

TOWARD PREDICTABLE PERFORMANCE IN SOFTWARE PACKET-PROCESSING PLATFORMS

Mihai Dobrescu, EPFL
Katerina Argyraki, EPFL
Sylvia Ratnasamy, UC Berkeley

Programmable Networks

2

- Industry/research community efforts
 - ▣ Easily deploy new services
 - ▣ Test research ideas

- Software packet processing
 - ▣ General purpose hardware
 - ▣ Familiar programming environment

Extensible network functionality

Problem: Unpredictable Performance

3

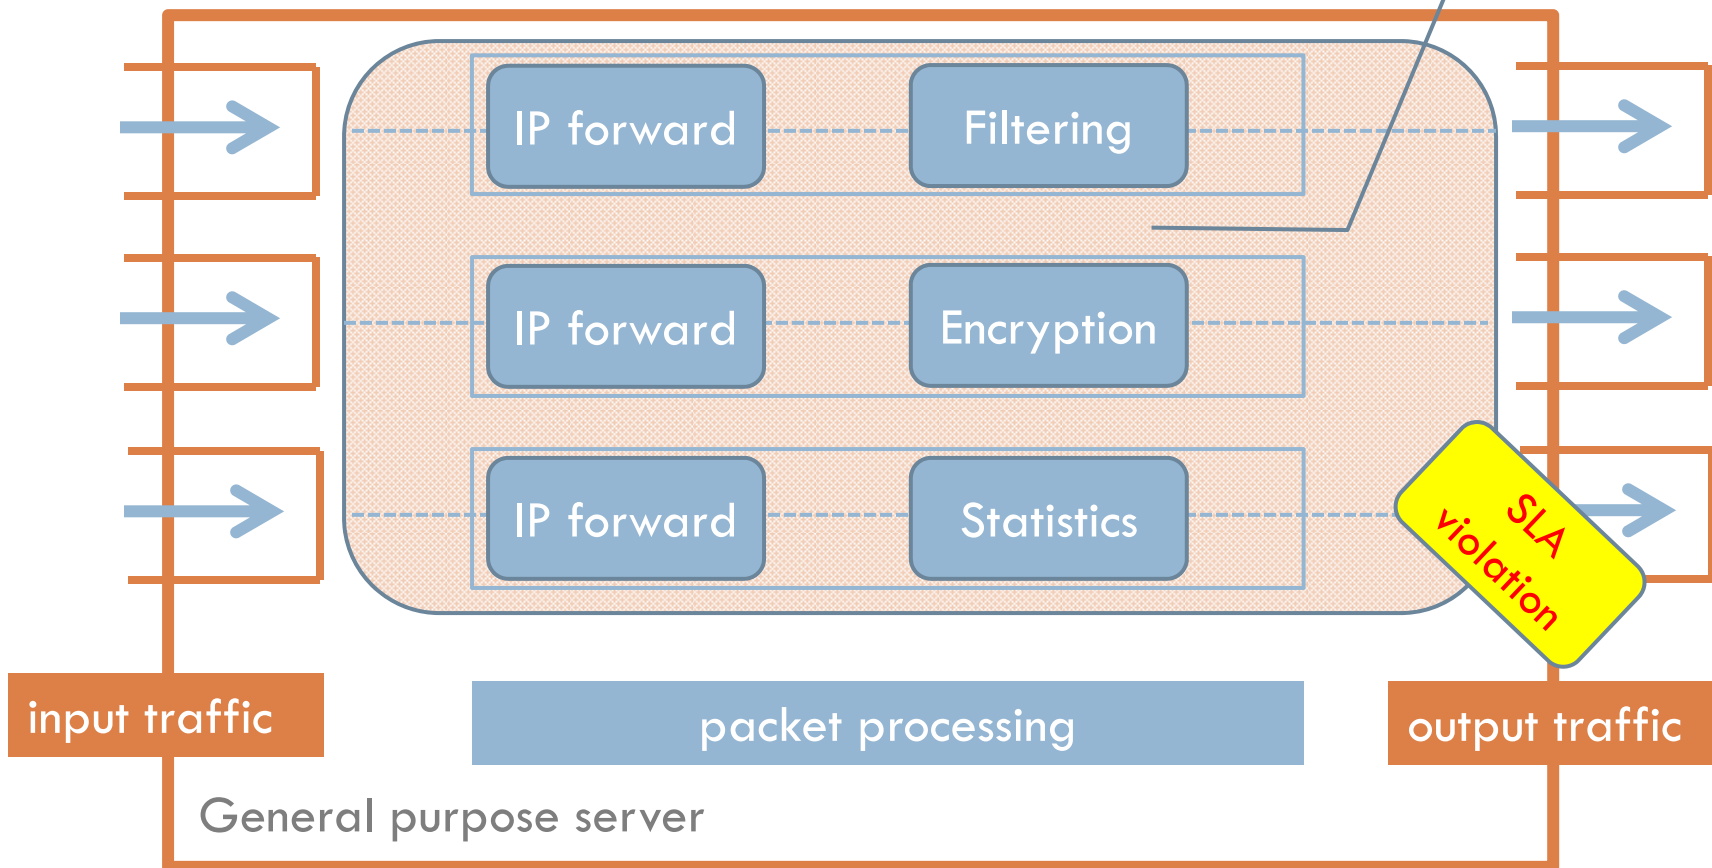
- Resource contention
 - ▣ Caches, memory controllers, buses
 - ▣ Performance interference
- Software packet-processing systems [Dobrescu'09, Han'10]
 - ▣ High performance
 - ▣ Same processing for all packets

Goal: software packet processing with predictable performance

System Overview

4

Contention for shared resources



Is This Hard?

5

- Yes, in general-purpose context
 - ▣ Math models to predict contention
 - ▣ Contention-aware job placement

- In packet-processing context?

