

Green-CM: Energy efficient contention management for Transactional Memory

Shady Alaa

Paolo Romano – INESC-ID/IST

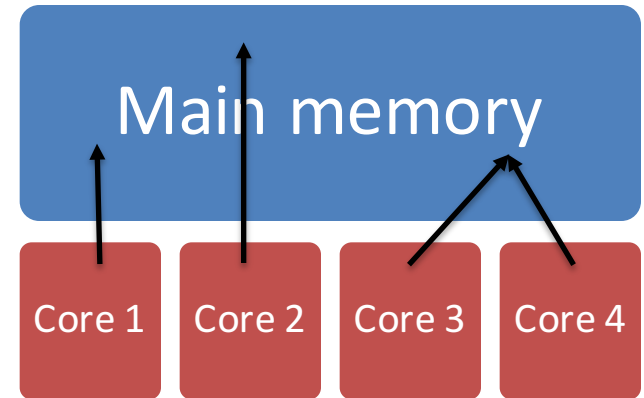
Mats Brorsson - KTH

Agenda

- Introduction
- Related work
- Architecture
- Green-CM
- Evaluation
- Conclusion

Introduction

- Multicores are everywhere
 - Complex programming
 - Locks
 - Deadlocks
 - Transactional memory
 - Atomics blocks
 - Transparent from programmer



```
atomic{
    if (bal > amount)
        withdraw (amount);
}
```

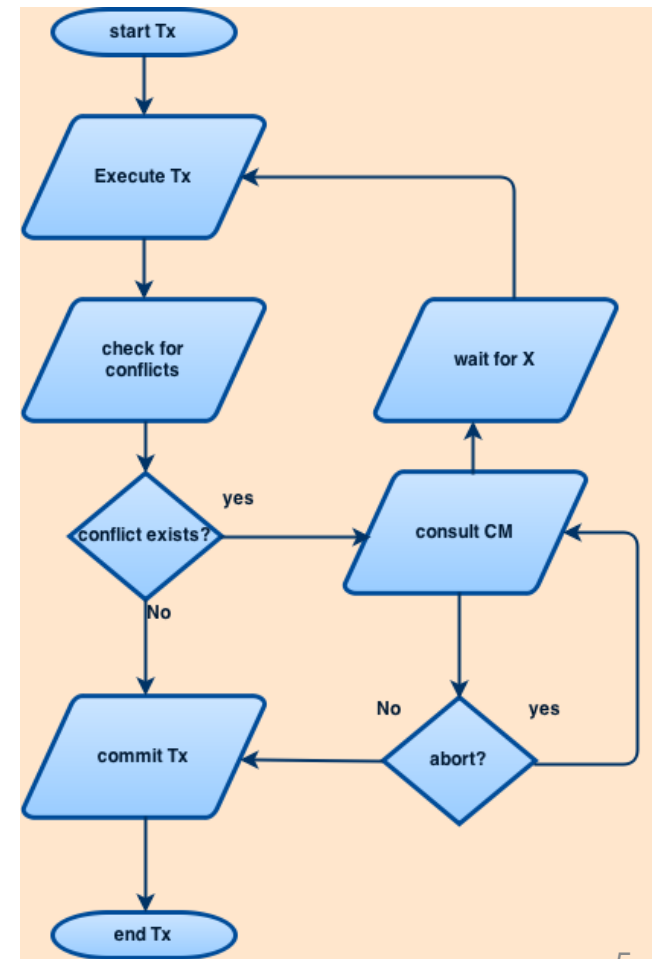
Introduction

- Energy efficiency
 - First order design choice
 - Battery based devices
 - Data centers
- Goal
 - Energy efficient transactional memory in terms of both energy and performance



Introduction

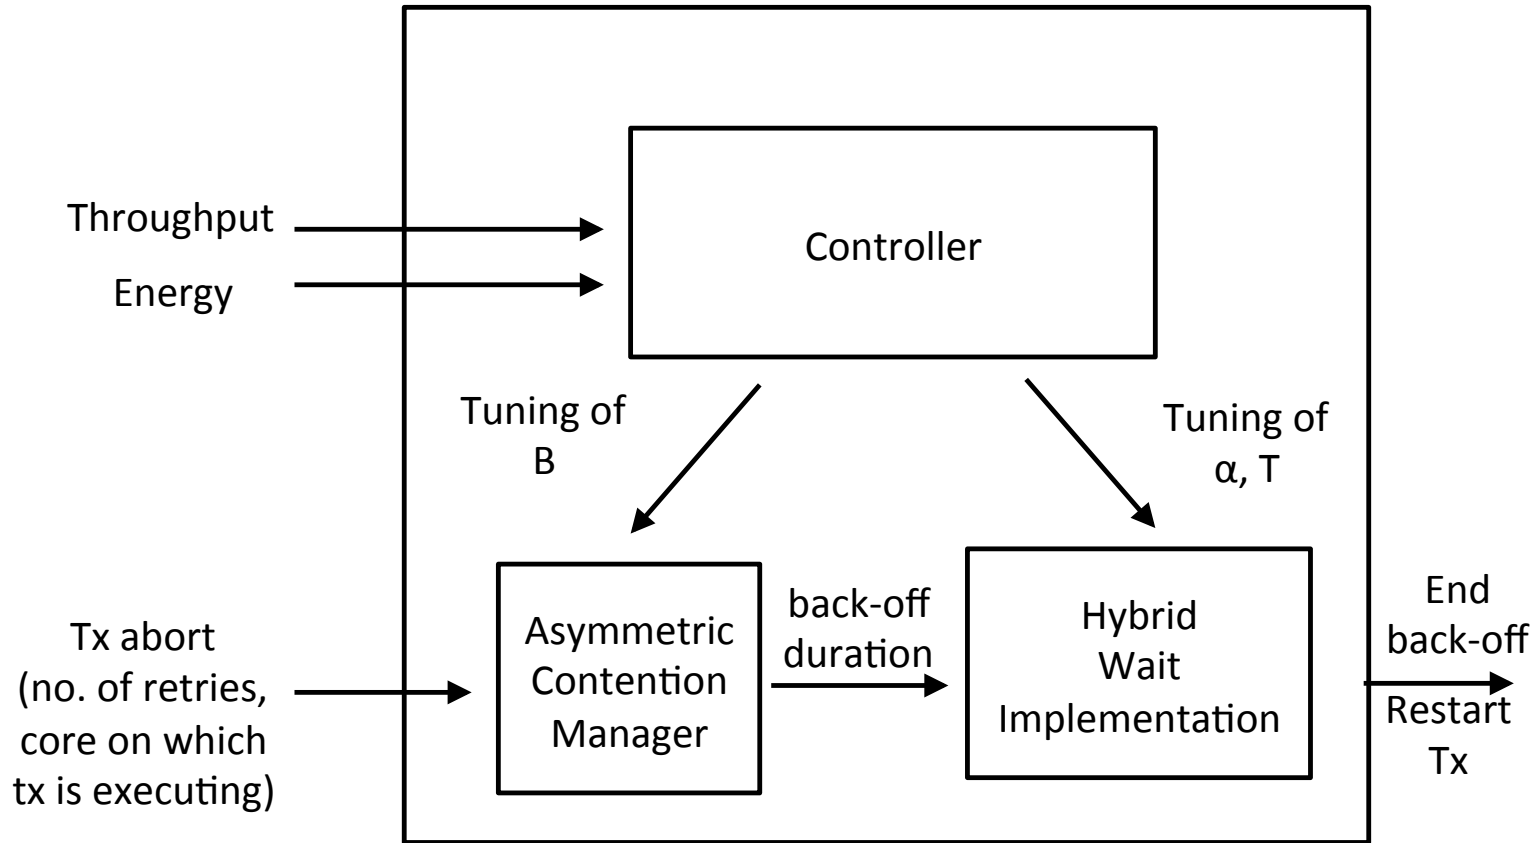
- Contention Manager
 - minimize contention
 - which transaction to abort
 - when to restart an aborted transaction
- Energy efficiency:
 - wait implementation
 - DVFS



