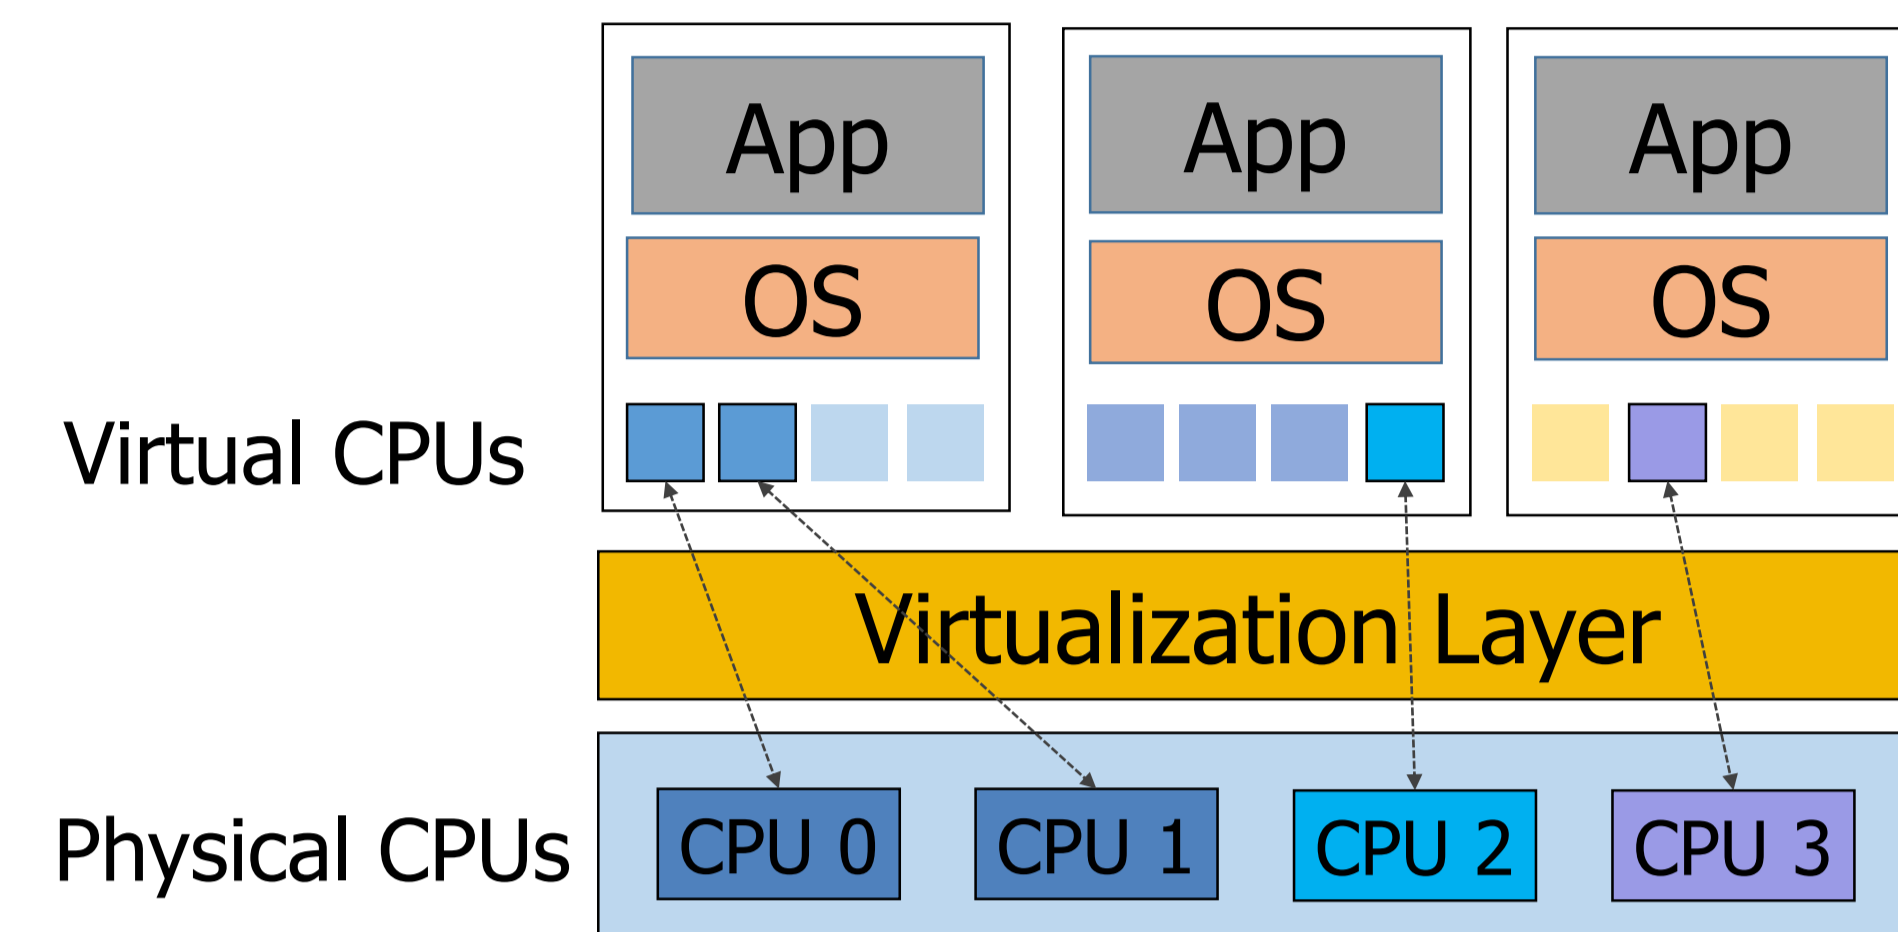


# Micro-sliced Virtual Processors

to Hide the Effect of Discontinuous CPU Availability for Consolidated Systems


Jeongseob Ahn, Chang Hyun Park, and Jaehyuk Huh