Our Contribution

6

1. It is feasible to build a packet-processing platform with predictable performance using simple techniques.

2. Contention-aware job placement does not bring significant benefit to the overall performance.

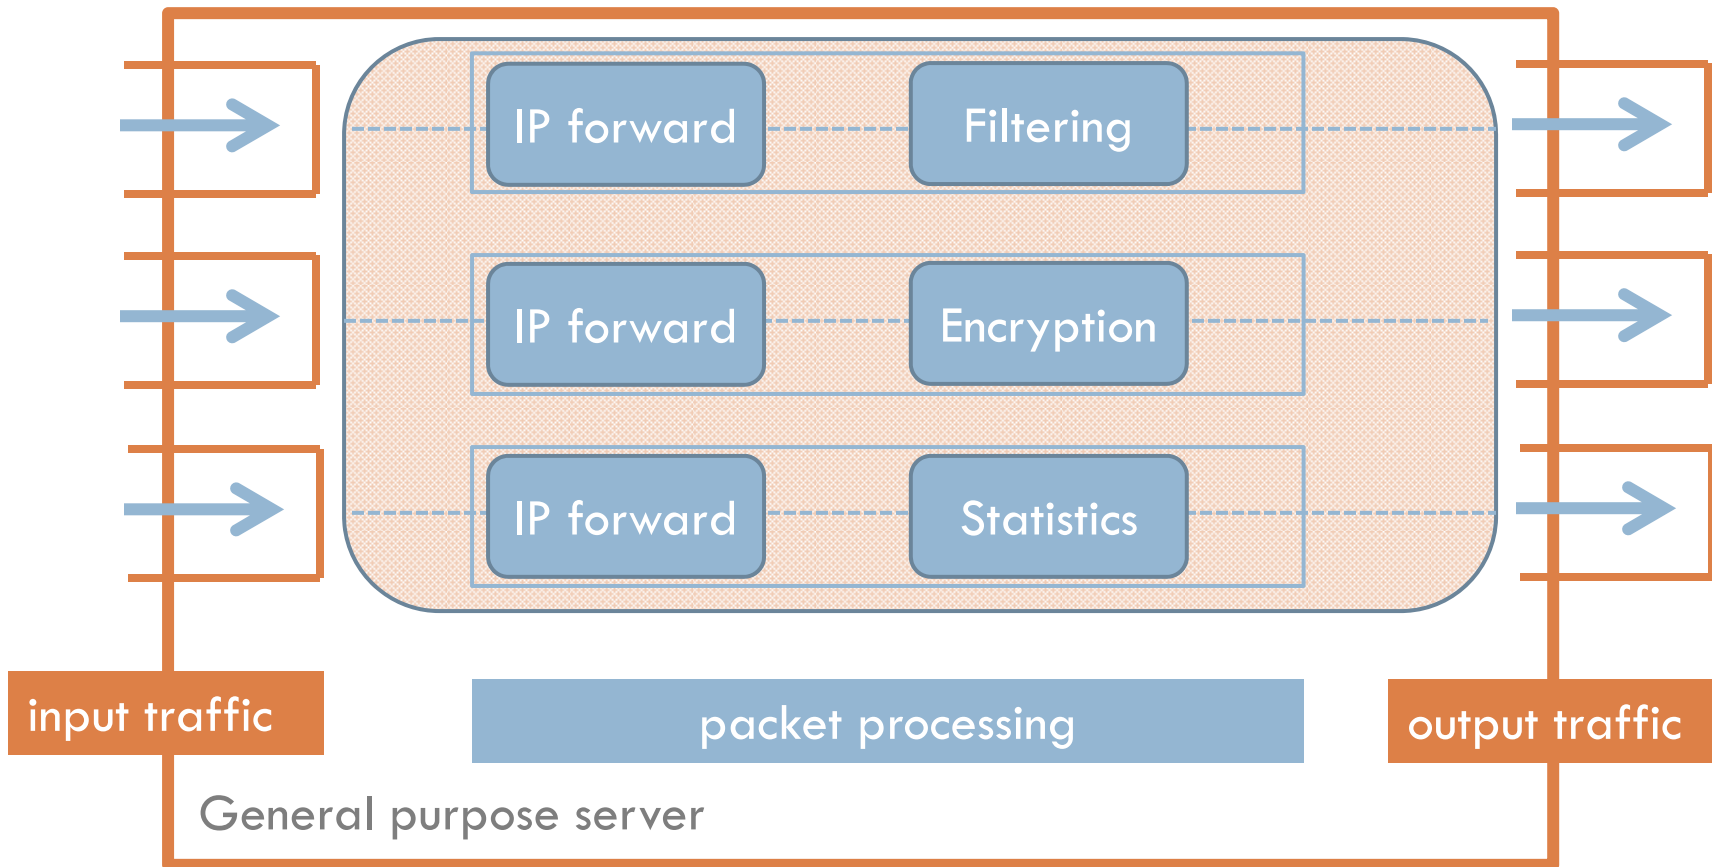
Outline

7

- System overview
- Contention factors
- Observations on application behavior
- A simple prediction method
- Intuition

System Overview

8



Workloads

9


| Application | Main functionality | Characteristics |
|-------------|--------------------------|--------------------|
| IP | IP routing, 128k entries | L3 cache intensive |
| MON | Monitoring, 100k flows | L3 cache intensive |
| FW | Firewall, 1000 rules | L2 cache intensive |
| RE | Redundancy elimination | Memory intensive |
| VPN | Encryption | CPU intensive |
| Synthetic | Random cache reads | Cache/memory/CPU |