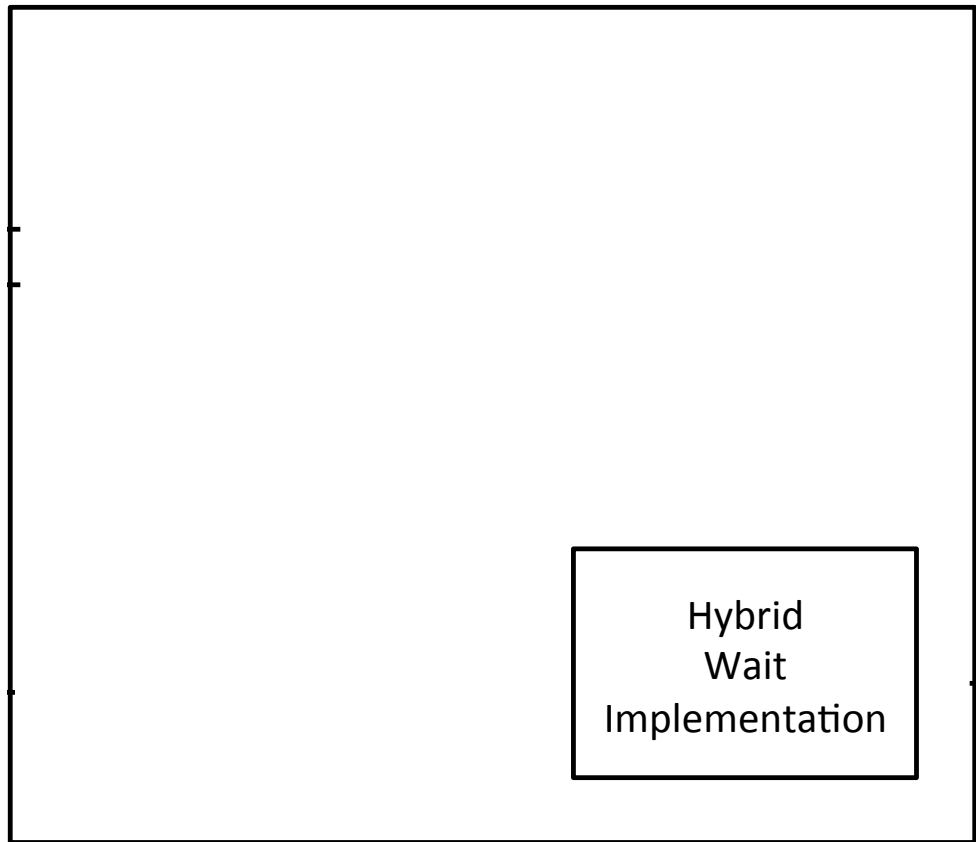
Related work

- Few work in literature
 - Mainly HTM
 - Clock gating processors upon abort
 - Lowering frequency upon abort
 - Using simulator
- Studies
 - HTM consume lower energy
 - Does not fit all workloads
 - Need for adaptability
- Using DVFS in TM
 - Fastlane
 - Designed for low number of threads

Architecture



Architecture



Implementing waits

- Building block for contention managers
- Drastic effect on energy consumption
- Can be implemented in two ways:
 - Busy waiting
 - sleeping

Implementing waits

- Busy waiting
 - Fine granularity
 - Similar to real actual work
- Sleeping
 - Coarse granularity
 - Low energy consumption
 - expensive



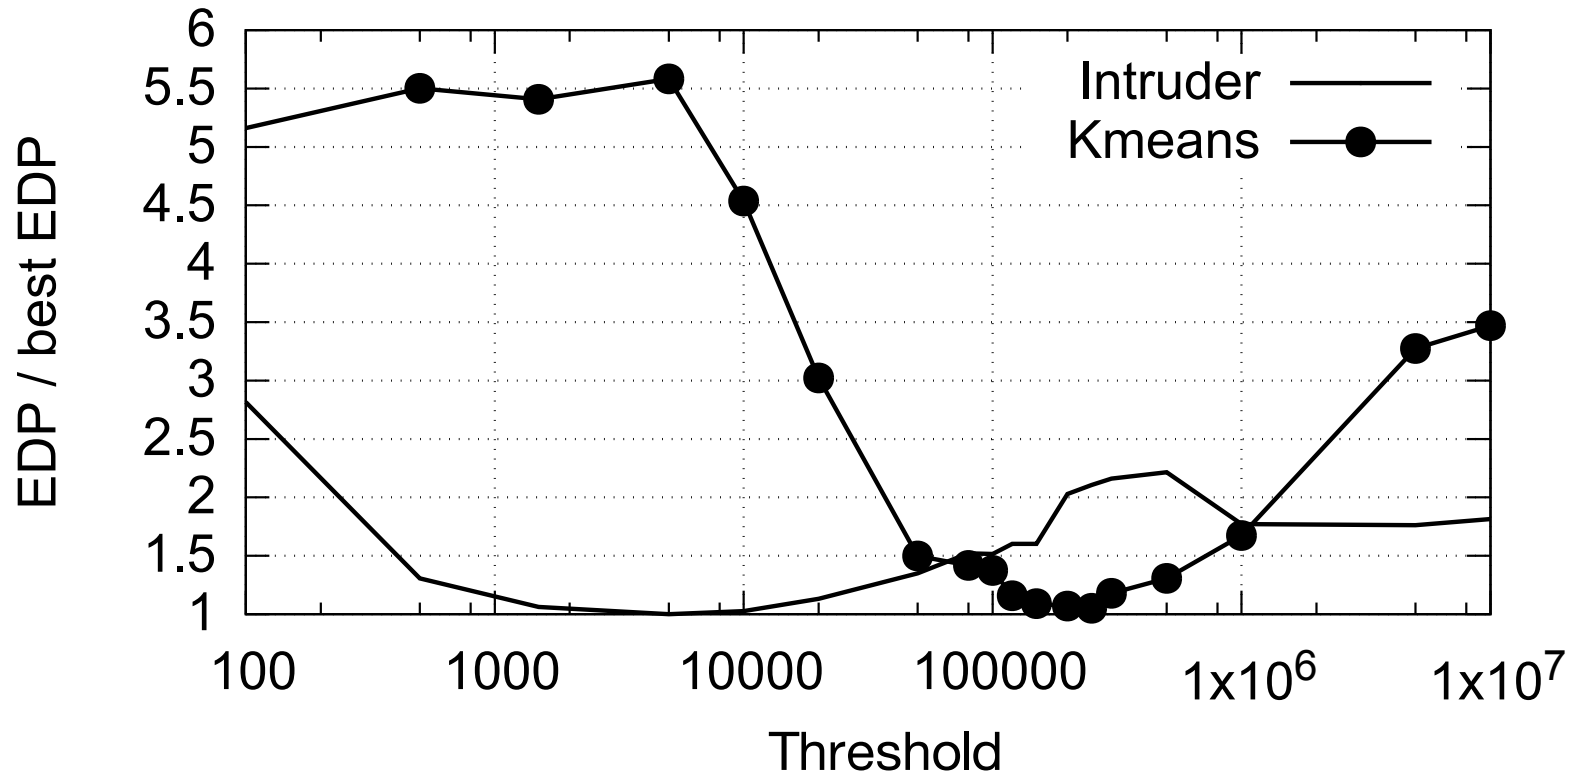
Implementing waits

- Hybrid approach
 - Either busy wait or sleep
 - Adaptive fashion
 - How to determine the **threshold**
 - Cost of sleep

