Runtime Measurements in the Cloud

Observing, Analyzing, and Reducing Variance

Jörg Schad, Jens Dittrich, and Jorge-Arnulfo Quiané-Ruiz

Information Systems Group
Saarland University

VLDB 2010
September 14th, Singapore

Motivation

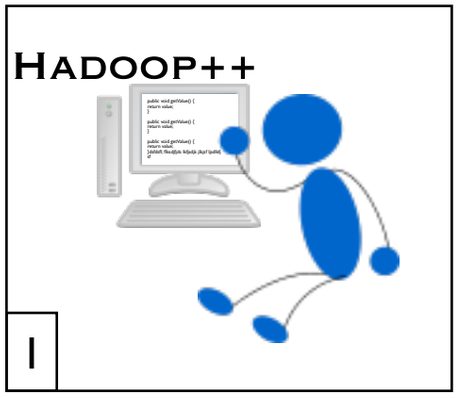

Motivation

SAARLAND UNIVERSITY
Information Systems Group

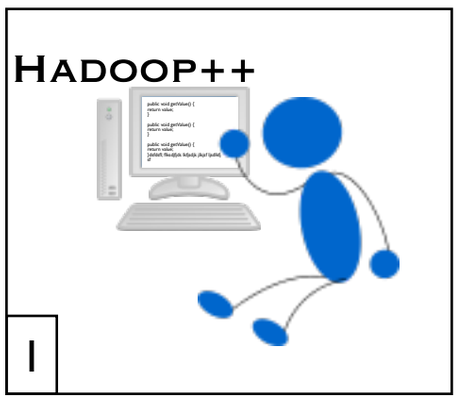

J. Dittrich et al., *Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)*

VLDB 2010

Presentation on Wednesday at 12:00

Research track: Cloud computing session

Motivation

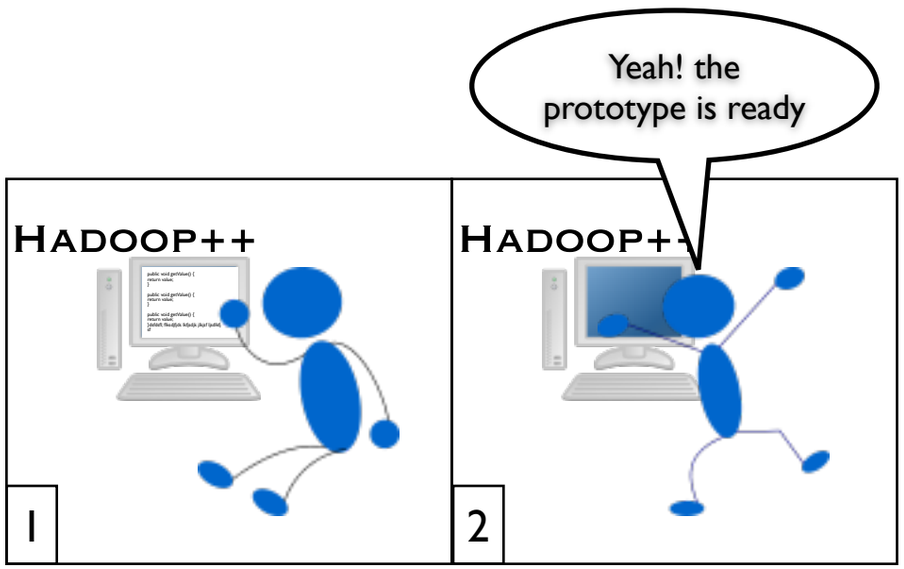

J. Dittrich et al., *Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)*
VLDB 2010
Presentation on Wednesday at 12:00
Research track: Cloud computing session

Motivation

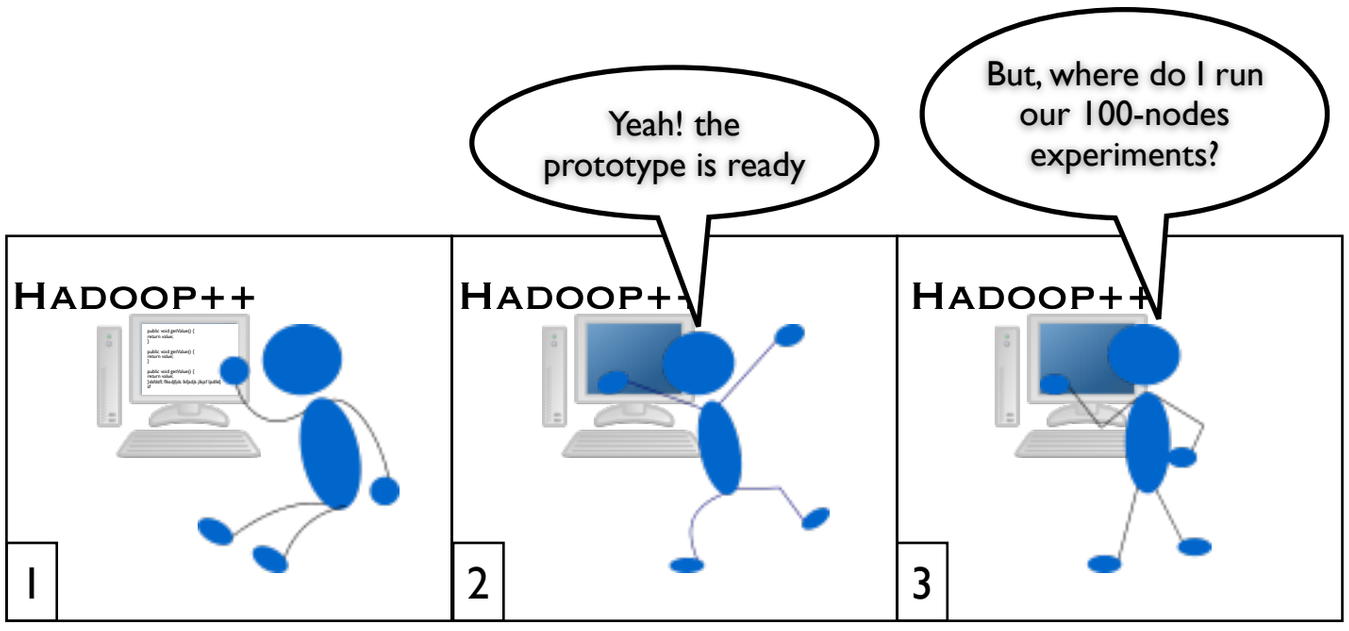

J. Dittrich et al., *Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)*
VLDB 2010
Presentation on Wednesday at 12:00
Research track: Cloud computing session

Motivation

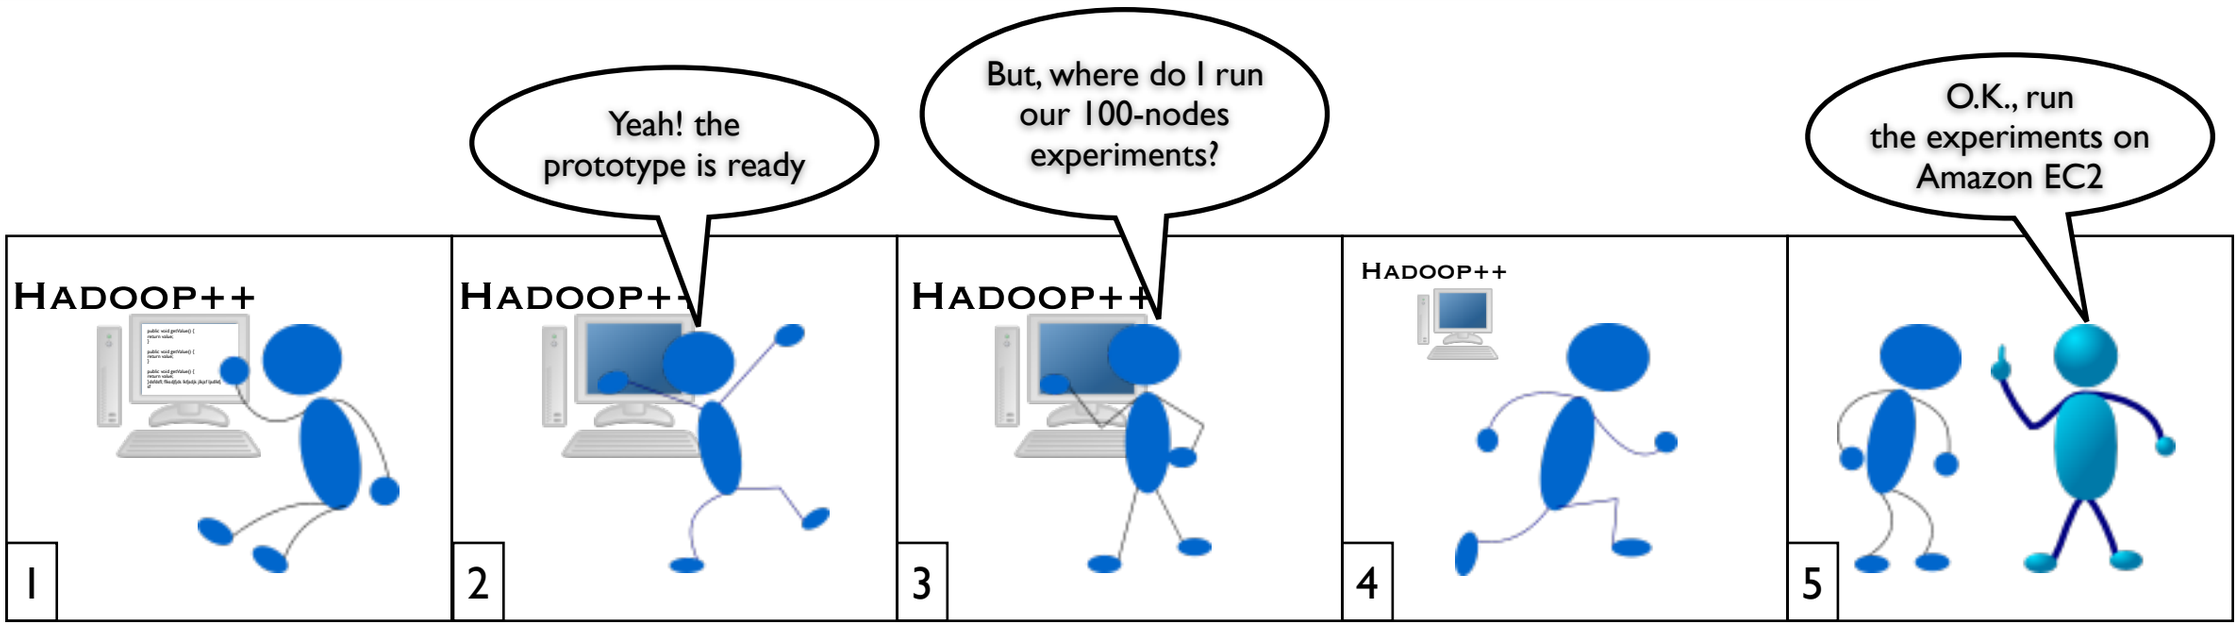

J. Dittrich et al., *Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)*
VLDB 2010
Presentation on Wednesday at 12:00
Research track: Cloud computing session

Motivation

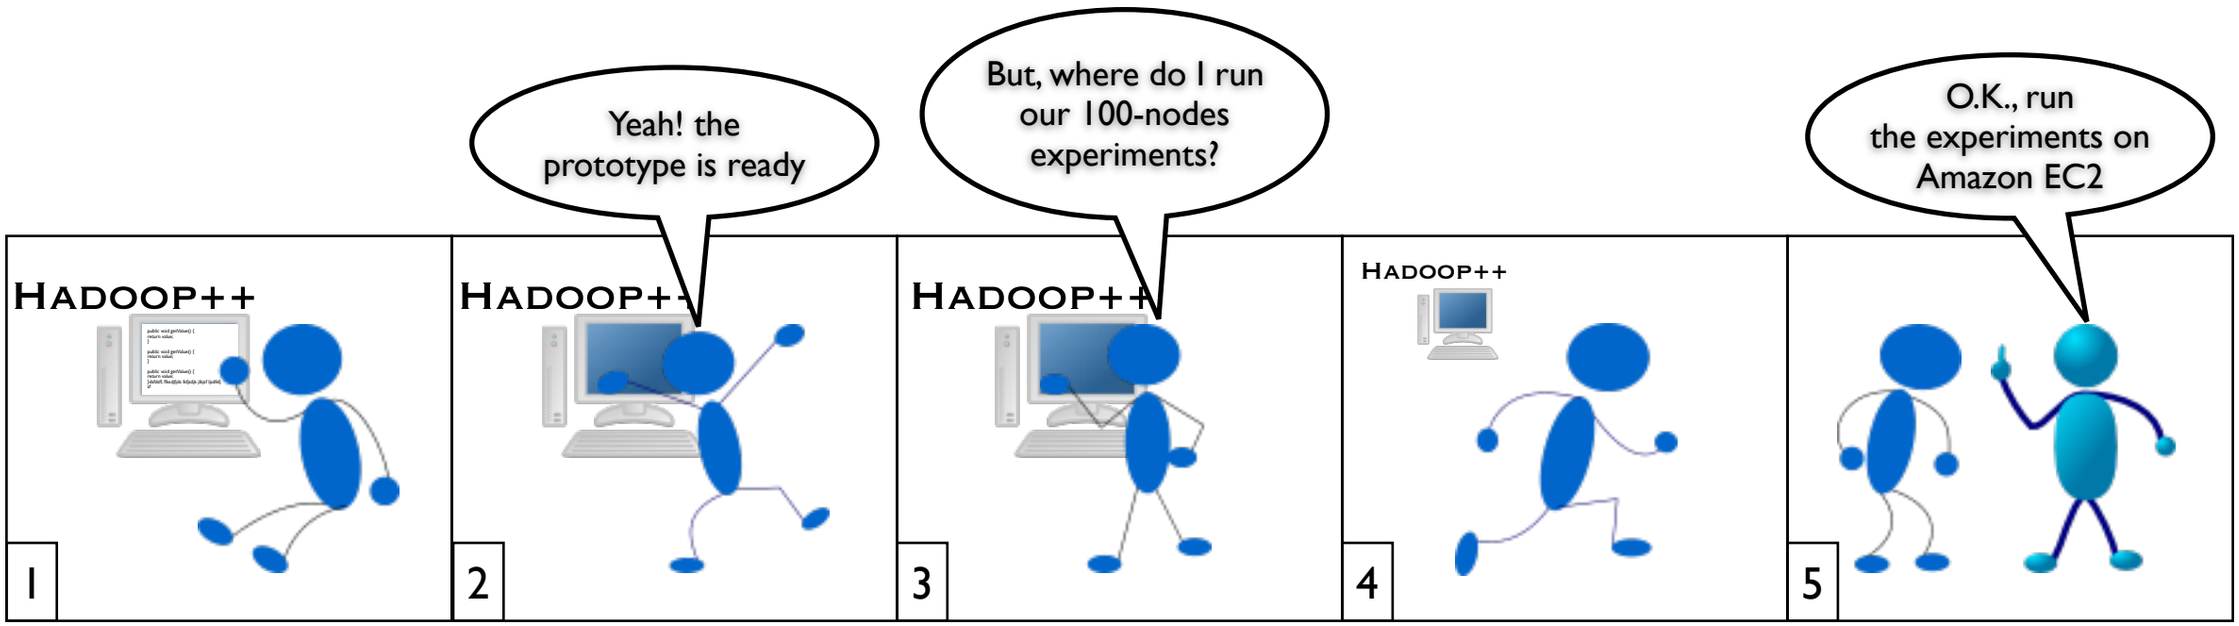

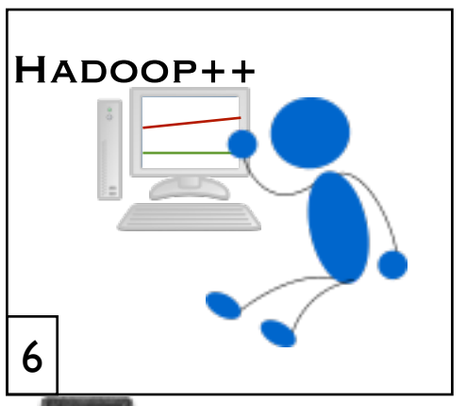

J. Dittrich et al., *Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)*
VLDB 2010
Presentation on Wednesday at 12:00
Research track: Cloud computing session

Motivation

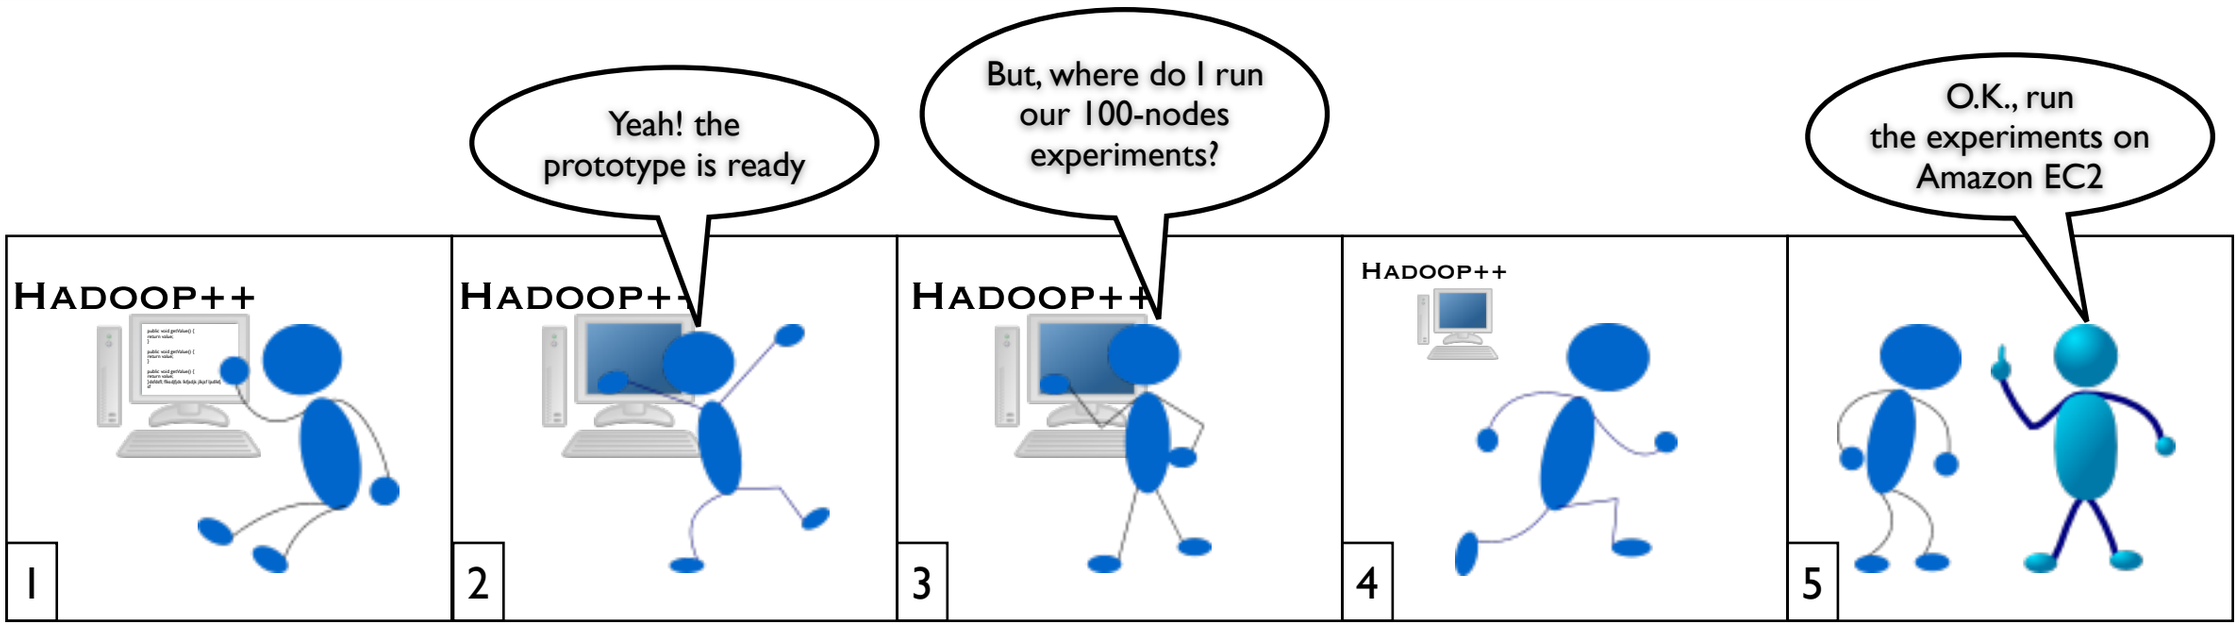

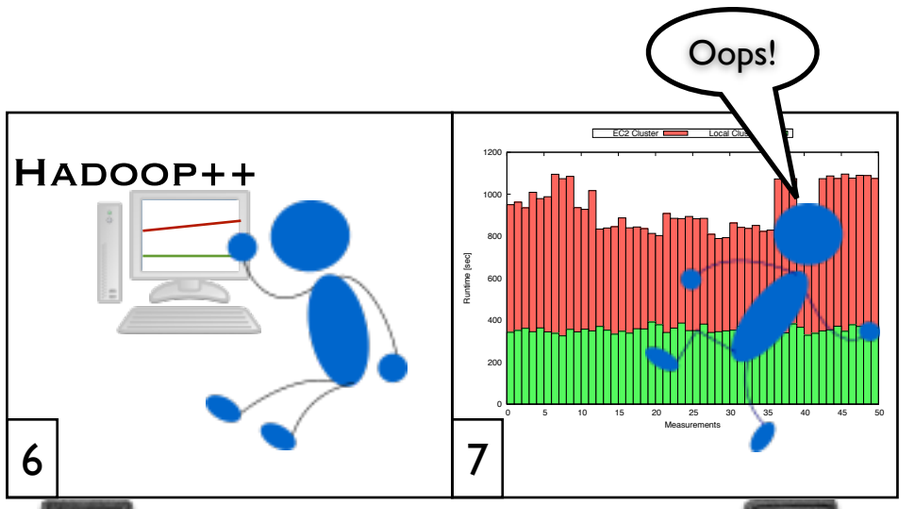

J. Dittrich et al., *Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)*
VLDB 2010
Presentation on Wednesday at 12:00
Research track: Cloud computing session

Motivation

SAARLAND UNIVERSITY
Information Systems Group

- **Application:** a MapReduce aggregation job
- **Number of virtual nodes:** 50
- **Repetitions:** once every hour

