


Erasing Belady's Limitations: In Search of Flash Cache Offline Optimality

Yue Cheng^{1,2}, Fred Douglass², Philip Shilane², Michael Trachtman², Grant Wallace², Peter Desnoyers³, Kai Li⁴
¹Virginia Tech, ²EMC Corporation, ³Northeastern University, ⁴Princeton University

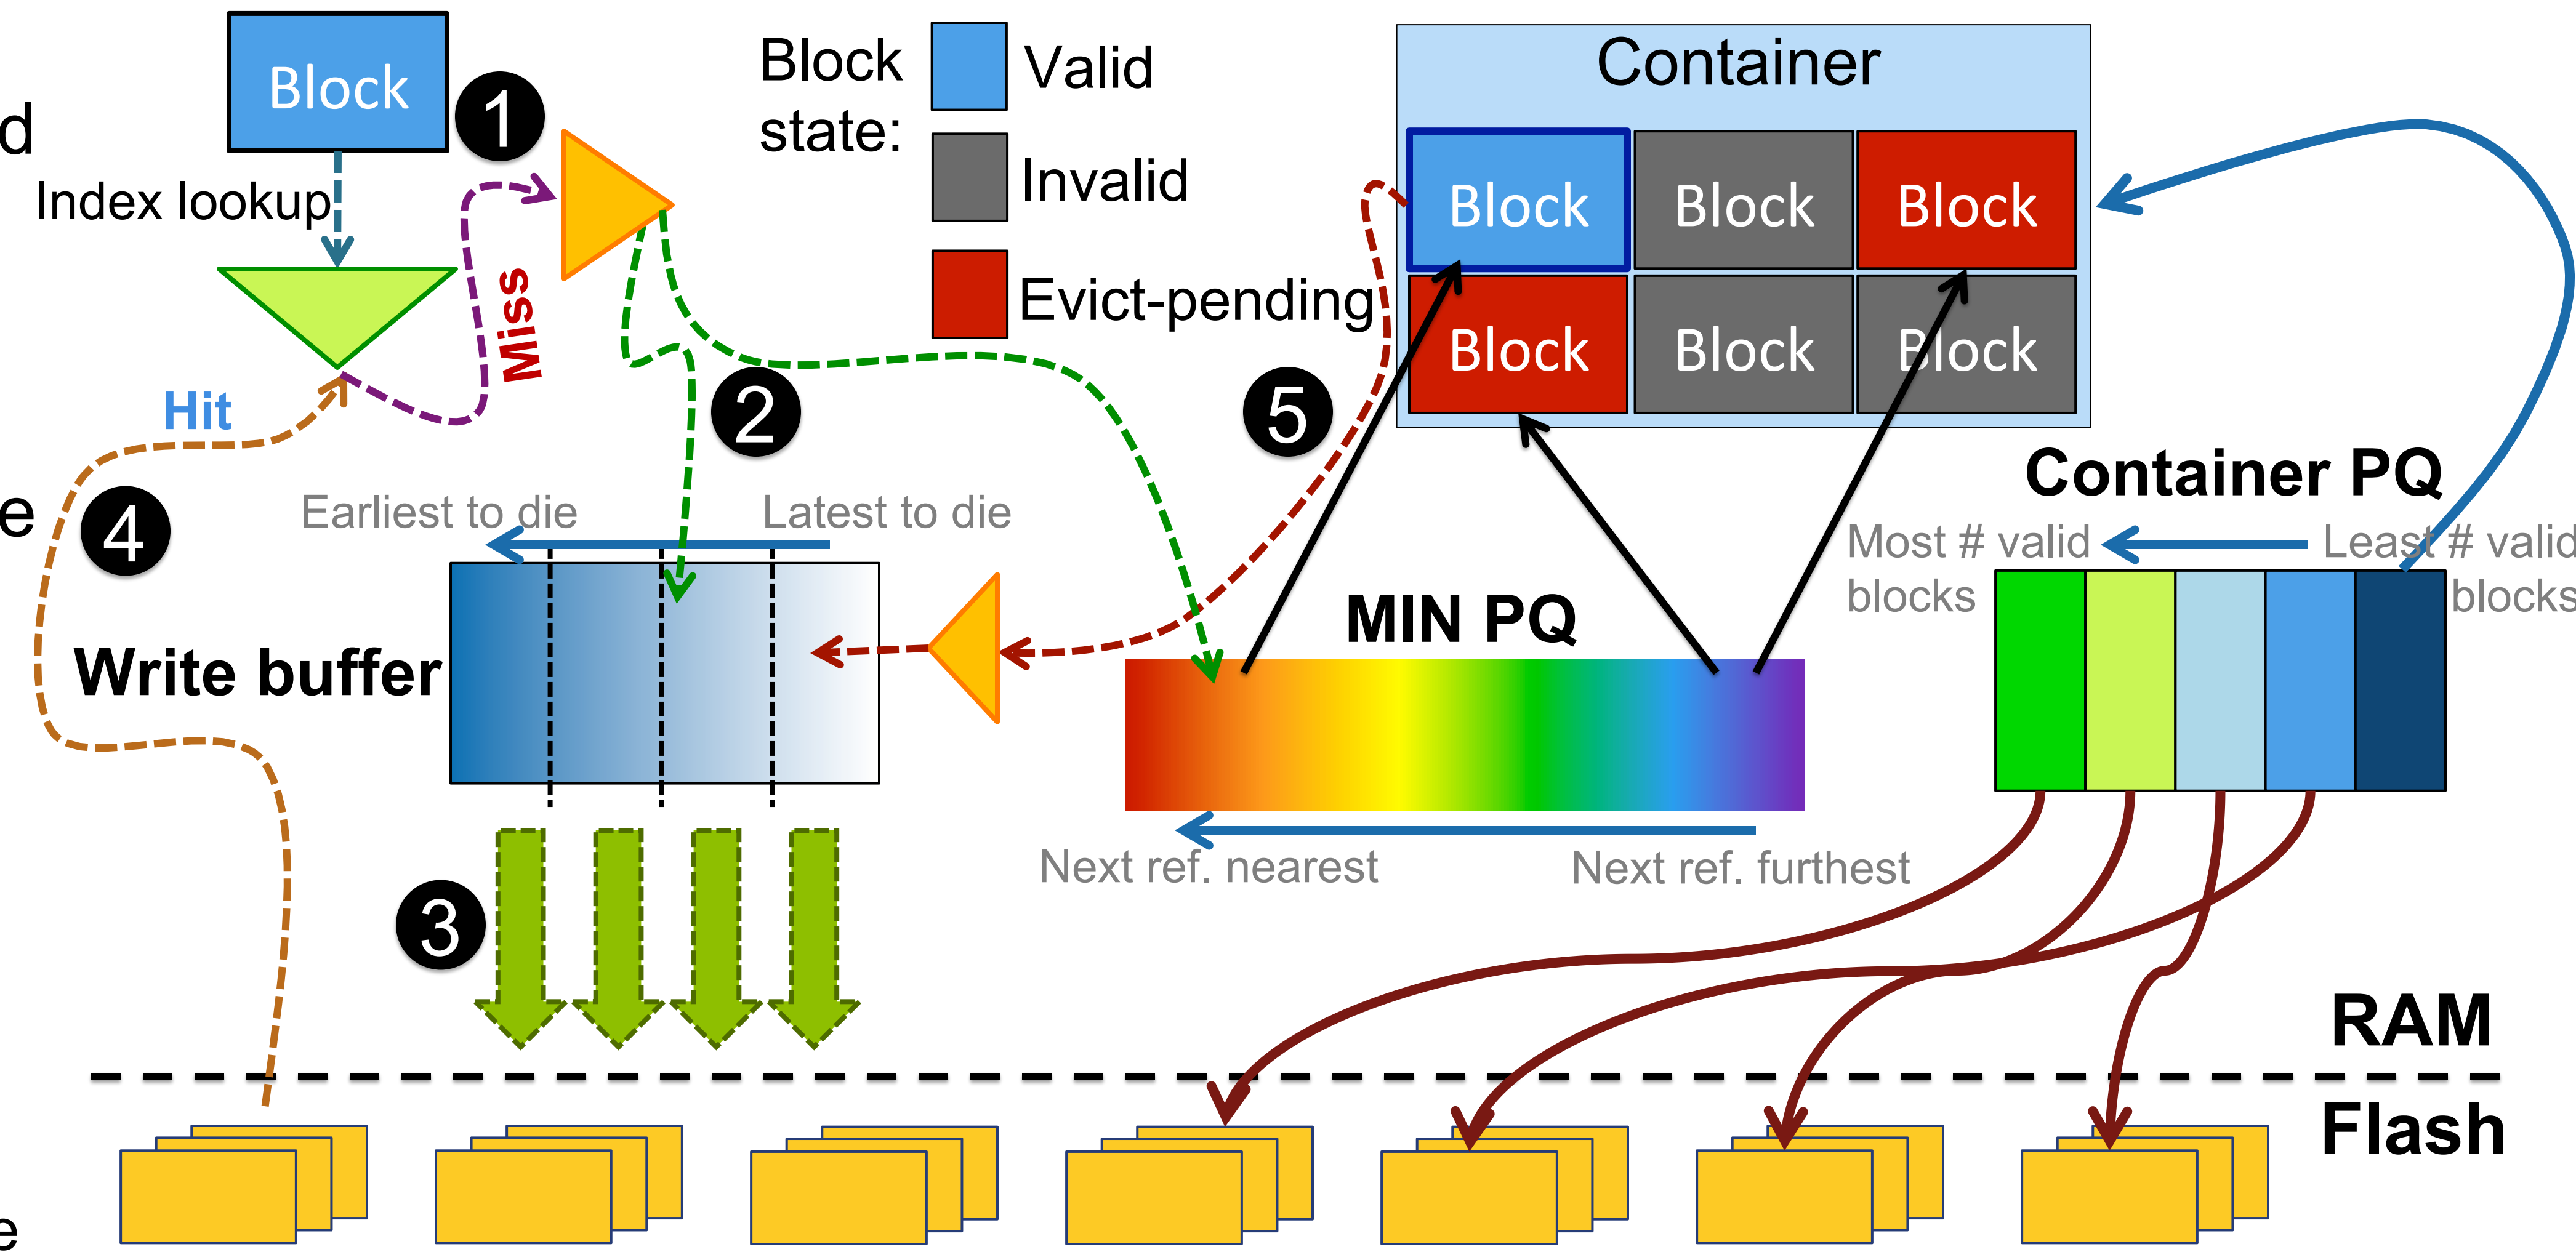
Motivation

- Using SSDs for an HDD cache
- SSD cache btw. DRAM & HDD
 - Goal: to balance performance against endurance
- Nitro [ATC'14], CacheDedup [FAST'16]
- RIPQ [FAST'15]
- Pannier [Middleware'15]
- Are we doing well?
 - Comparison against an offline "best case"
 - But what is the offline optimal for flash cache?