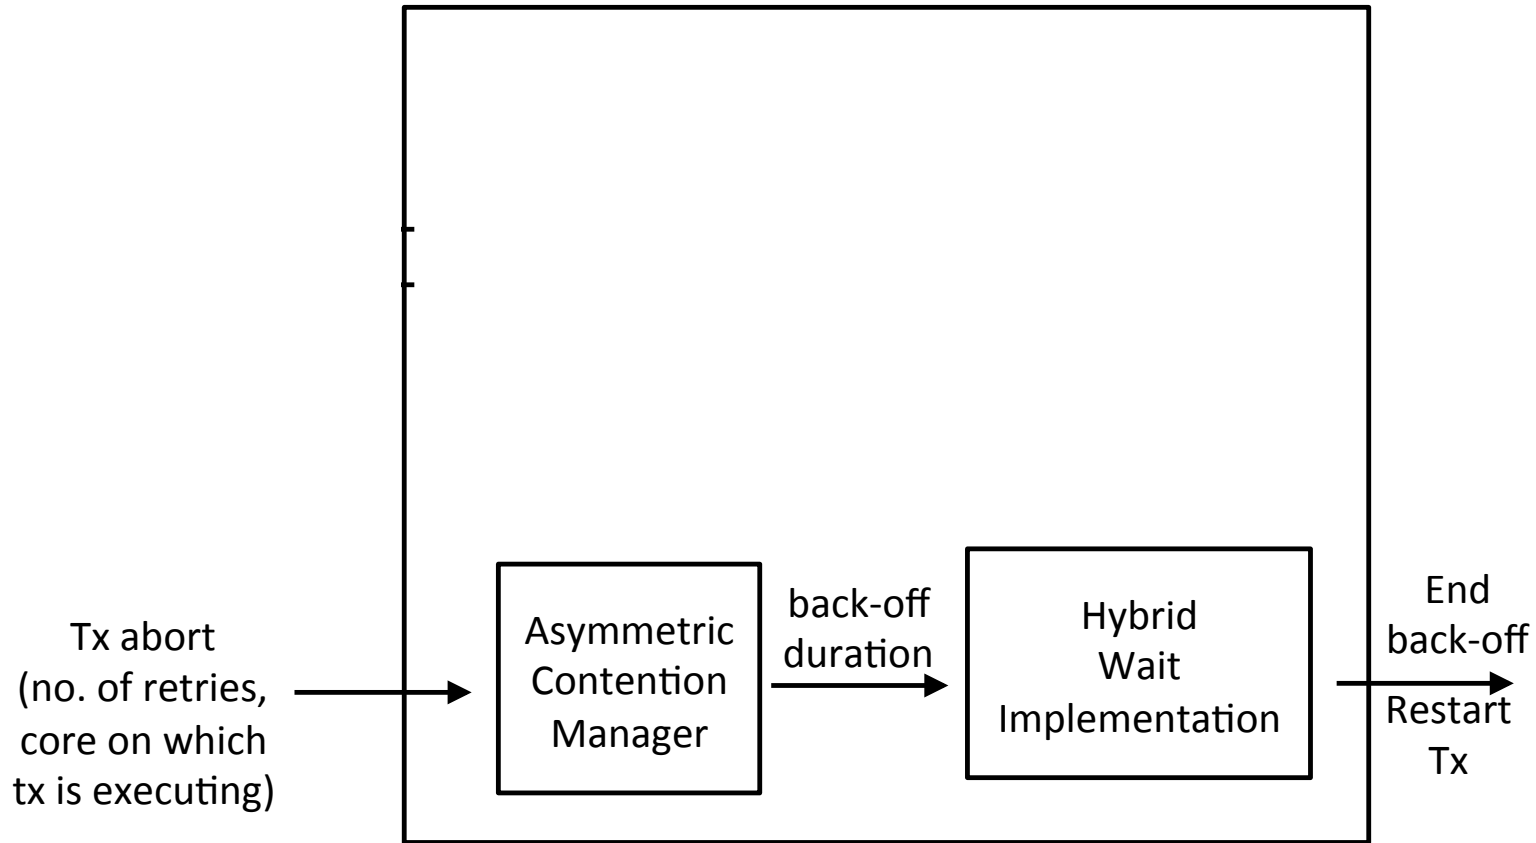
Implementing waits

No one
size fits all

Static Thresholds



Architecture



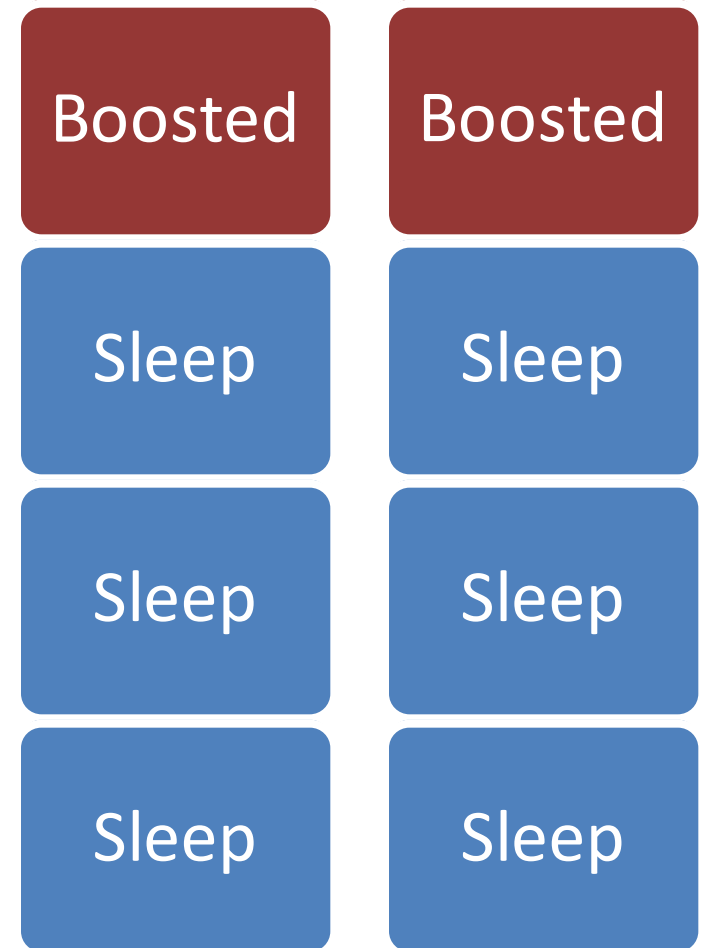
Asymmetric CM

- DVFS
 - Variable operating frequency
- Exploiting DVFS
 - Boosting active threads
 - Reducing freq. of backing off threads
- Enabling DVFS
 - Manual control is expensive
 - How to favor automatic boosting

P0	3.0 GHz
P1	2.4 GHz
P2	2.2 GHz
P3	2.0 GHz
P4	1.8 GHz
P5	1.6 GHz
P6	1.4 GHz

Asymmetric CM

- Linear backoff cores:
 - Shorter backoff periods
 - Mainly busy waiting backoffs
- Exp. Backoff cores:
 - Longer backoff periods
 - Mainly sleep waiting
- Favor boosting
 - When enough cores are in sleep states



8 core
processor

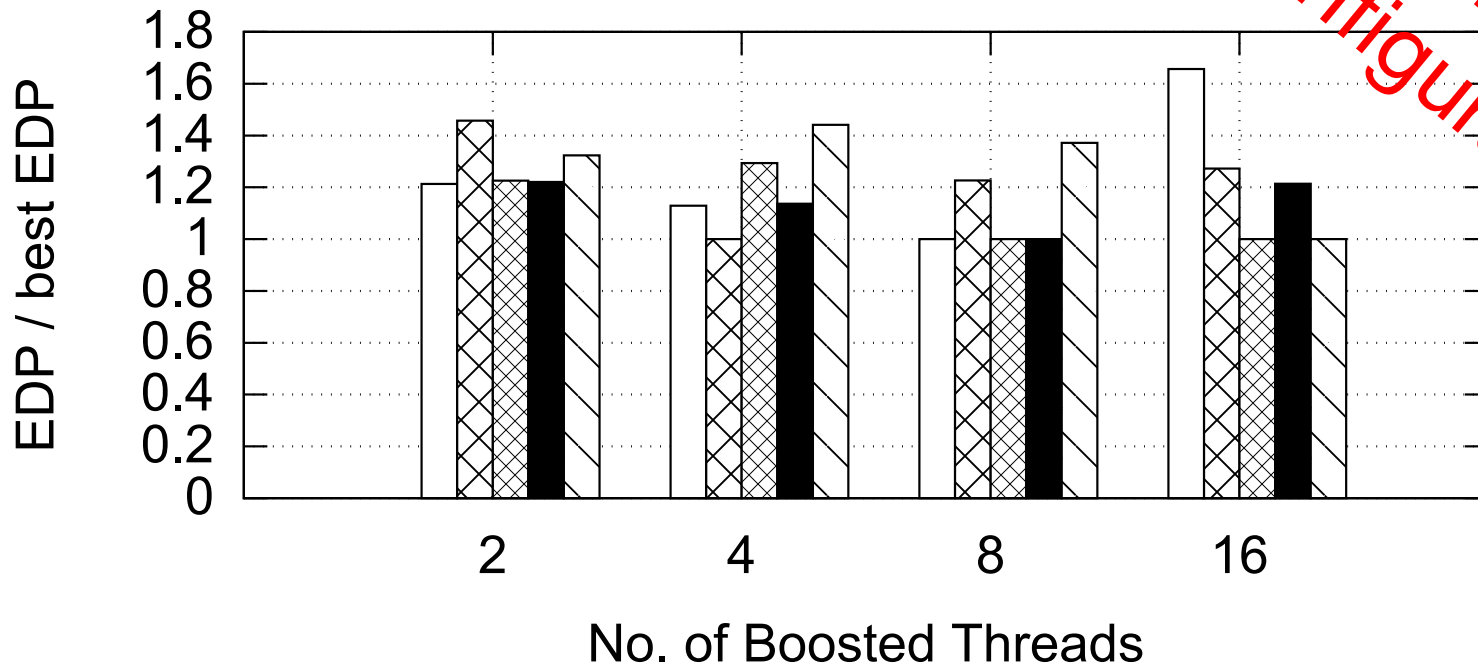
Asymmetric CM

- Increased contention?
 - Cores not backing off exponentially
- Control **number of cores to be boosted**