## CPU Scheduling in Consolidated Systems



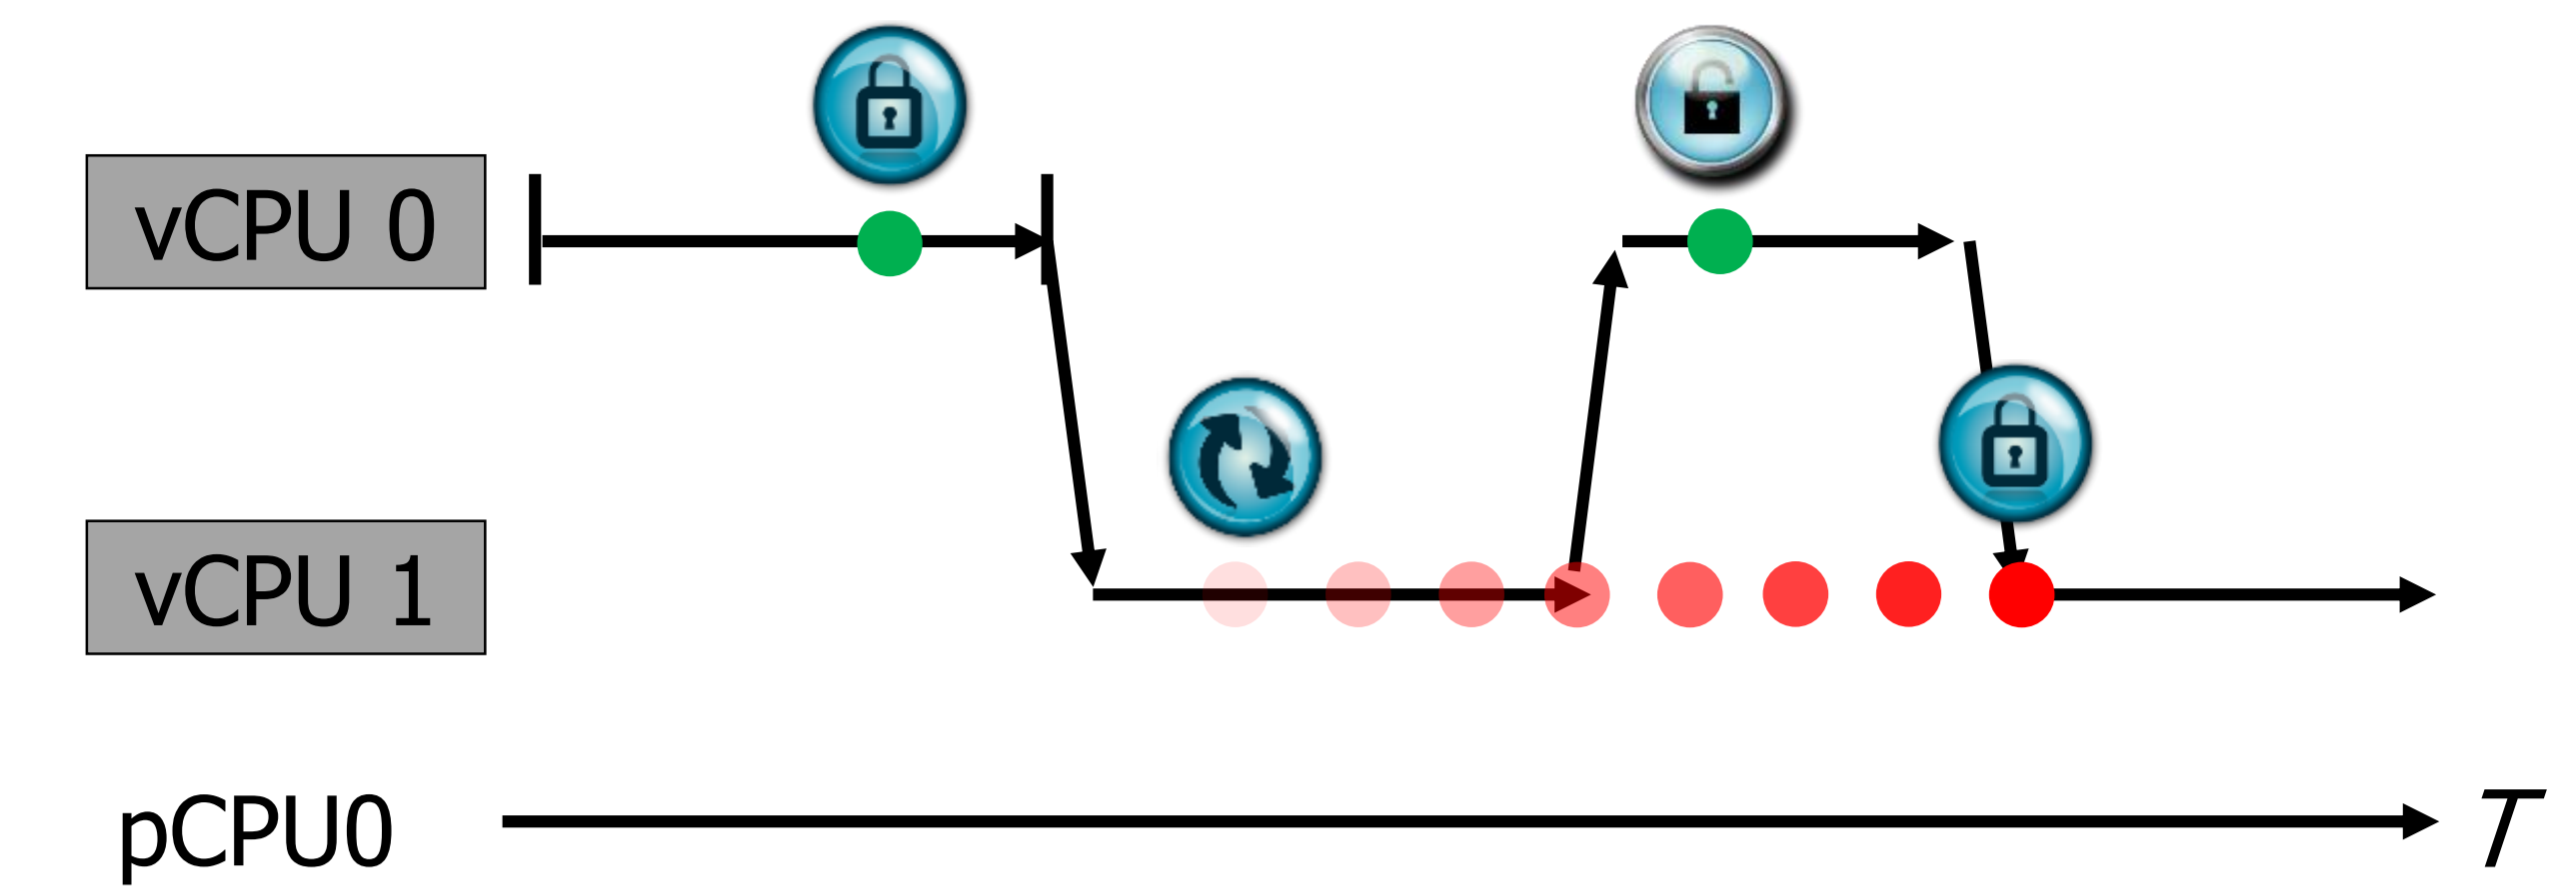
- Virtual machines **share** the physical CPUs by time sharing
- Virtual CPUs are not always running on physical CPUs