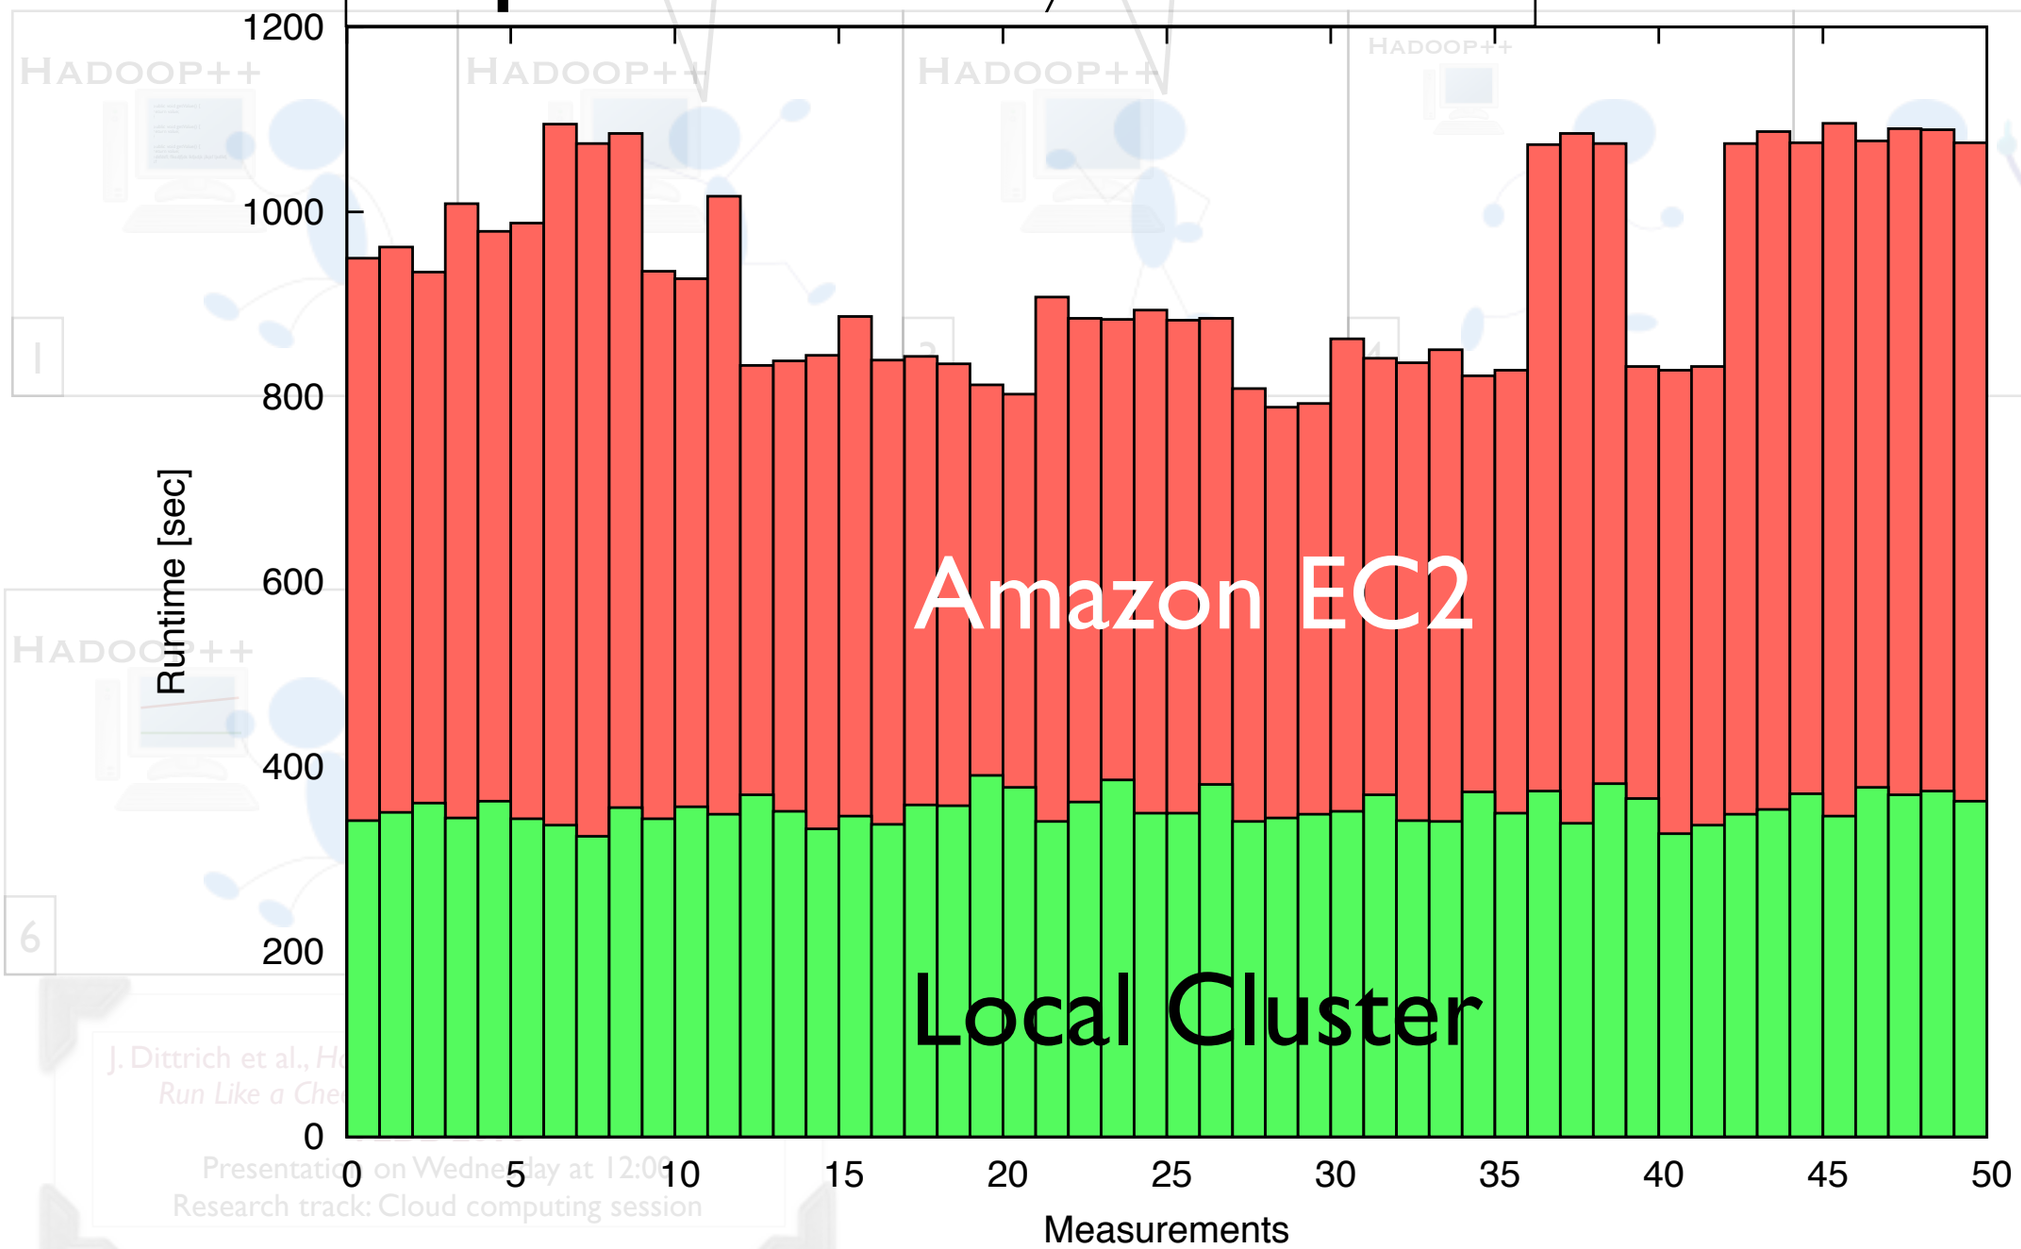

Motivation

SAARLAND UNIVERSITY
Information Systems Group

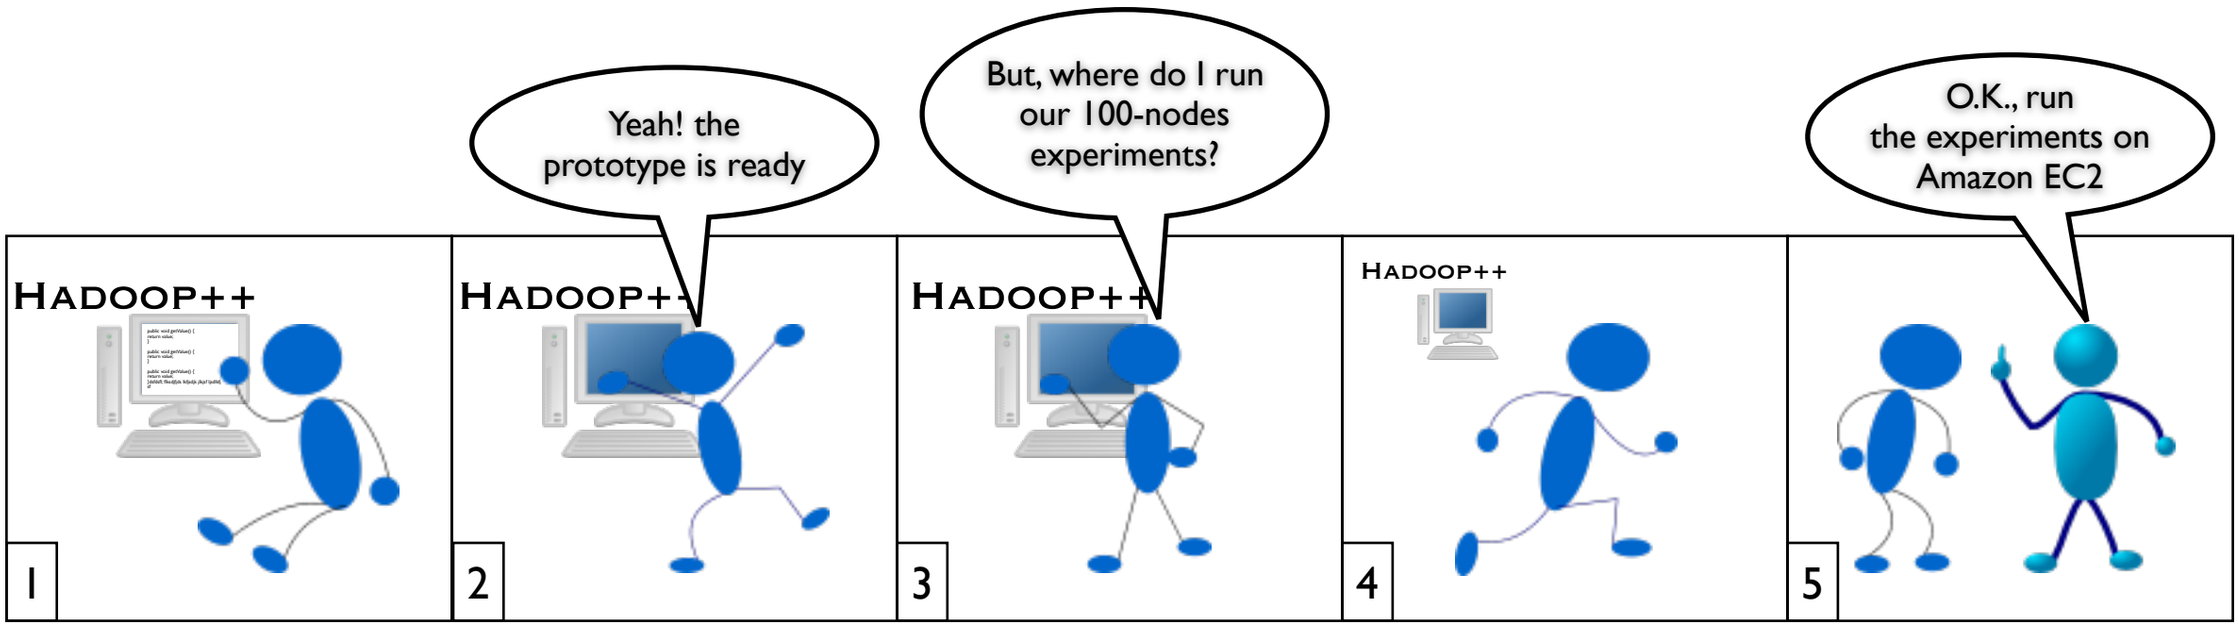

J. Dittrich et al., *Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)*
VLDB 2010
Presentation on Wednesday at 12:00
Research track: Cloud computing session

Motivation

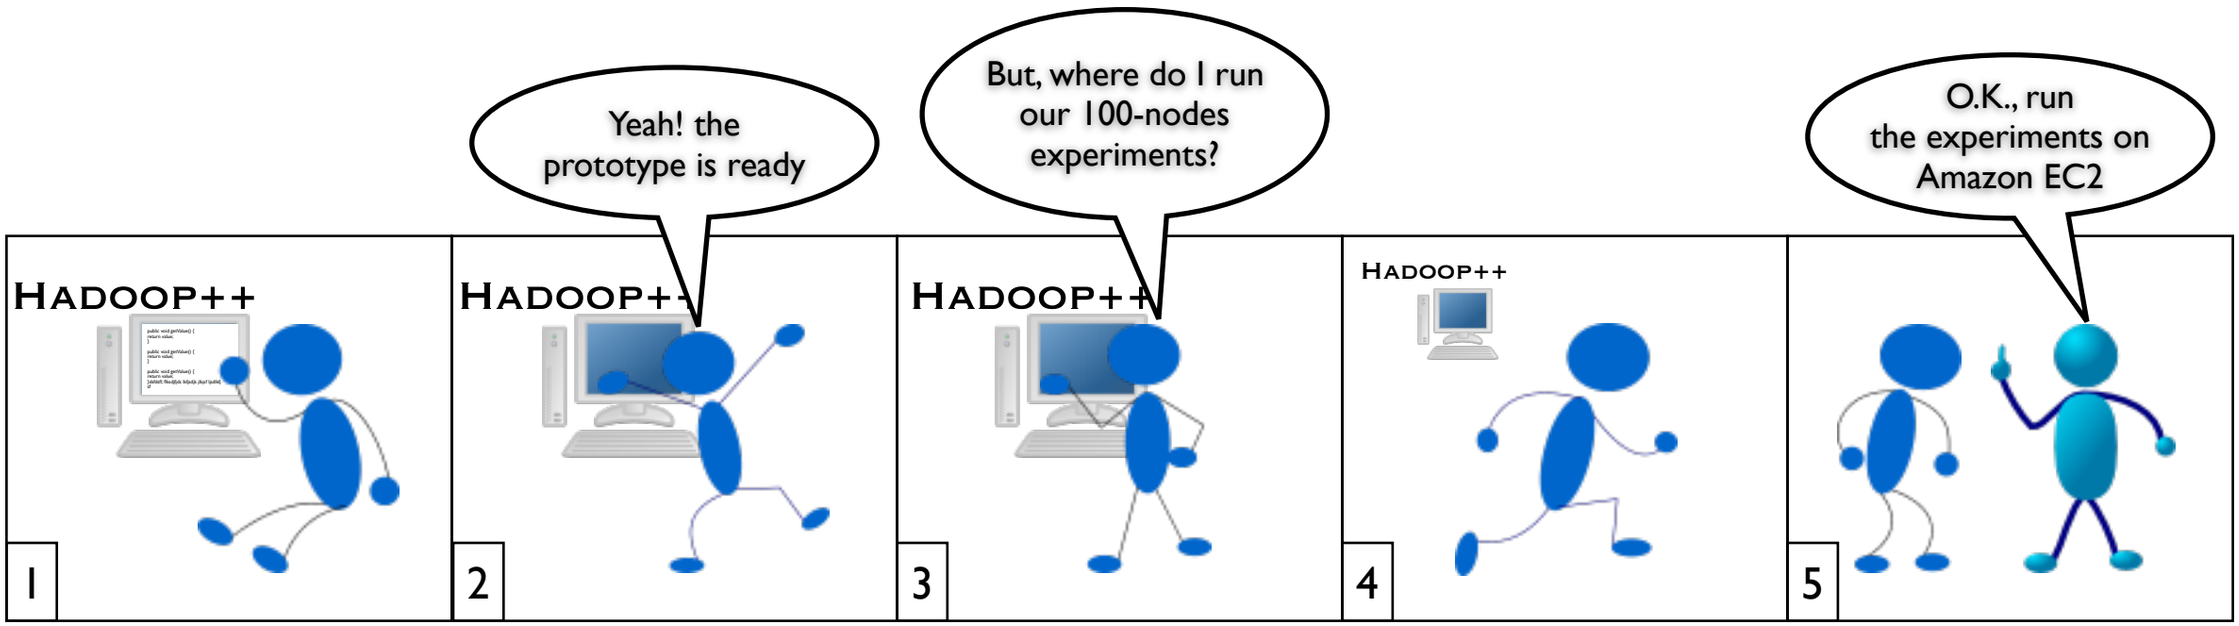

J. Dittrich et al., *Hadoop++: Making a Yellow Elephant Run Like a Cheetah (Without It Even Noticing)*
VLDB 2010
Presentation on Wednesday at 12:00
Research track: Cloud computing session

Related Work

SAARLAND UNIVERSITY
Information Systems Group

Related Work

SAARLAND UNIVERSITY
Information Systems Group

M. Armbrust et al., Above the Clouds: A Berkeley View of Cloud Computing. UCB Technical Report, 2009.

Summary: performance unpredictability is mentioned as one of the major obstacles for Cloud computing.

Related Work

SAARLAND UNIVERSITY
Information Systems Group

Research Challenges

M. Armbrust et al., Above the Clouds: A Berkeley View of Cloud Computing. UCB Technical Report, 2009.

Summary: performance unpredictability is mentioned as one of the major obstacles for Cloud computing.

S. Ostermann et al., A Performance Analysis of EC2 Cloud computing Services for Scientific Computing. Cloudcomp, 2009.

Summary: evaluation of different Cloud services of Amazon in terms of cost and performance.

Related Work

SAARLAND UNIVERSITY
Information Systems Group

Research Challenges

M. Armbrust et al., Above the Clouds: A Berkeley View of Cloud Computing. UCB Technical Report, 2009.

Summary: performance unpredictability is mentioned as one of the major obstacles for Cloud computing.

Absolute Performance

S. Ostermann et al., A Performance Analysis of EC2 Cloud computing Services for Scientific Computing. Cloudcomp, 2009.

Summary: evaluation of different Cloud services of Amazon in terms of cost and performance.

D. Kossmann et al., An Evaluation of Alternative Architectures for Transaction Processing in the Cloud. SIGMOD, 2010.

Summary: cost and performance evaluation of different distributed databases architectures and cloud providers.

[Appeared after VLDB'10 deadline]

Related Work

SAARLAND UNIVERSITY
Information Systems Group

Research Challenges

M. Armbrust et al., Above the Clouds: A Berkeley View of Cloud Computing. UCB Technical Report, 2009.

Summary: performance unpredictability is mentioned as one of the major obstacles for Cloud computing.

Absolute Performance

S. Ostermann et al., A Performance Analysis of EC2 Cloud computing Services for Scientific Computing. Cloudcomp, 2009.

Summary: evaluation of different Cloud services of Amazon in terms of cost and performance.

Application Performance

D. Kossmann et al., An Evaluation of Alternative Architectures for Transaction Processing in the Cloud. SIGMOD, 2010.

Summary: cost and performance evaluation of different distributed databases architectures and cloud providers.

Variability in Performance

?

Related Work

SAARLAND UNIVERSITY
Information Systems Group

Research Challenges

M. Armbrust et al., Above the Clouds: A Berkeley View of Cloud Computing. UCB Technical Report, 2009.

Summary: performance unpredictability is mentioned as one of the major obstacles for Cloud computing.

Absolute Performance

S. Ostermann et al., A Performance Analysis of EC2 Cloud computing Services for Scientific Computing. Cloudcomp, 2009.

Summary: evaluation of different Cloud services of Amazon in terms of cost and performance.

Application Performance

D. Kossmann et al., An Evaluation of Alternative Architectures for Transaction Processing in the Cloud. SIGMOD, 2010.

Summary: cost and performance evaluation of different distributed databases architectures and cloud providers.

Variability in Performance

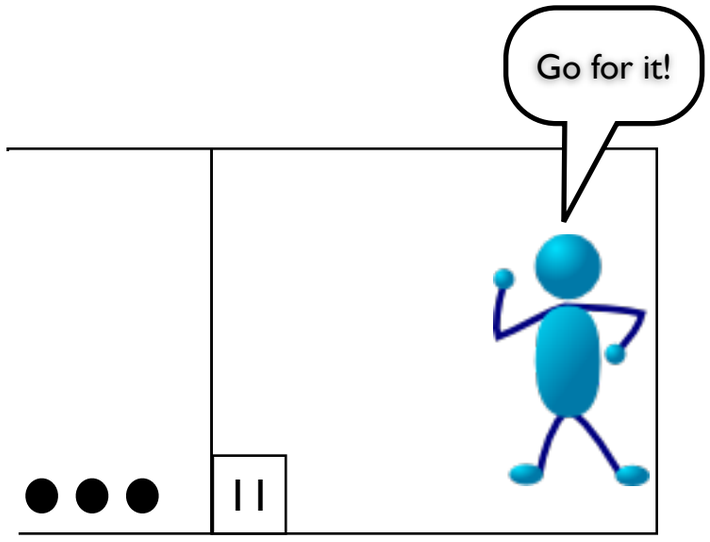

Agenda

SAARLAND UNIVERSITY
Information Systems Group

- Motivation
- Related Work
- Background
- Methodology
- Results & Analysis

Amazon EC2

SAARLAND UNIVERSITY
Information Systems Group

- Most **popular** Cloud infrastructure
- Three **locations**: US, EU, and ASIA [after VLDB'10 deadline]
- Different **availability zones** for US
- Linux-based virtual machines (**instances**)
- Five EC2 **Instance types**: standard, micro [from September 9th], high-memory, high-cpu, and cluster-compute [after VLDB'10 deadline]

Standard Instances

SAARLAND UNIVERSITY
Information Systems Group

- **Small size instance**
  - 1.7 GB of main memory
  - 1 EC2 Compute Unit
  - 160 GB of local storage
- **Large size instance**
  - 7.5 GB of main memory
  - 4 EC2 Compute Units
  - 850 GB of local storage
- **Extra Large size instance**
  - 15 GB of main memory
  - 8 EC2 Compute Unit
  - 1690 GB of local storage

Standard Instances

- **Small size instance**
  - 1.7 GB of main memory
  - 1 EC2 Compute Unit
  - 160 GB of local storage
- **Large size instance**
  - 7.5 GB of main memory
  - 4 EC2 Compute Units
  - 850 GB of local storage
- **Extra Large size instance**
  - 15 GB of main memory
  - 8 EC2 Compute Unit
  - 1690 GB of local storage

> *"one EC2 compute unit (ECU) provides the **equivalent** CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor."*

Agenda

SAARLAND UNIVERSITY
Information Systems Group

- Motivation
- Related Work
- Background
- Methodology
- Results & Analysis

What to Measure?

SAARLAND UNIVERSITY
Information Systems Group

What to Measure?

SAARLAND UNIVERSITY
Information Systems Group

Microbenchmarks

What to Measure?

SAARLAND UNIVERSITY
Information Systems Group

Microbenchmarks

- CPU performance
- Memory performance

What to Measure?

SAARLAND UNIVERSITY
Information Systems Group

Microbenchmarks

- CPU performance
- Memory performance
- Disk I/O (sequential and random)

What to Measure?

SAARLAND UNIVERSITY
Information Systems Group

Microbenchmarks

In this talk:

- CPU performance
- Memory performance
- Disk I/O (sequential and random)
- Internal network bandwidth
- External network bandwidth
- Instance startup

How to Measure?

SAARLAND UNIVERSITY
Information Systems Group

Microbenchmarks

In this talk:

- CPU performance:
- Memory performance:
- Disk I/O (sequential and random):
- Internal network bandwidth
- External network bandwidth
- Instance startup

How to Measure?

SAARLAND UNIVERSITY
Information Systems Group

Microbenchmarks

In this talk:

- CPU performance: **Ubench**
- Memory performance: **Ubench**
- Disk I/O (sequential and random): **Bonnie++**
- Internal network bandwidth
- External network bandwidth
- Instance startup

Goal of Our Study

SAARLAND UNIVERSITY
Information Systems Group

Goal of Our Study

SAARLAND UNIVERSITY
Information Systems Group

- Do different **Instance types** have different variations in performance?

Goal of Our Study

SAARLAND UNIVERSITY
Information Systems Group

- Do different **Instance types** have different variations in performance?
- Do different **locations** or **availability zones** impact performance?

Goal of Our Study

SAARLAND UNIVERSITY
Information Systems Group

In this talk:

- Do different **Instance types** have different variations in performance?
- Do different **locations** or **availability zones** impact performance?
- Does performance depend on the **time** of the day, weekday, or week?

