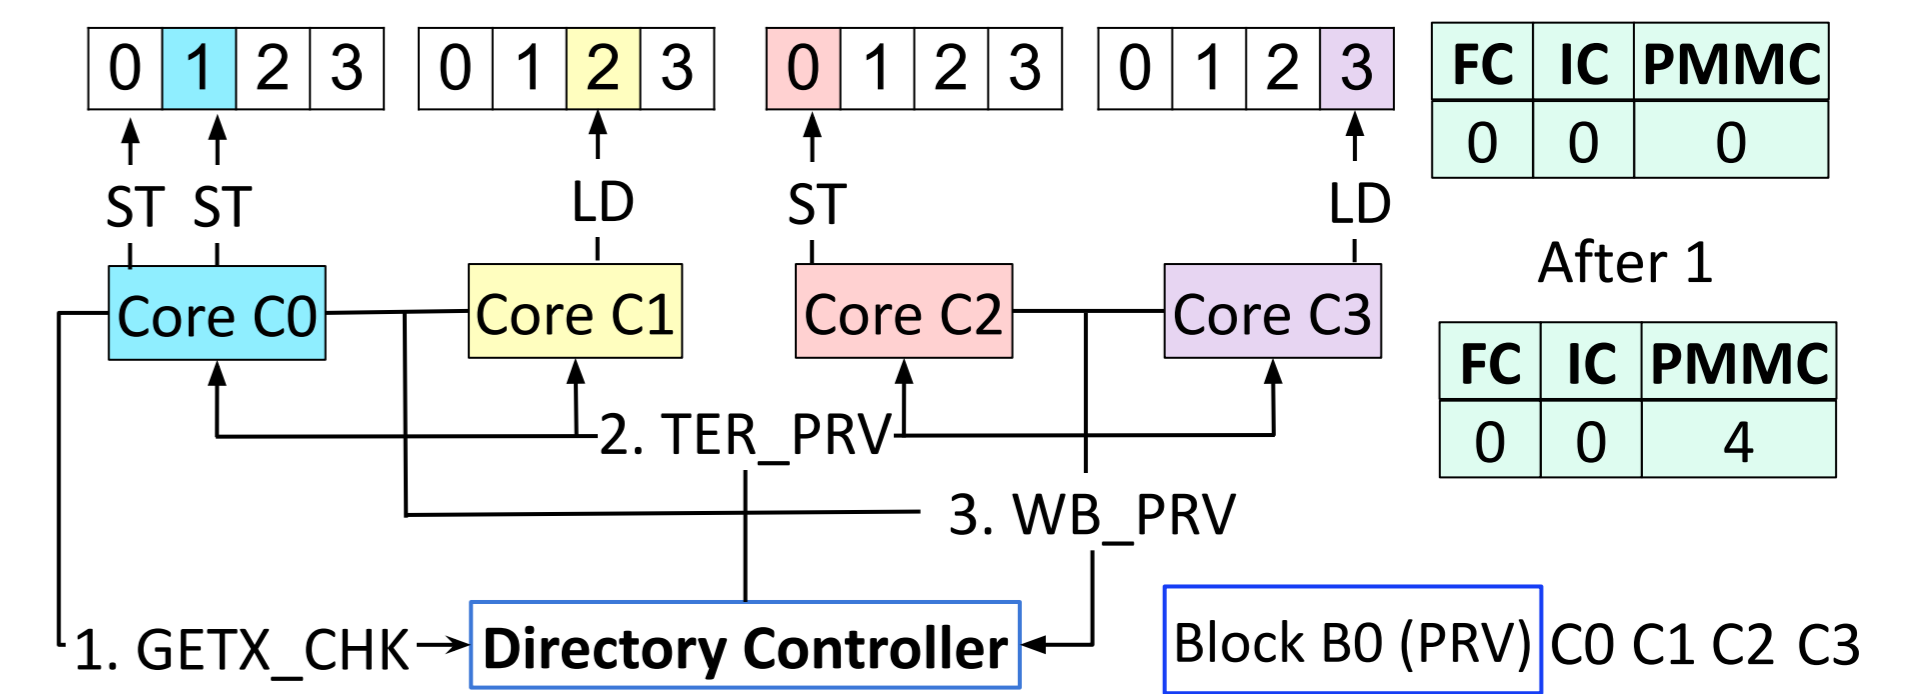
## FSLite: Privatization of falsely shared blocks