## Interrupt Handling Delay

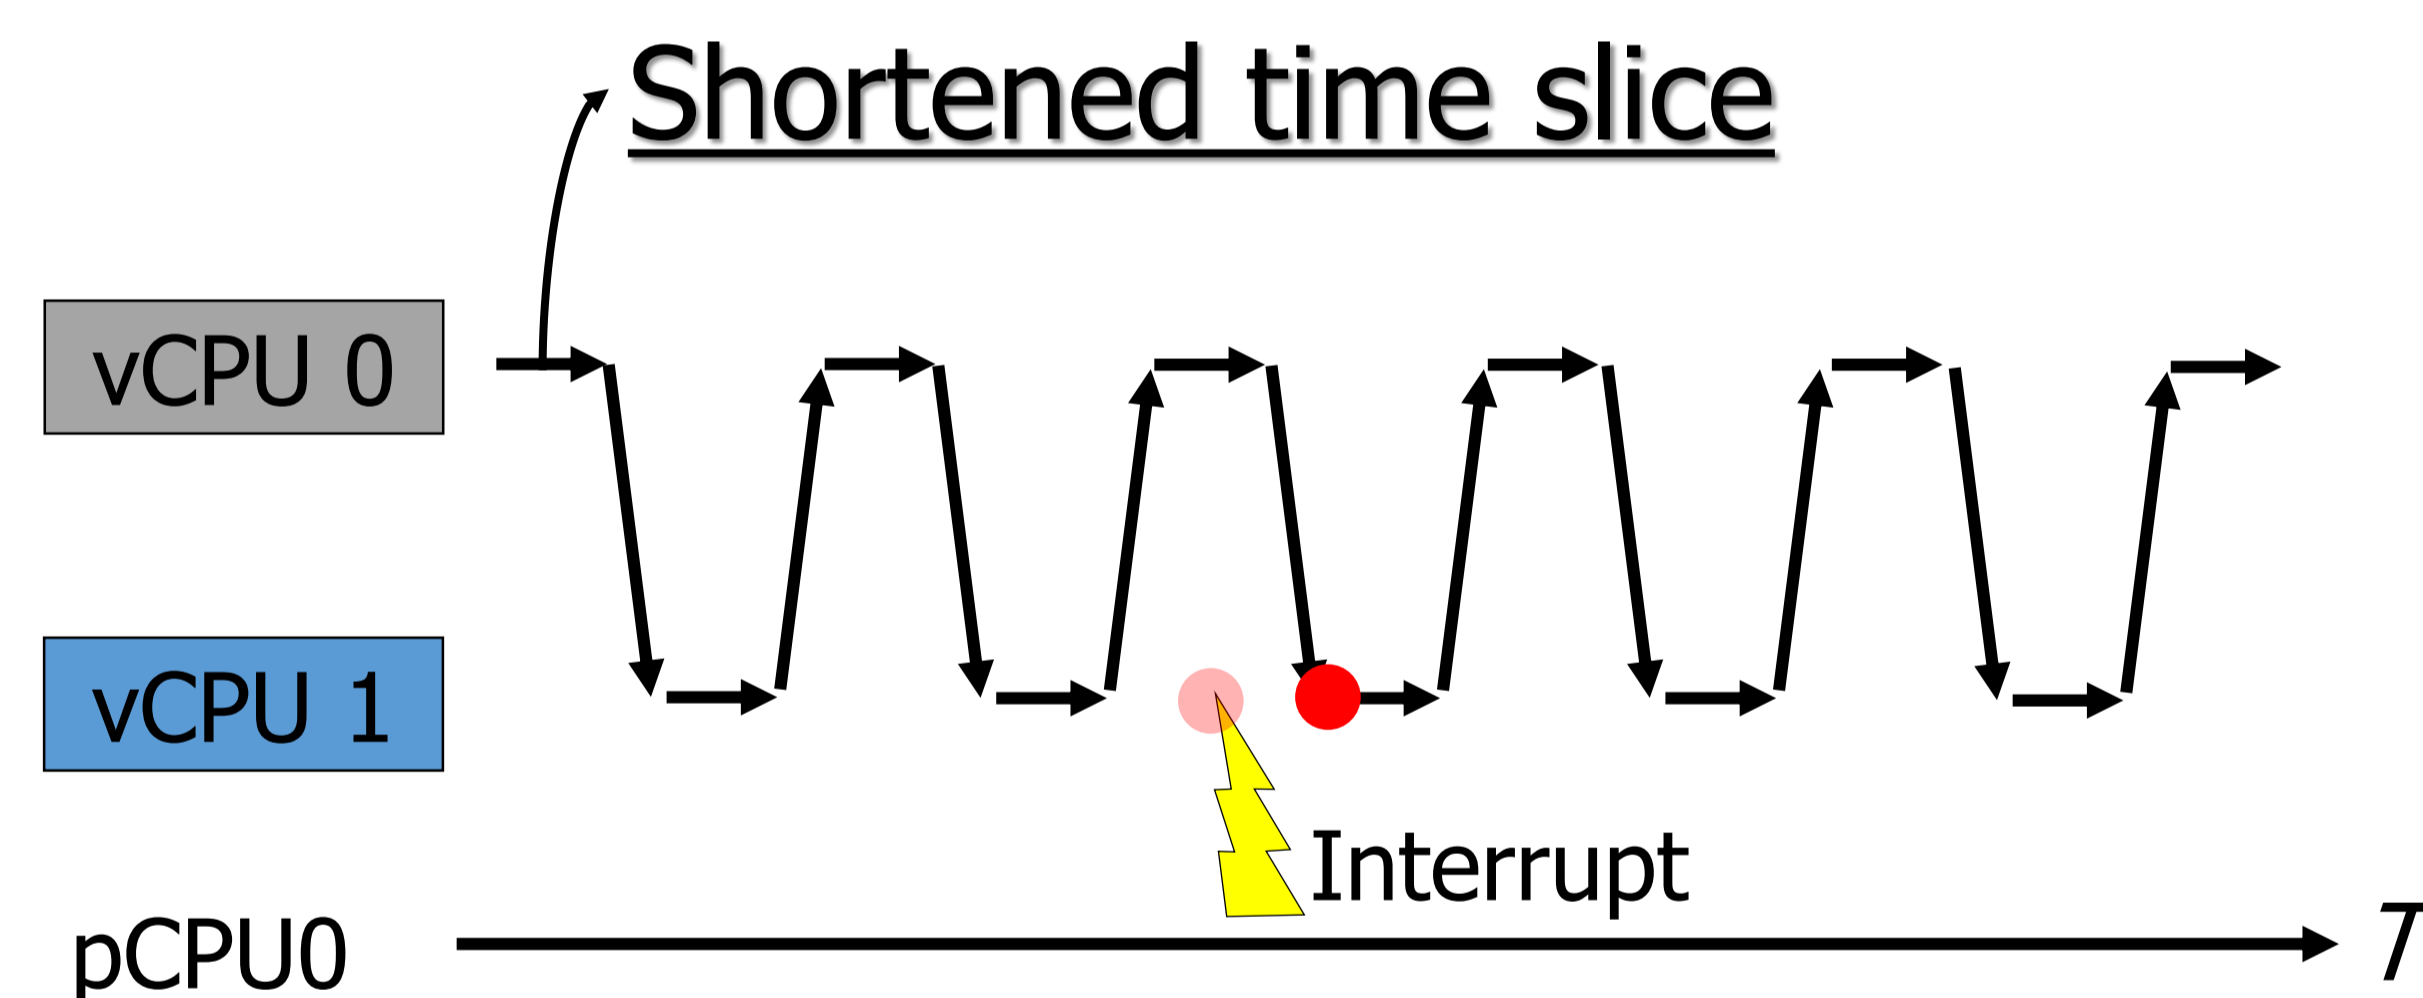


- Guest OS assumes the continuous execution of its vCPUs, but each vCPU is periodically scheduled in and out