Setup

SAARLAND UNIVERSITY
Information Systems Group

- **Small** and **large** Instances in **US** and **EU** locations
- **Default** settings for Ubench and Bonnie++
- Results reported in **CET** time

- **Baseline**: our team's cluster at Saarland University
  - 50 Xeon-based virtual nodes
  - 2.66 GHz Quad Core Xeon CPU
  - 16 GB of main memory
  - 6x750 GB SATA hard disks

Setup

SAARLAND UNIVERSITY
Information Systems Group

[**ASIA** location was introduced after VLDB deadline]

- **Small** and **large** Instances in **US** and **EU** locations
- **Default** settings for Ubench and Bonnie++
- Results reported in **CET** time

- **Baseline**: our team's cluster at Saarland University
  - 50 Xeon-based virtual nodes
  - 2.66 GHz Quad Core Xeon CPU
  - 16 GB of main memory
  - 6x750 GB SATA hard disks

Methodology

SAARLAND UNIVERSITY
Information Systems Group

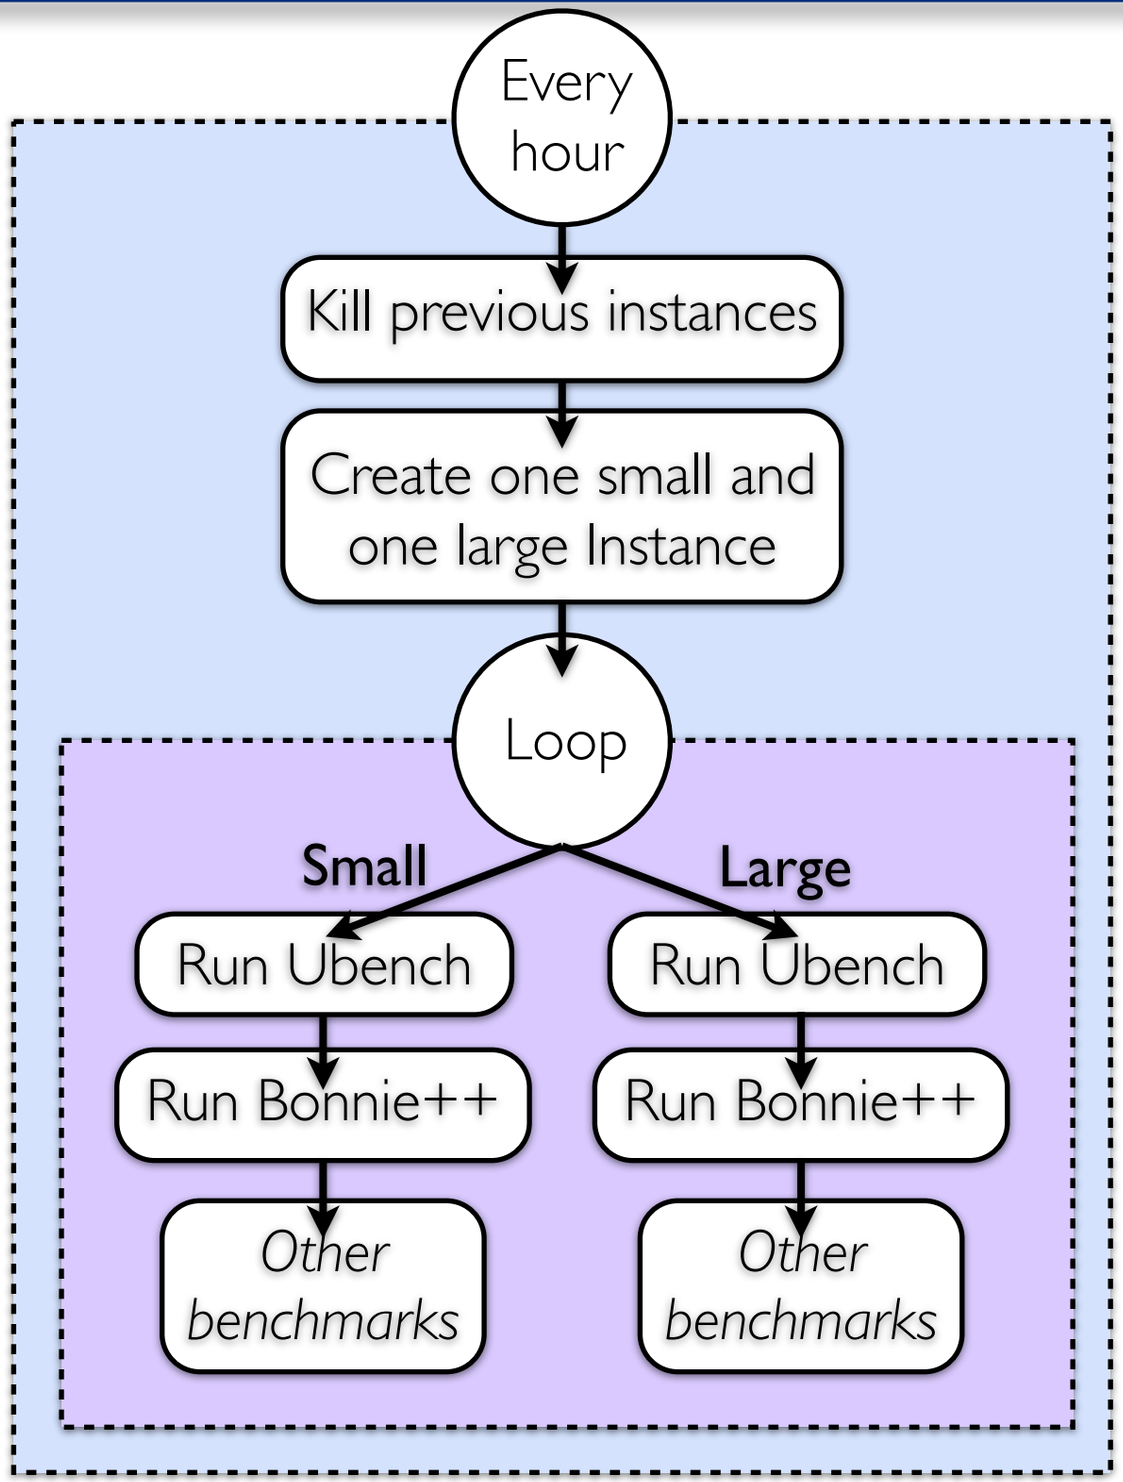

Methodology

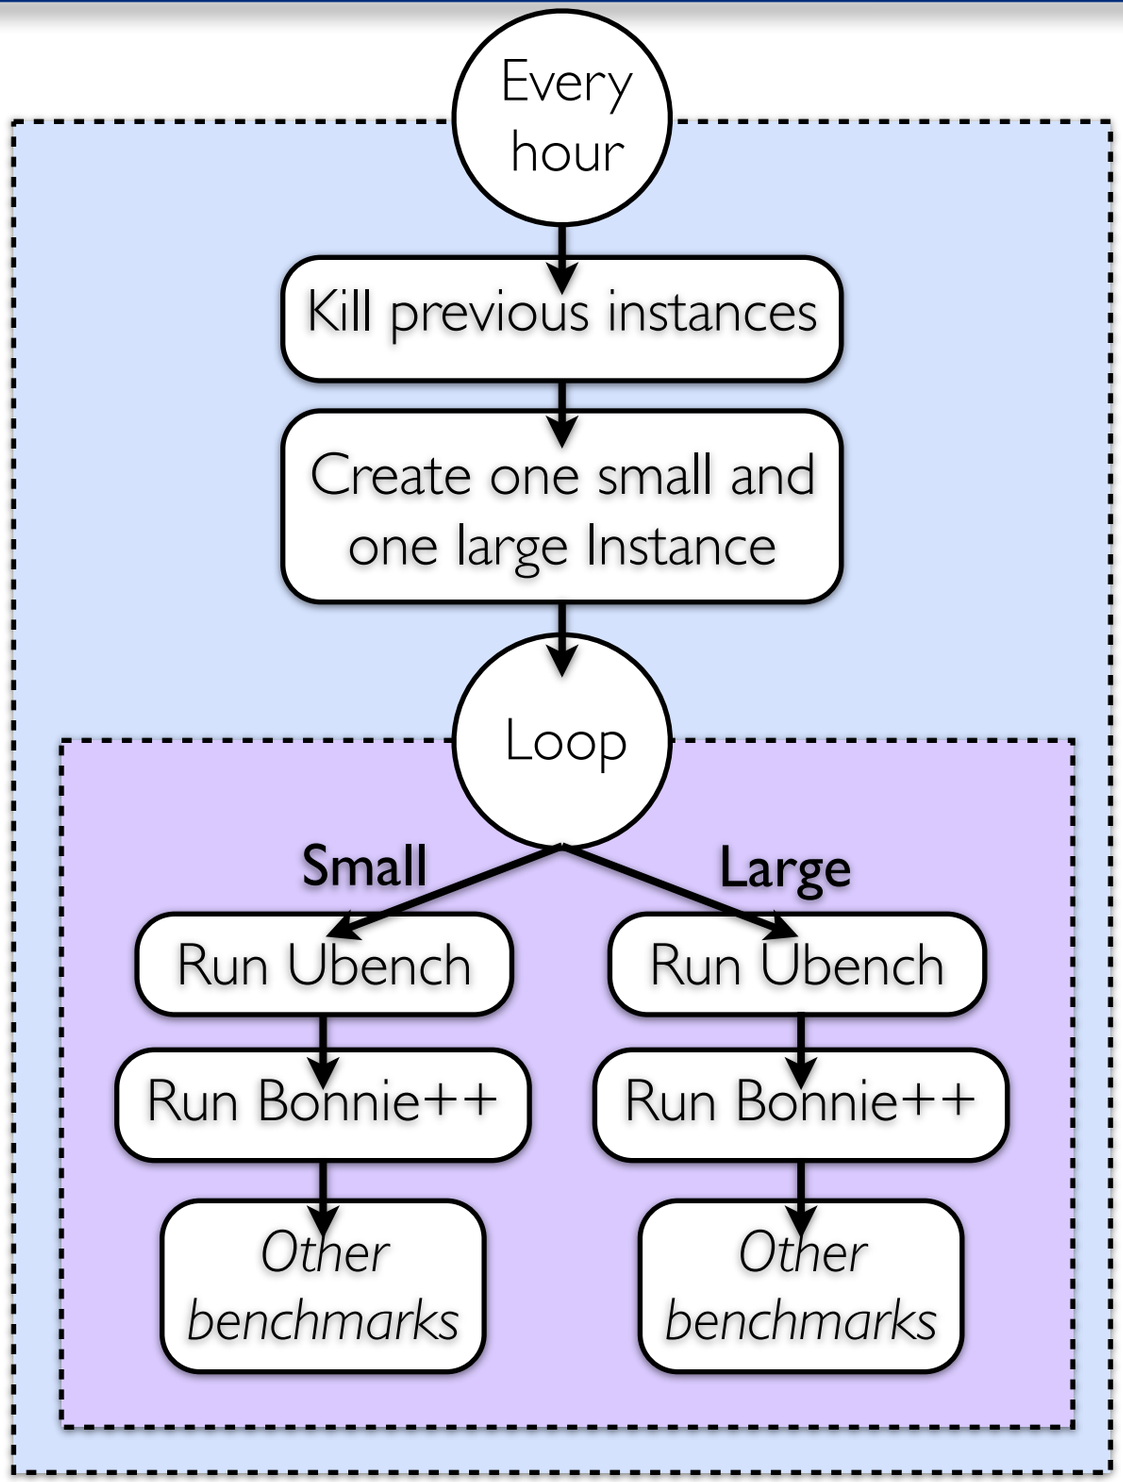

Start: December 14, 2009
End: January 12, 2010
Duration: 31 days

[Results for one additional month, but without any additional pattern]

Measure of Variation

Measure of Variation

SAARLAND UNIVERSITY
Information Systems Group

- Different ones: **range**, **variance**, **standard deviation**, ...

Measure of Variation

SAARLAND UNIVERSITY
Information Systems Group

- Different ones: **range**, **variance**, **standard deviation**, ...
- Need to **compare** data series in different **scales**

Measure of Variation

SAARLAND UNIVERSITY
Information Systems Group

- Different ones: **range**, **variance**, **standard deviation**, ...
- Need to **compare** data series in different **scales**
- **Coefficient of Variation** (COV): ratio of the standard deviation to the mean