Asymmetric CM

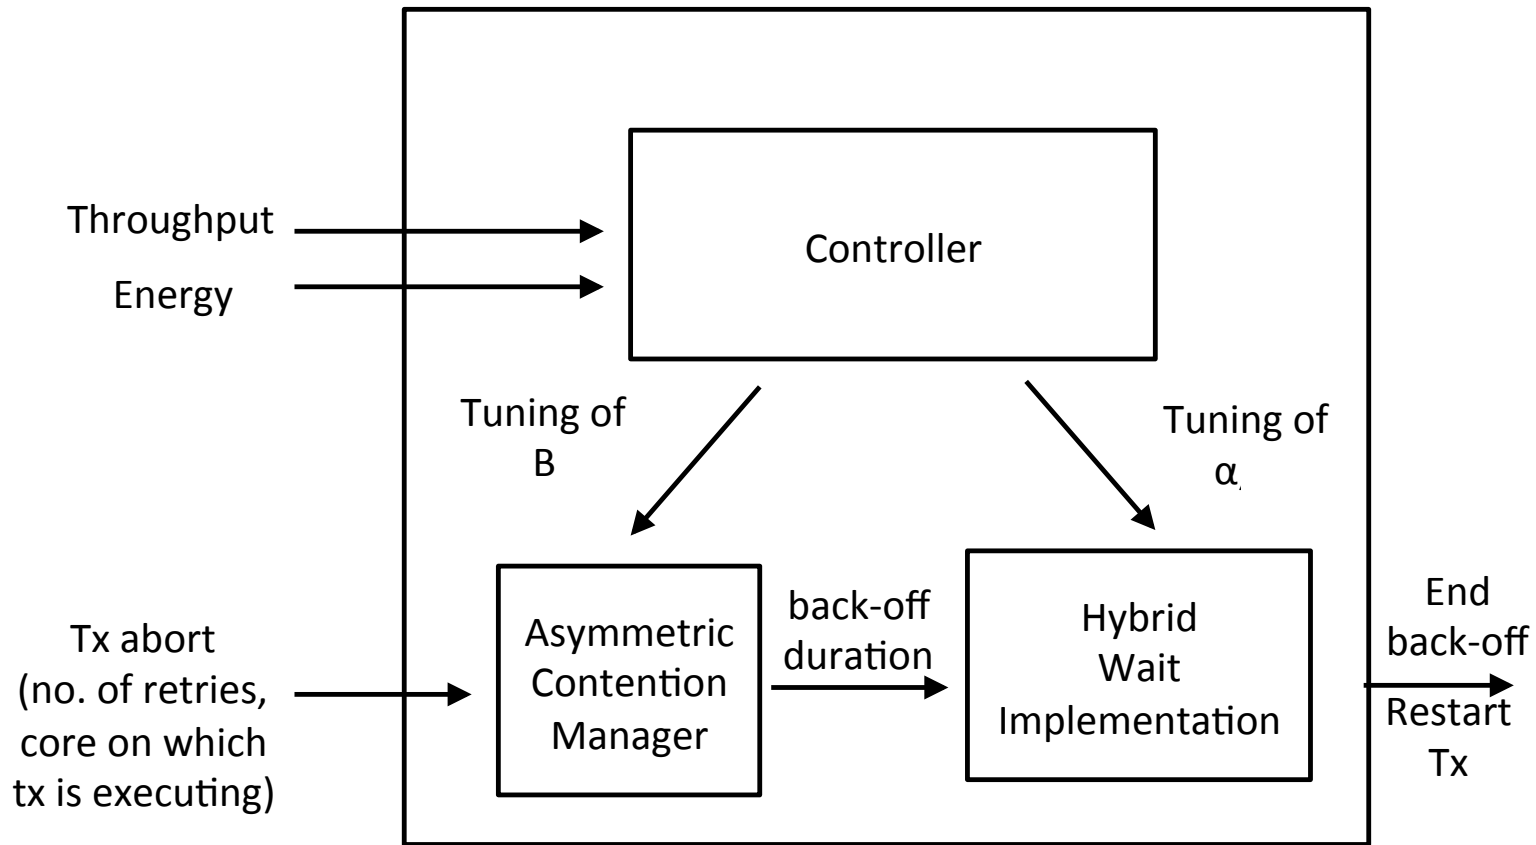
Intruder  Genome  STM7 
Kmeans  Memcached 

Static No. of Boosted Threads



No universal configuration

Architecture

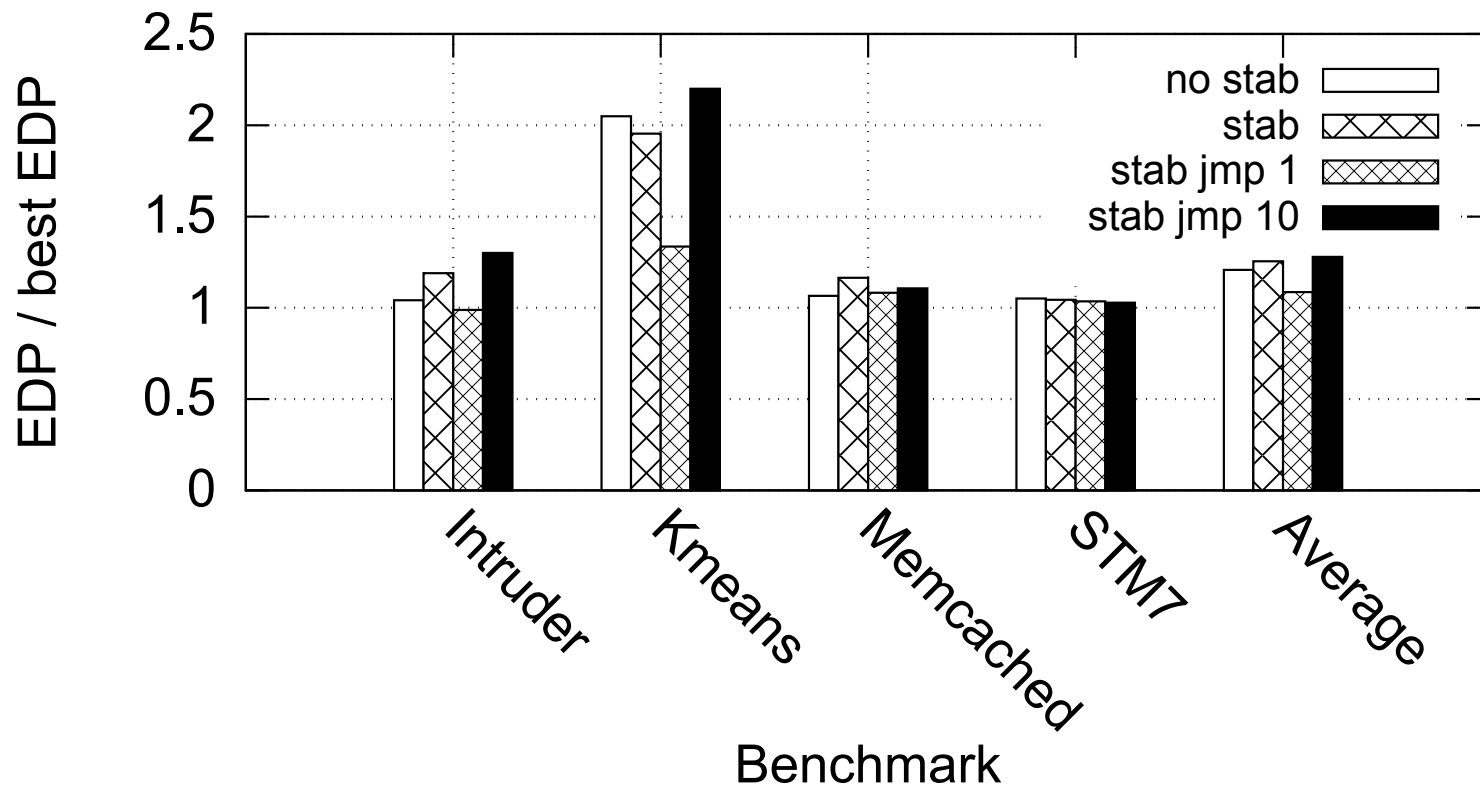


Controller

- Online, lightweight
- Hill climbing
- Challenges:
 - Collection of energy
 - Multi dimensional
- Different exploration strategies
 - Stabilization
 - Random jumps