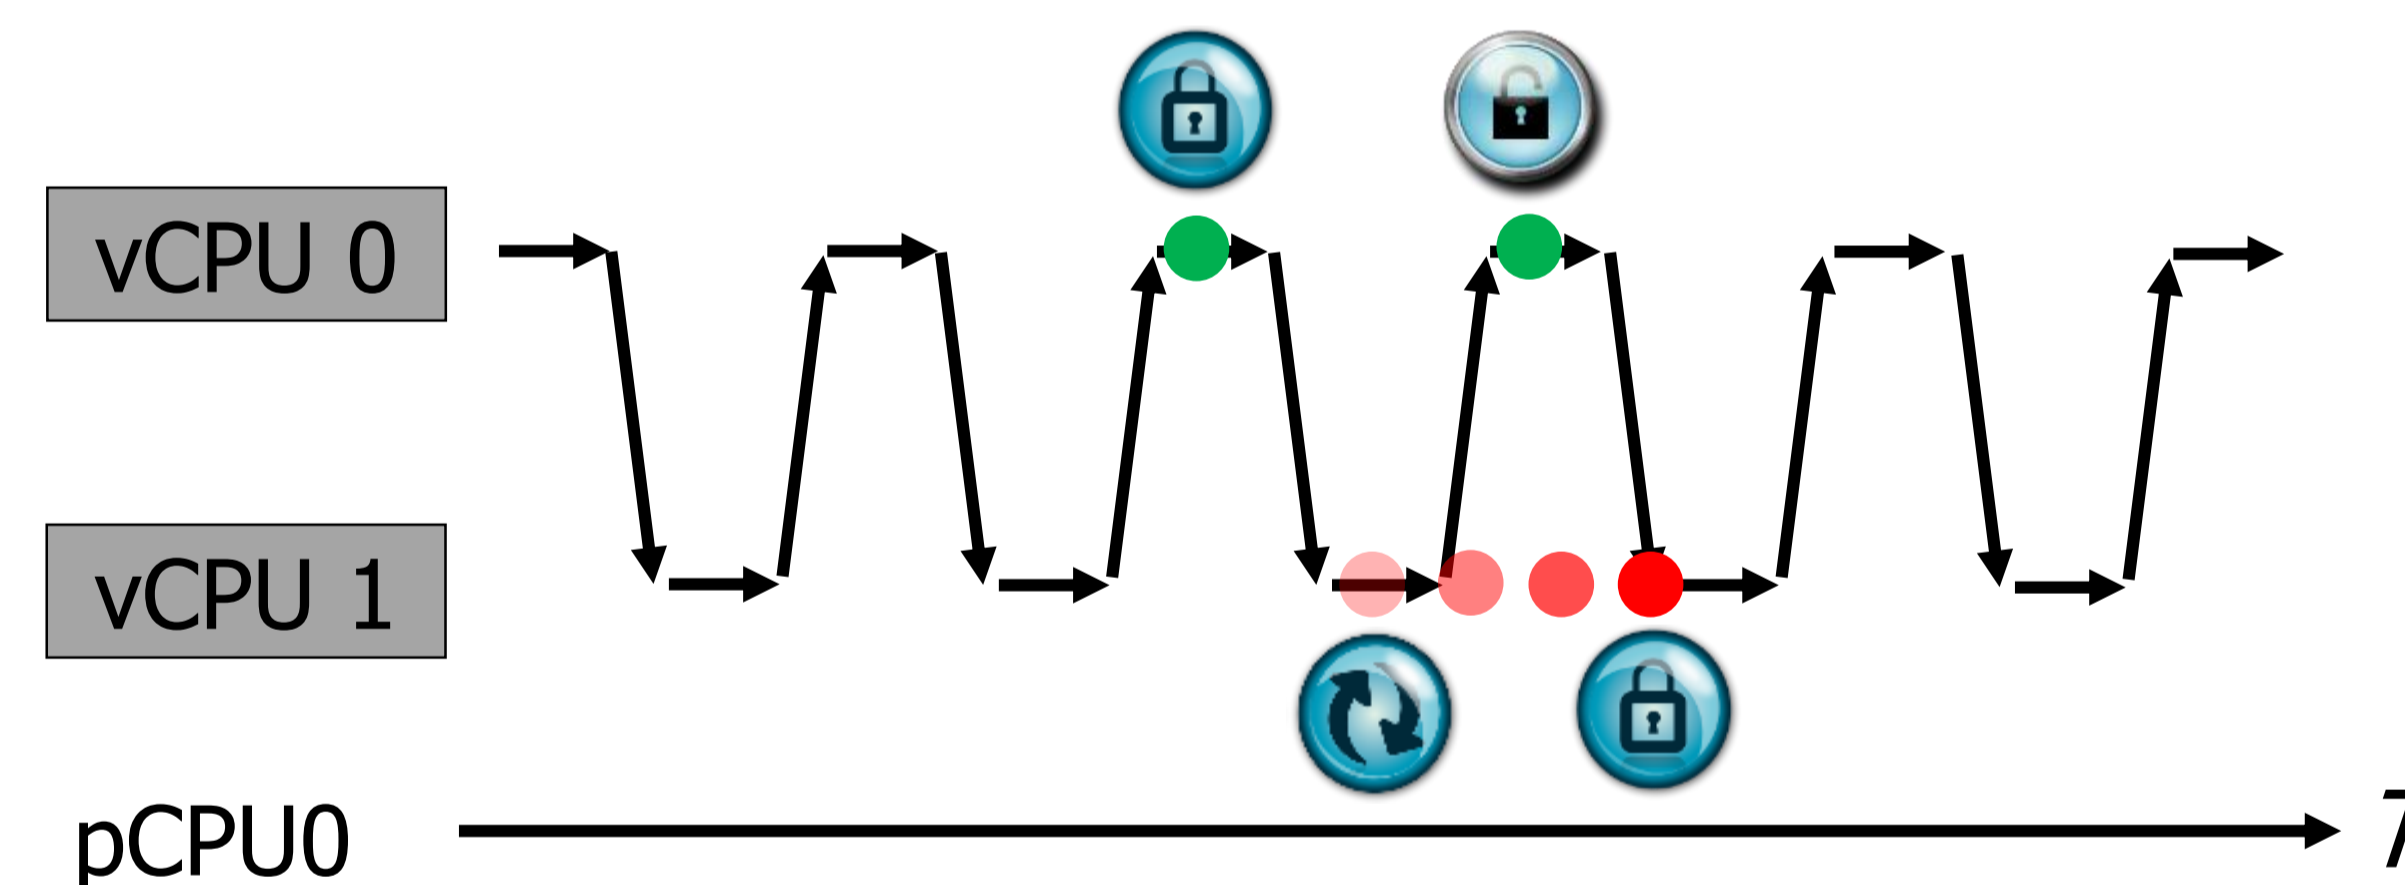
## Lock Holder Preemption