Representative set of realistic applications

Setup

10

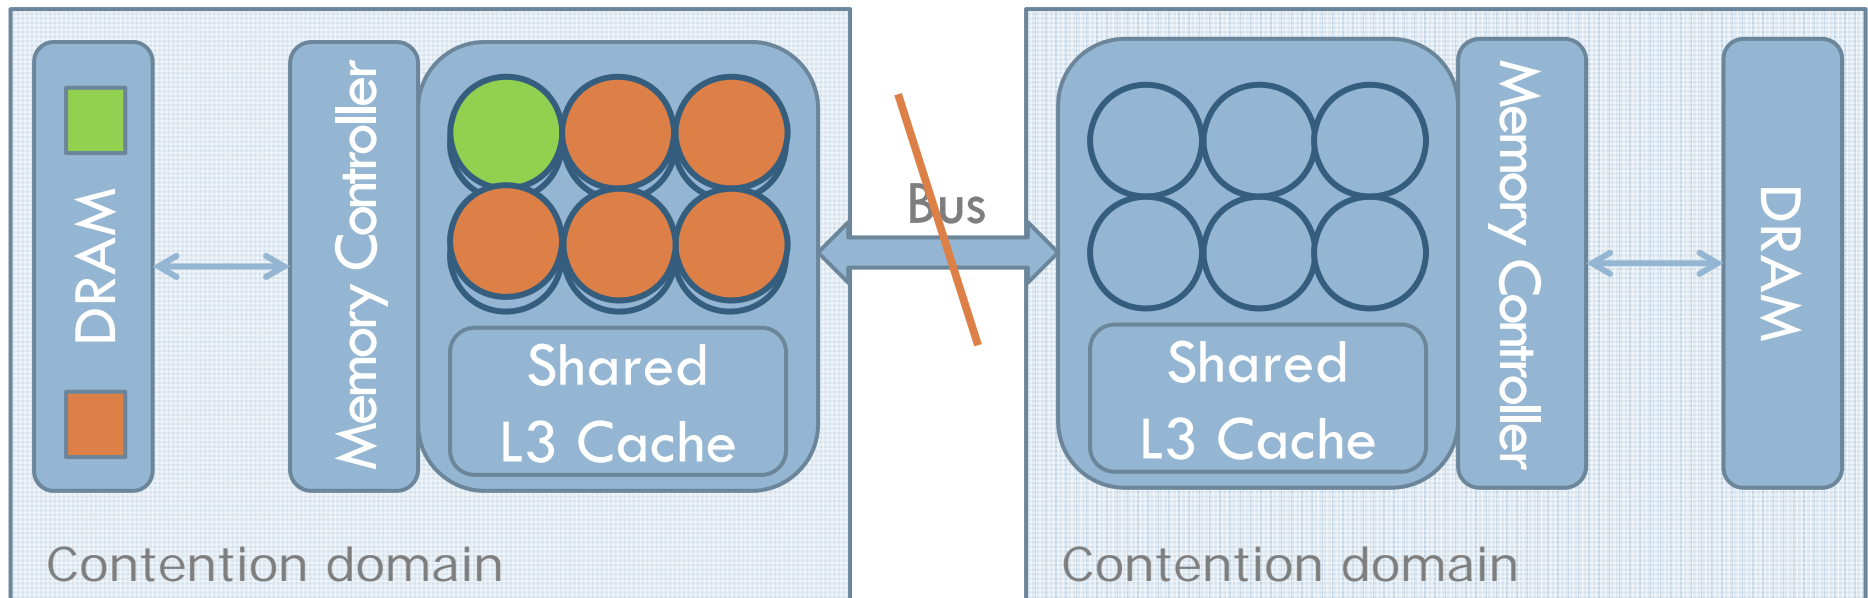
- Linux + Click
- Commodity Intel Xeon server



Basic Configuration

11

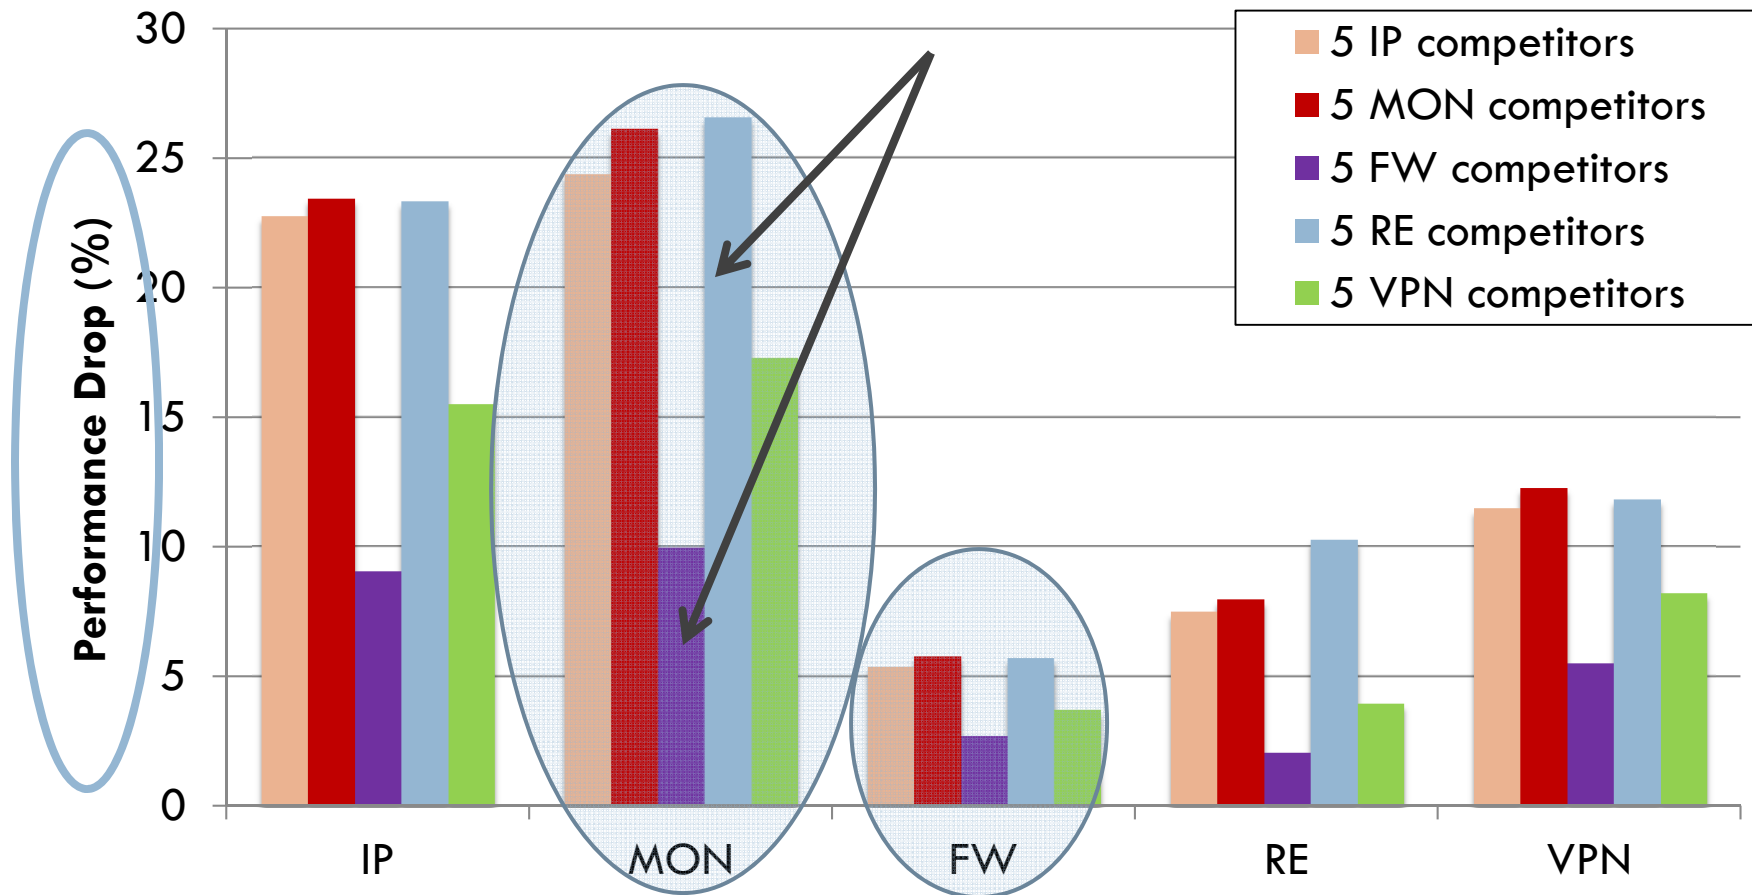
- One application per core
- NUMA-aware memory allocation