Container-optimized Heuristic

- On a **read miss**, check if the block will be read before evicted
 - Pack a block into in-RAM write buffer and insert it to MIN PQ*
 - Disperse the write buffer into containers and write them to the flash cache
 - On a **read hit**, read the block from the flash cache
 - At garbage collection, copy-forward (CF) the valid block to the write buffer by checking if it will be reread before evicted
- *PQ=Priority queue



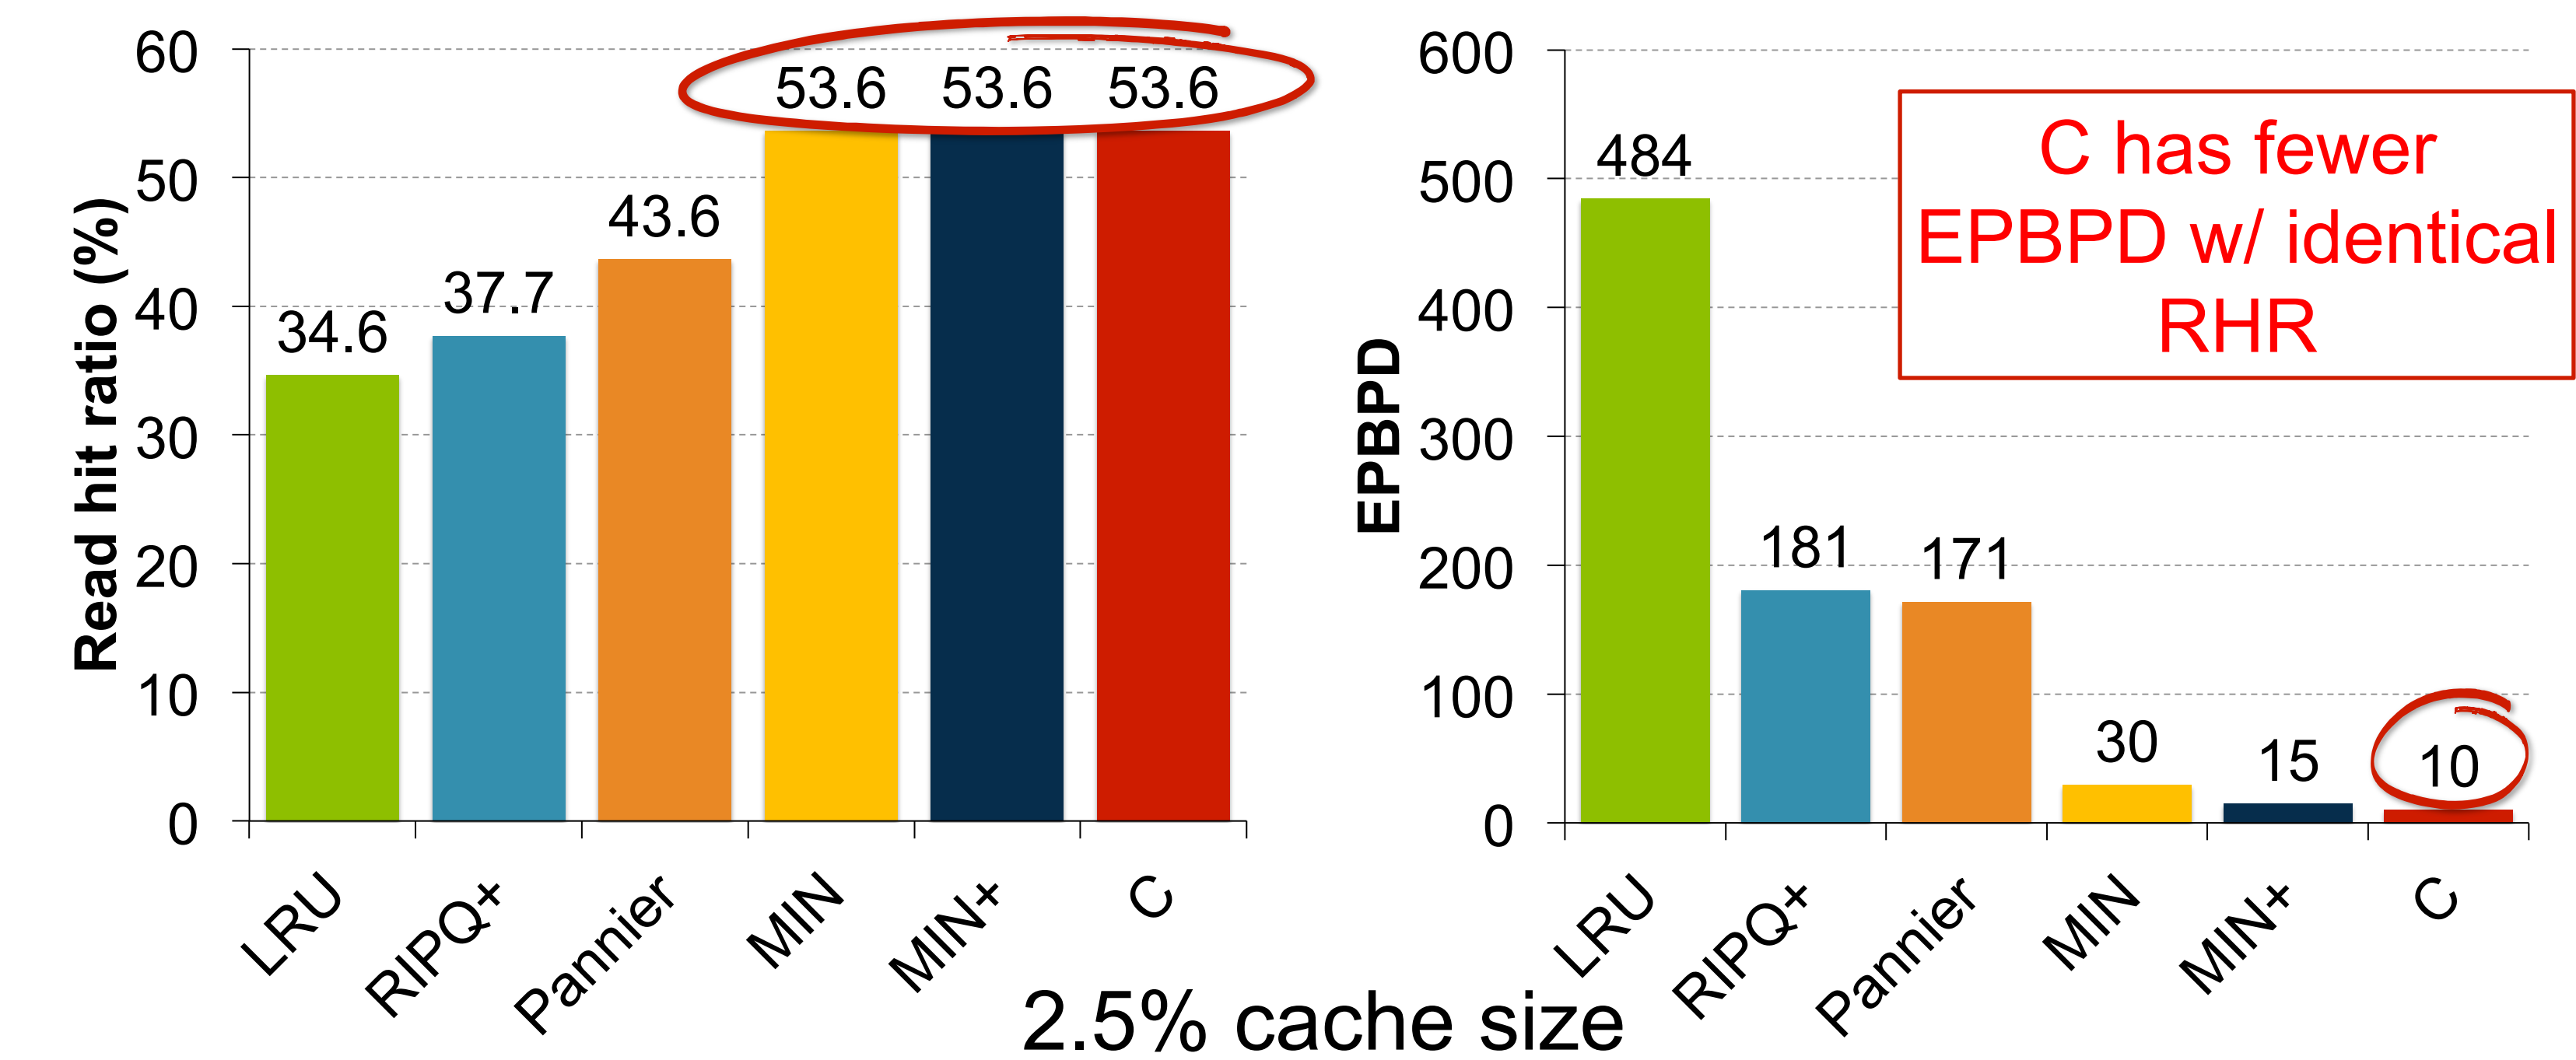
Background: Offline Optimality

- Belady's MIN: A simple offline caching alg. **when the next access is known**
 - Evicts the item that will be used furthest in the future
 - Yields the optimal read hit ratio (RHR)
 - Ignorant of minimizing erasures/block/day (EPBPD)
- MIN is not able to provide the optimal erasures in the context of flash caching
 - MIN inserts items that **won't actually be read**

Comparing Algorithms

Policy	Description	O*	C+
LRU	Least recently used	X	X
RIPQ+	Static web content	X	X
Pannier	Handles divergent containers	X	✓
MIN	Don't insert data w/ furthest next ref	✓	X
MIN+	Don't insert data evicted w/o read	✓	✓
C	Our container-optimized heuristic	✓	✓