## Our Approach "Short, but more Frequent Runs"



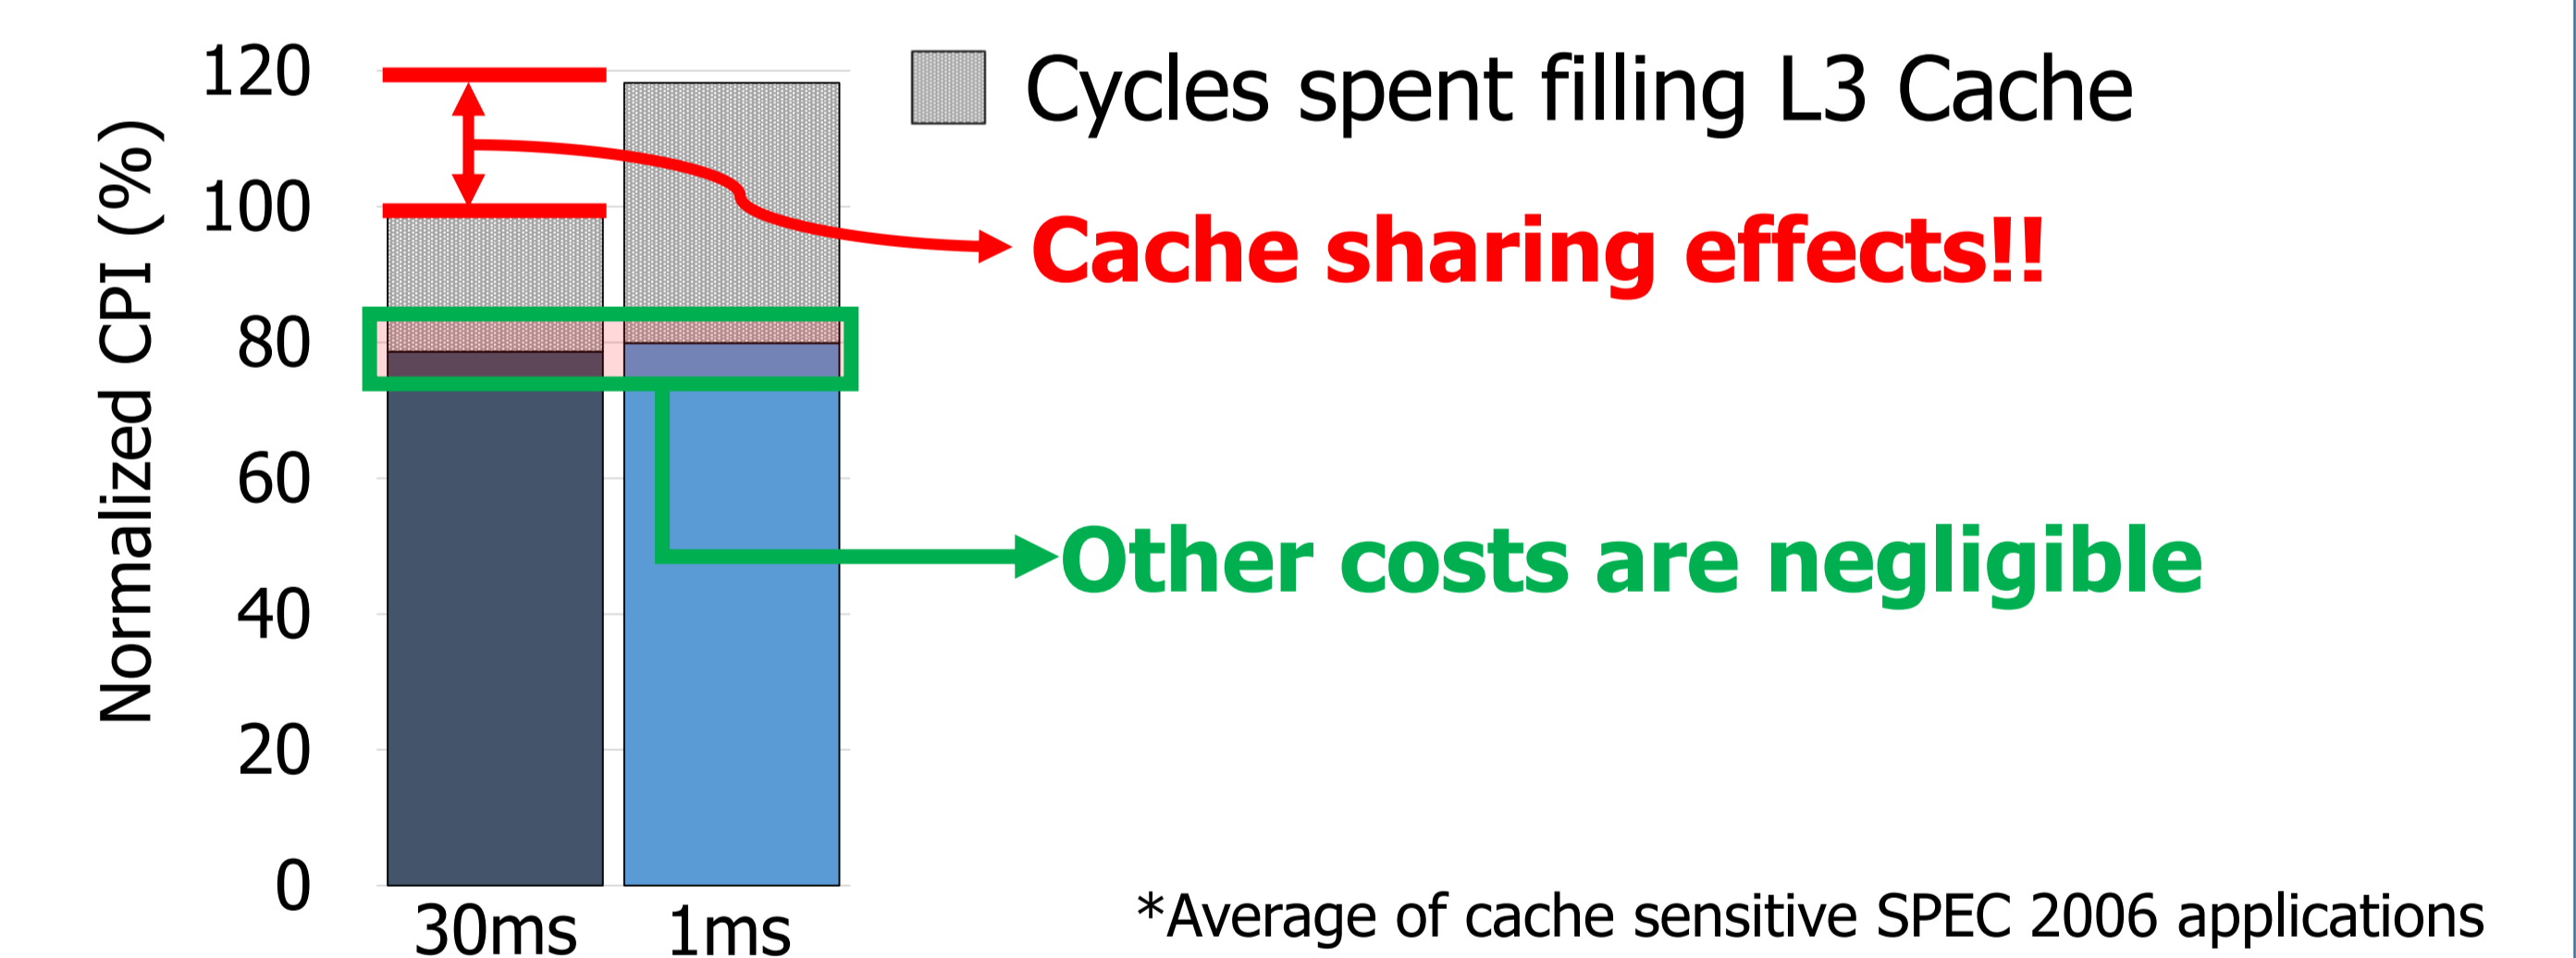
- Interrupt is handled within a short time