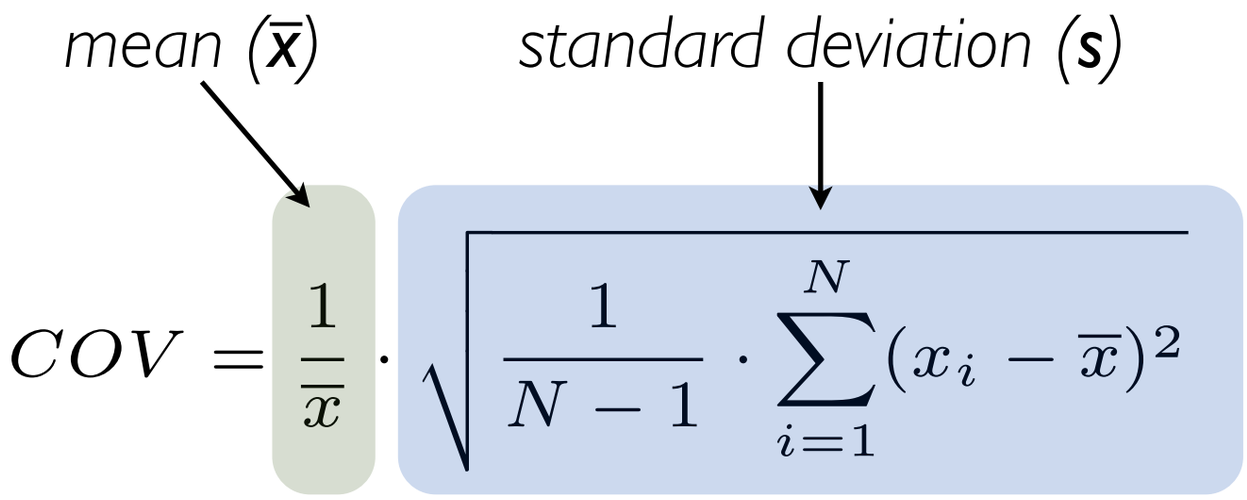

Agenda

SAARLAND UNIVERSITY
Information Systems Group

- Motivation
- Related Work
- Background
- Methodology
- Results & Analysis

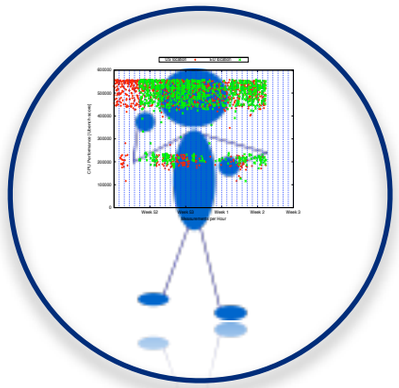

CPU Performance

SAARLAND UNIVERSITY
Information Systems Group

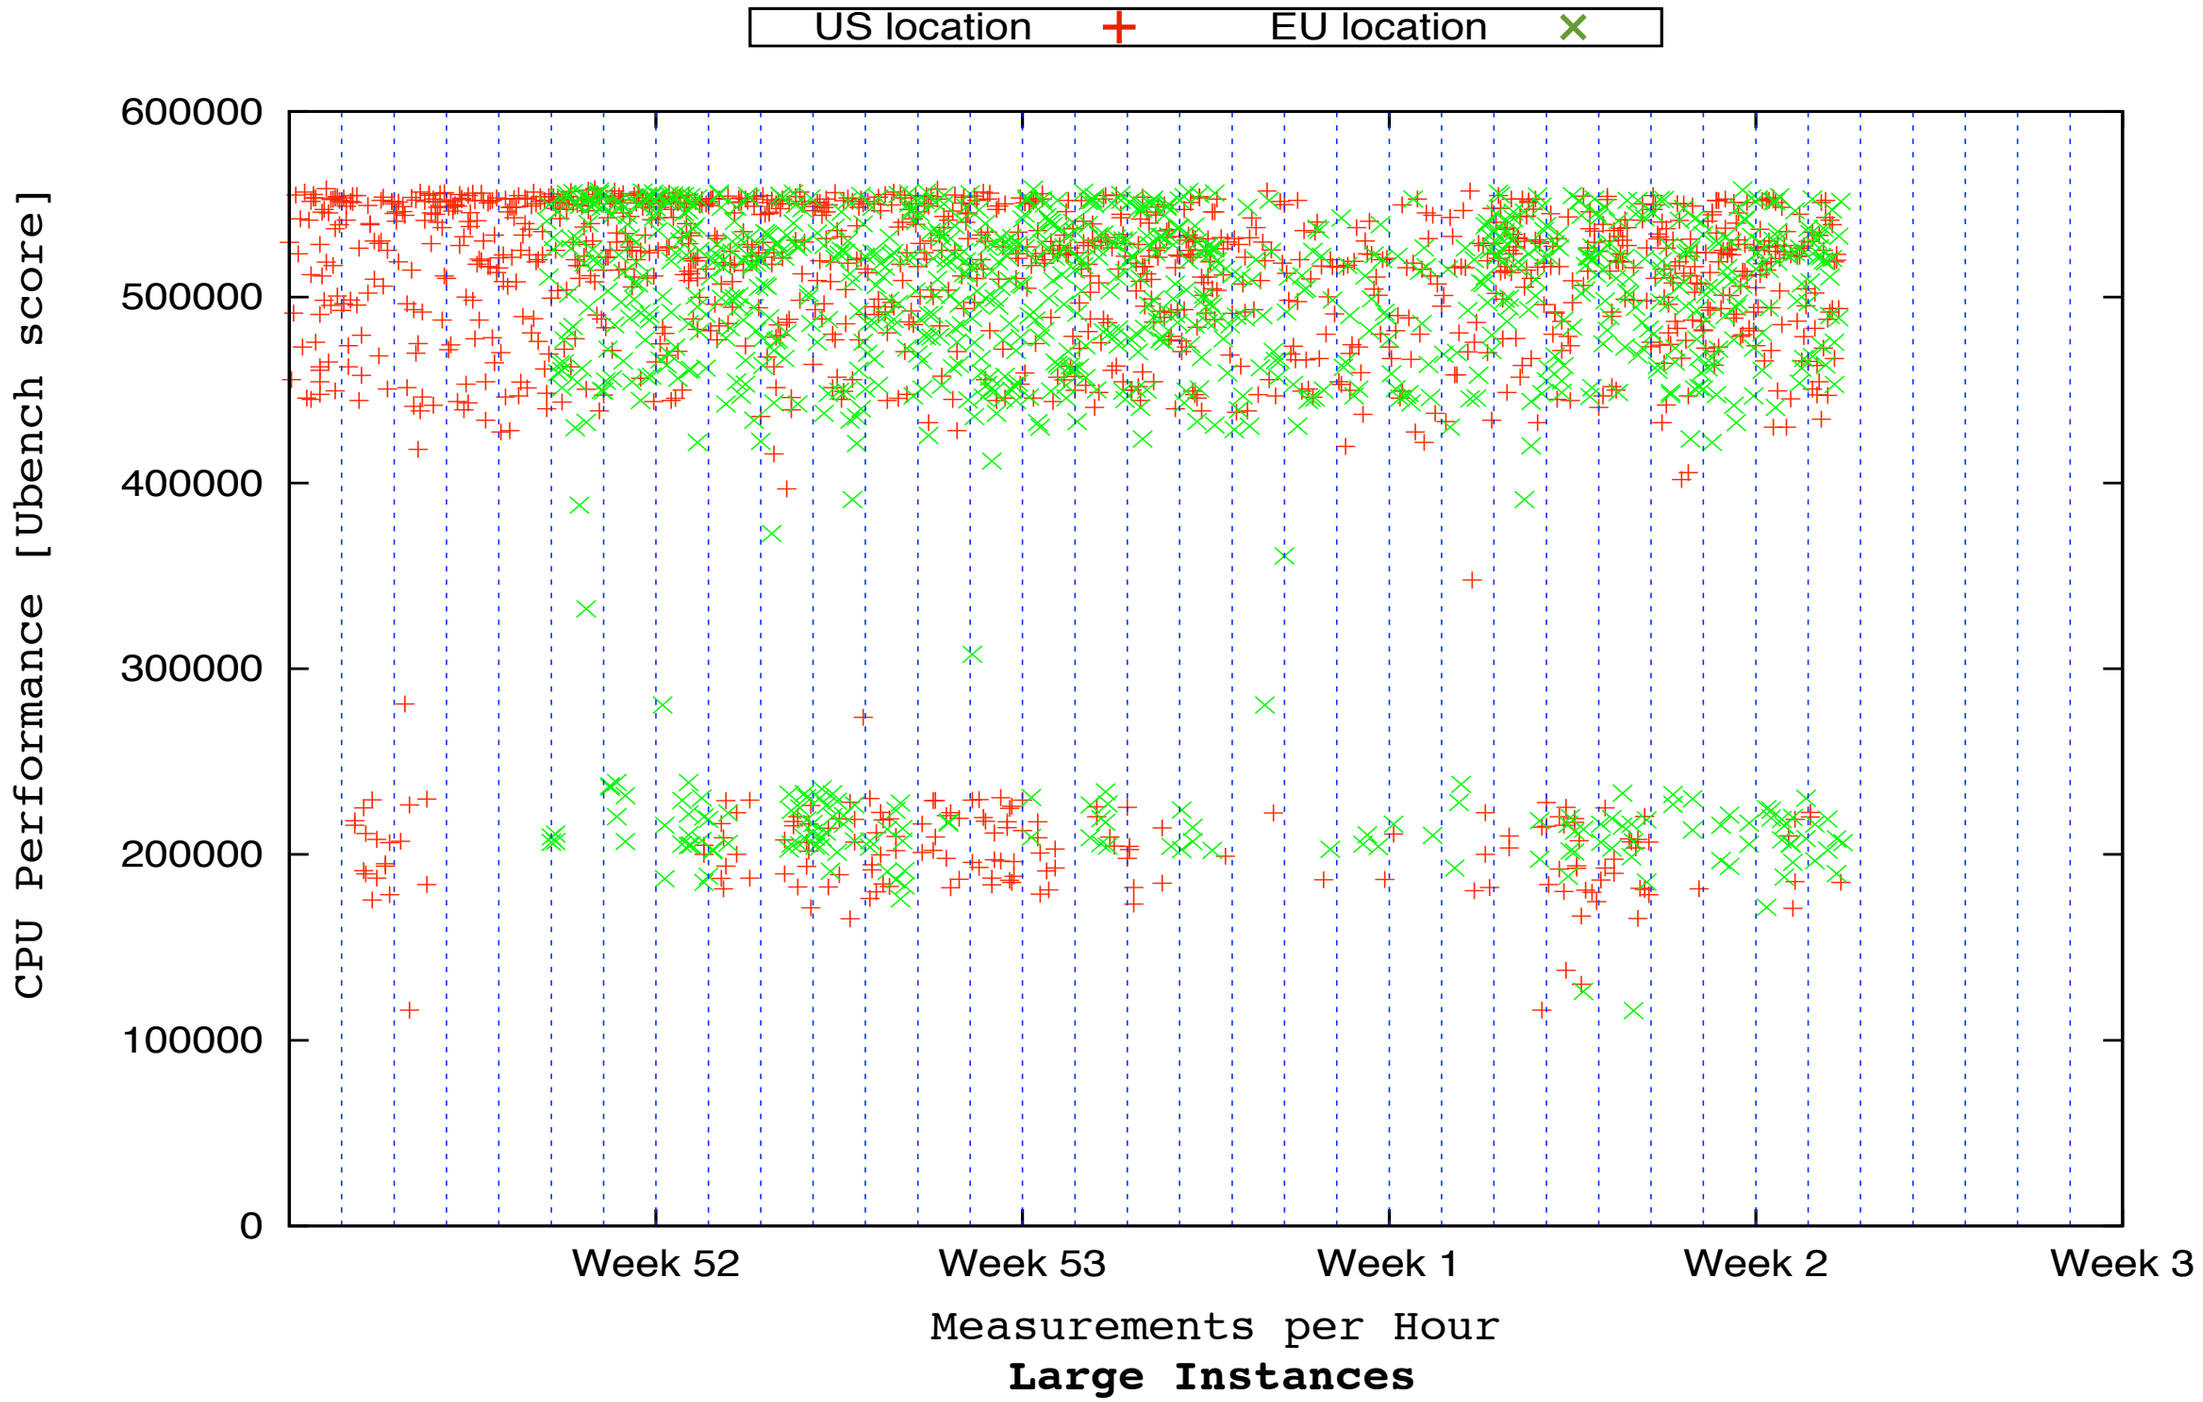

CPU Performance

SAARLAND UNIVERSITY