## Handling Request during Privatization

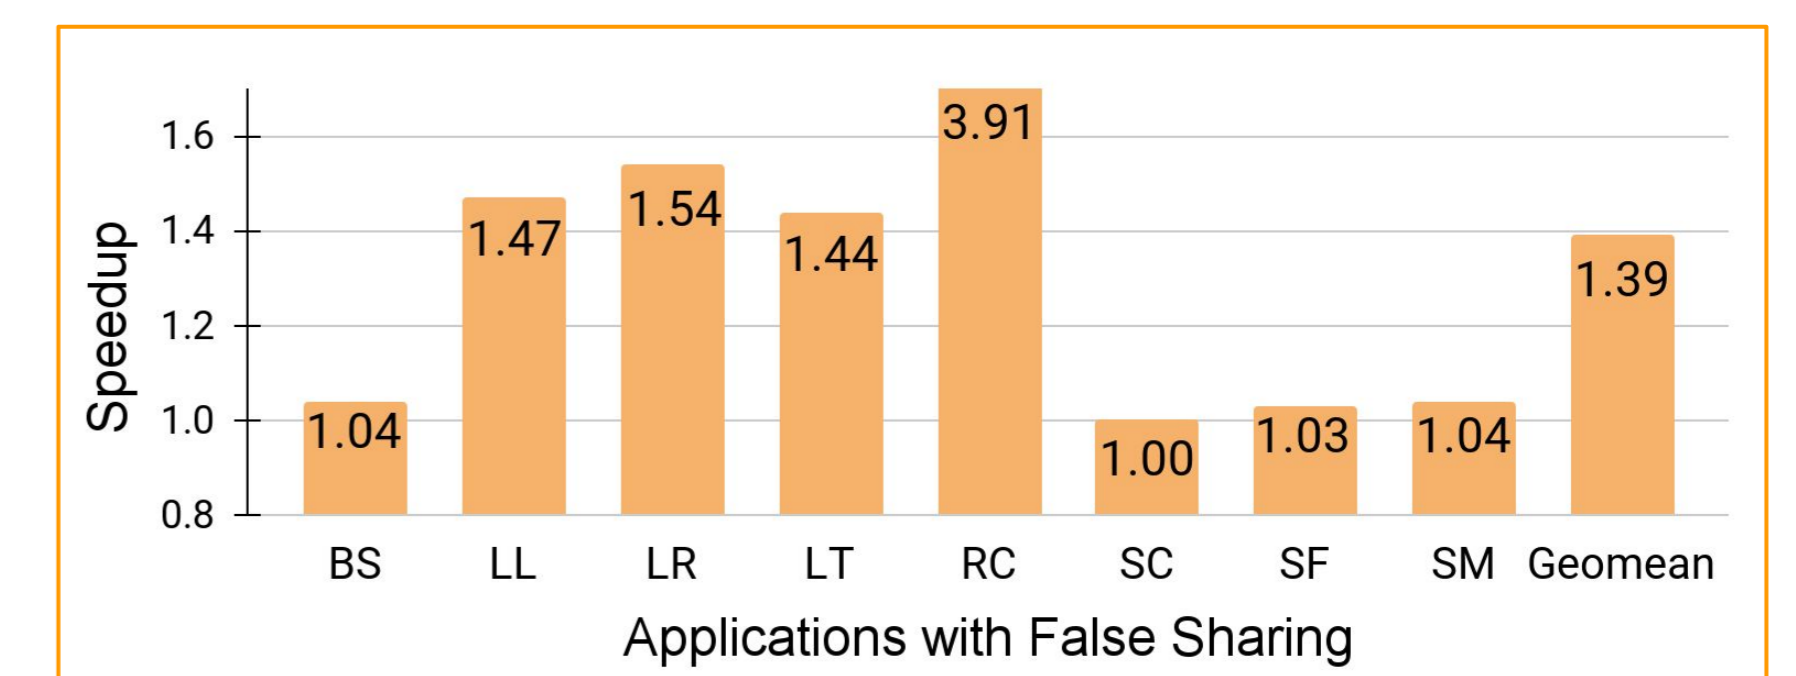


## Termination of Privatization



- Core C1 access (store) Byte 1 of the block B0
- C1 sends an access conflict detection request (GETX\_CHK) to the directory
- Directory identifies the conflict with C2 and sends termination requests (TER\_PRV) to all the sharers
- Each core responds with write back message (WB\_PRV) to the directory
- Directory updates the LLC copy of the block by merging the bytes modified by each core using **Write-Mask**