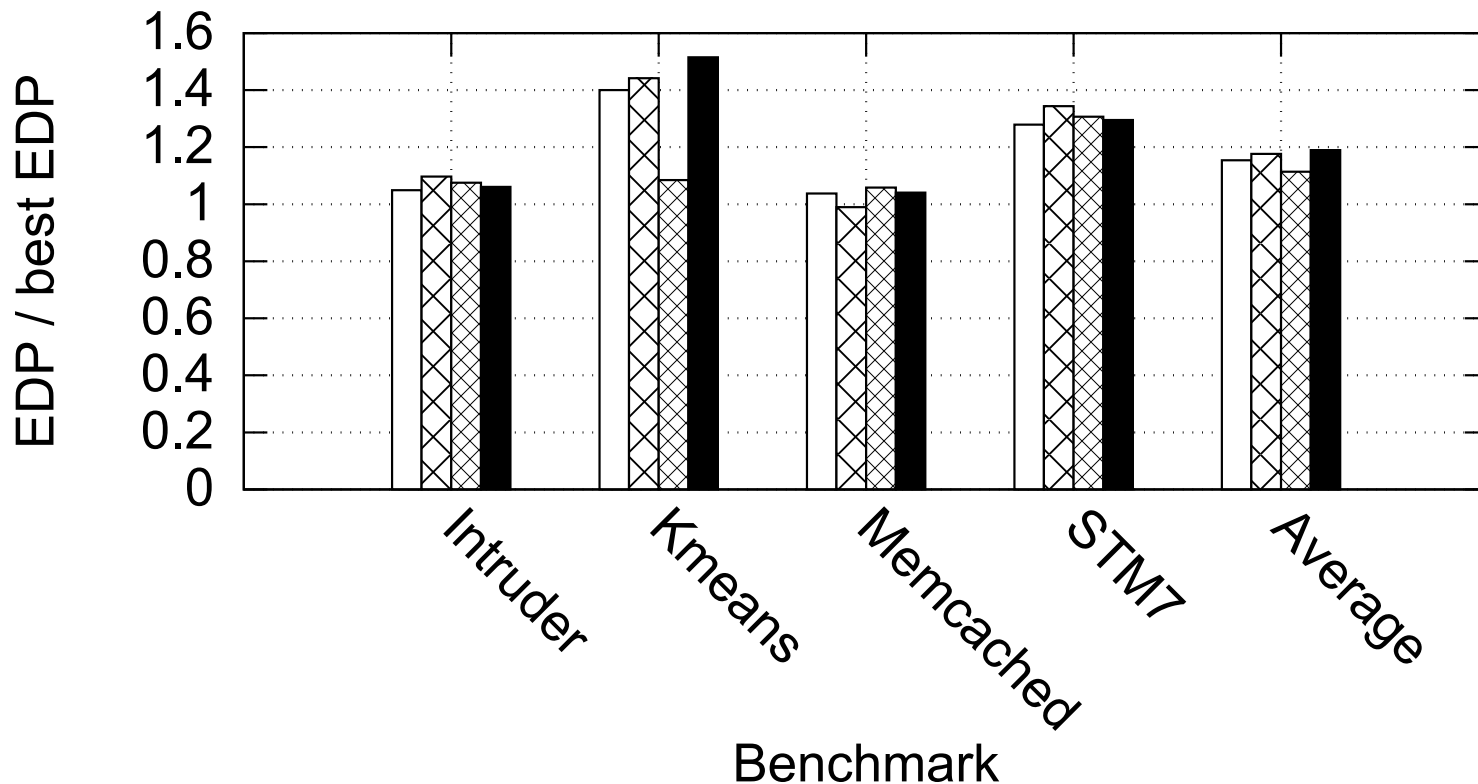
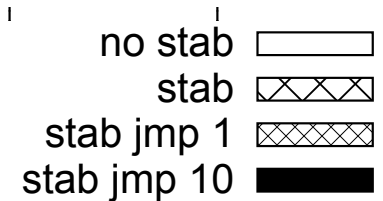
Controller

- Tuning α (threshold for hybrid)