Information Systems Group

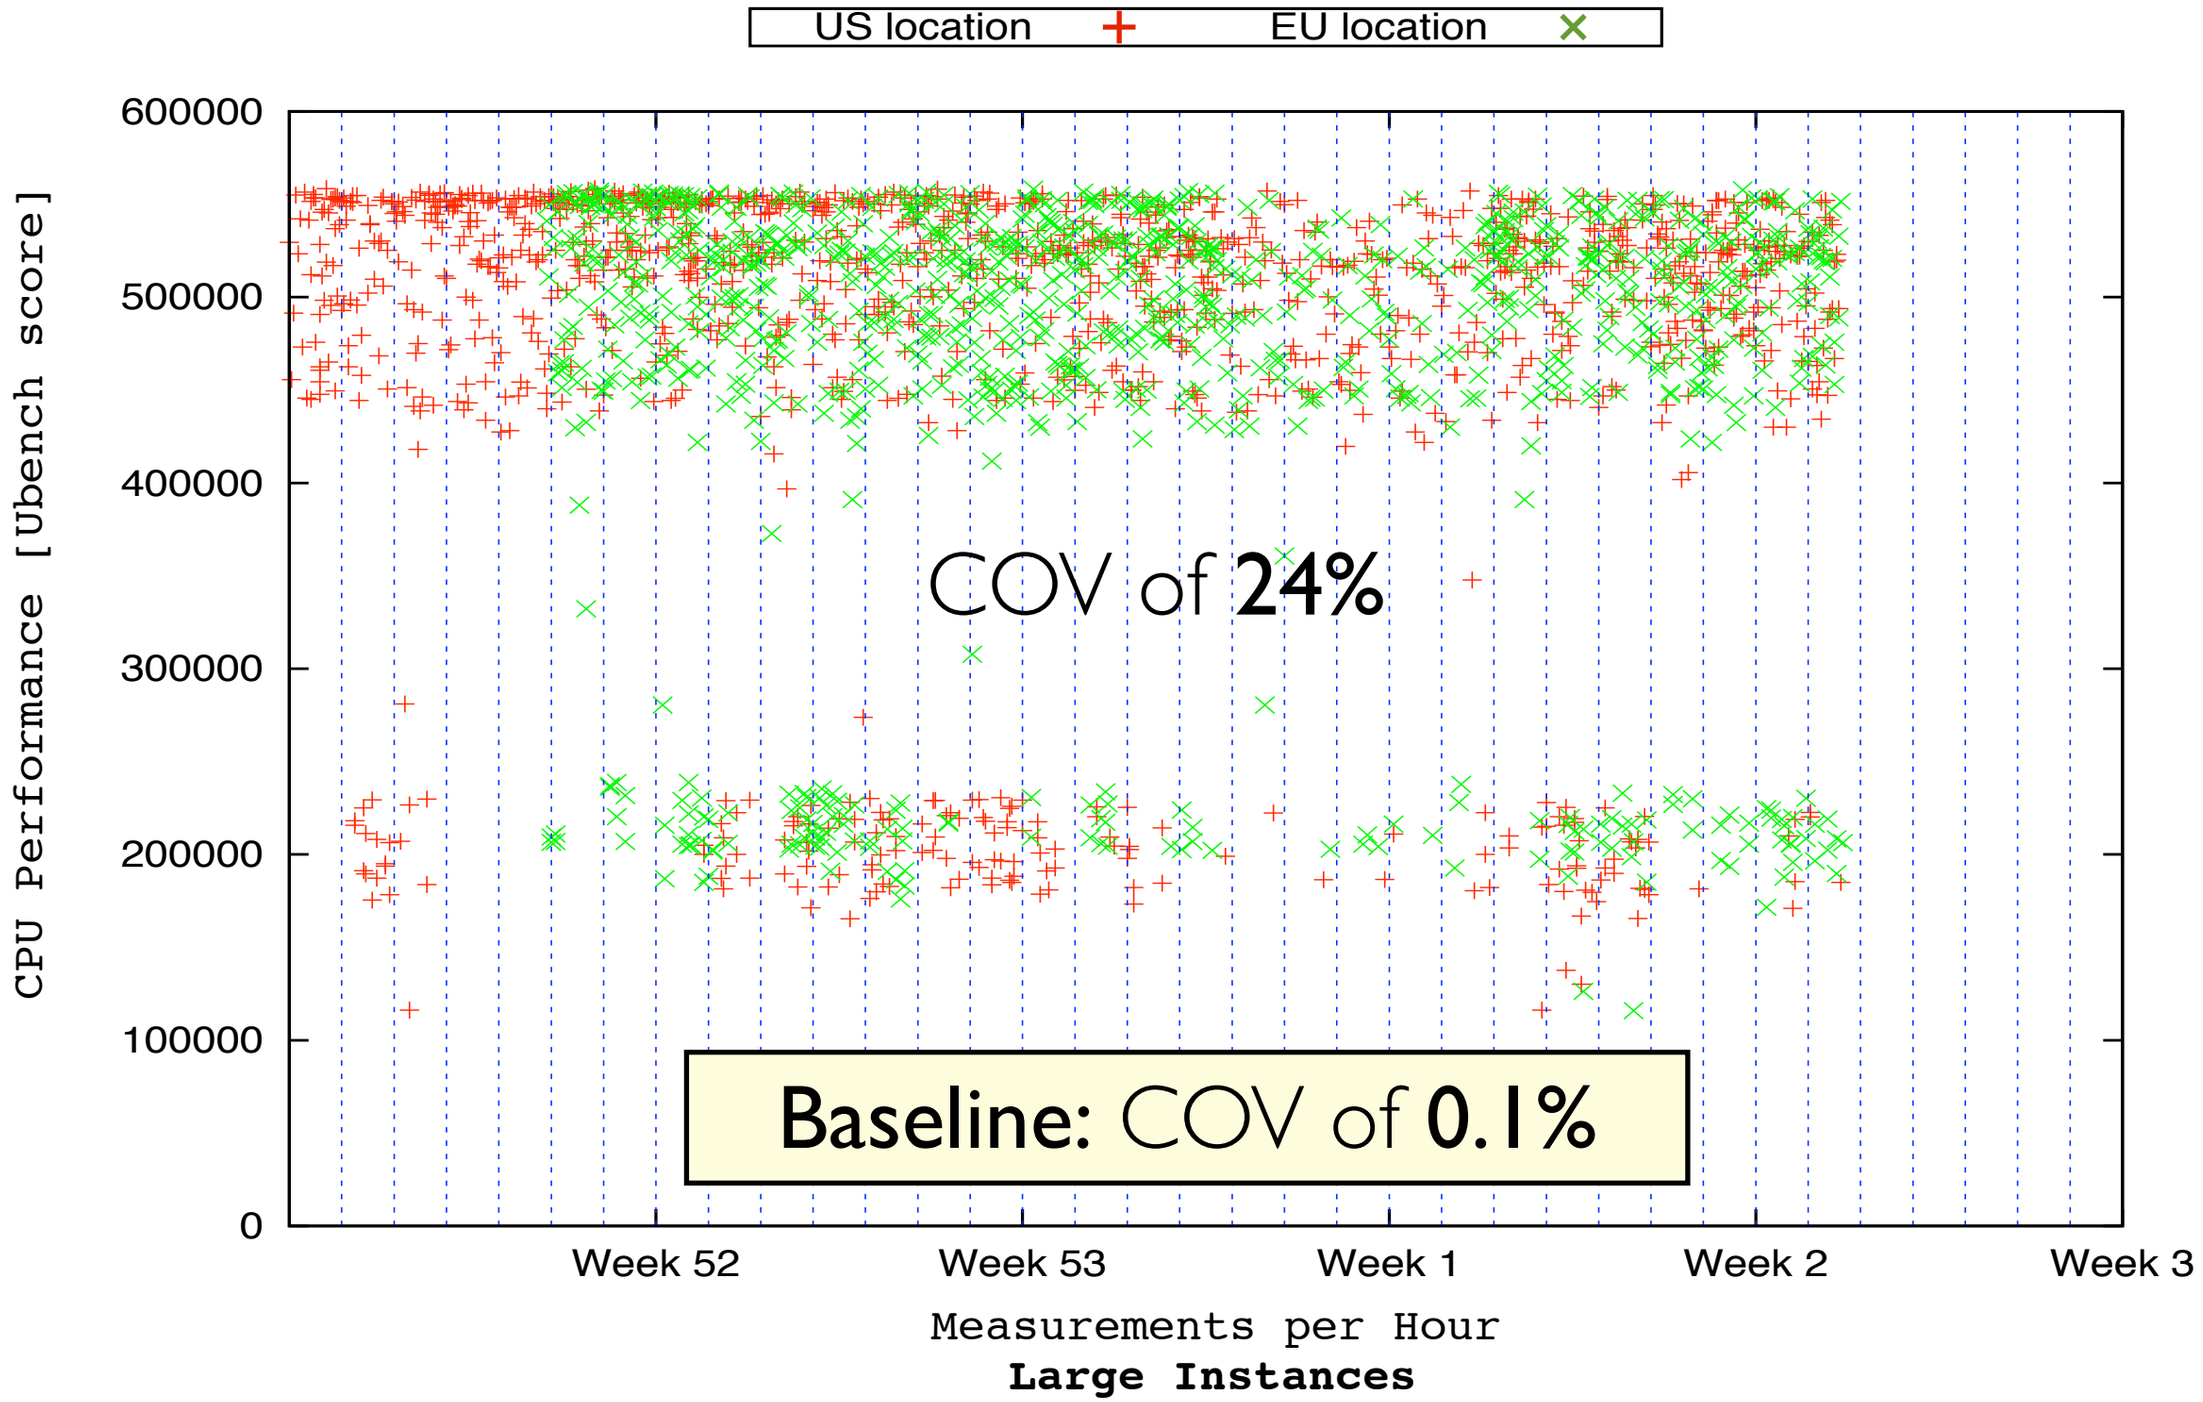

CPU Performance

SAARLAND UNIVERSITY
Information Systems Group

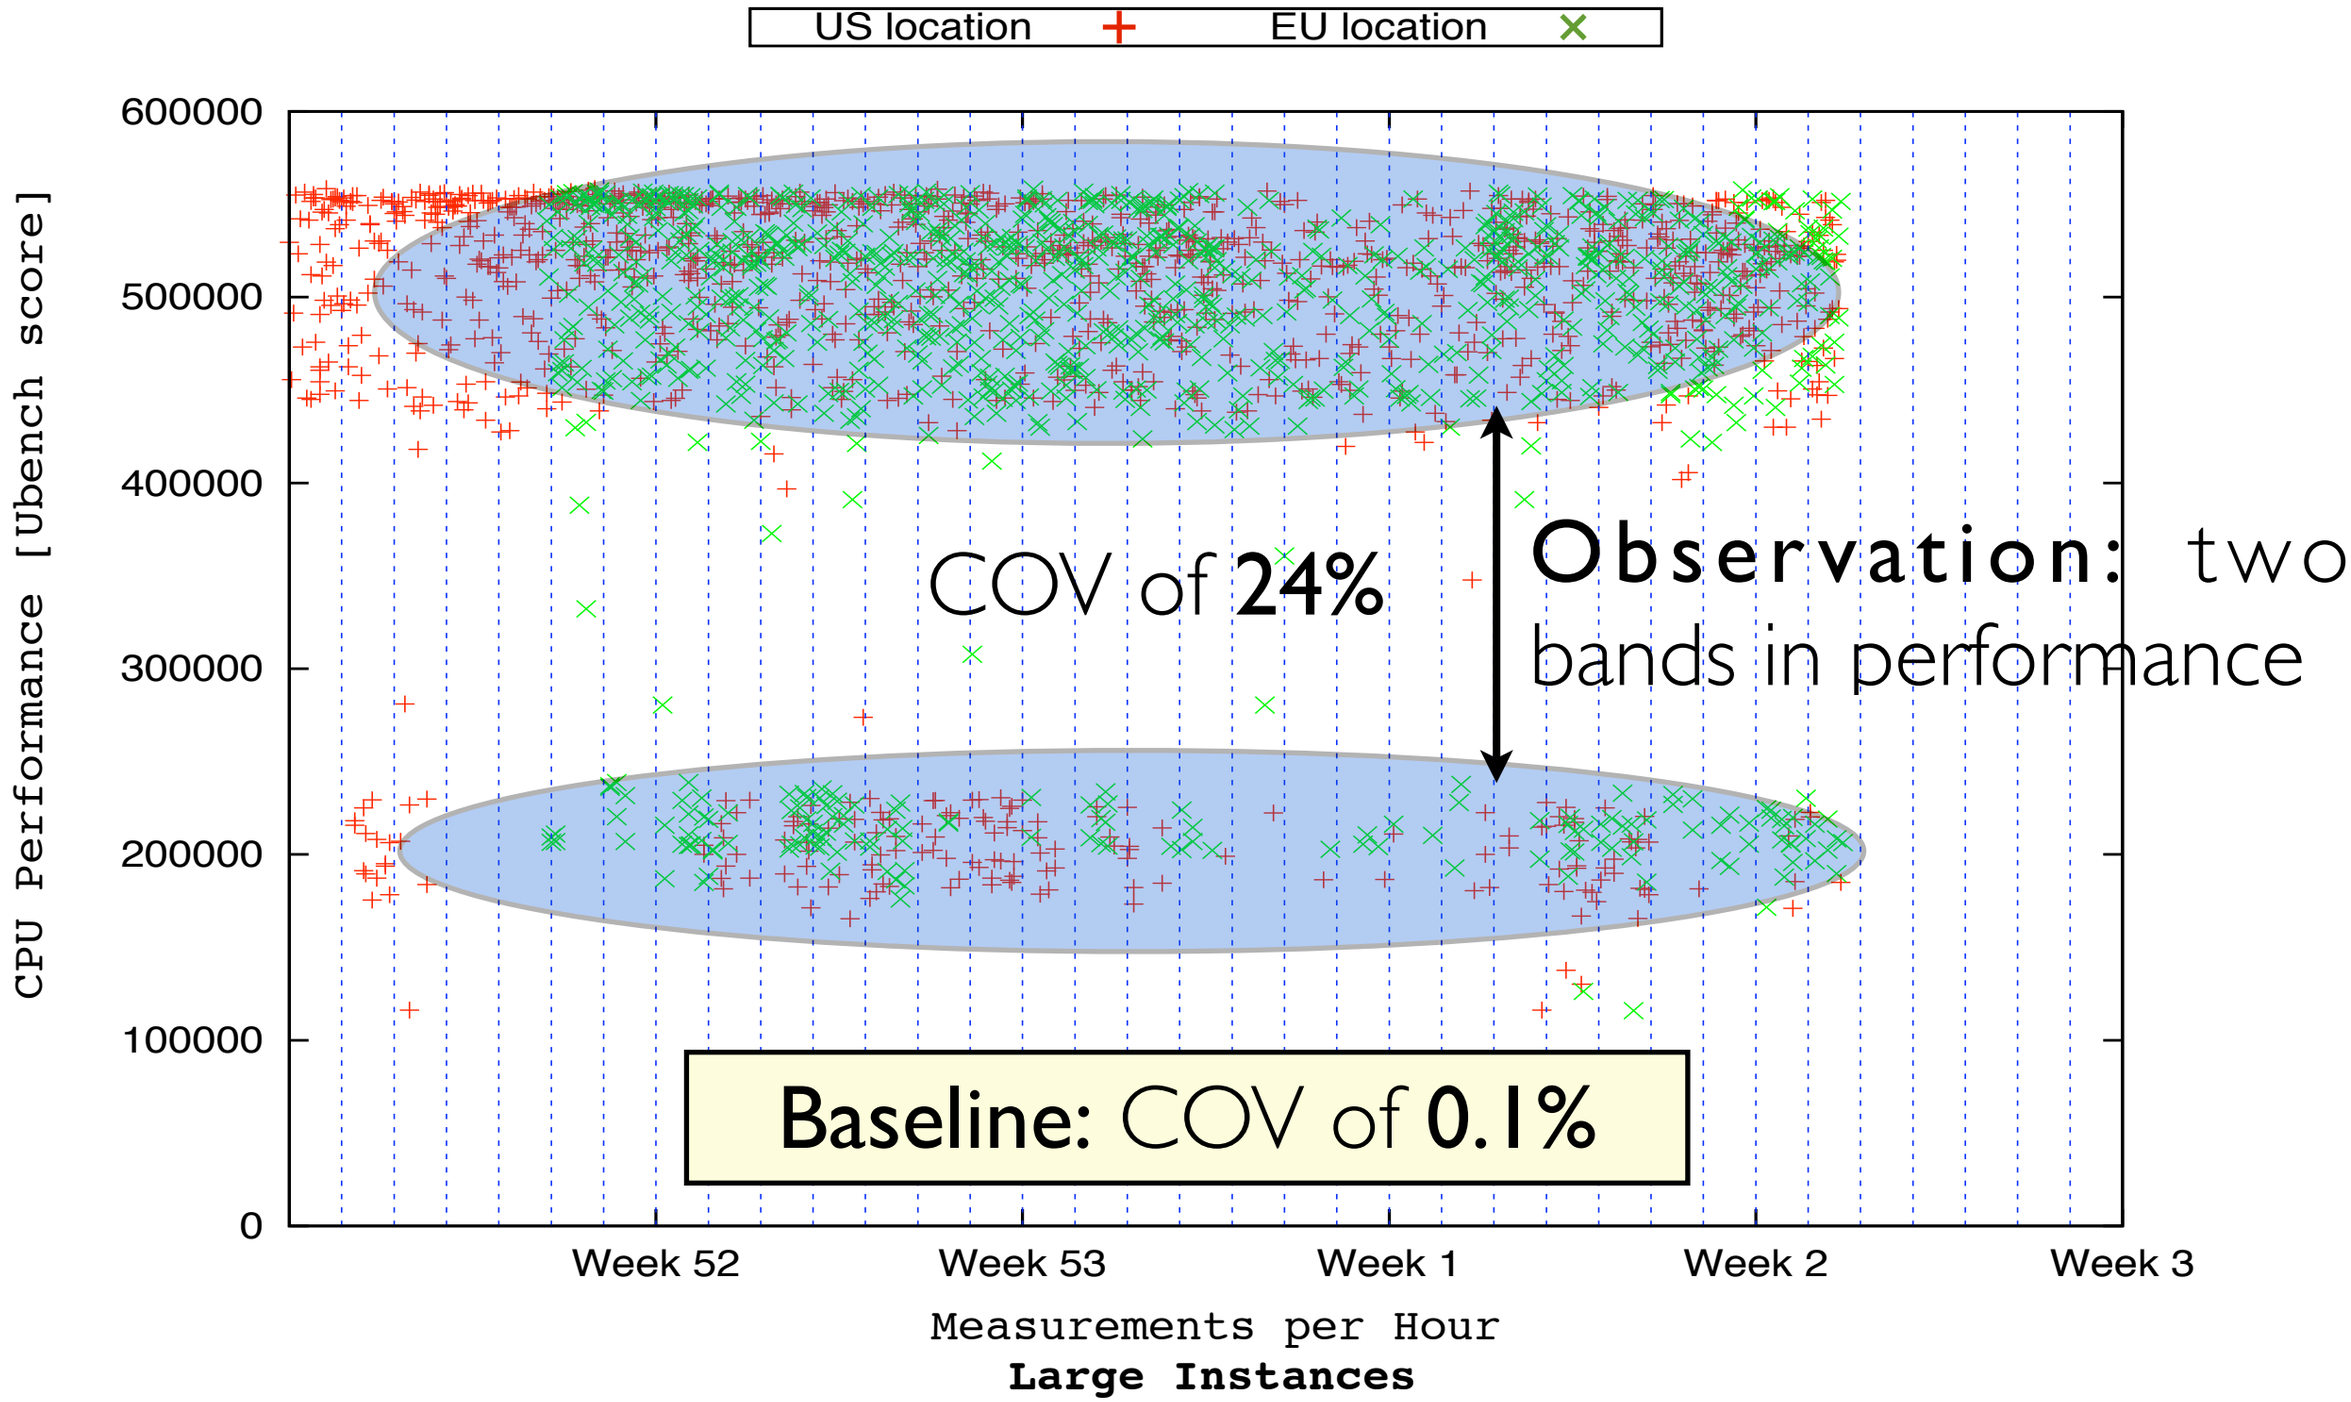

Memory Performance

SAARLAND UNIVERSITY
Information Systems Group

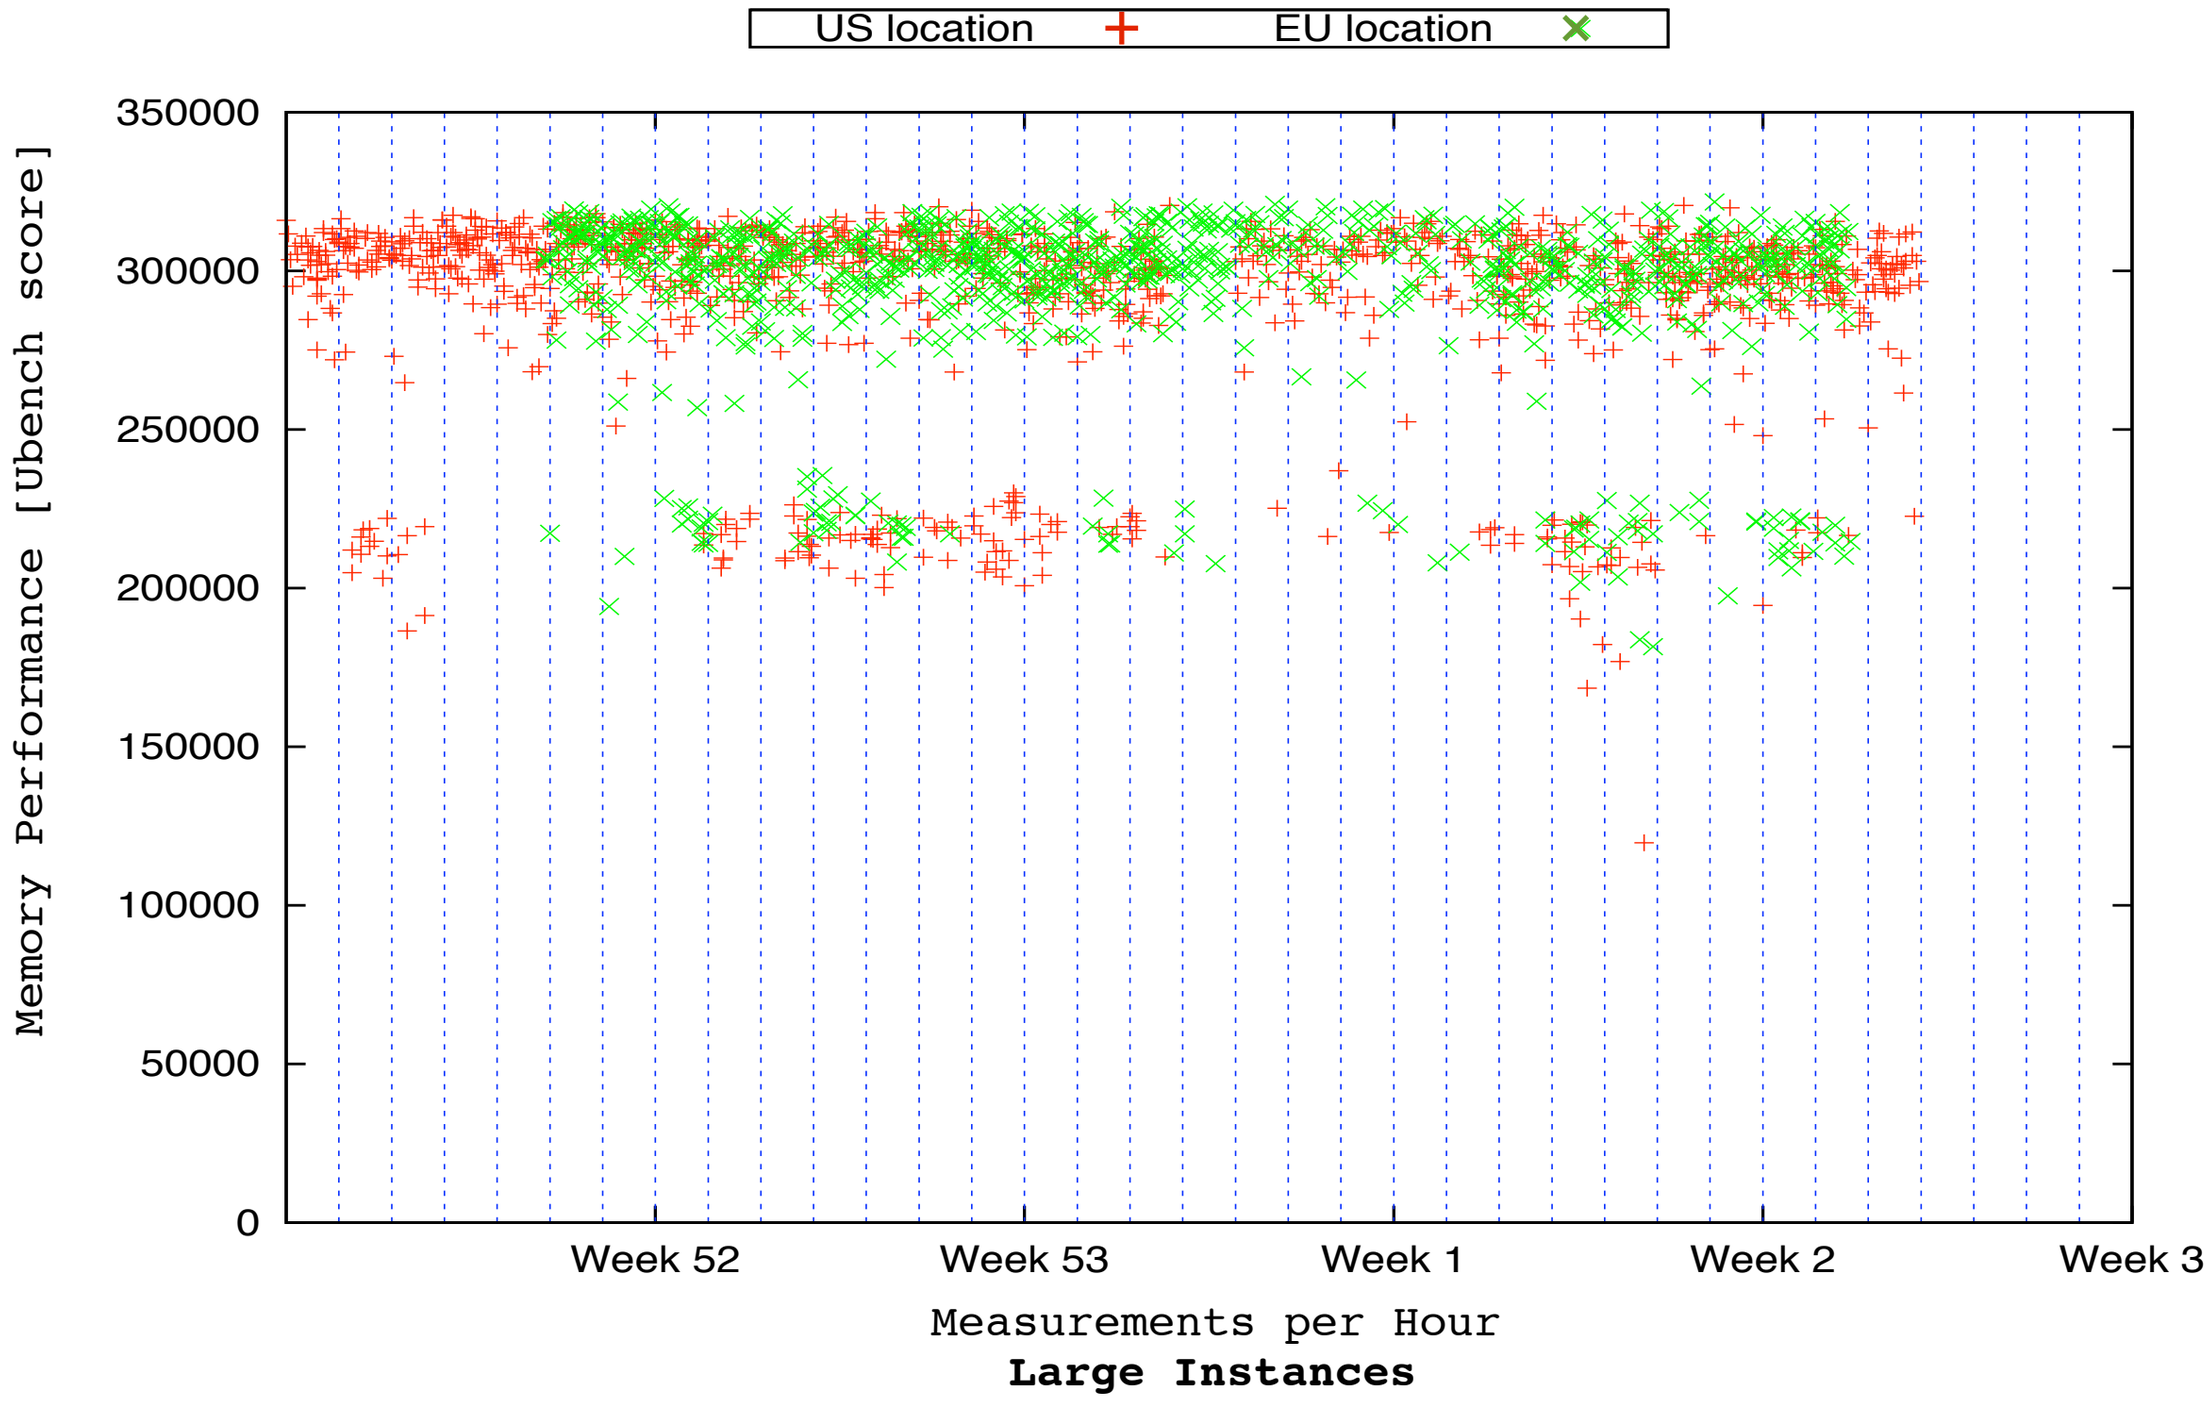

Memory Performance

SAARLAND UNIVERSITY
Information Systems Group