*O=Offline, +C=Container-optimized

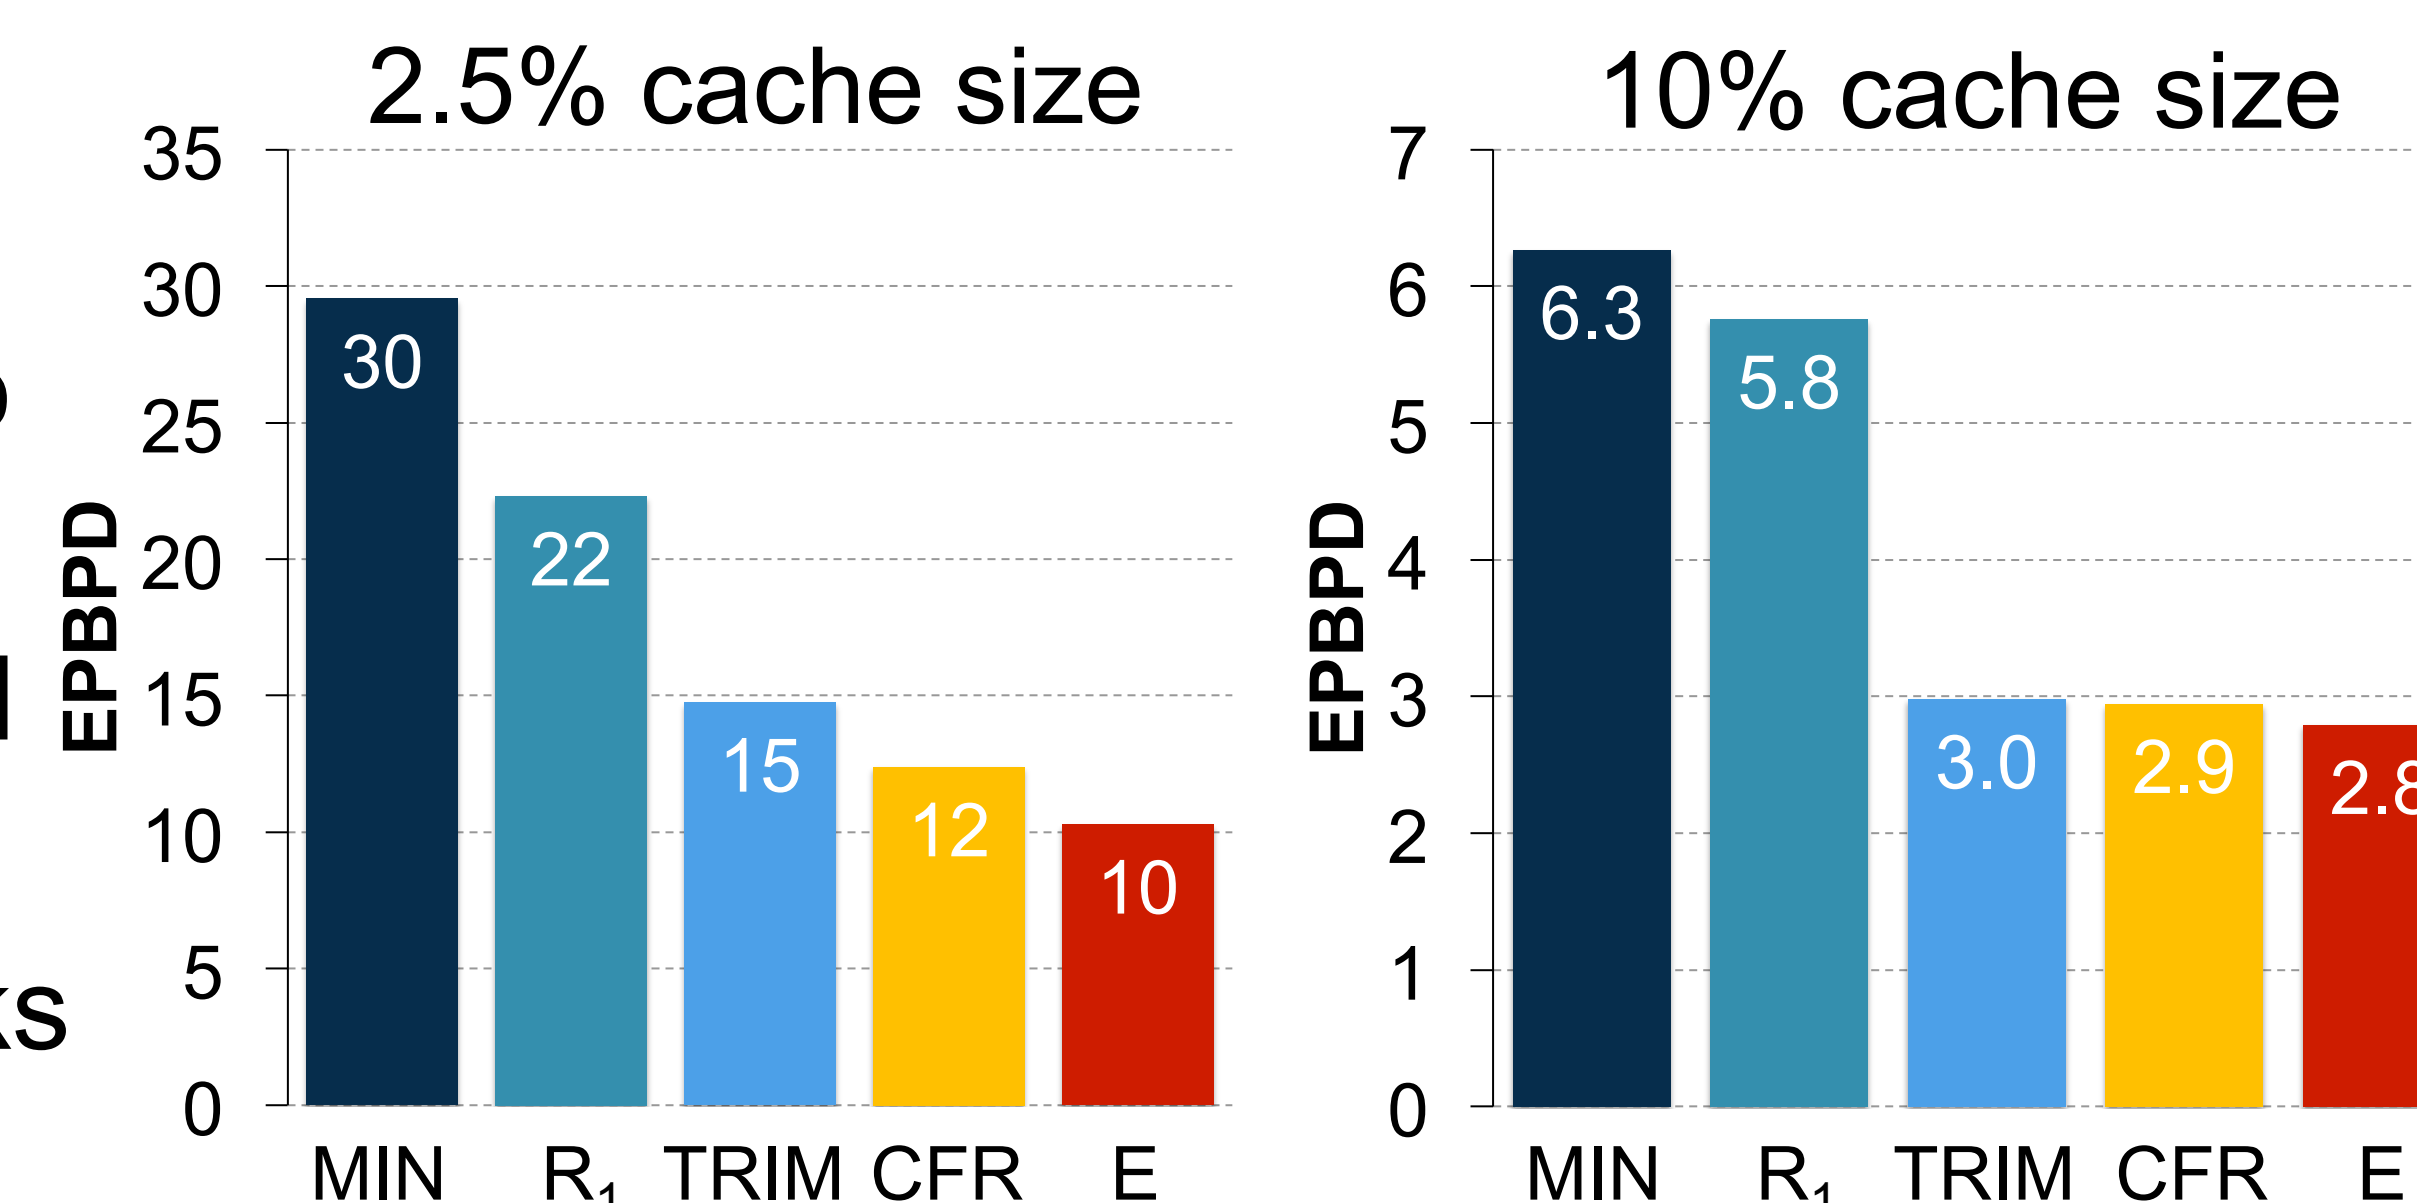


Offline Flash Caching

- Objectives
 - Minimize erasures s.t. maximal RHR
 - Never insert items if it does not increase RHR
 - Other objectives possible as discussed in the paper
- True optimal vs. Heuristic
 - Complexity of true optimal?
 - Approximation** is the focus of this work

Evaluating Optimization Techniques

- R₁**: Omit insertions w/ no reread
- TRIM**: Notify GC to omit dead blocks
- CFR**: Avoid wasted CF blocks
- E**: Segregate blocks by evict timestamp



Conclusion

- Important to have a baseline for the offline optimal considering **both RHR and endurance**
- Additional optimizations may be possible to move our heuristic to the **true optimal**