Contended resources: cache *and* memory controller

Resource Contention Effects

12



Outline

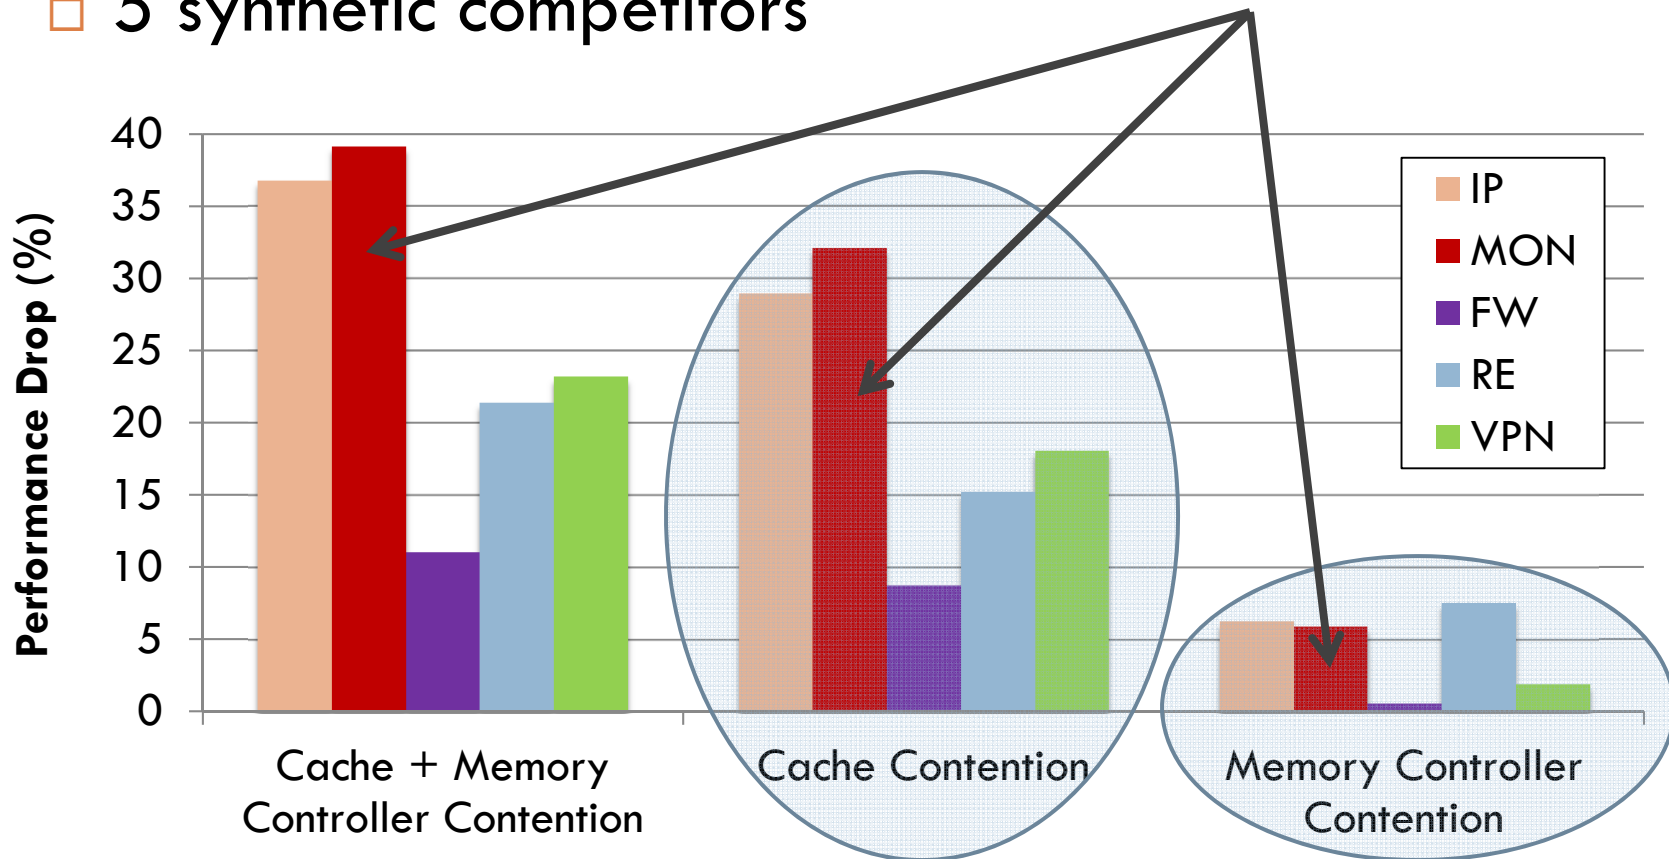
13

- System overview
- **Contention factors**
- Observations on application behavior
- A simple prediction method
- Intuition