- Lock is acquired sooner

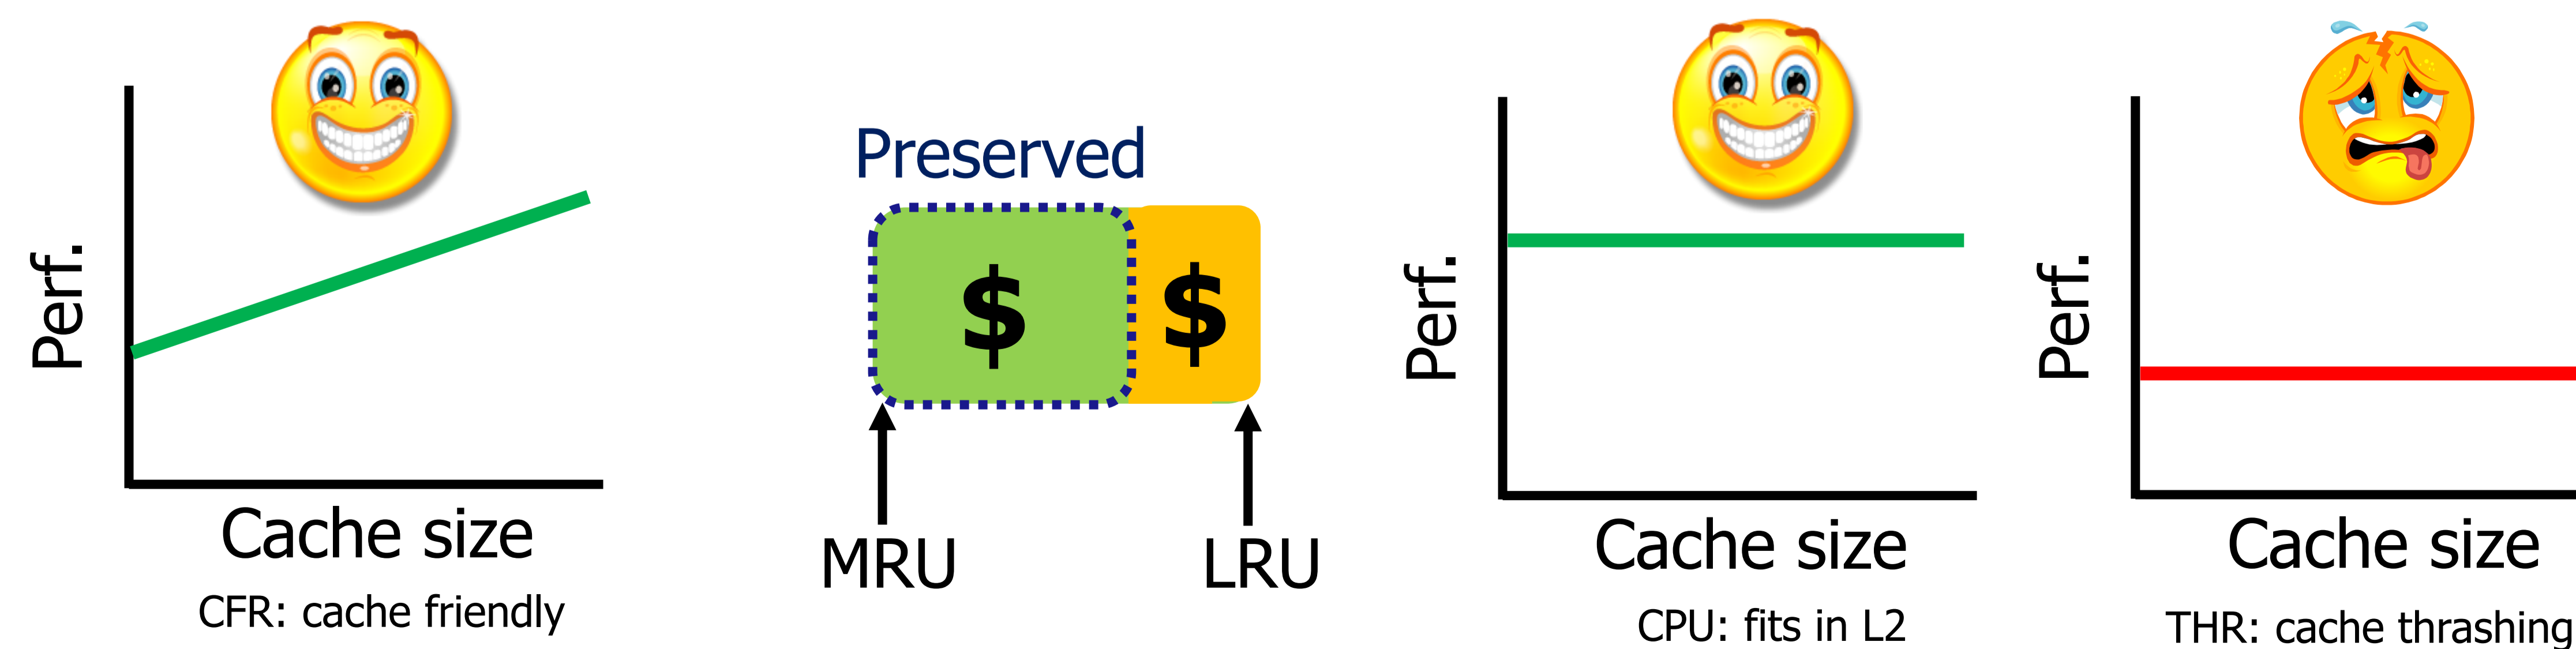
## What are the Negative Effects ?

- Frequent context switching cost (save & restore registers)
- Pollution of architectural structures (Caches, TLBs, etc.)