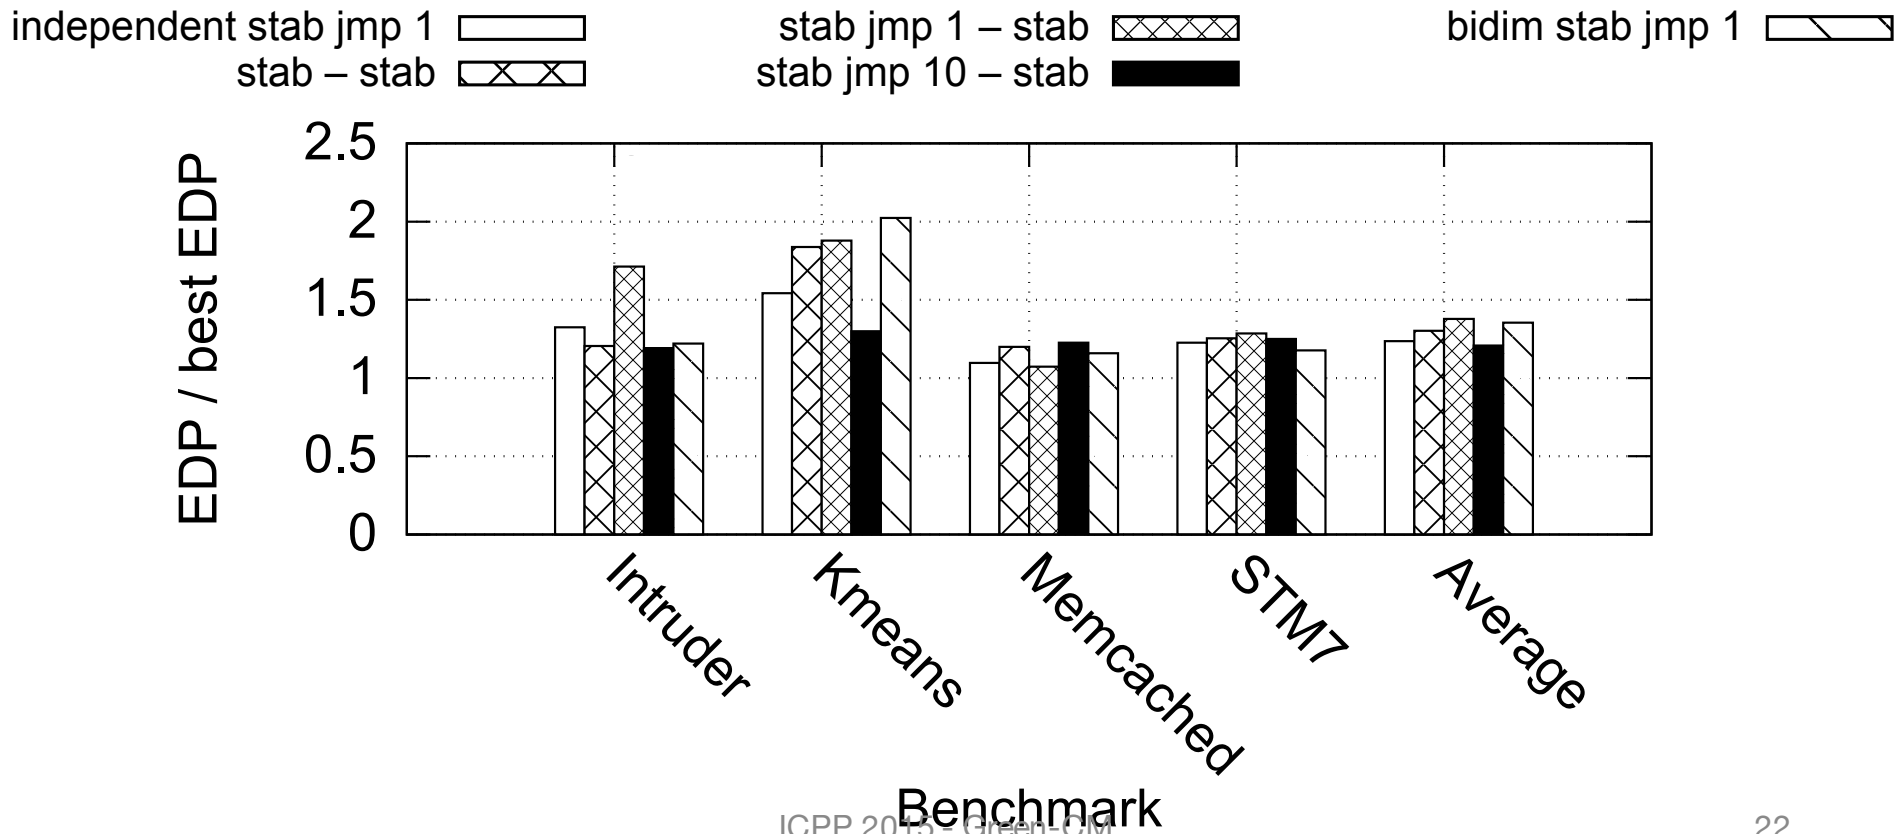
Controller

- Tuning β (no. of boosted threads)