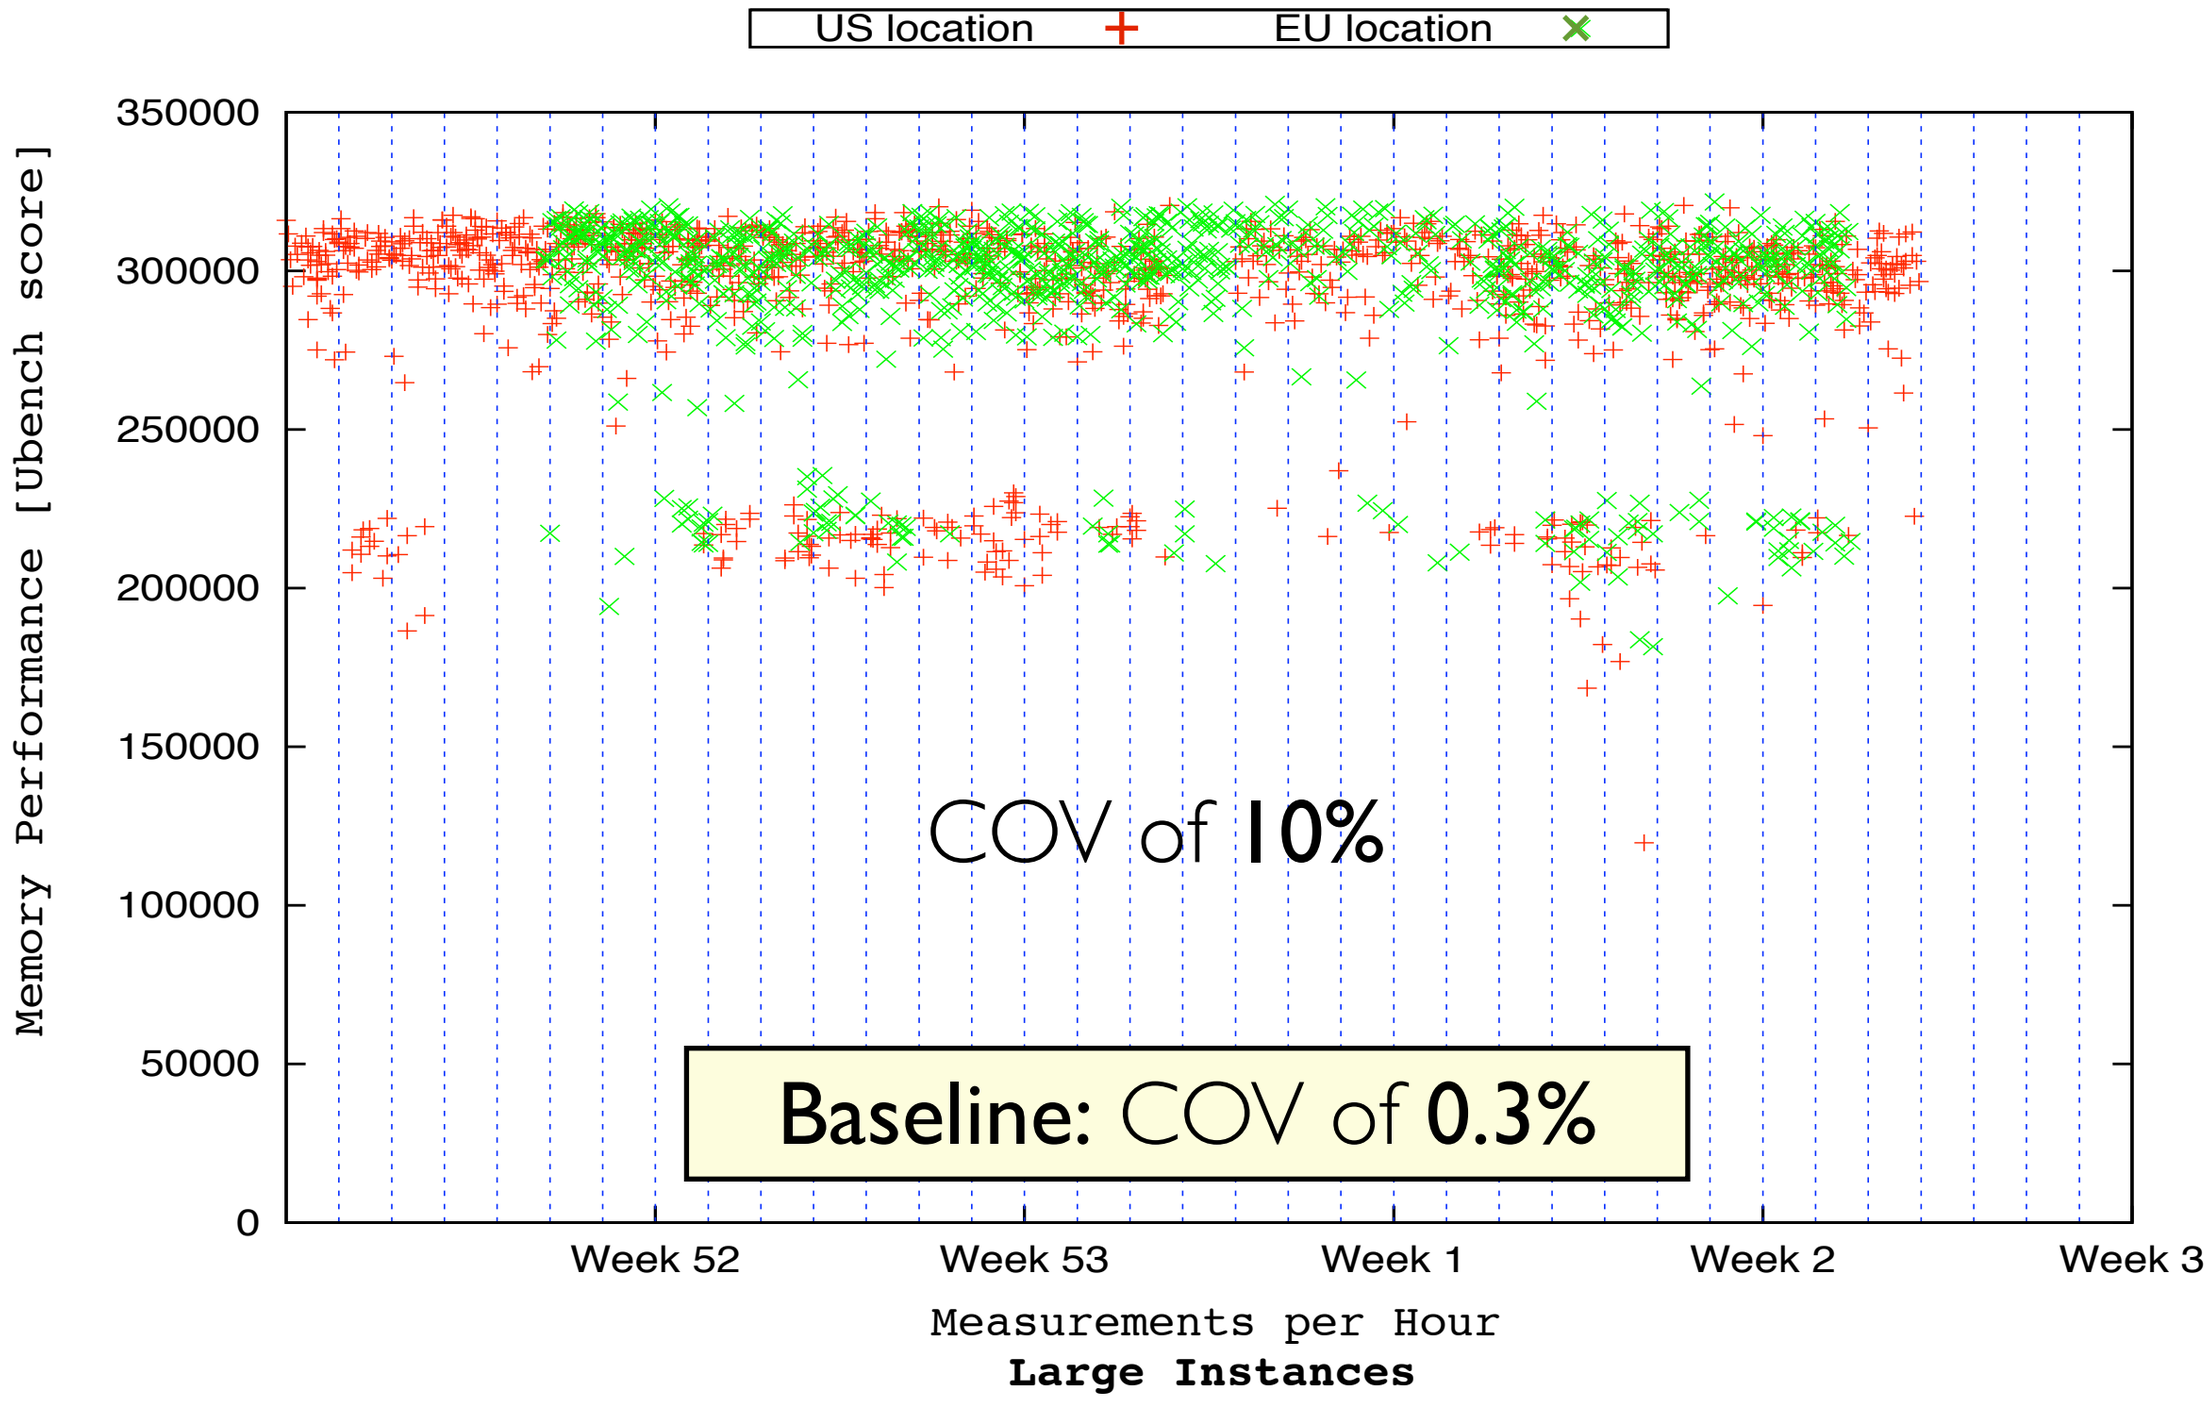

Memory Performance

SAARLAND UNIVERSITY
Information Systems Group

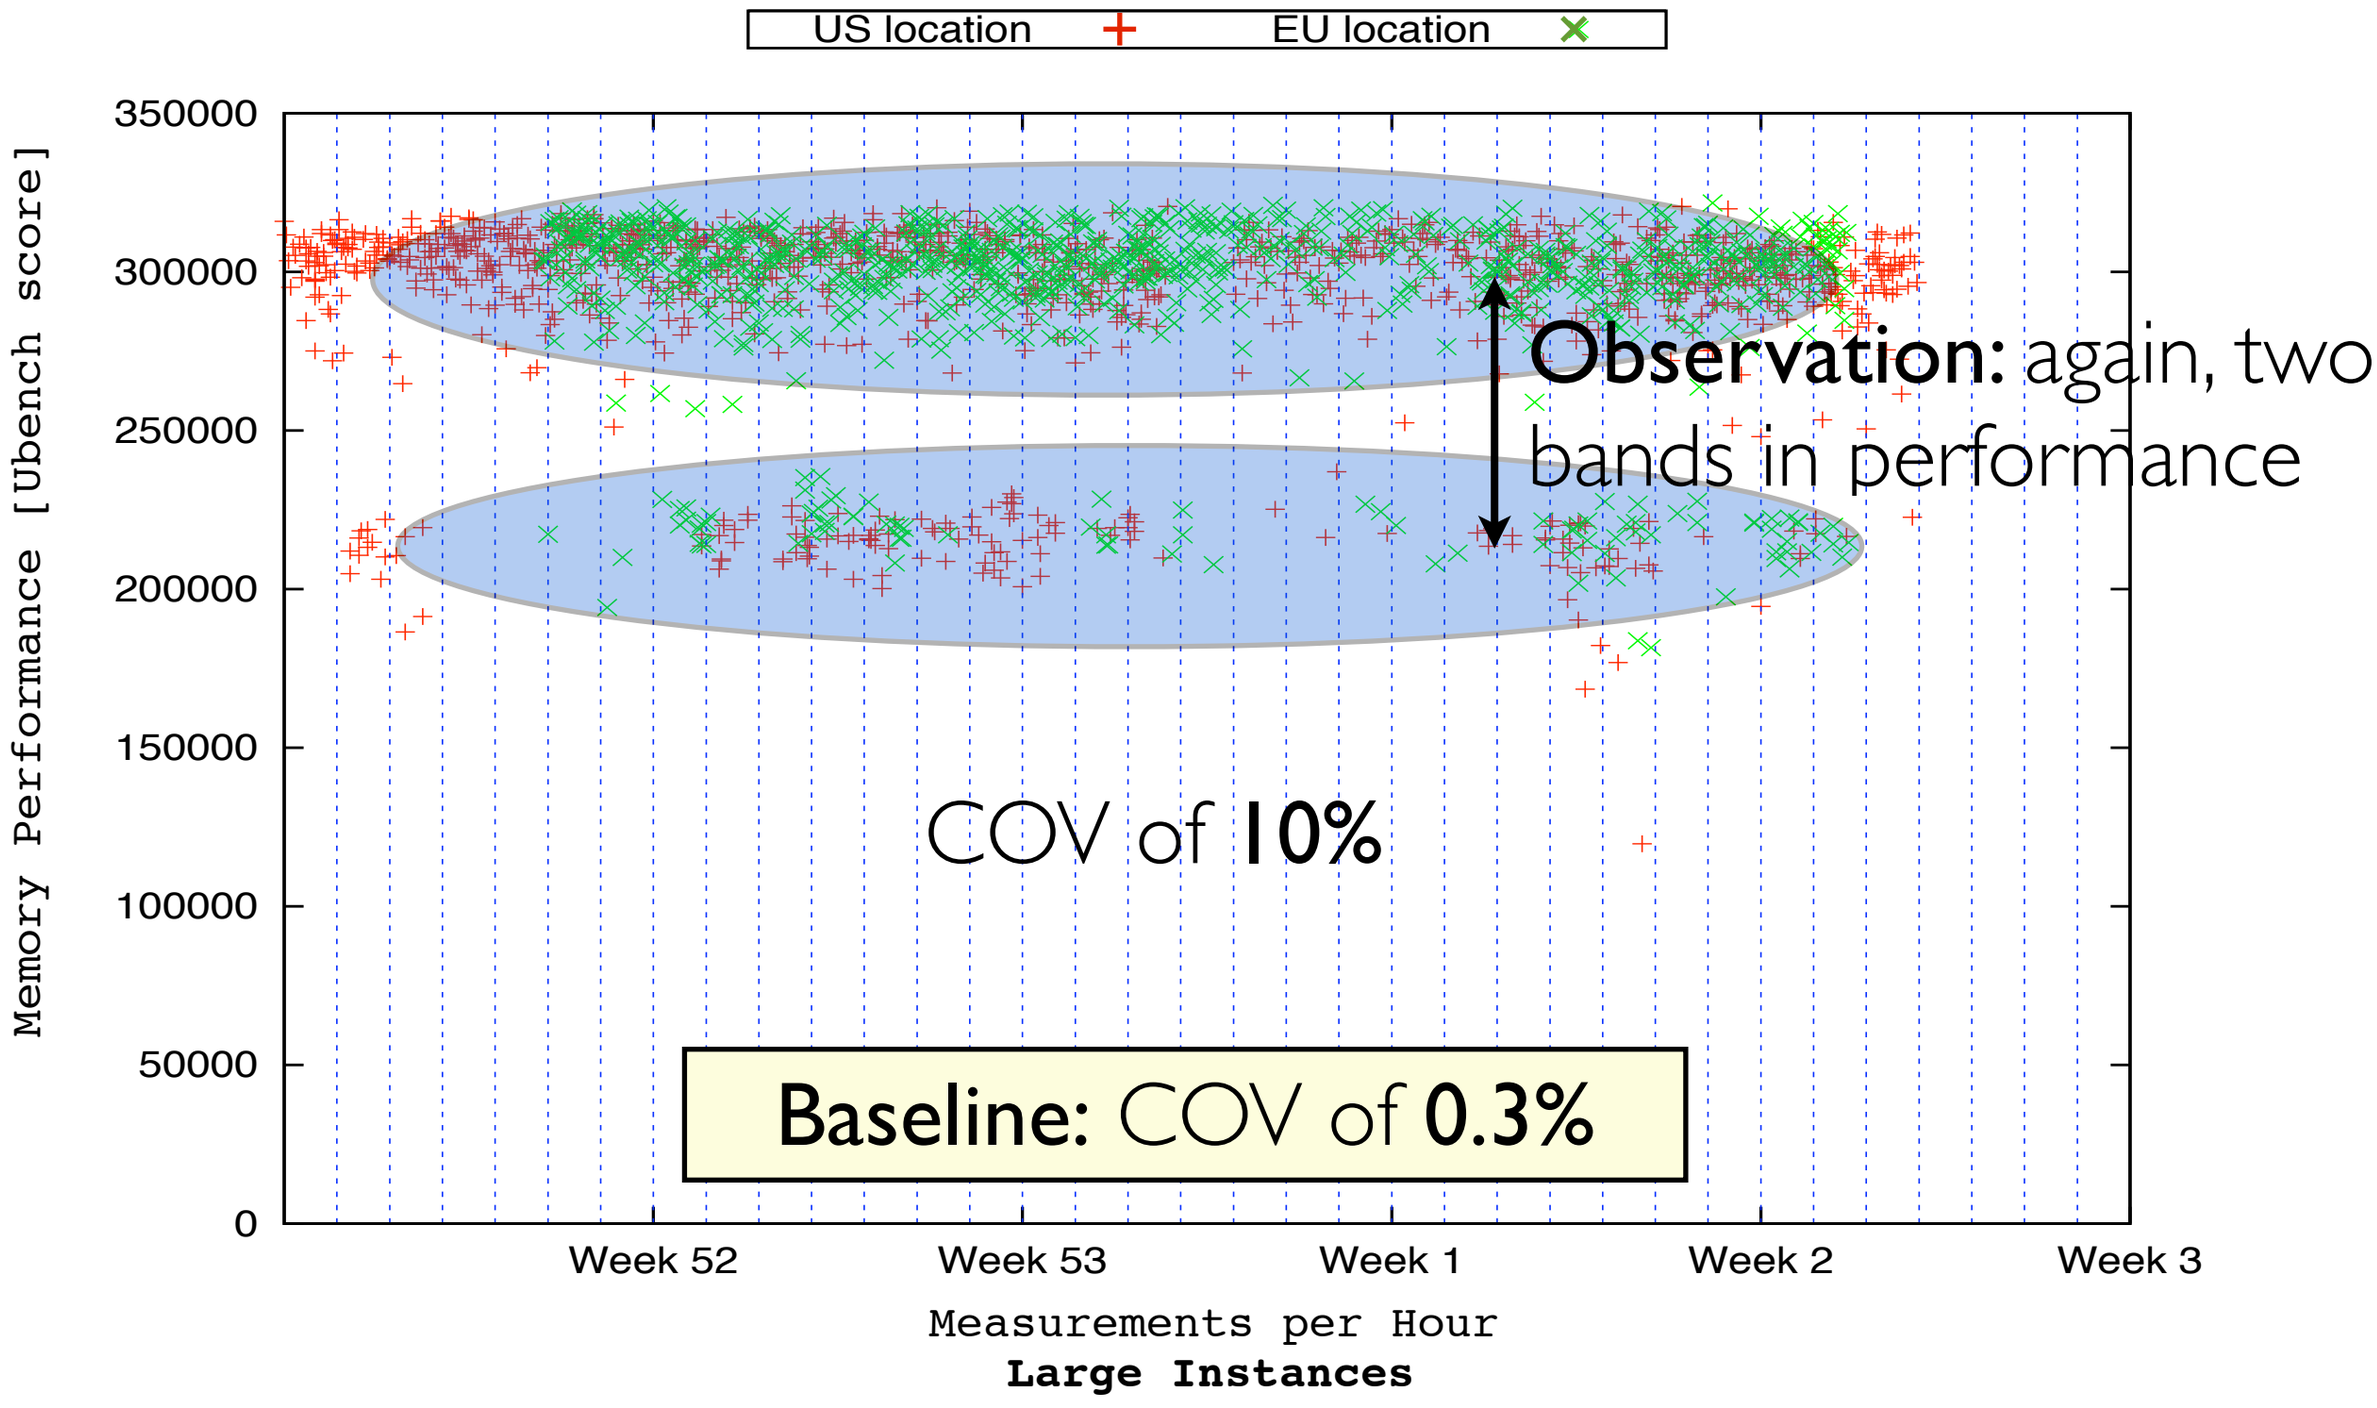

Random I/O Performance

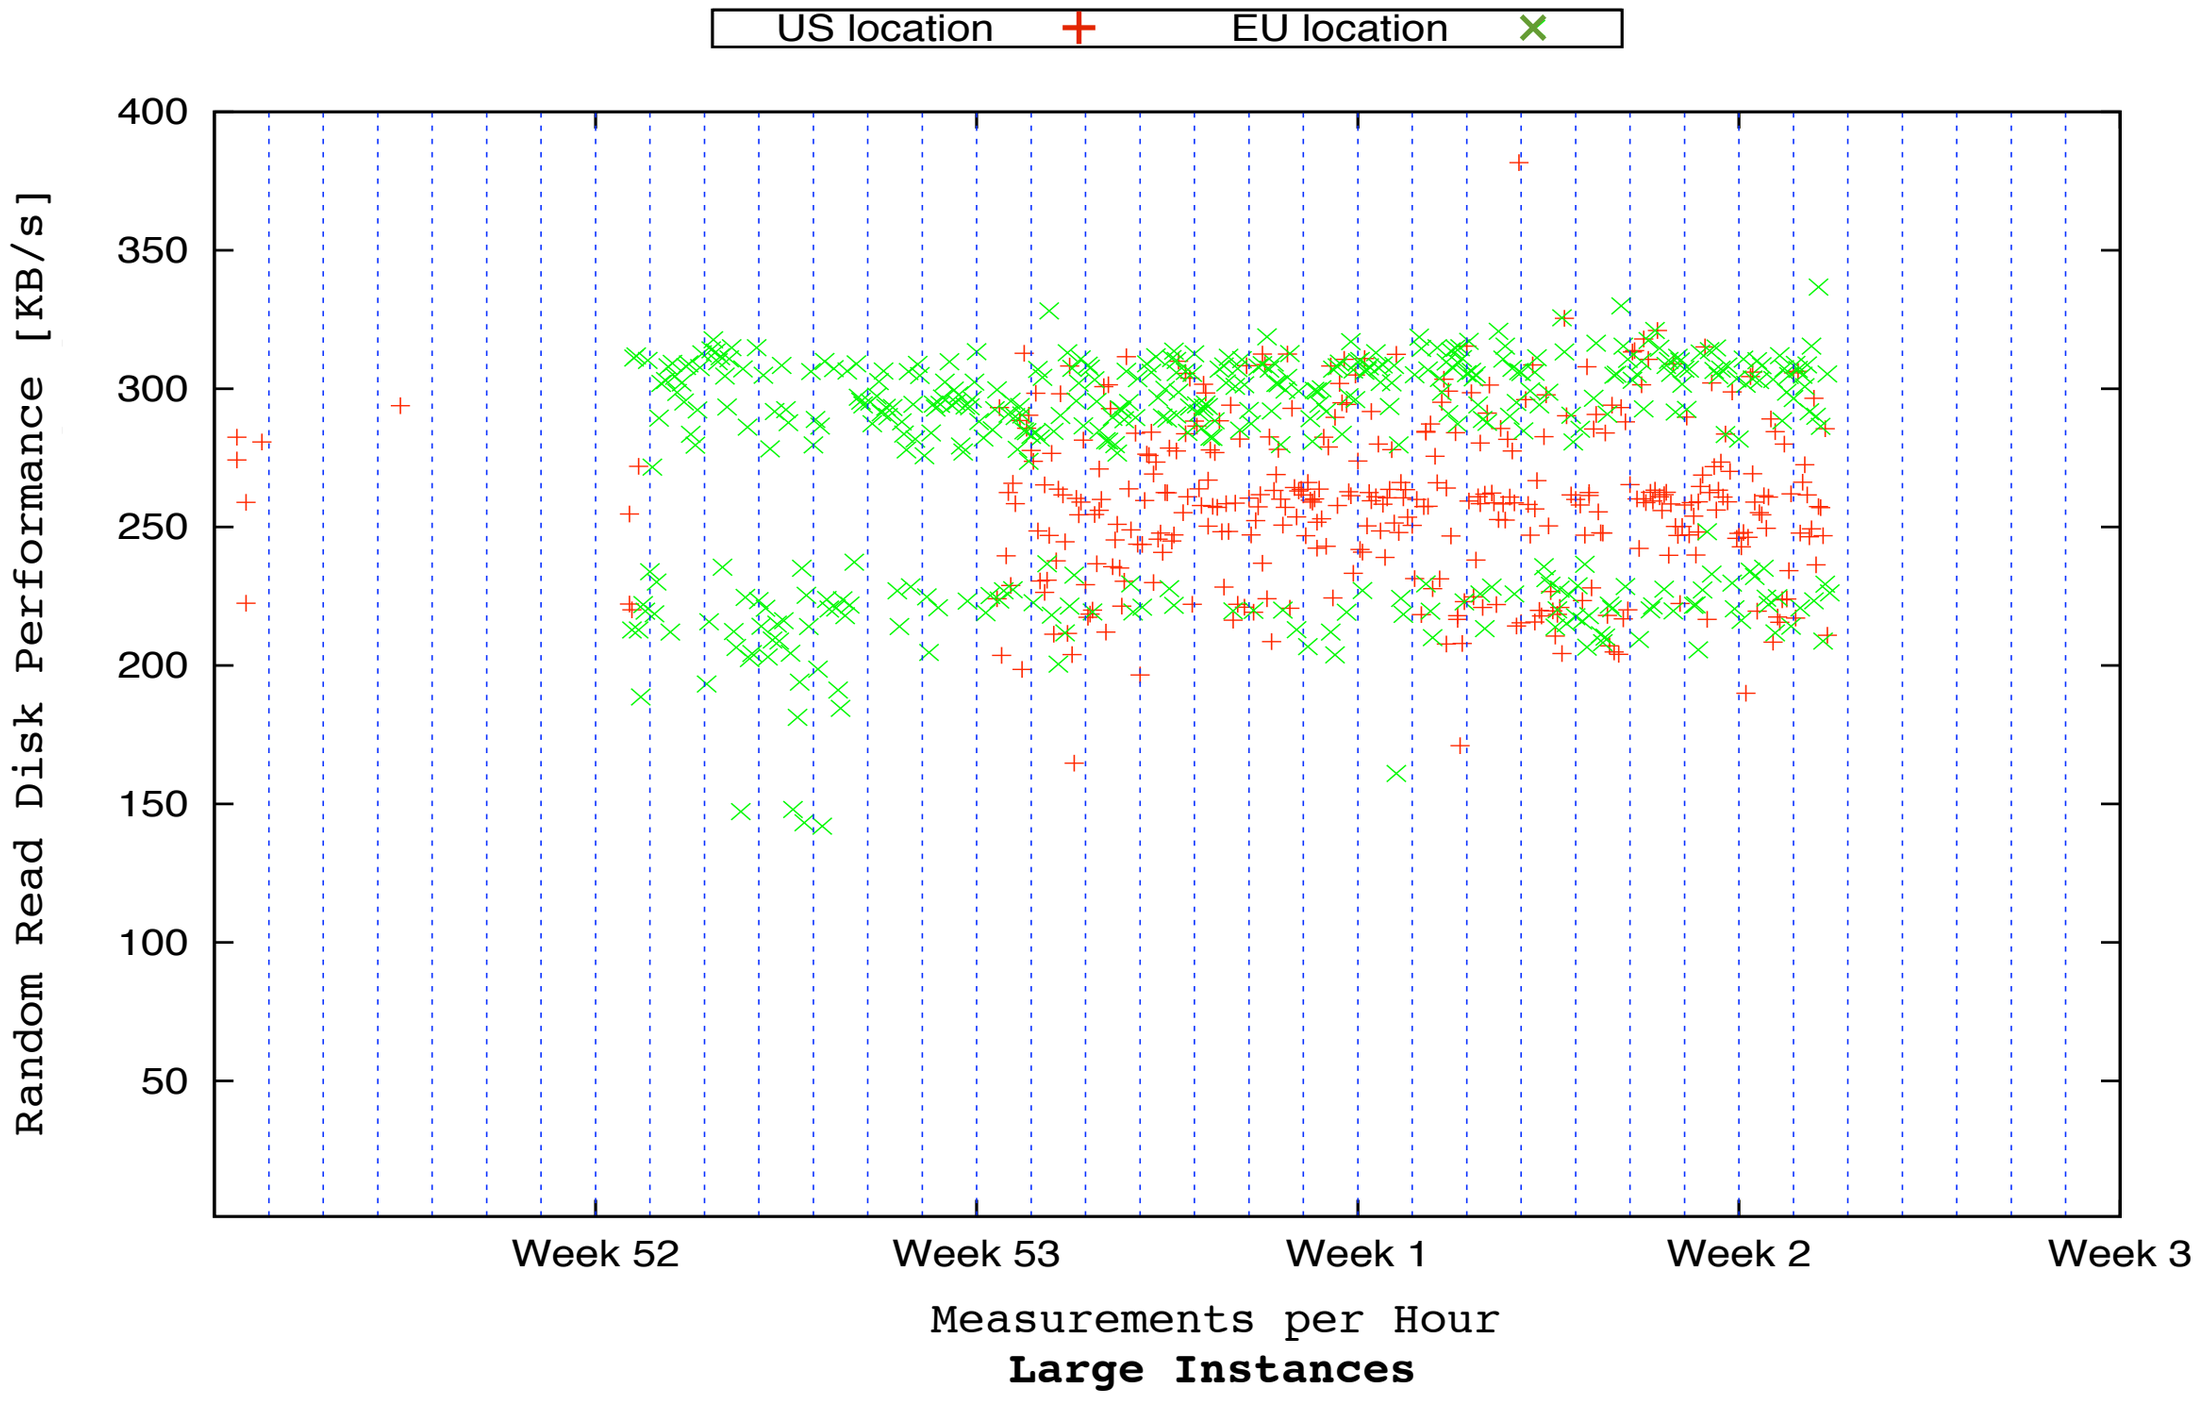

Random I/O Performance

SAARLAND UNIVERSITY
Information Systems Group

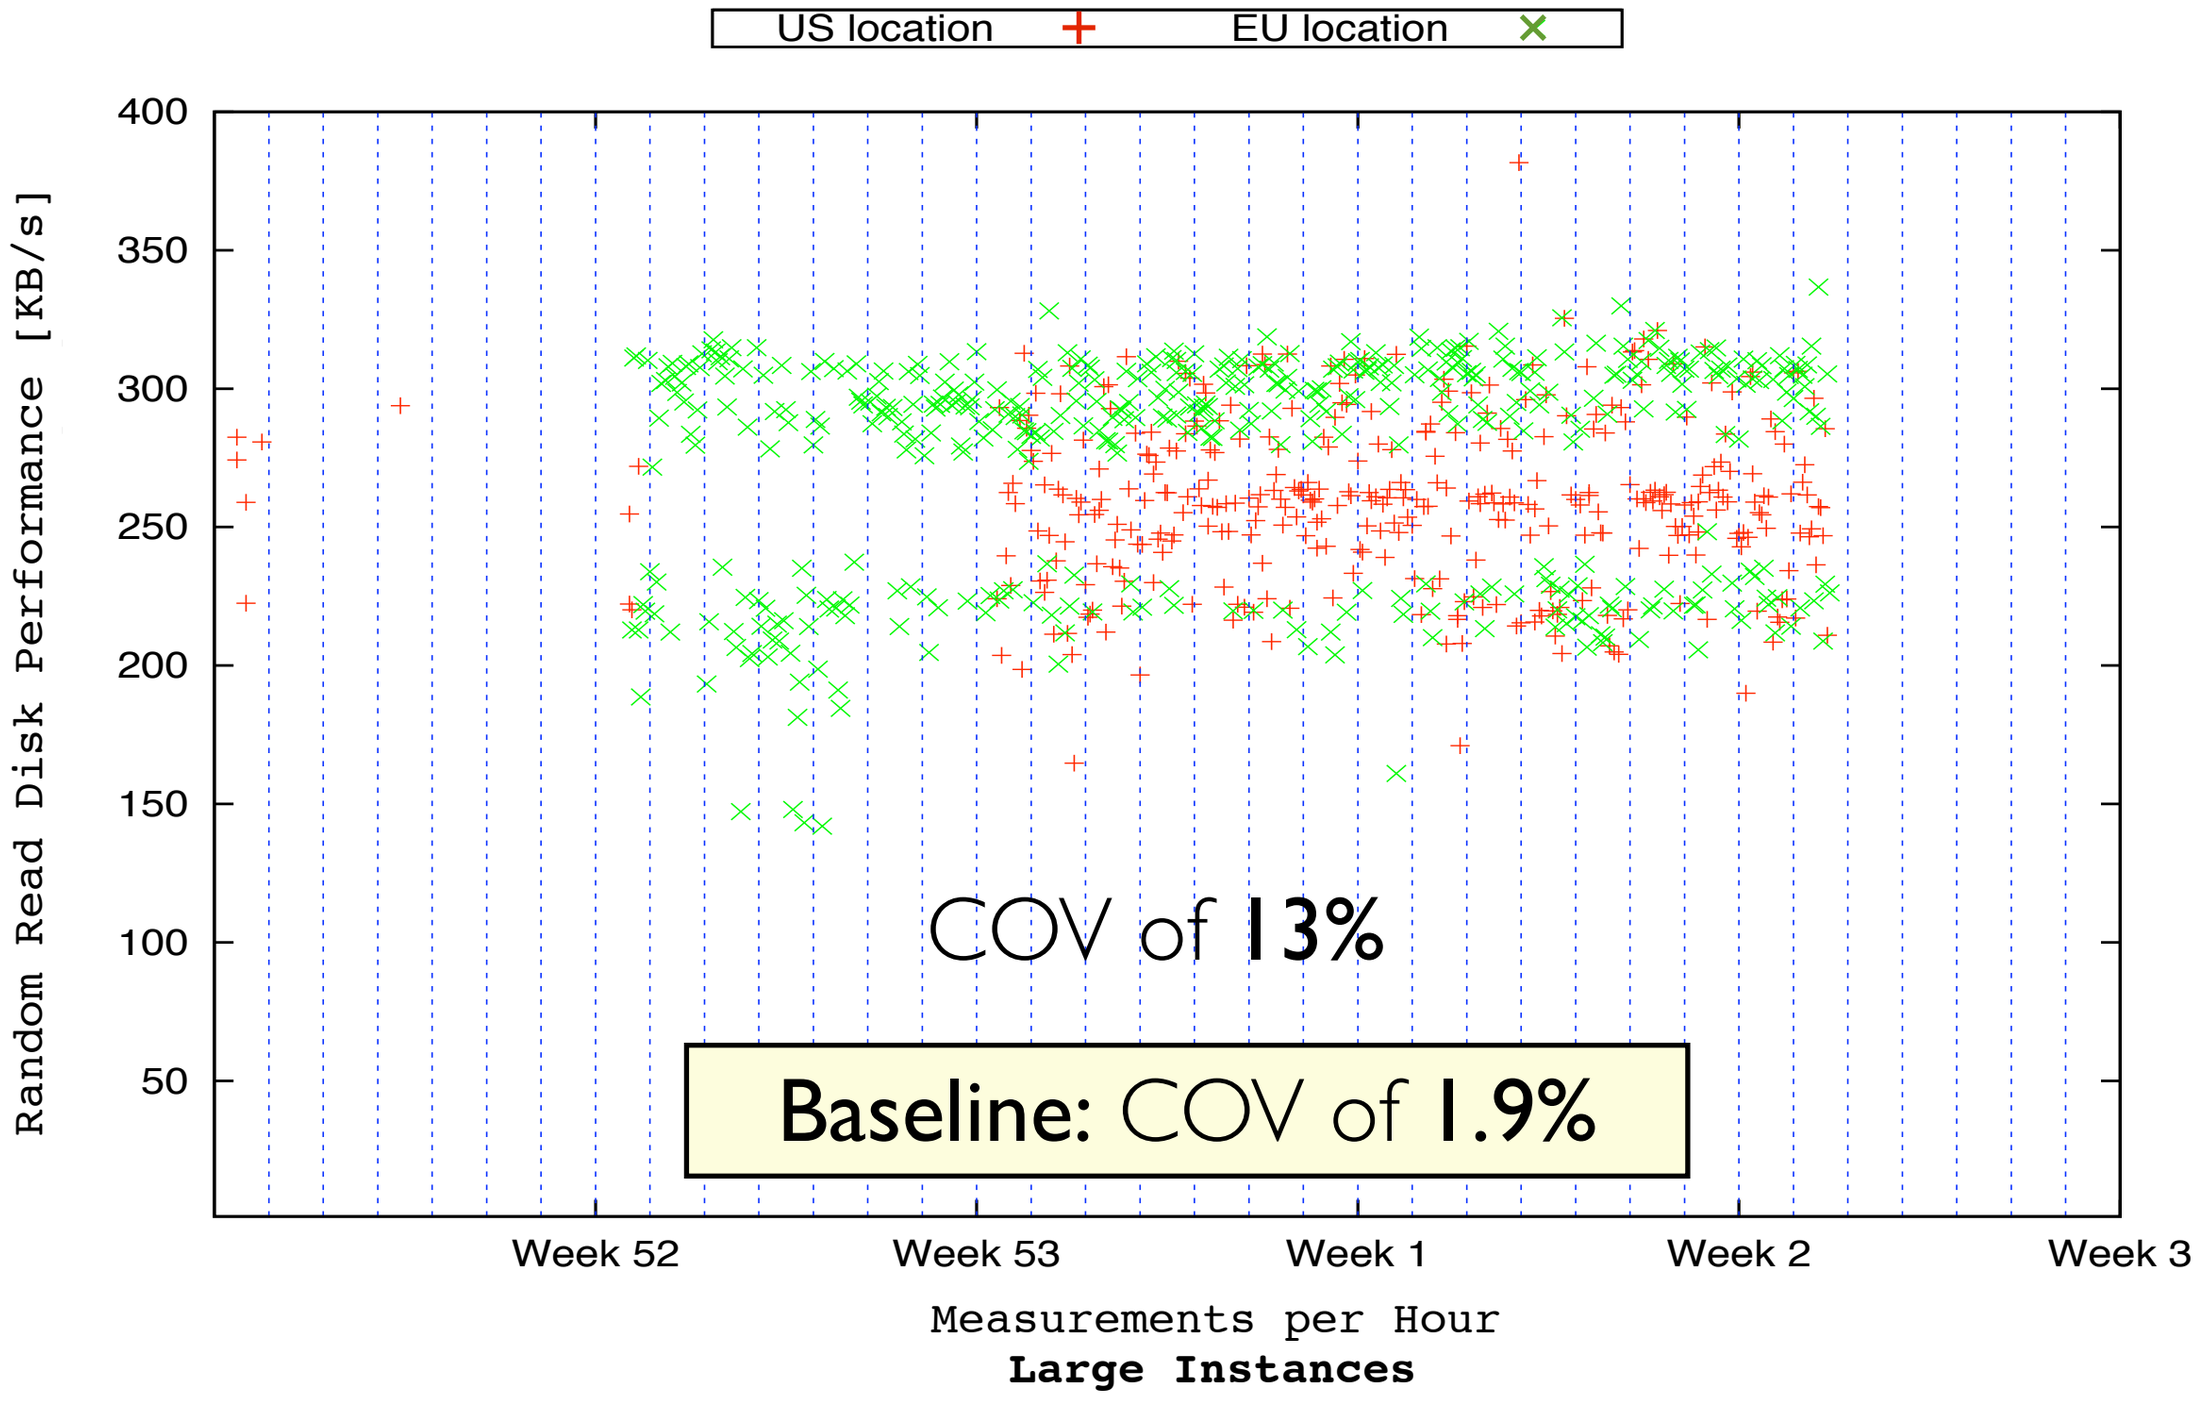

Random I/O Performance

SAARLAND UNIVERSITY