Contention Factors

14

- 5 synthetic competitors



Cache is the dominant contention factor

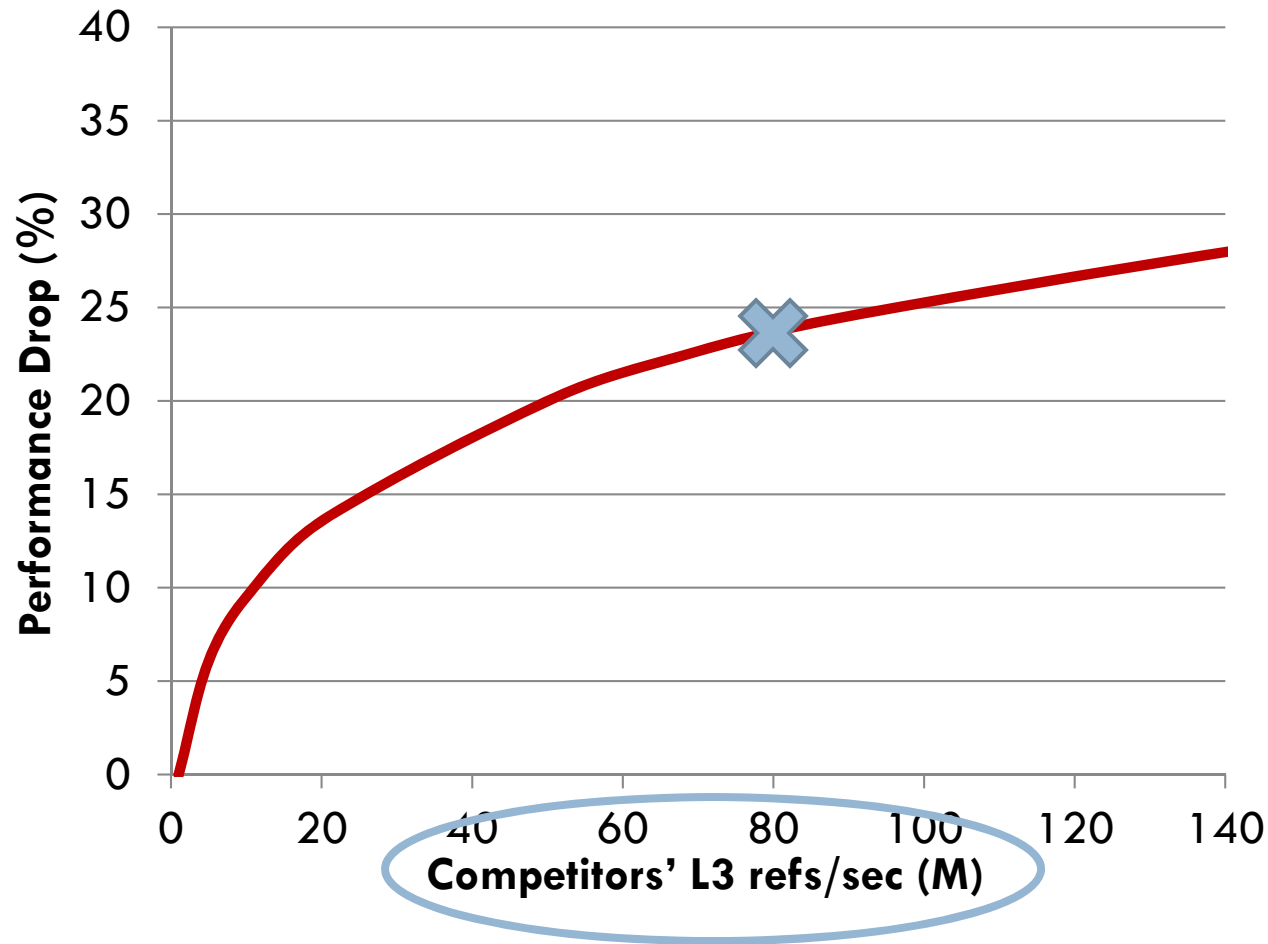
Outline

15

- System overview
- Contention factors
- **Observations on application behavior**
- A simple prediction method
- Intuition

Characterize Application Behavior


16



continuous curves:
synthetic competitors

Characterize Application Behavior

17




continuous curves:
synthetic competitors

individual points:
realistic competitors

Characterize Application Behavior

18




continuous curves:
synthetic competitors

individual points:
realistic competitors

Obs. #1: competitors' cache refs/sec determine drop

Characterize Application Behavior

19



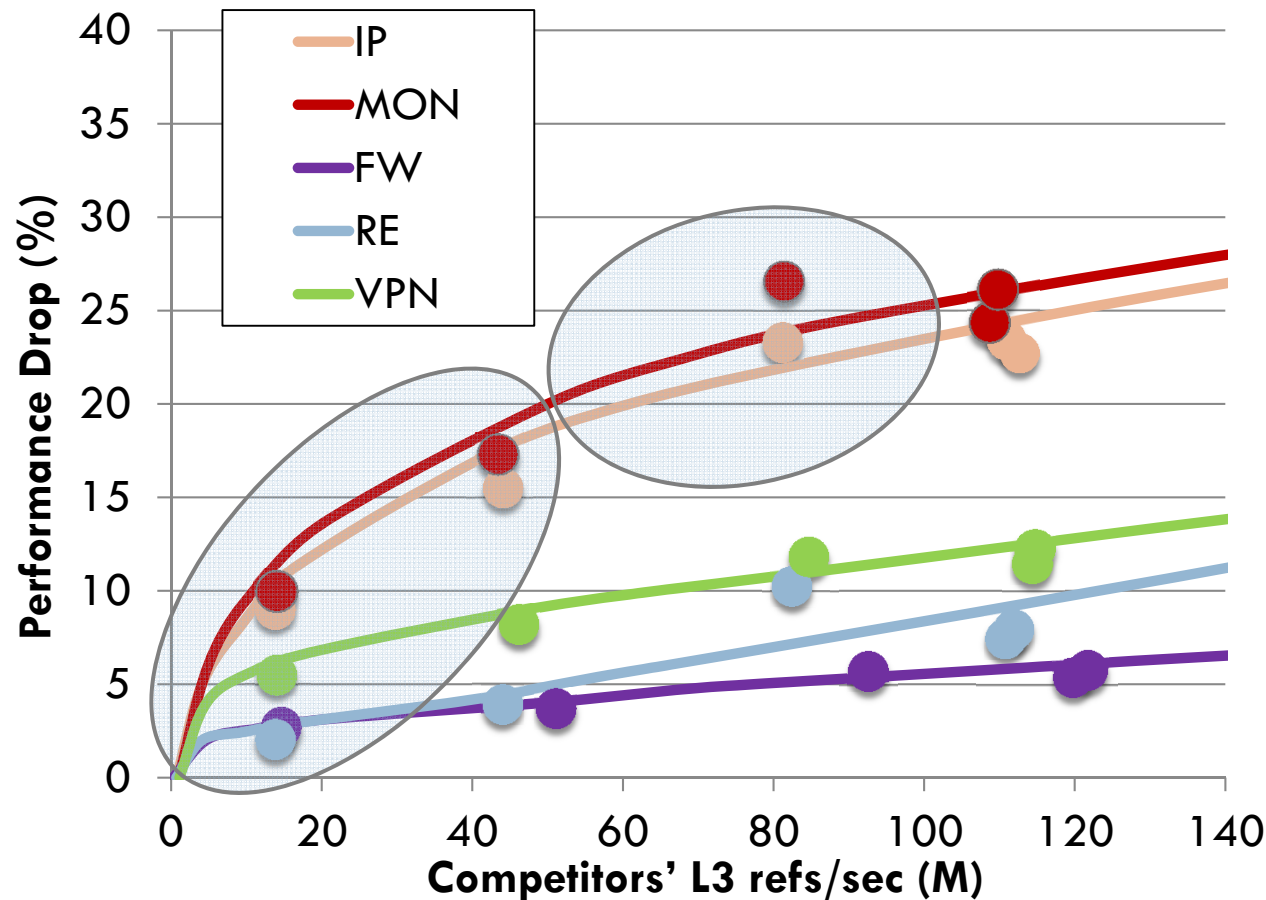
continuous curves:
synthetic competitors