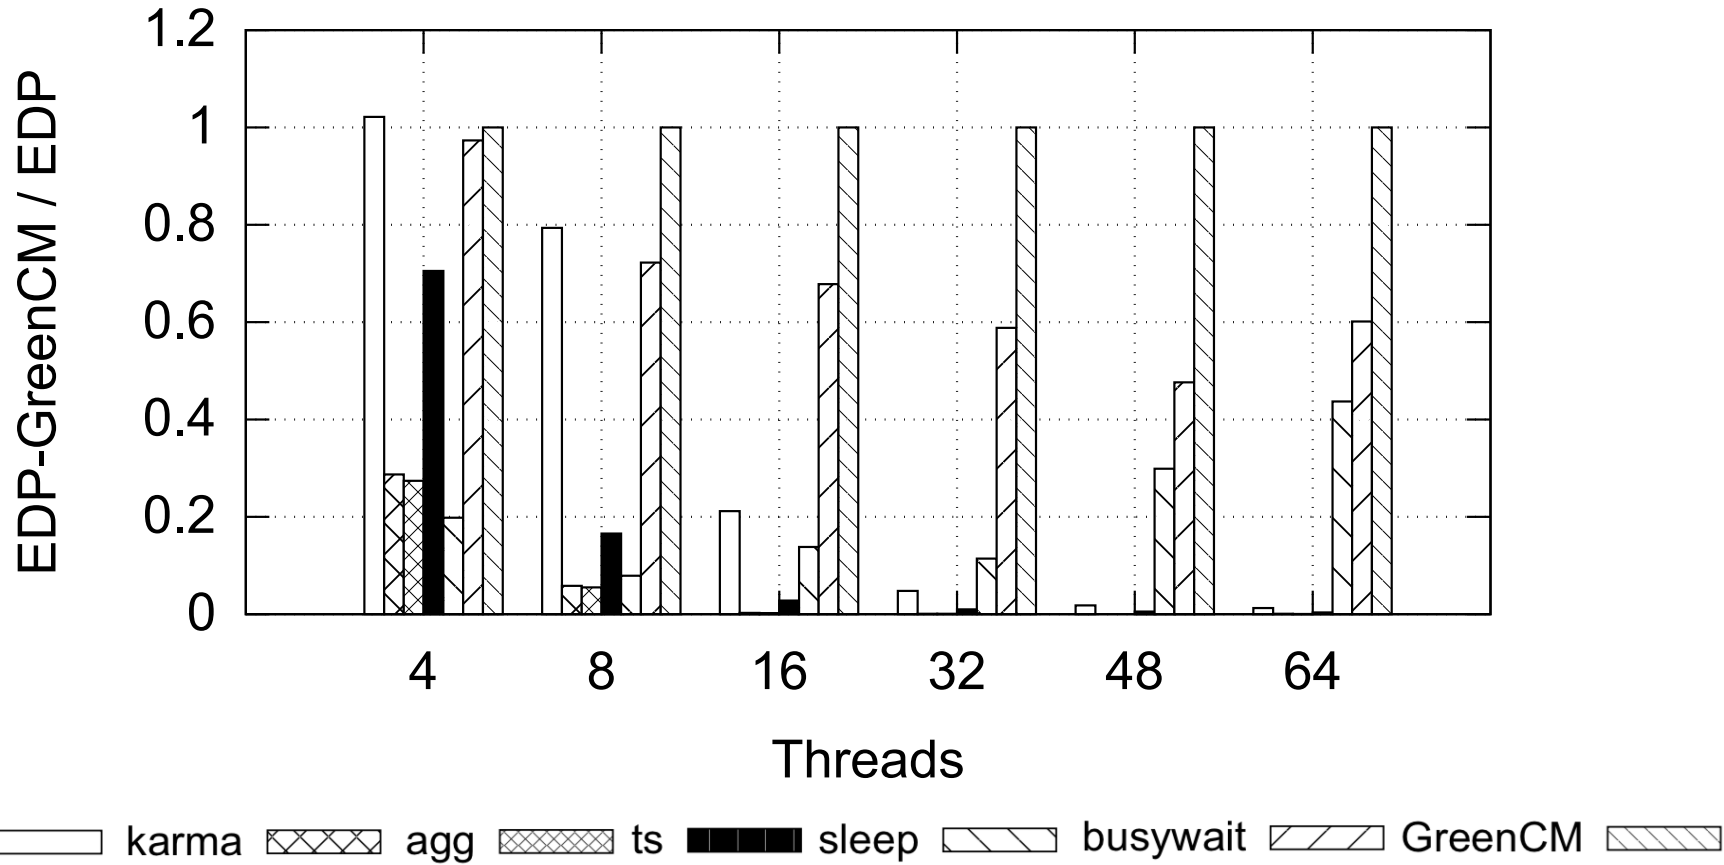
Controller

- Merging the learners



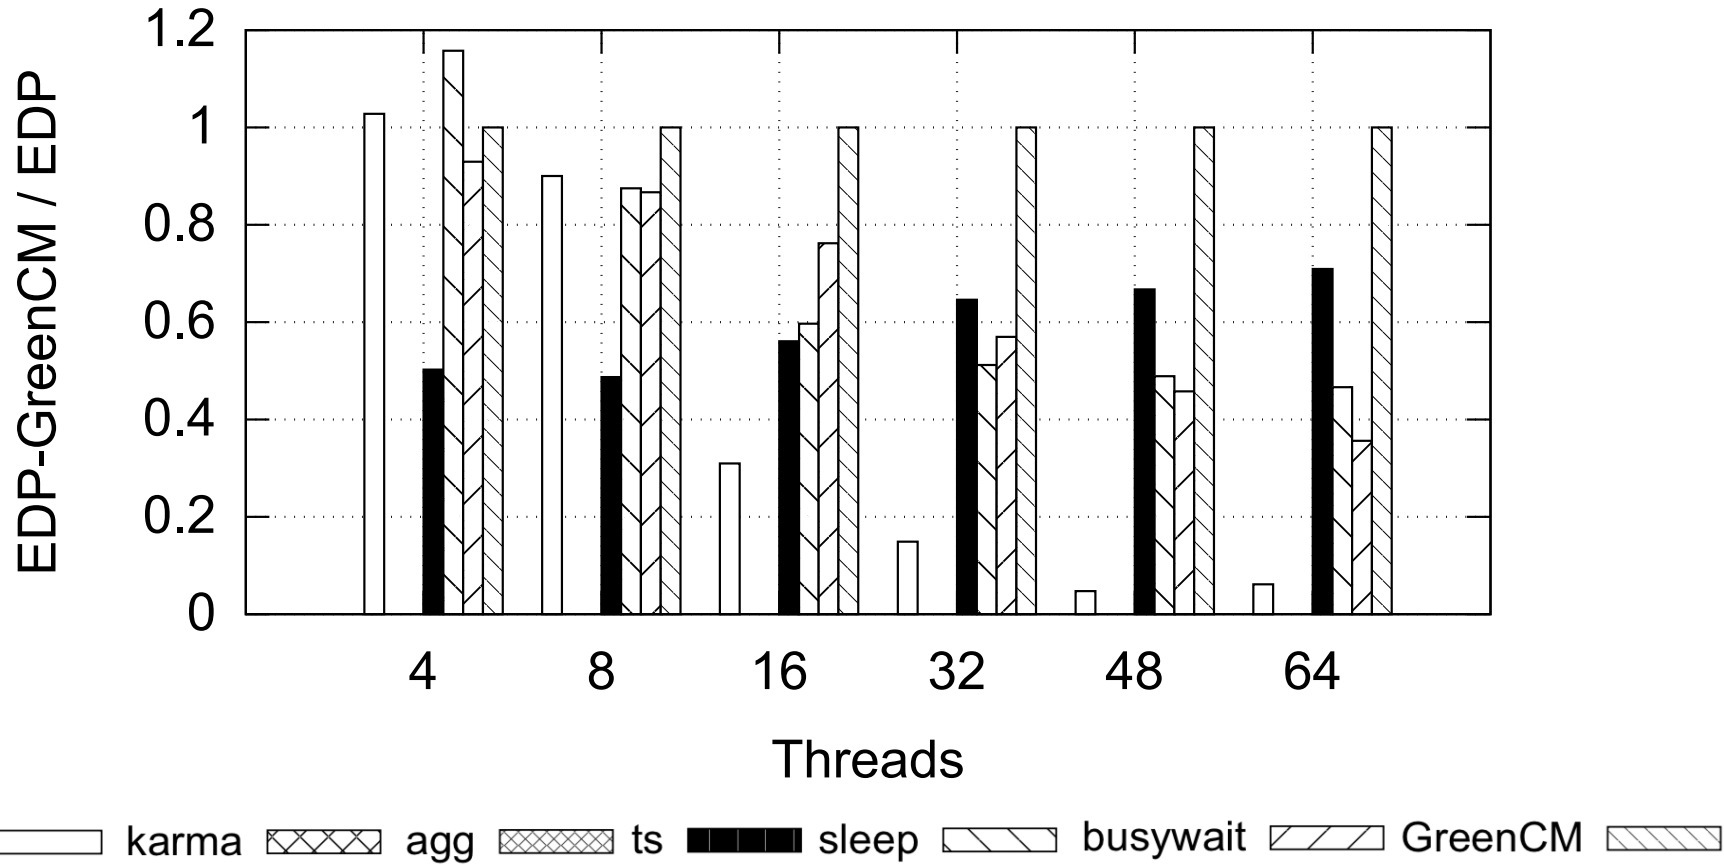
Evaluation

Intruder



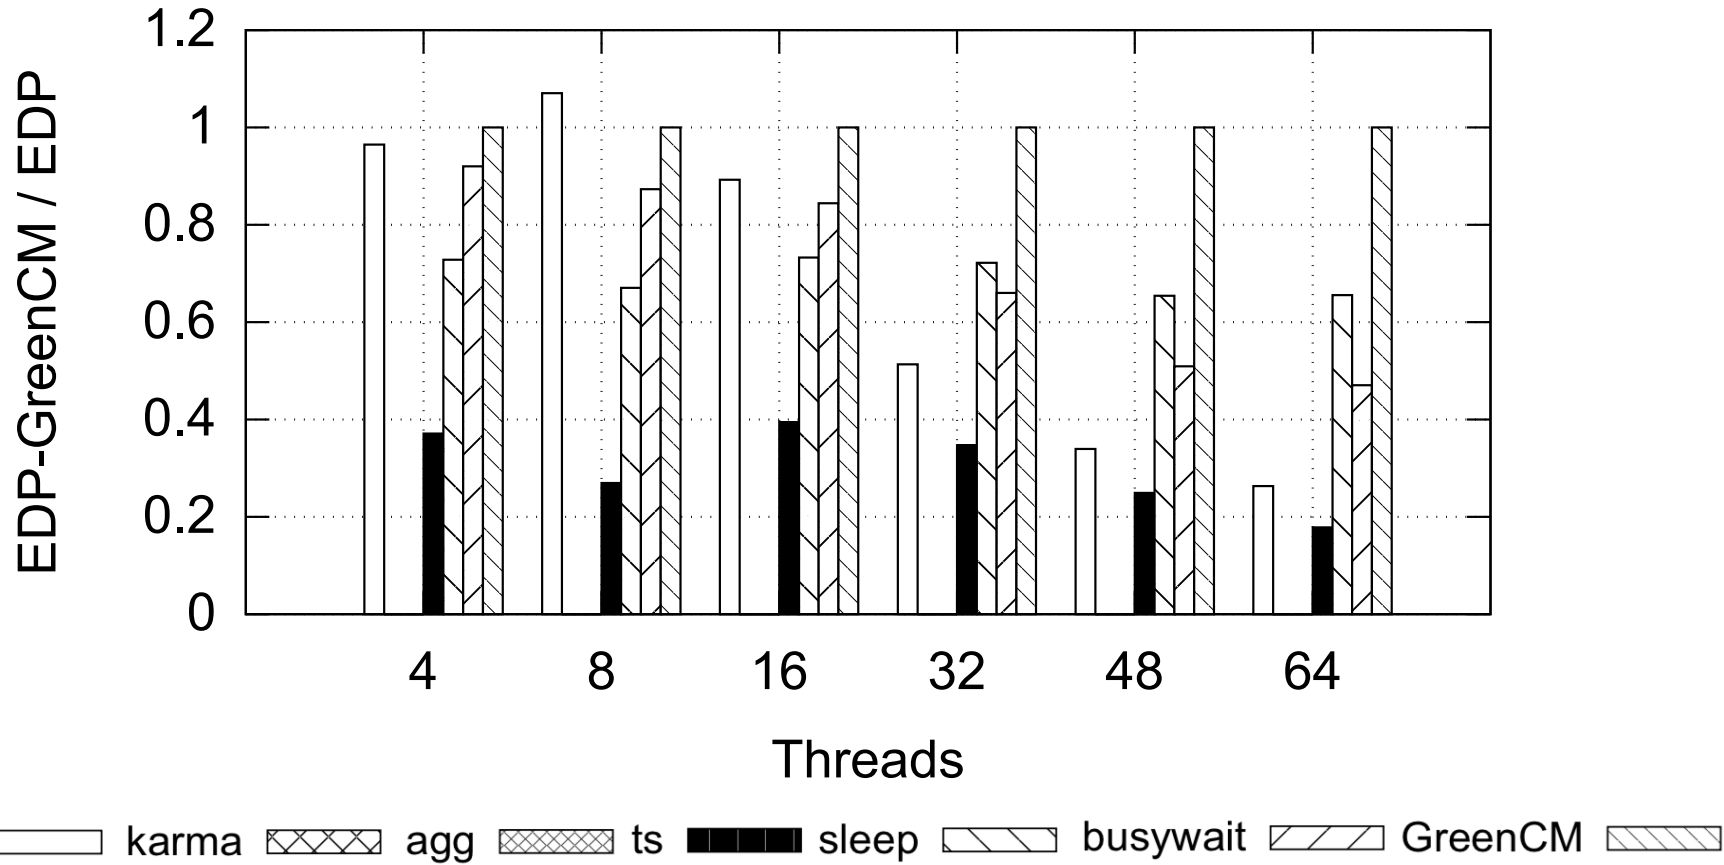
Evaluation

STM7



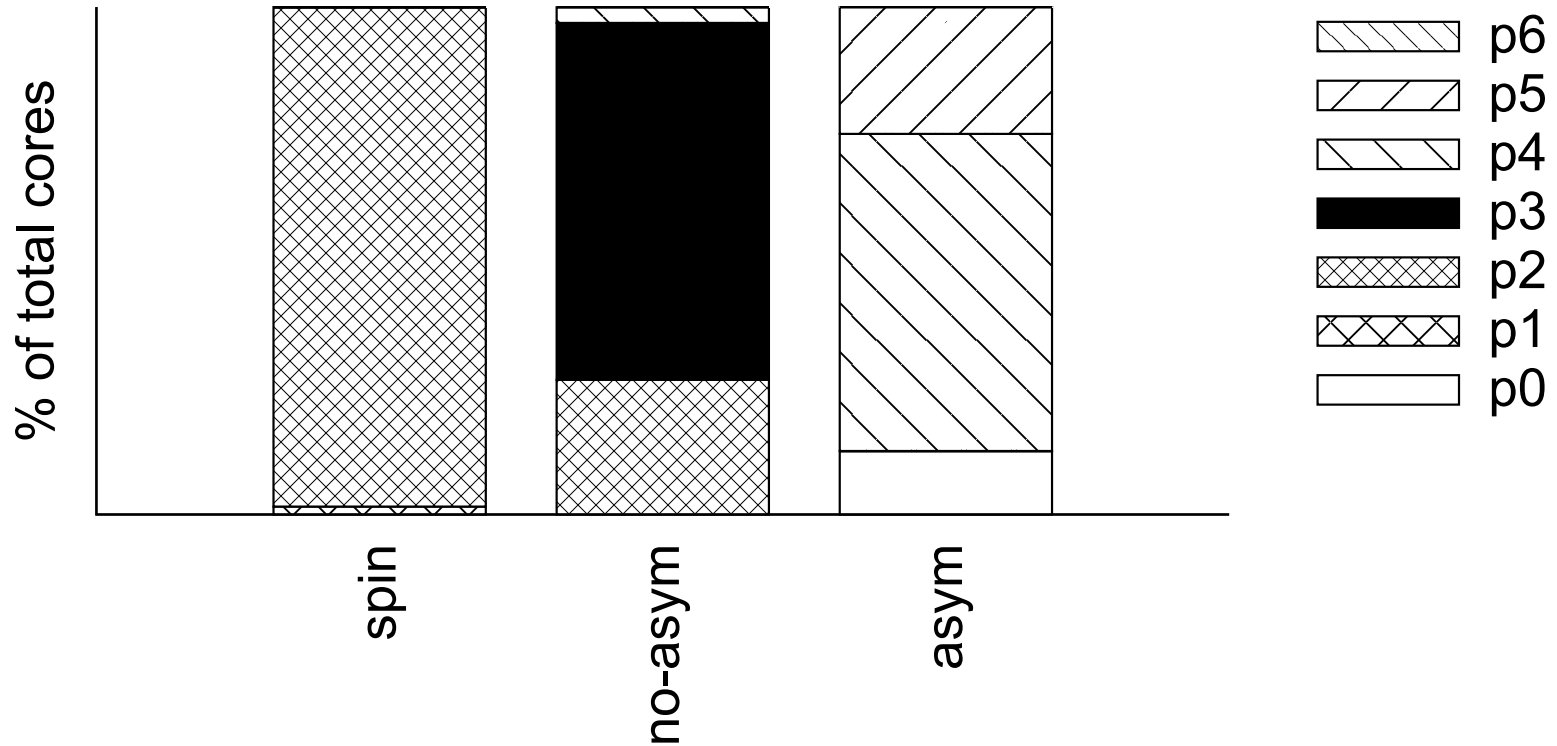
Evaluation

Memcached



Evaluation

Intruder, 64 threads



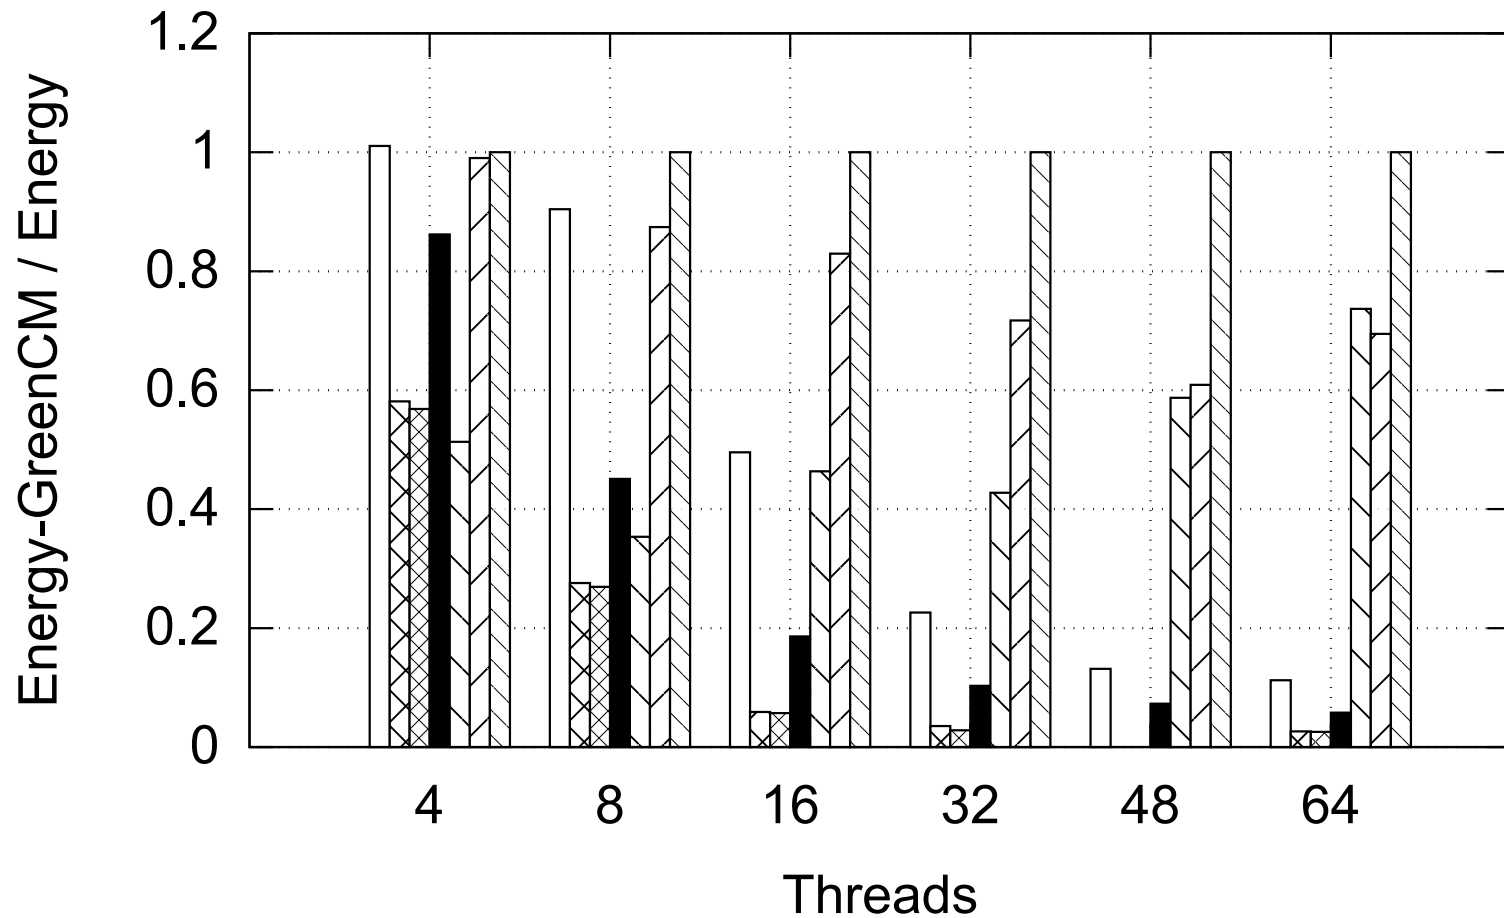
Conclusion

- Implementation of waits has a significant impact on energy efficiency
- Experimental results (obtained on real system) contradict previously published ones based on simulation
- Exploiting DVFS enhances energy efficiency
- Self-tuning is needed to adapt to different workloads

THANK YOU

Evaluation

Intruder



Evaluation

Intruder