Information Systems Group

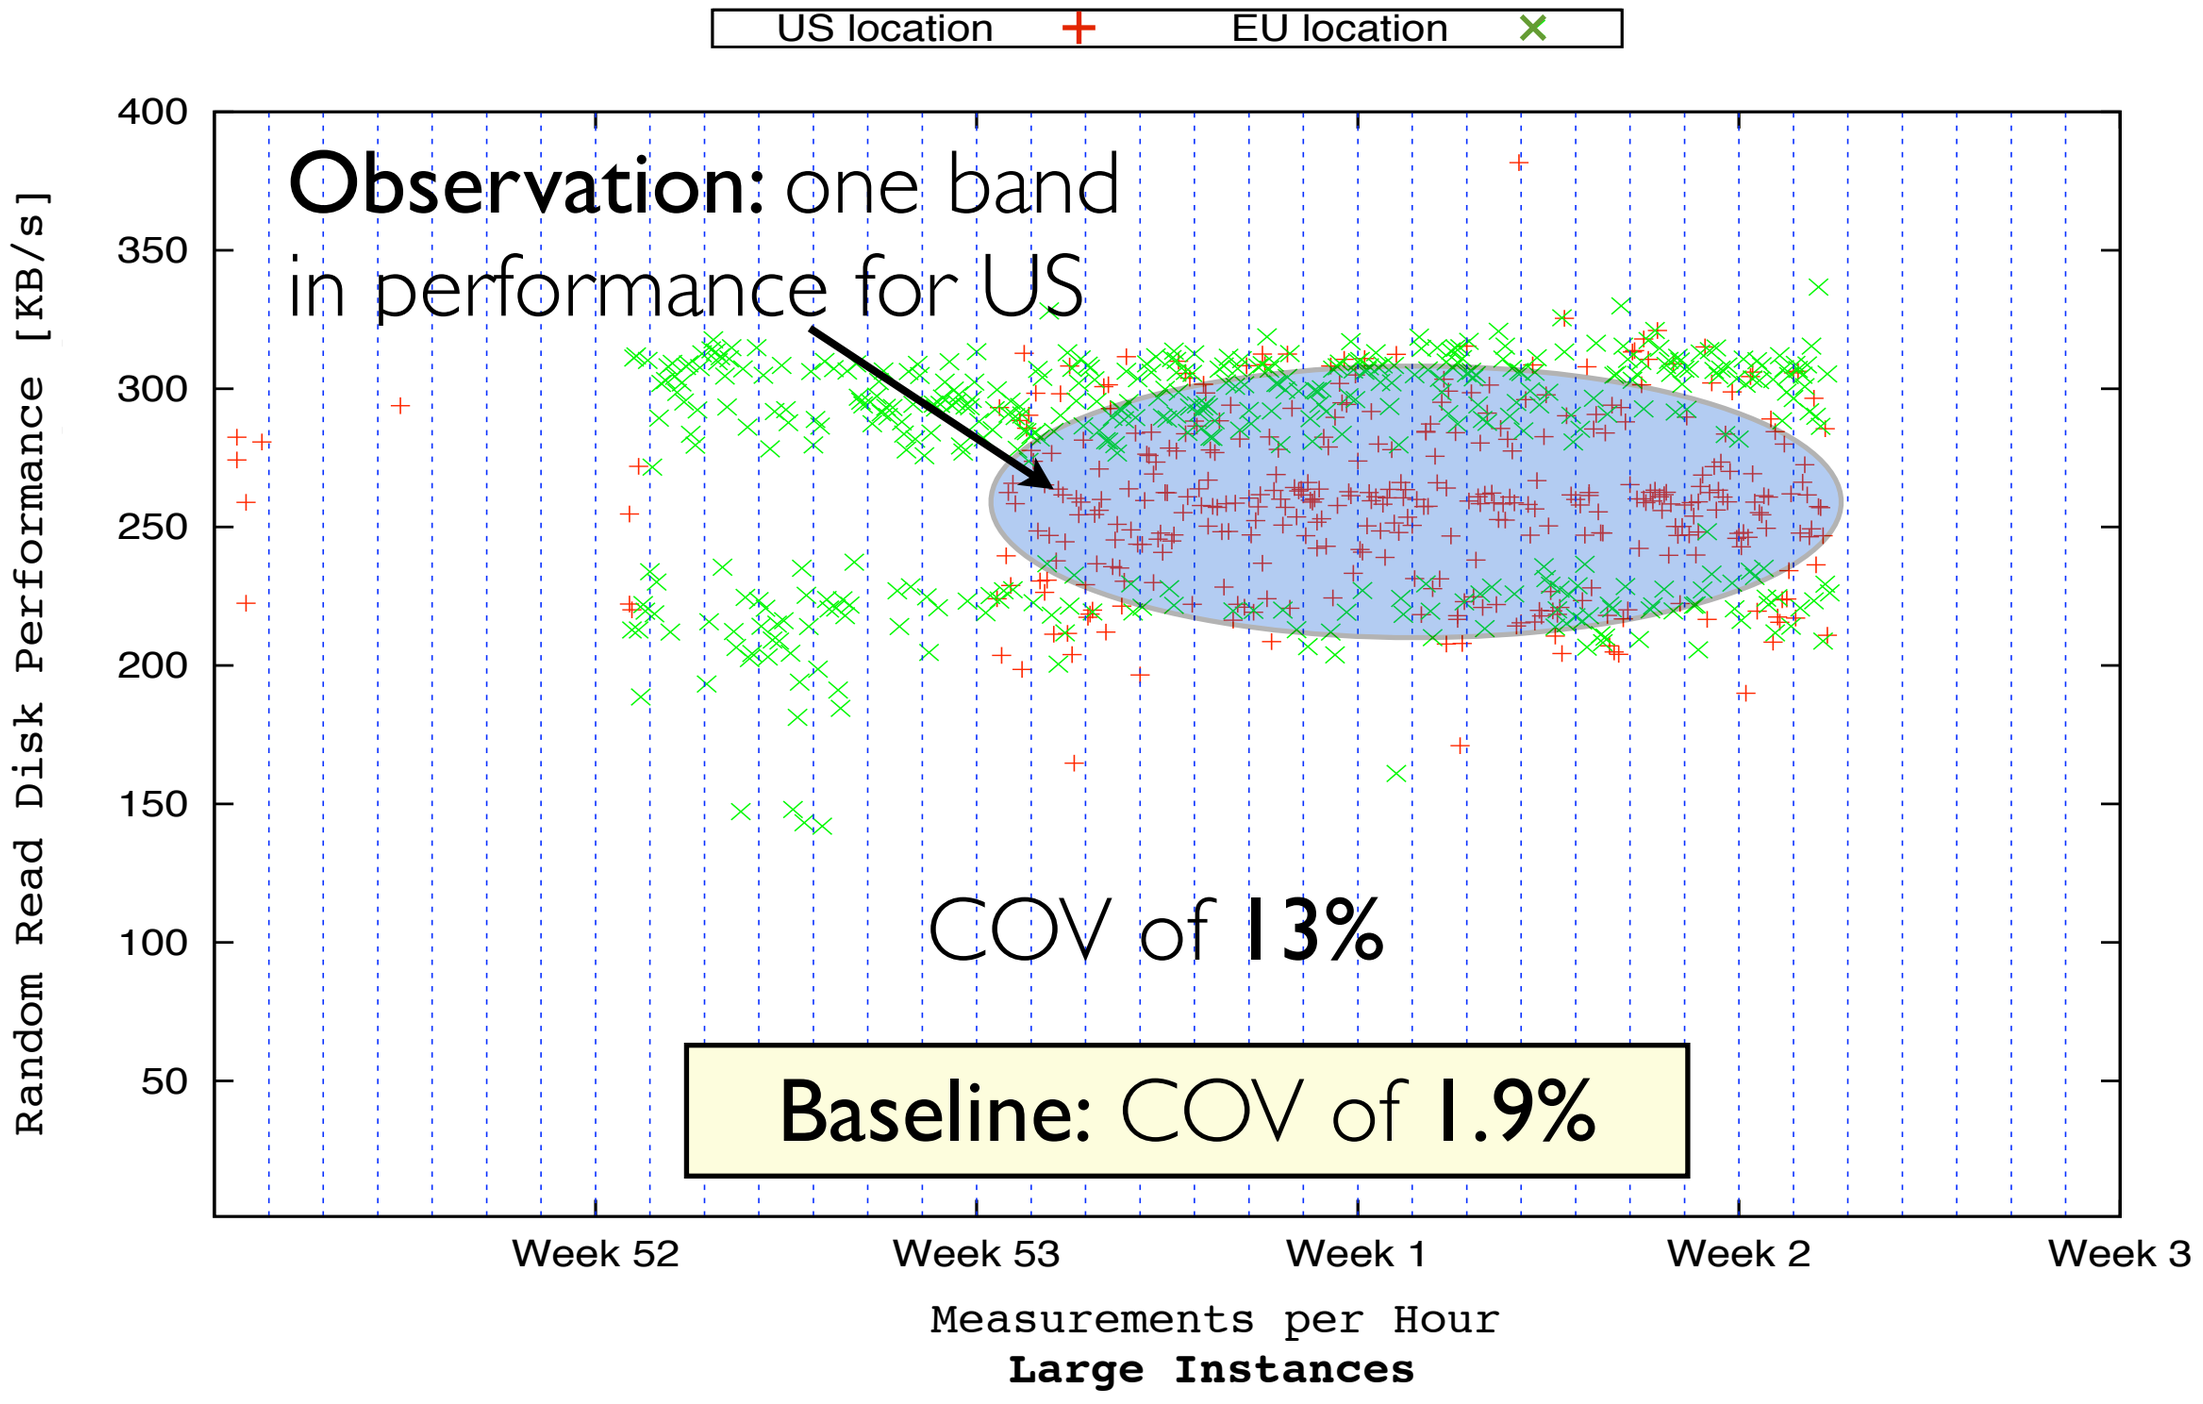

Availability Zones

SAARLAND UNIVERSITY
Information Systems Group

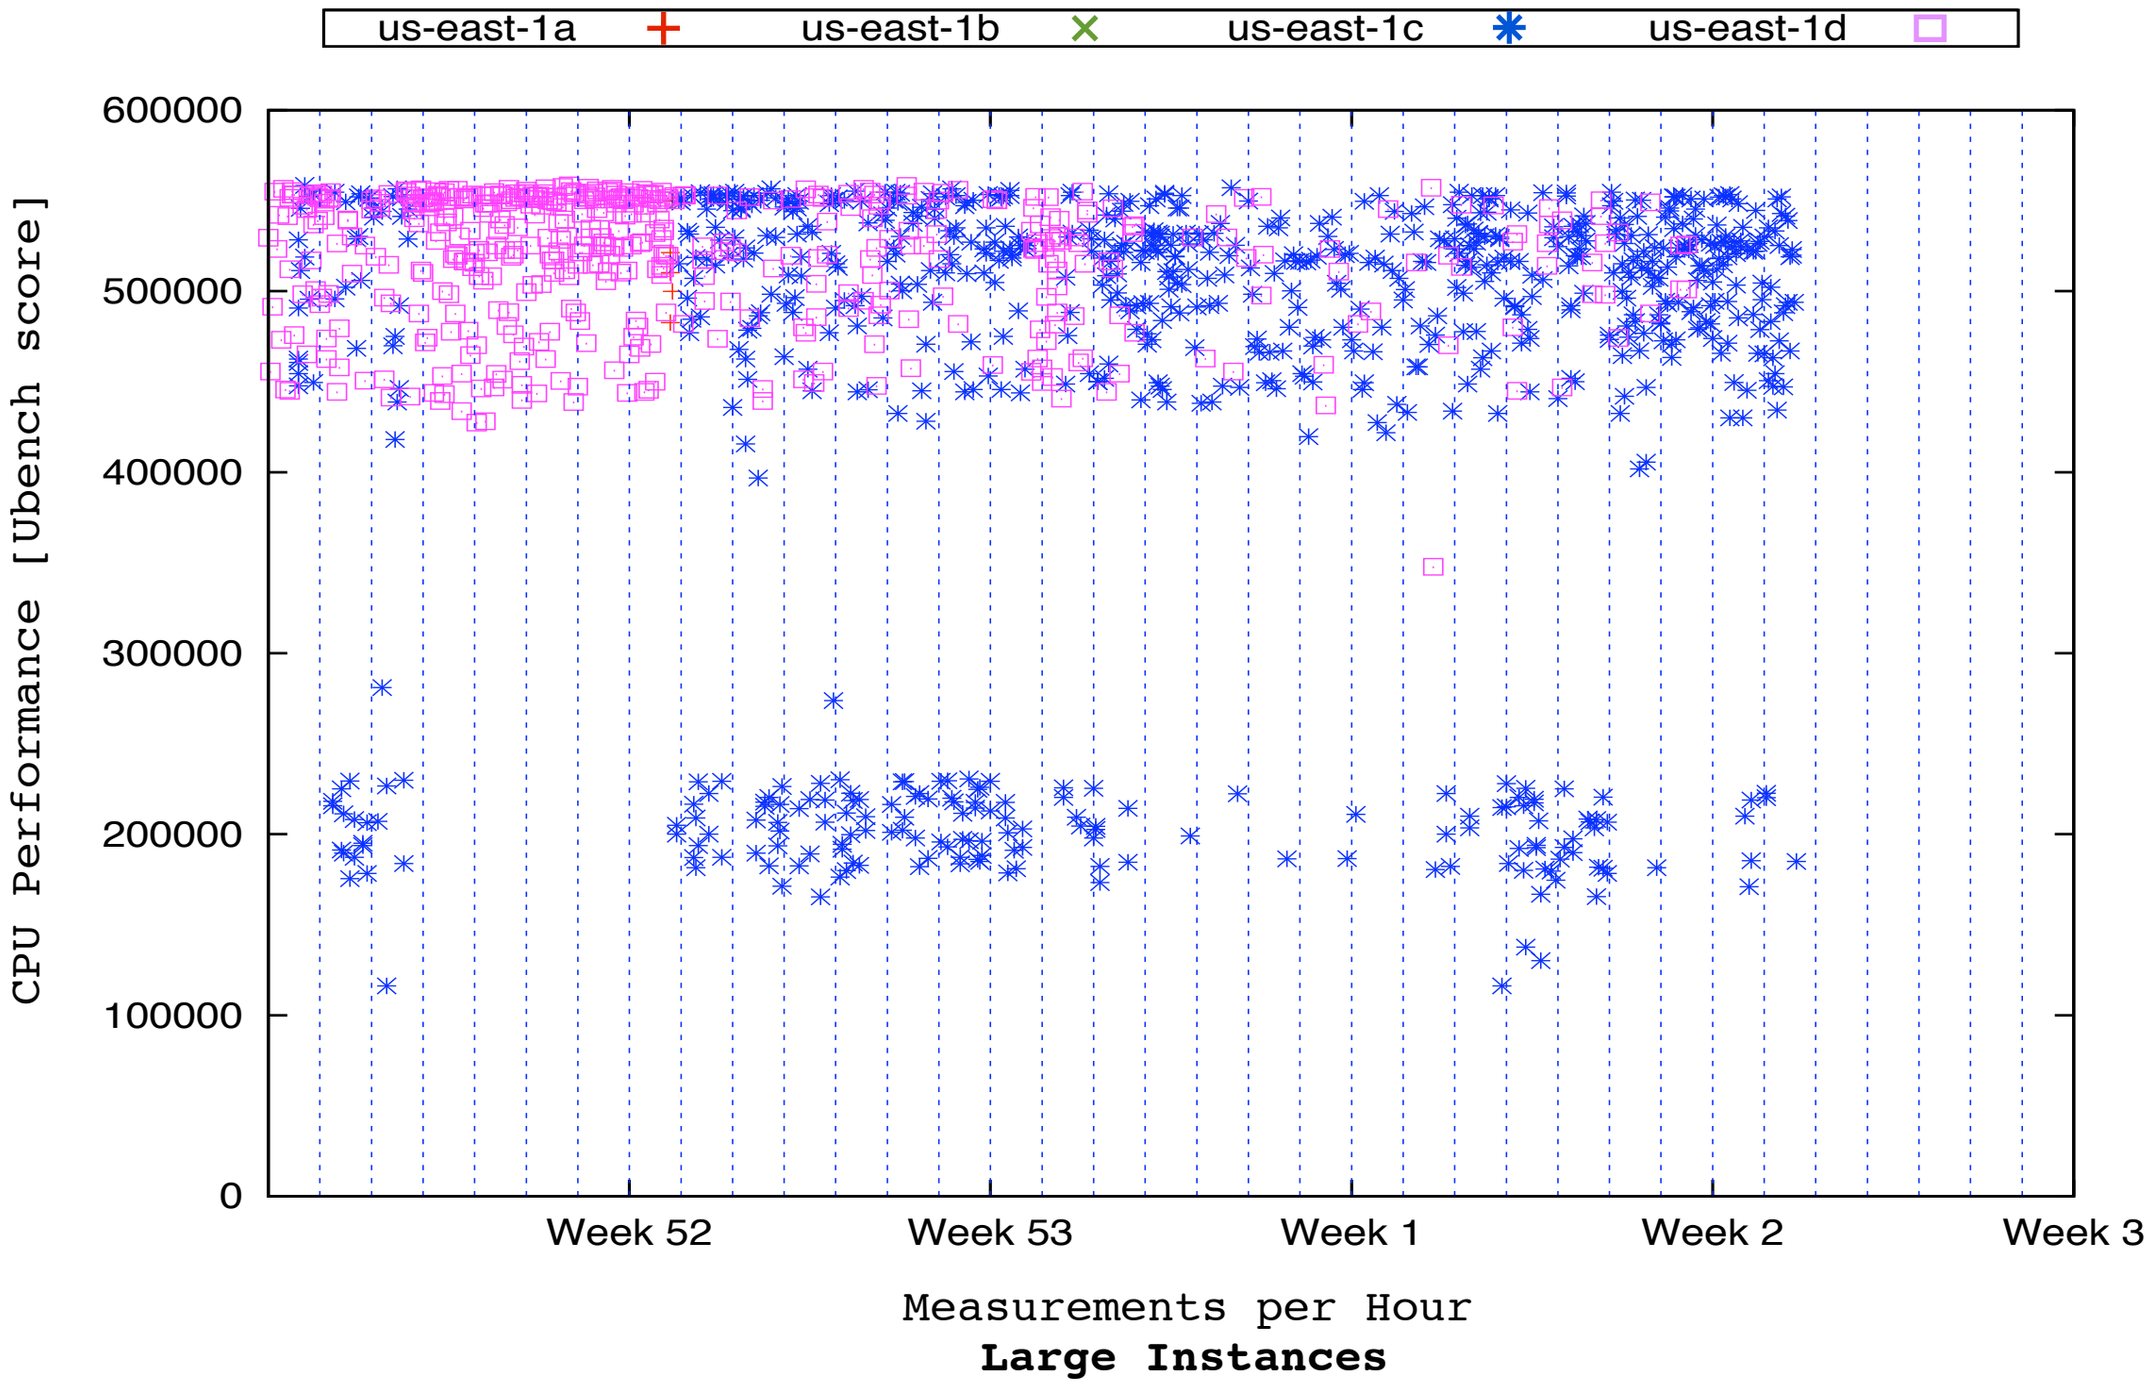

Availability Zones

SAARLAND UNIVERSITY
Information Systems Group

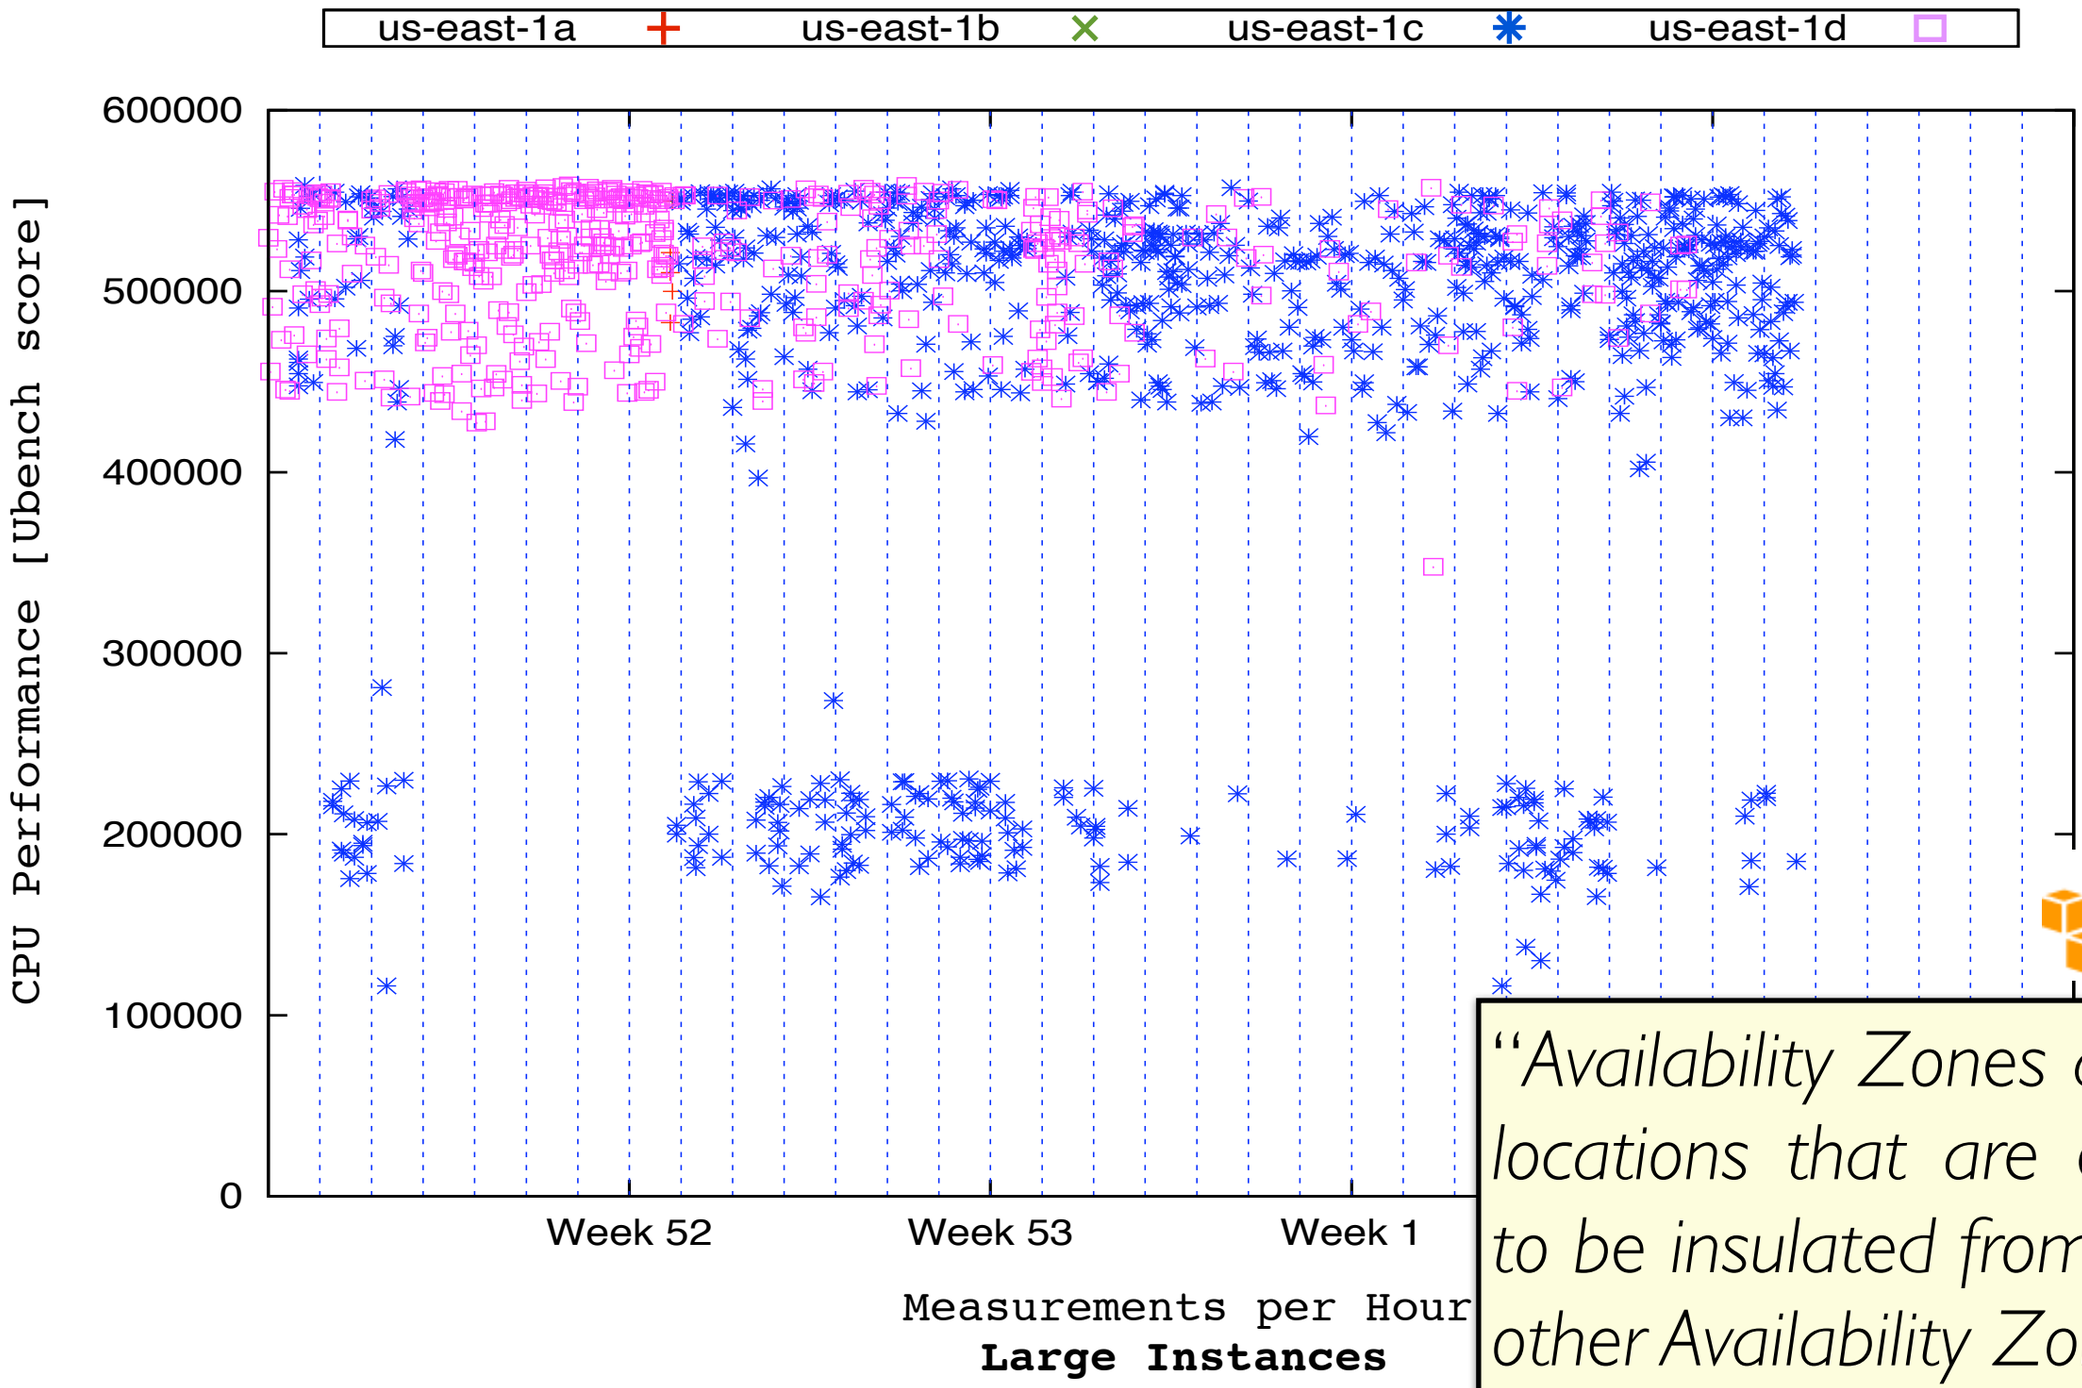

amazon web services™

"*Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones...*"