individual points:
realistic competitors

Obs. #1: competitors' cache refs/sec determine drop

Characterize Application Behavior

20



continuous curves:
synthetic competitors

individual points:
realistic competitors

Obs. #2: drop curve grows slowly after certain point

Outline

21

- System overview
- Contention factors
- Observations on application behavior
- **A simple prediction method**
- Intuition

Contention Effects Prediction

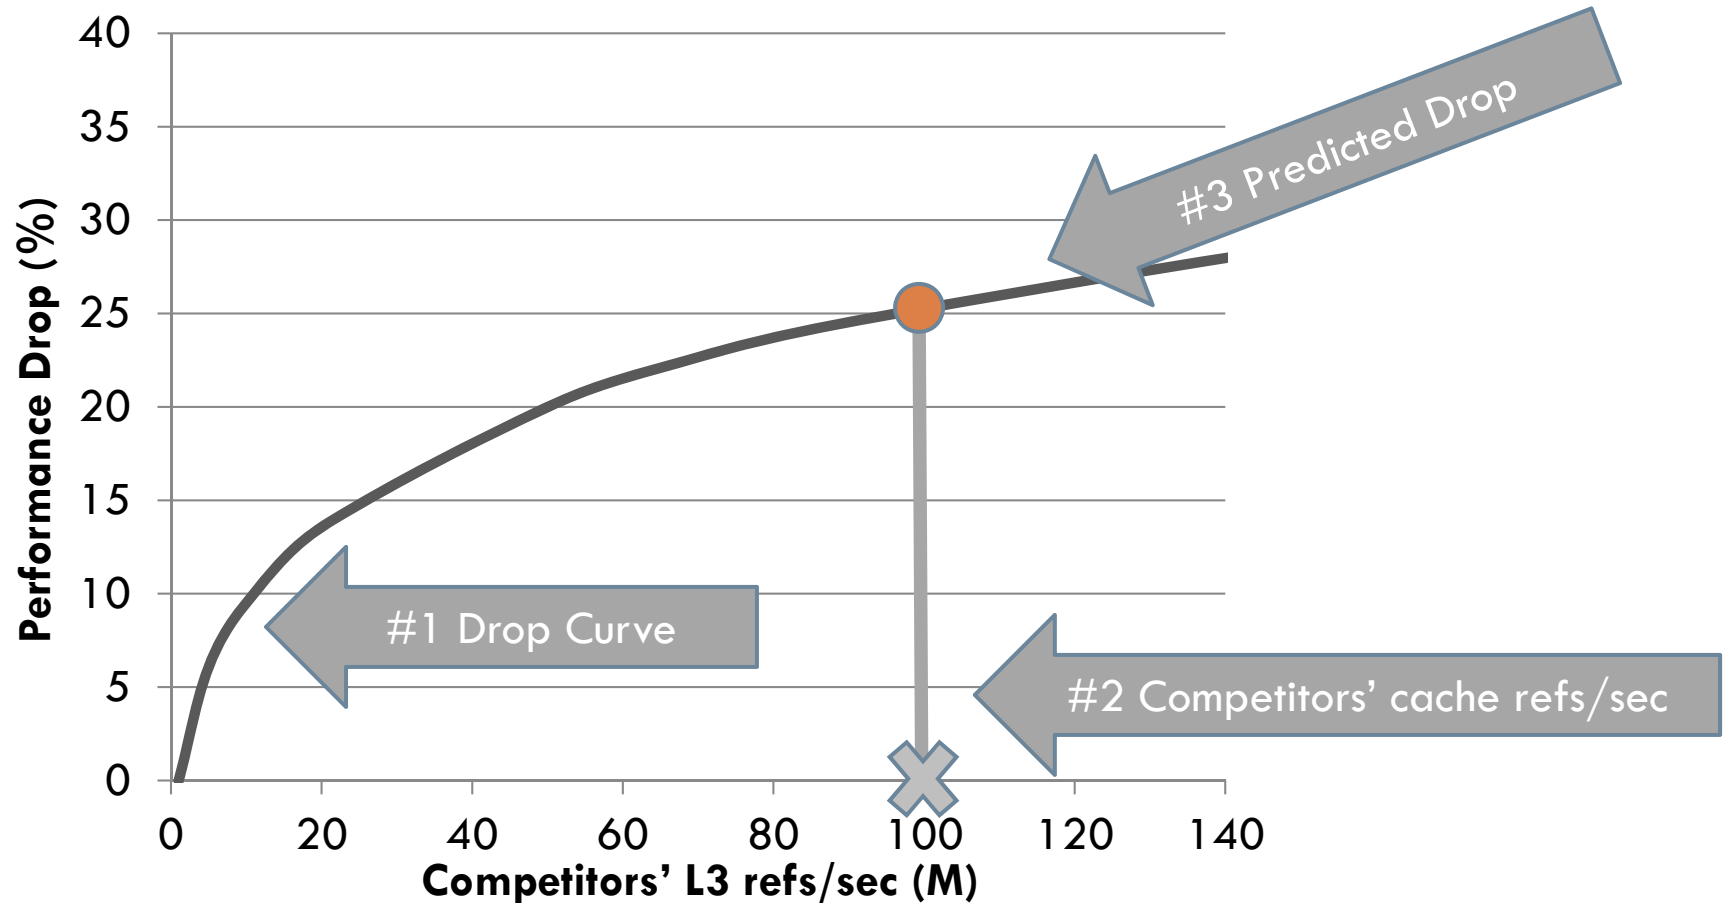
22

- Step#1: performance drop curve for each app
 - Synthetic competitors – random cache reads
 - Vary competitors' cache refs/sec
- Step#2: cache refs/sec for each app running alone
- Step#3: predicted drop equals the value of the drop curve corresponding to the competing cache refs/sec