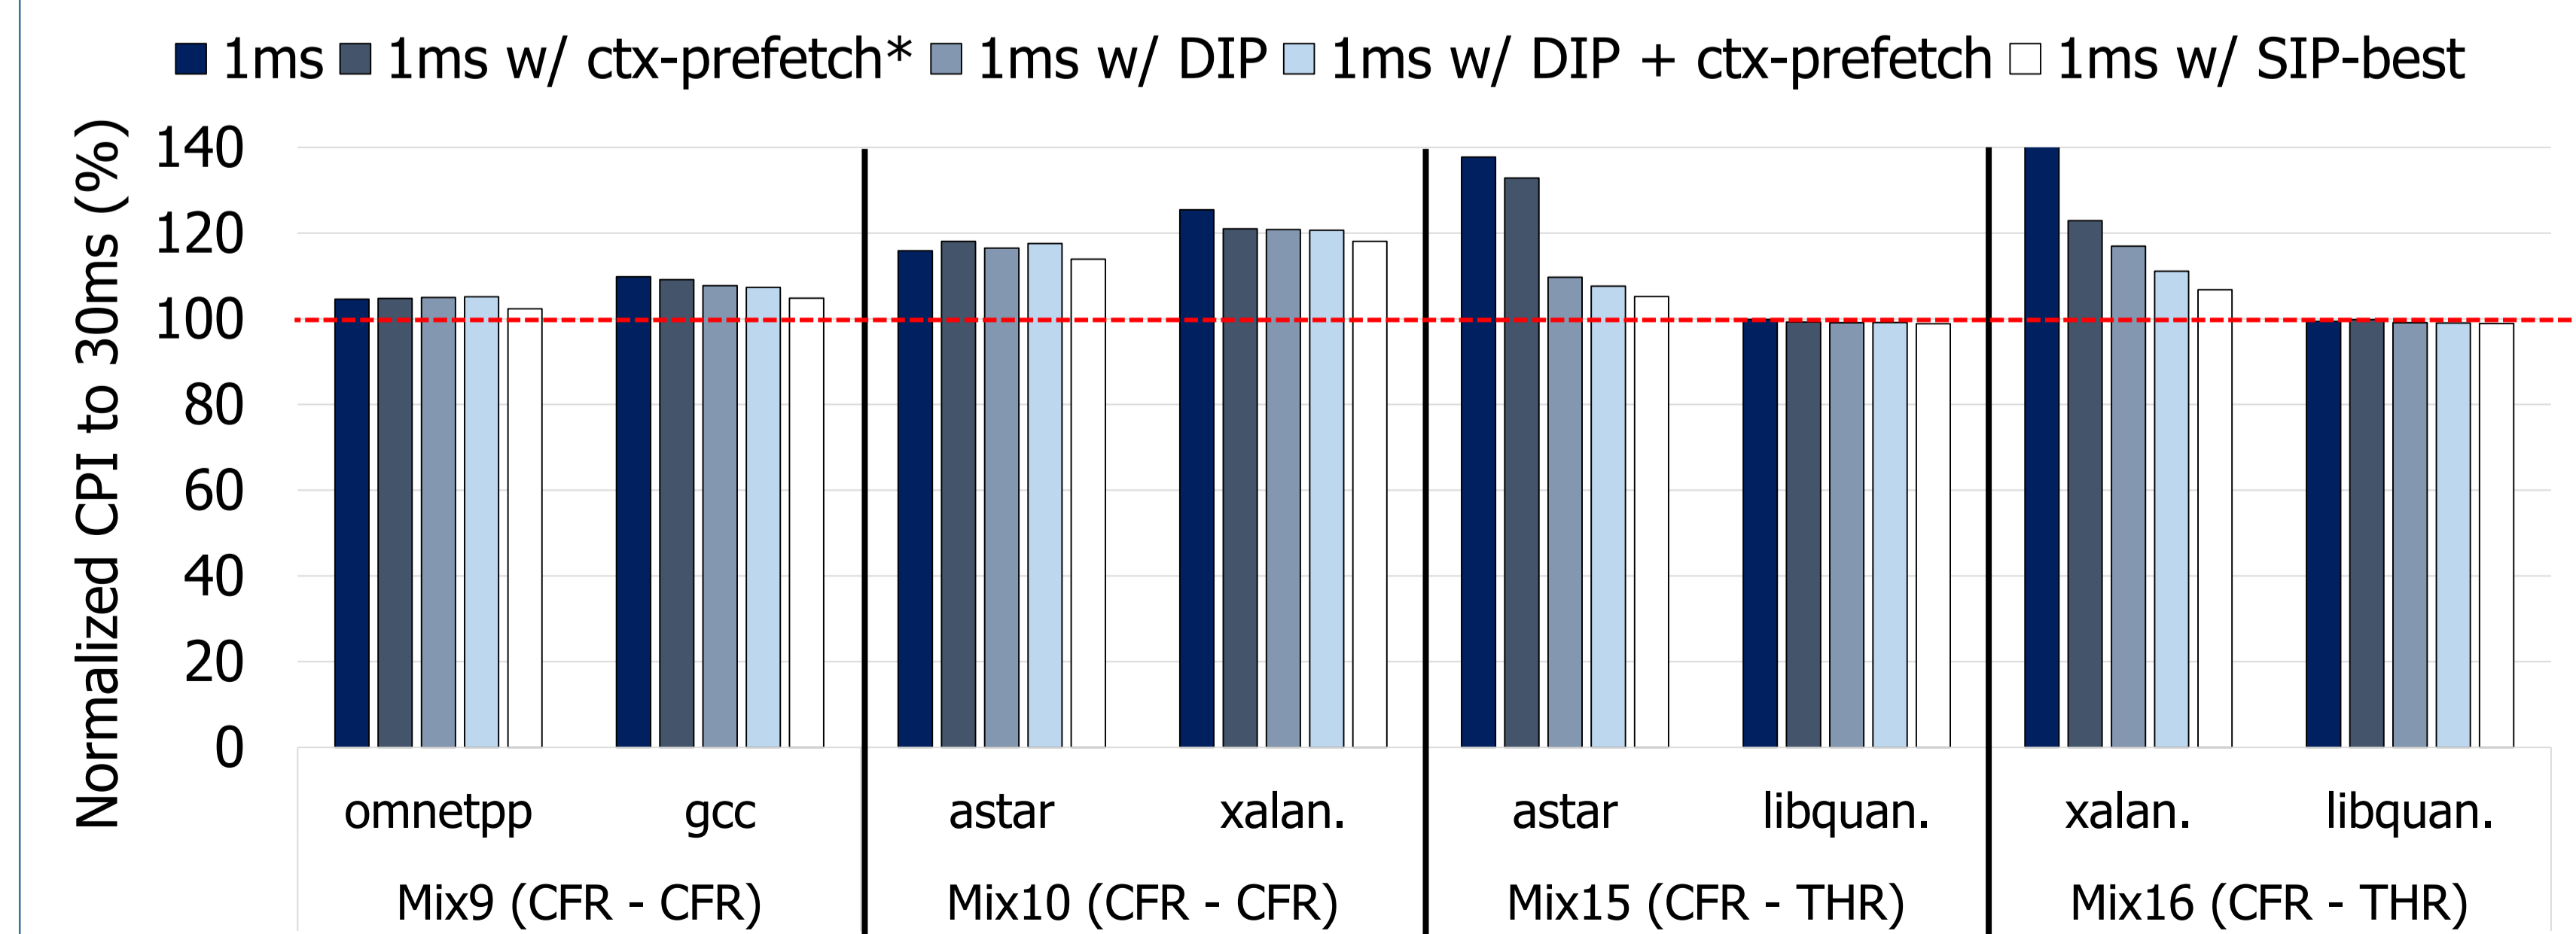


## Context Preservation with Dual Insertion Policy

- For *CFR*, a new cache block is inserted to the MRU
- For *CPU* and *THR*, a new cache block is inserted to the LRU



## Result Highlights



\* [Zebchuk et al., HPCA '13]