## Performance Improvement with FSLite



FSLite achieves on avg **1.39x** speedup and a max speedup **3.91x**

**27%** energy savings and **89%** less interconnect messages

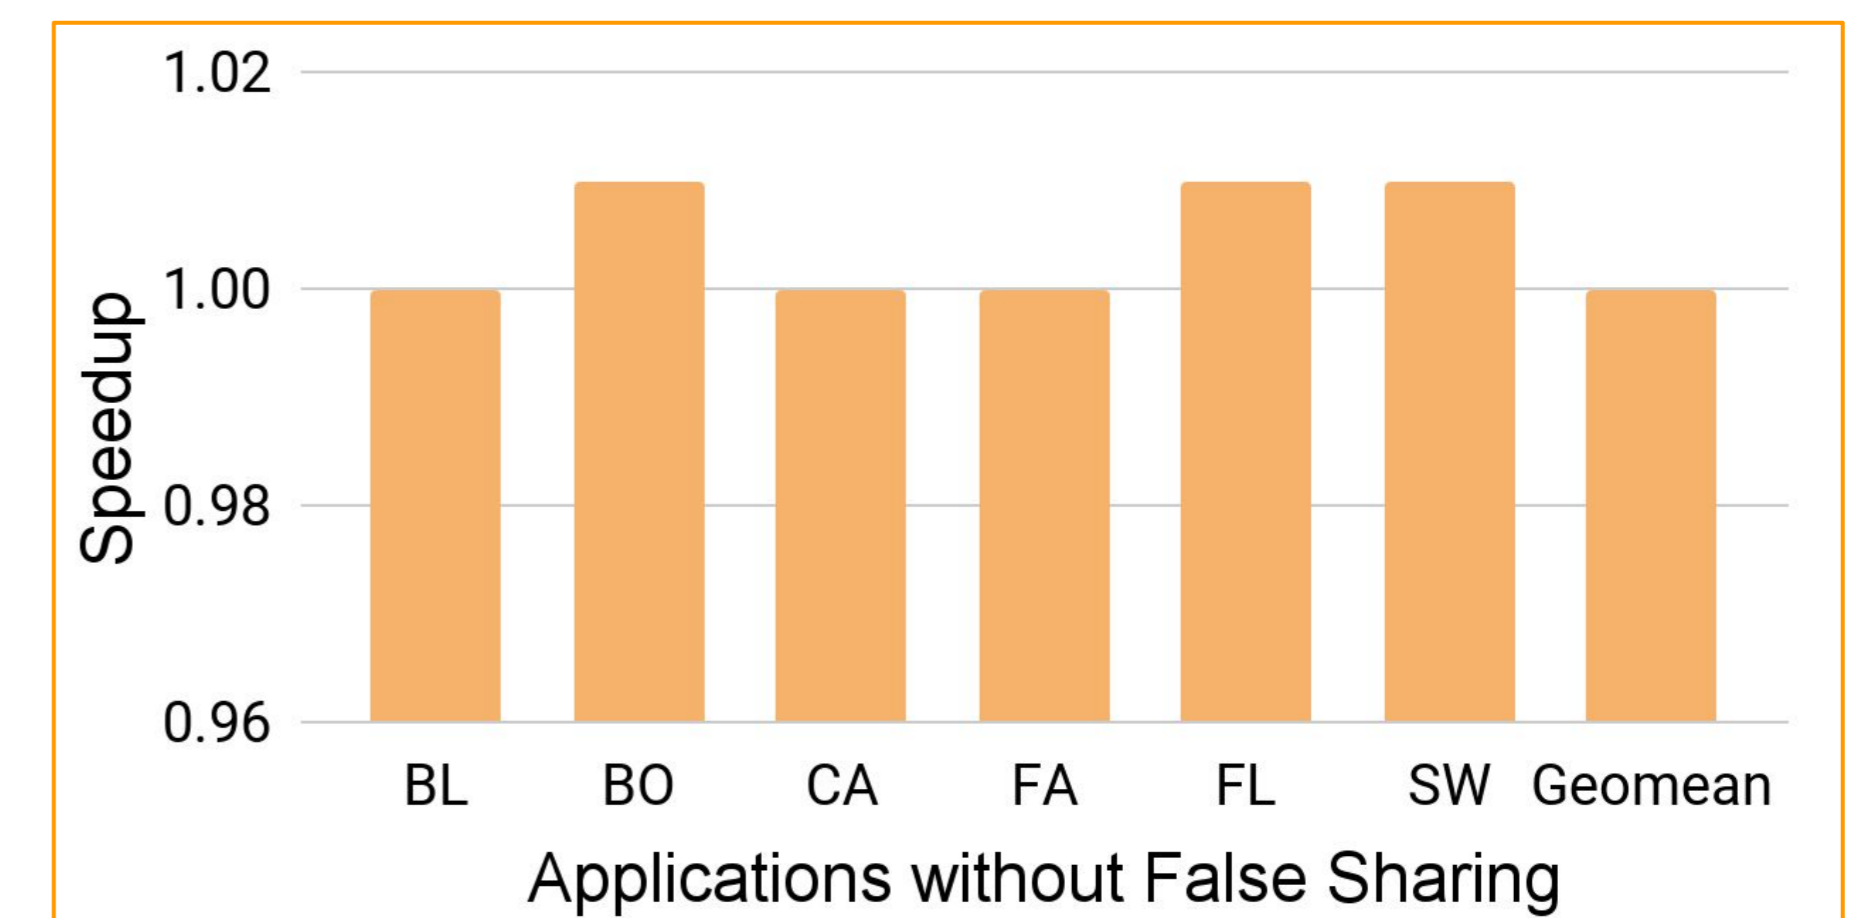
## Comparison with Manual Fix



Padding **inflates the memory footprint**

Padding **introduces additional instructions**

## Impact of FSLite on Applications w/o False Sharing



**No slowdown for applications without false sharing**

## Additional Results

- FSLite achieves an average speedup of **1.63x** for O-o-O cores
- FSLite achieves an average speedup of **1.21x** over a baseline with 128KB L1D cache across all applications

## References

- Emery D. Berger et al. "Hoard: A Scalable Memory Allocator for Multithreaded Applications." ASPLOS'00.
- Mihir Nanavati et al. "Whose Cache Line Is It Anyway? Operating System Support for Live Detection and Repair of False Sharing." EuroSys'13.
- Liang Luo et al. "LASER: Light, Accurate Sharing dETection and Repair" HPCA'16.