Simple offline profiling

Step by Step Prediction

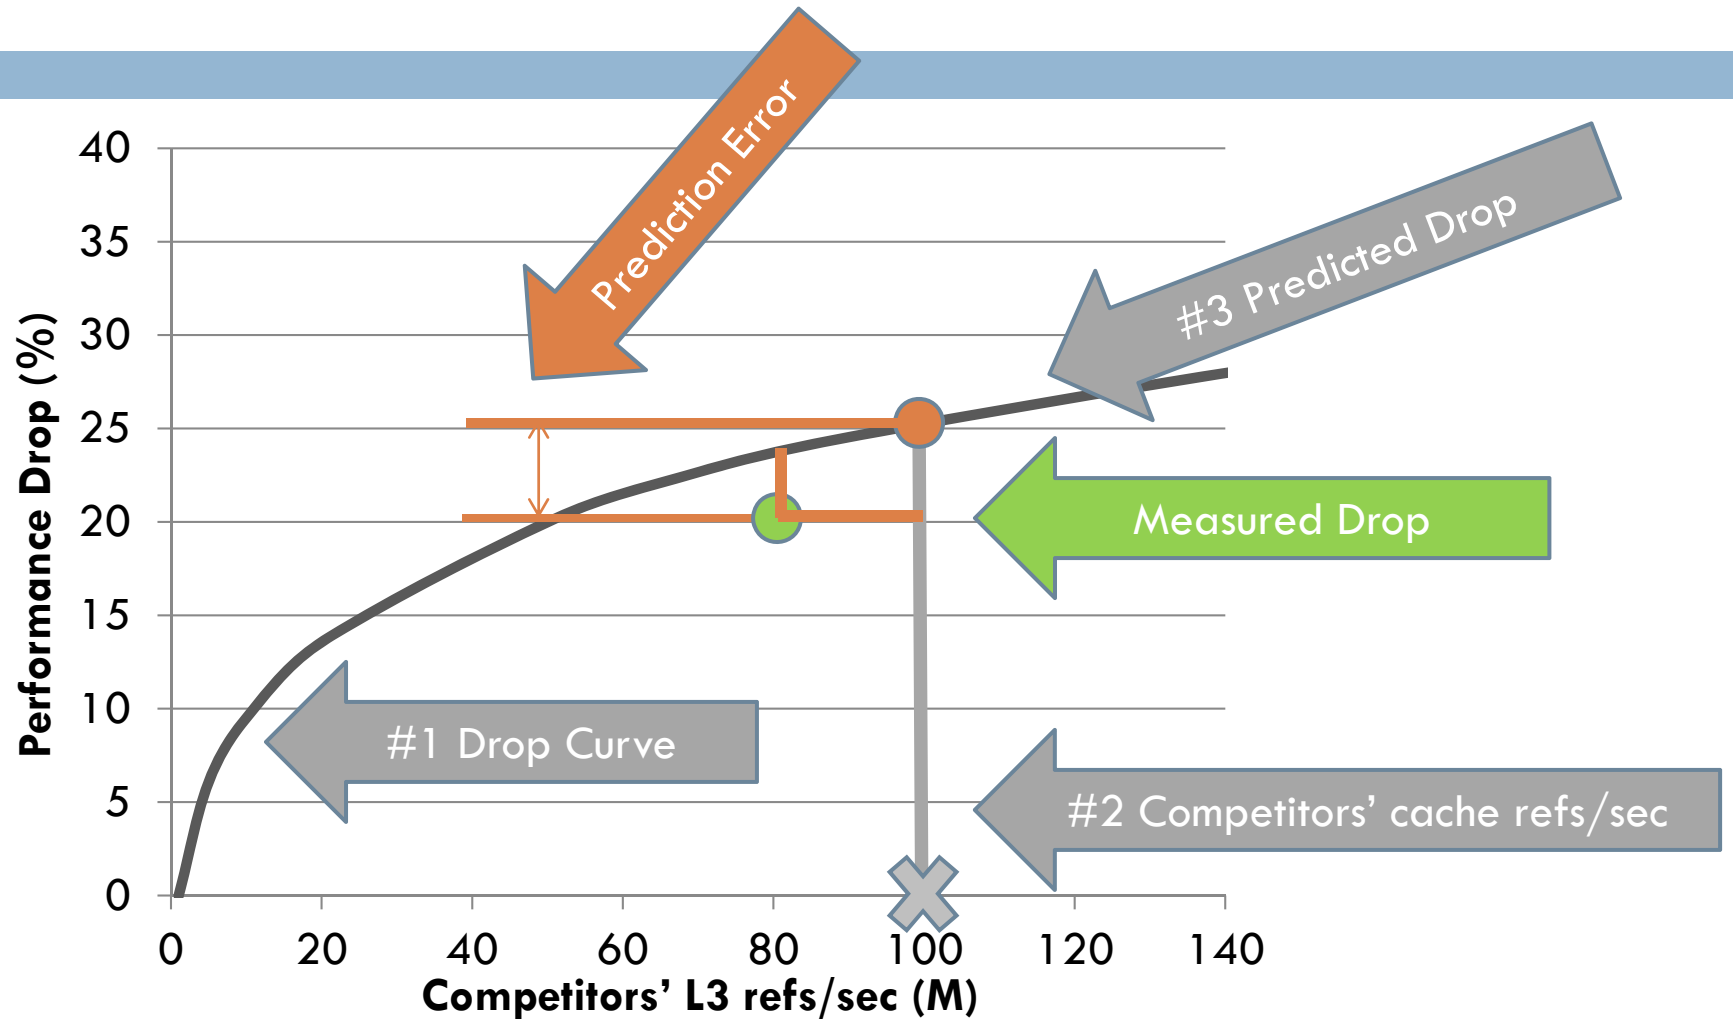
23



Simple offline profiling

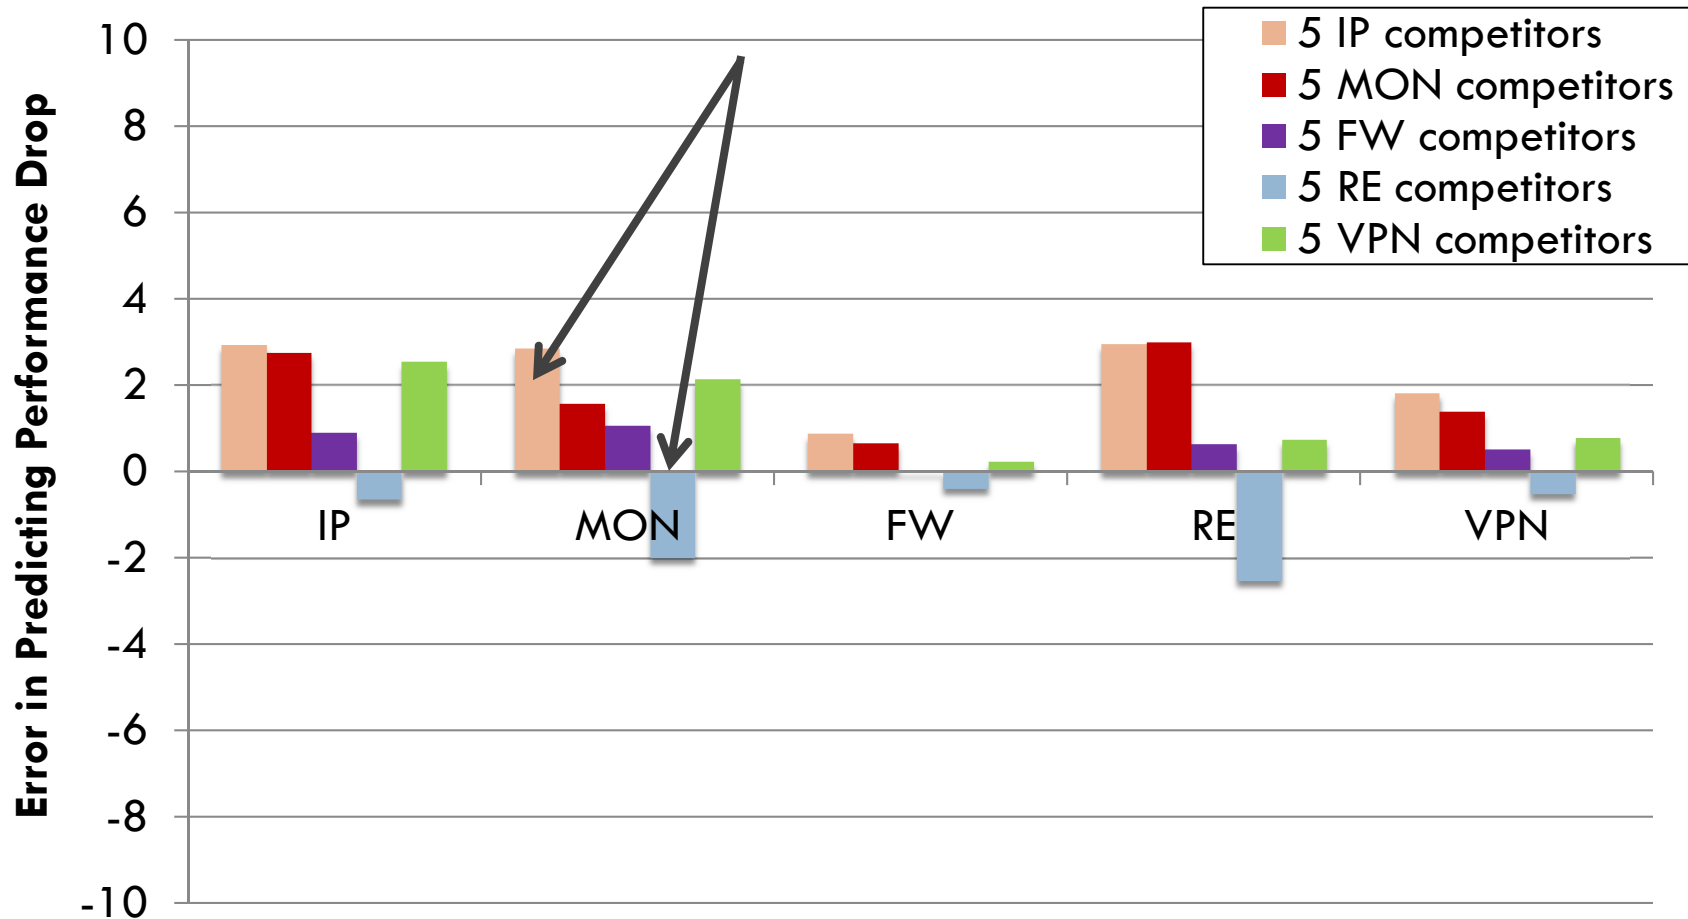
Prediction Errors

24



Evaluation

25



Contention effects are predictable

Outline

26

- System overview
- Contention factors
- Observations on application behavior
- A simple prediction method
- **Intuition**

The Intuition

27

- Obs. #1: competitors' cache refs/sec determine drop
 - Aggregate data exceeds cache size
 - 3MB shared cache/core

The Intuition

28

- Obs. #1: competitors' cache refs/sec determine drop
 - Aggregate data exceeds cache size
 - 3MB shared cache/core

- Obs. #2: drop curve grows slowly after certain point
 - Most damage happens early on
 - Simple probabilistic analysis

Conclusion

29

- It is feasible to build a packet-processing platform with predictable performance using simple techniques
 - ▣ 3% prediction error

- Contention-aware job placement does not bring significant benefit to the overall performance
 - ▣ 2% potential improvement