Availability Zones

SAARLAND UNIVERSITY
Information Systems Group

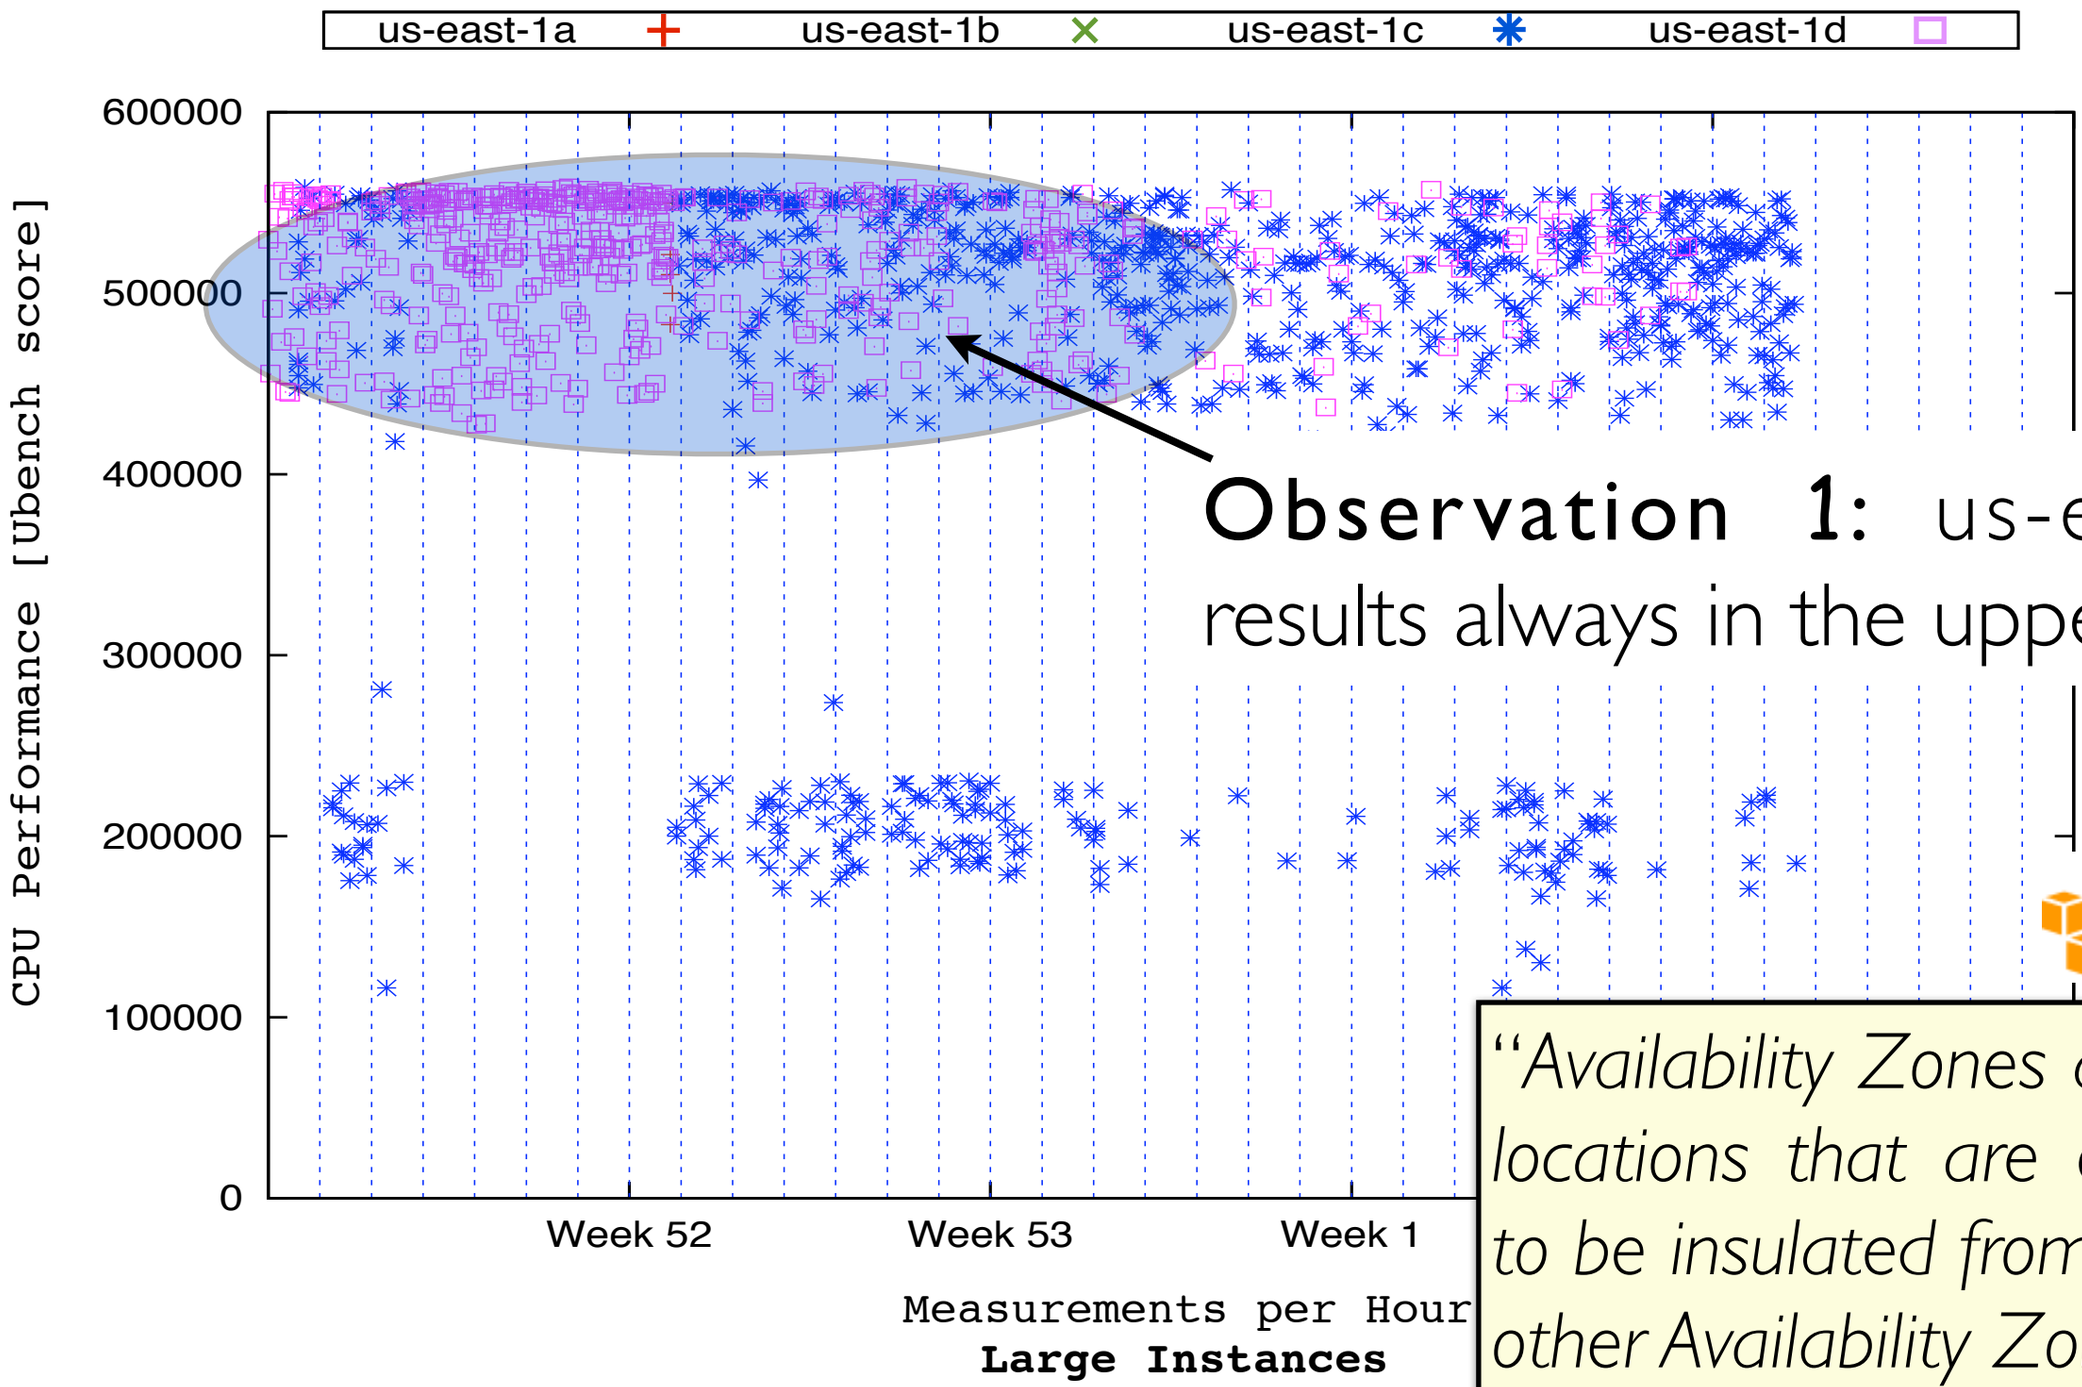

Observation 1: us-east-1d results always in the upper band

amazon web services™

"Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones..."

Availability Zones

SAARLAND UNIVERSITY
Information Systems Group

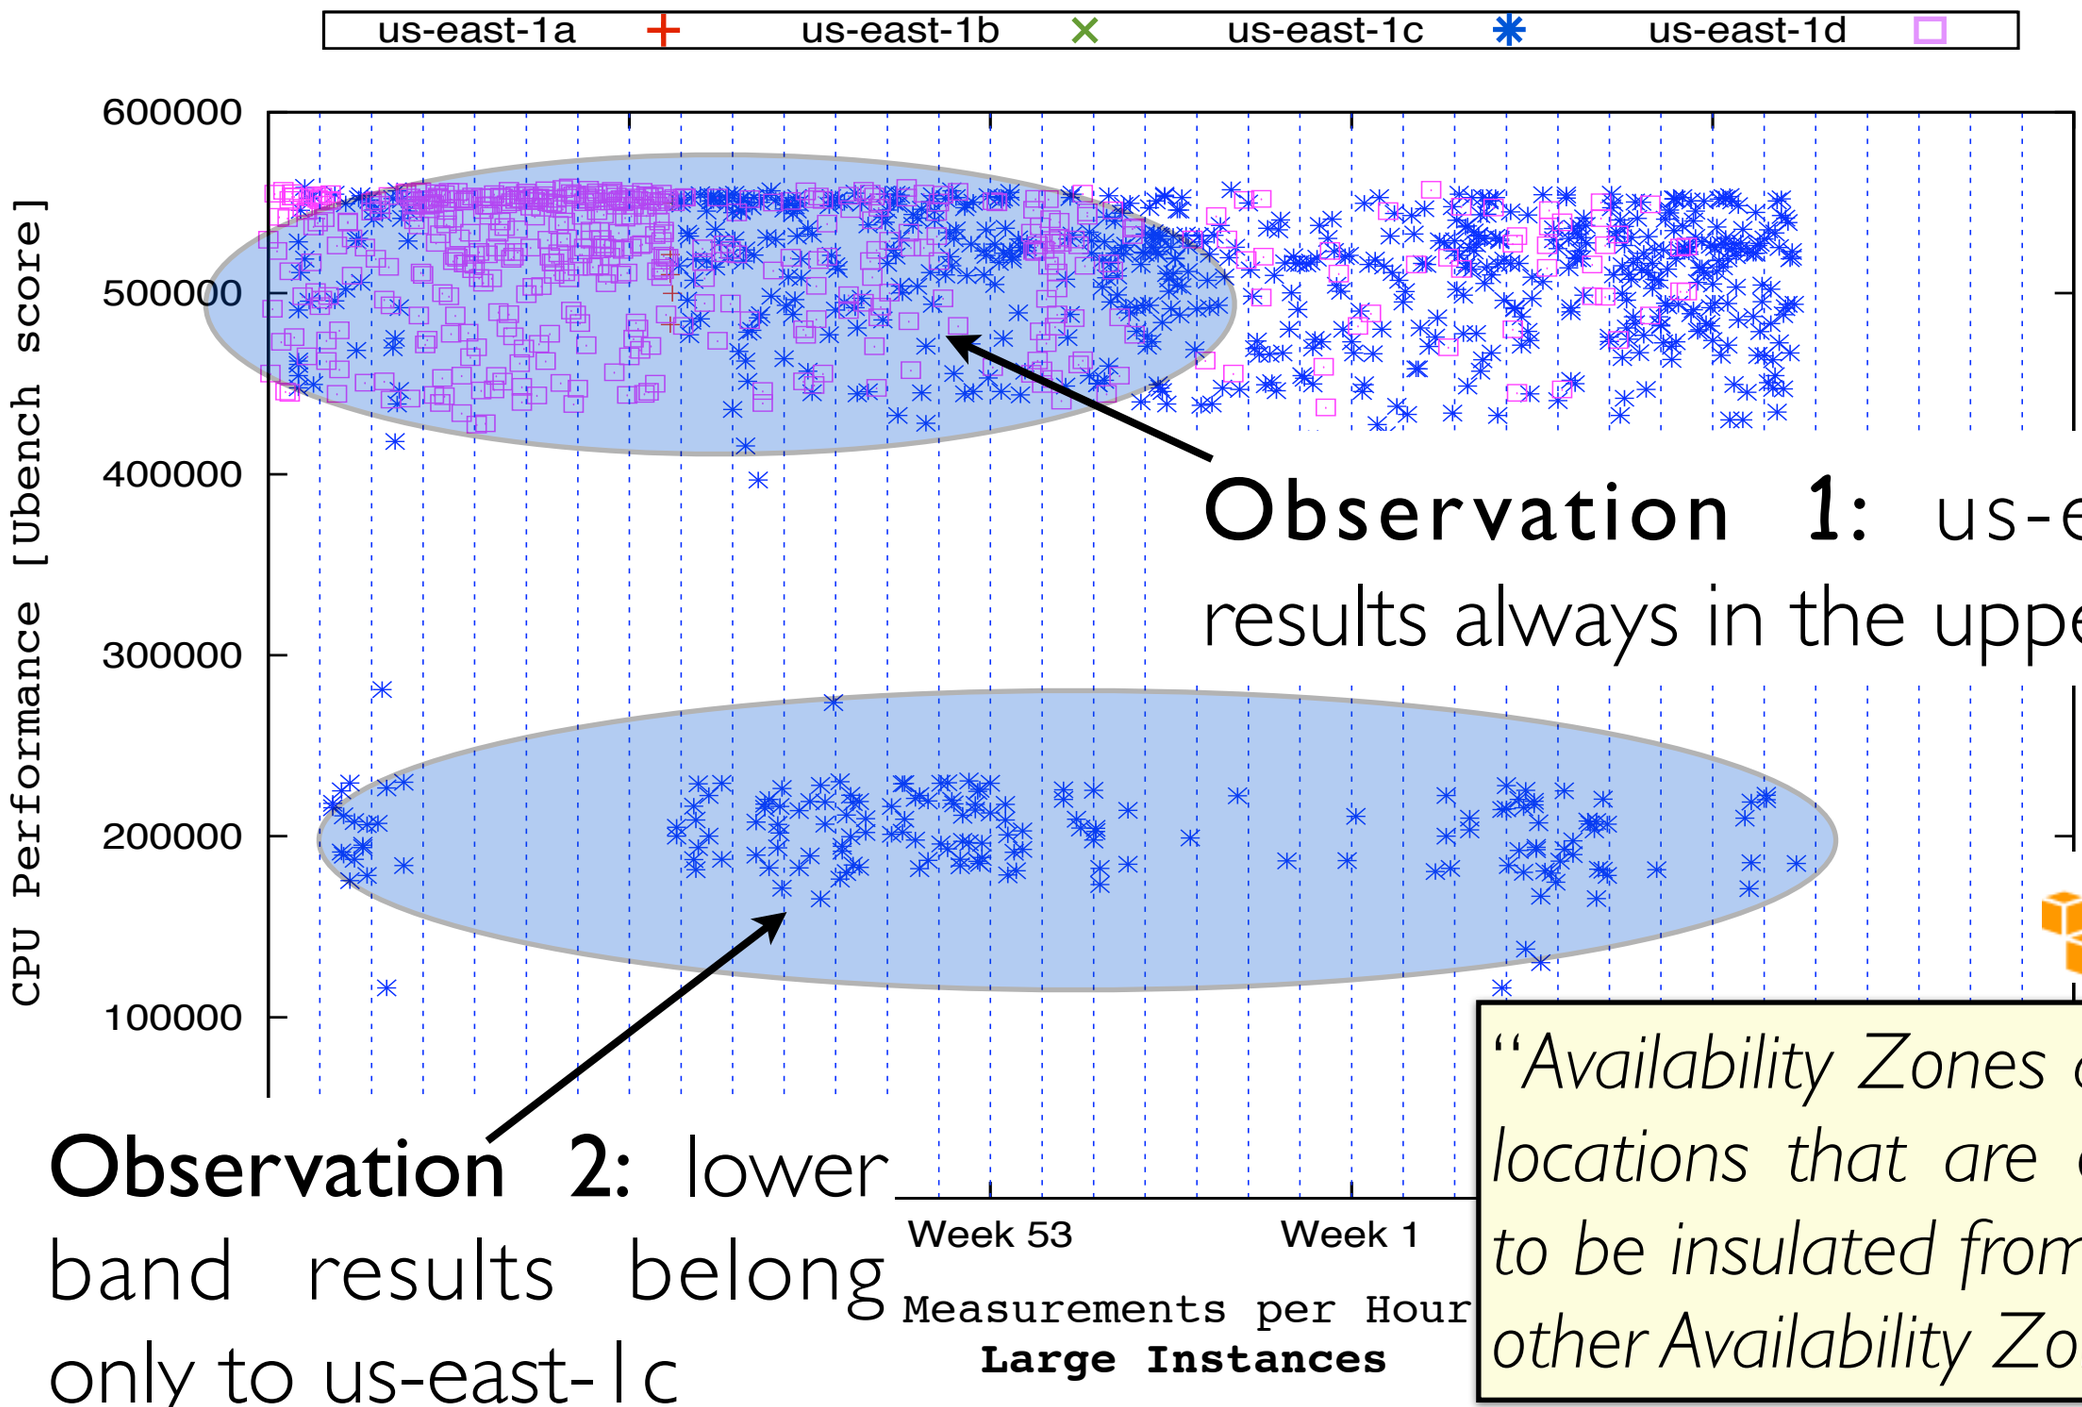

Observation 1: us-east-1d results always in the upper band

Observation 2: lower band results belong only to us-east-1c

amazon web services™

"Availability Zones are distinct locations that are engineered to be insulated from failures in other Availability Zones..."

CPU Distribution (1/2)

SAARLAND UNIVERSITY
Information Systems Group

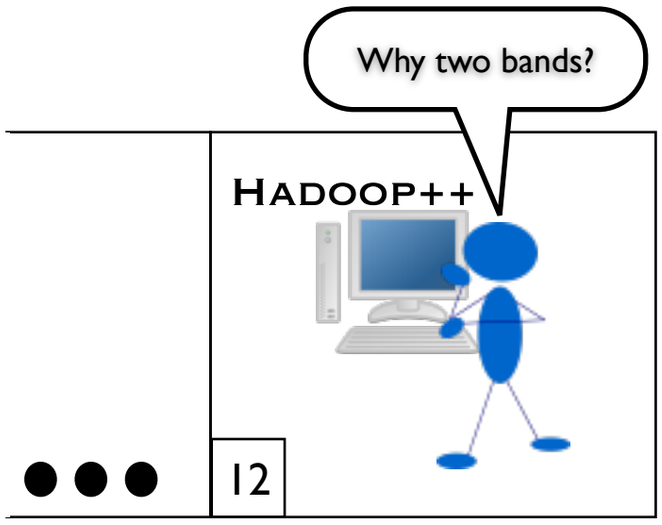

CPU Distribution (1/2)

SAARLAND UNIVERSITY
Information Systems Group

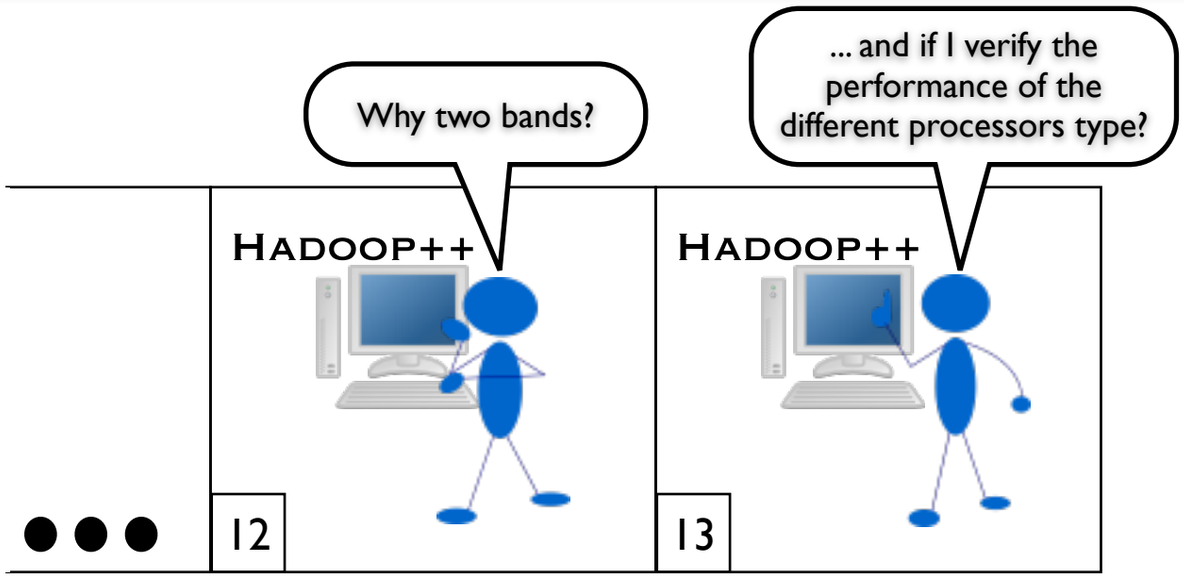

CPU Distribution (1/2)

SAARLAND UNIVERSITY
Information Systems Group

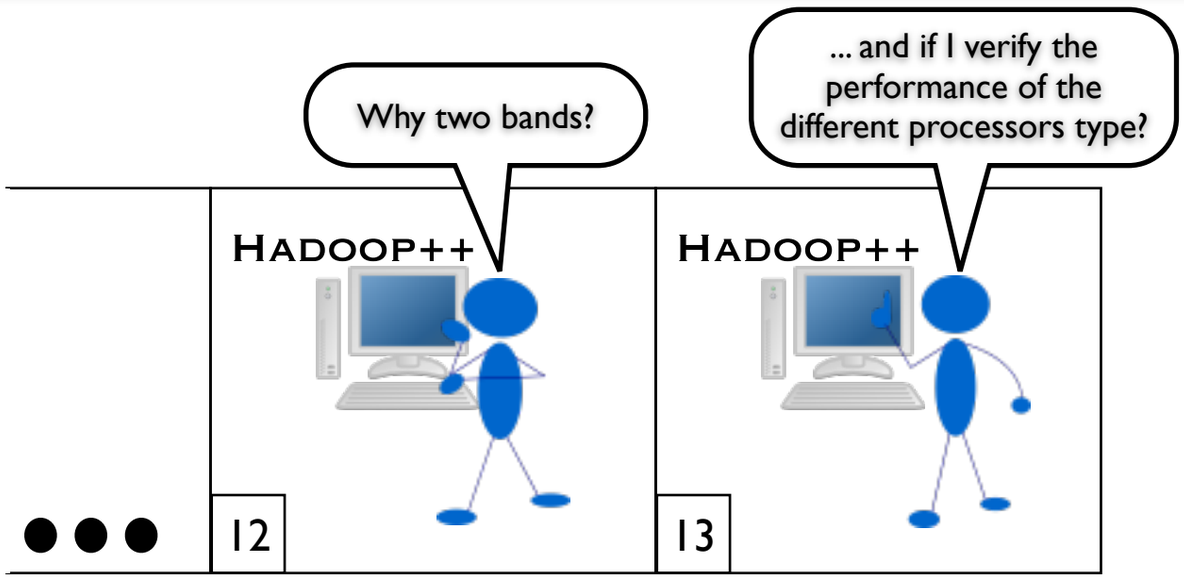

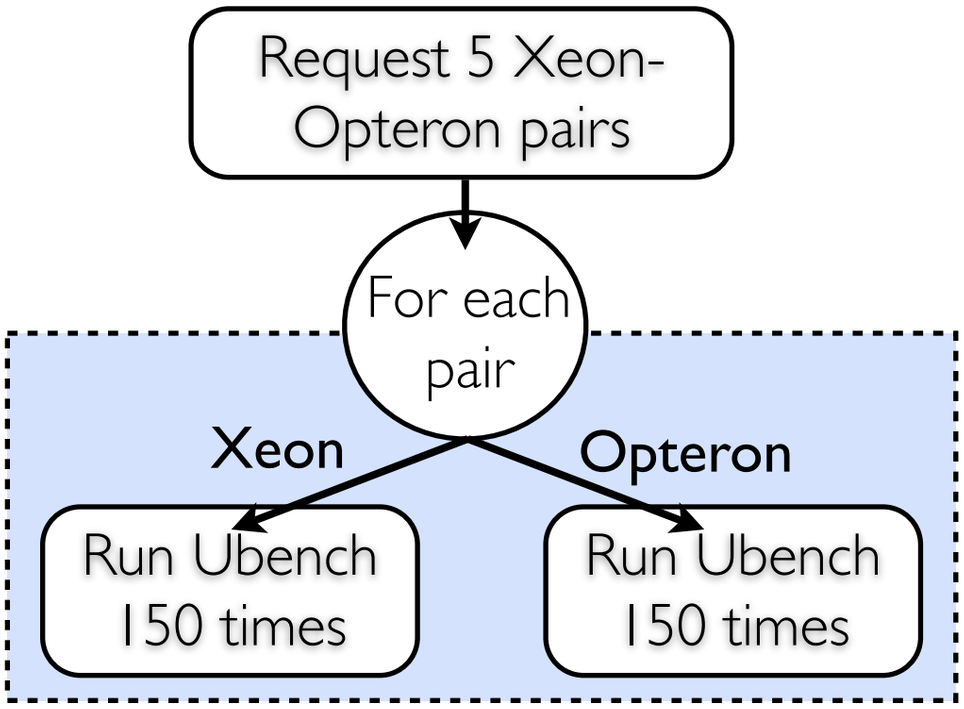

CPU Distribution (1/2)

SAARLAND UNIVERSITY
Information Systems Group

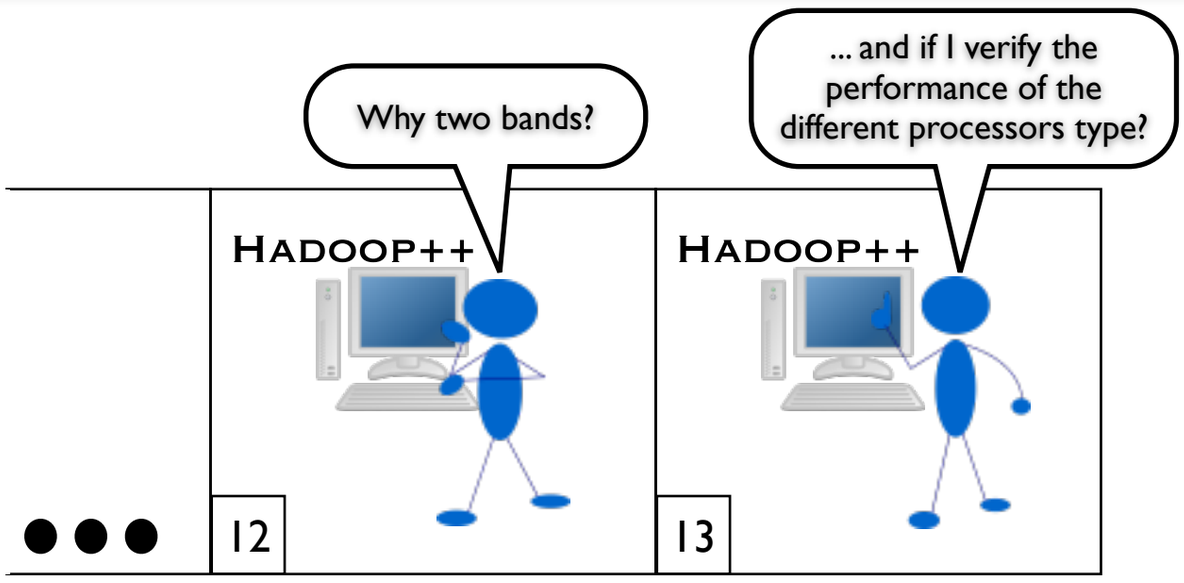

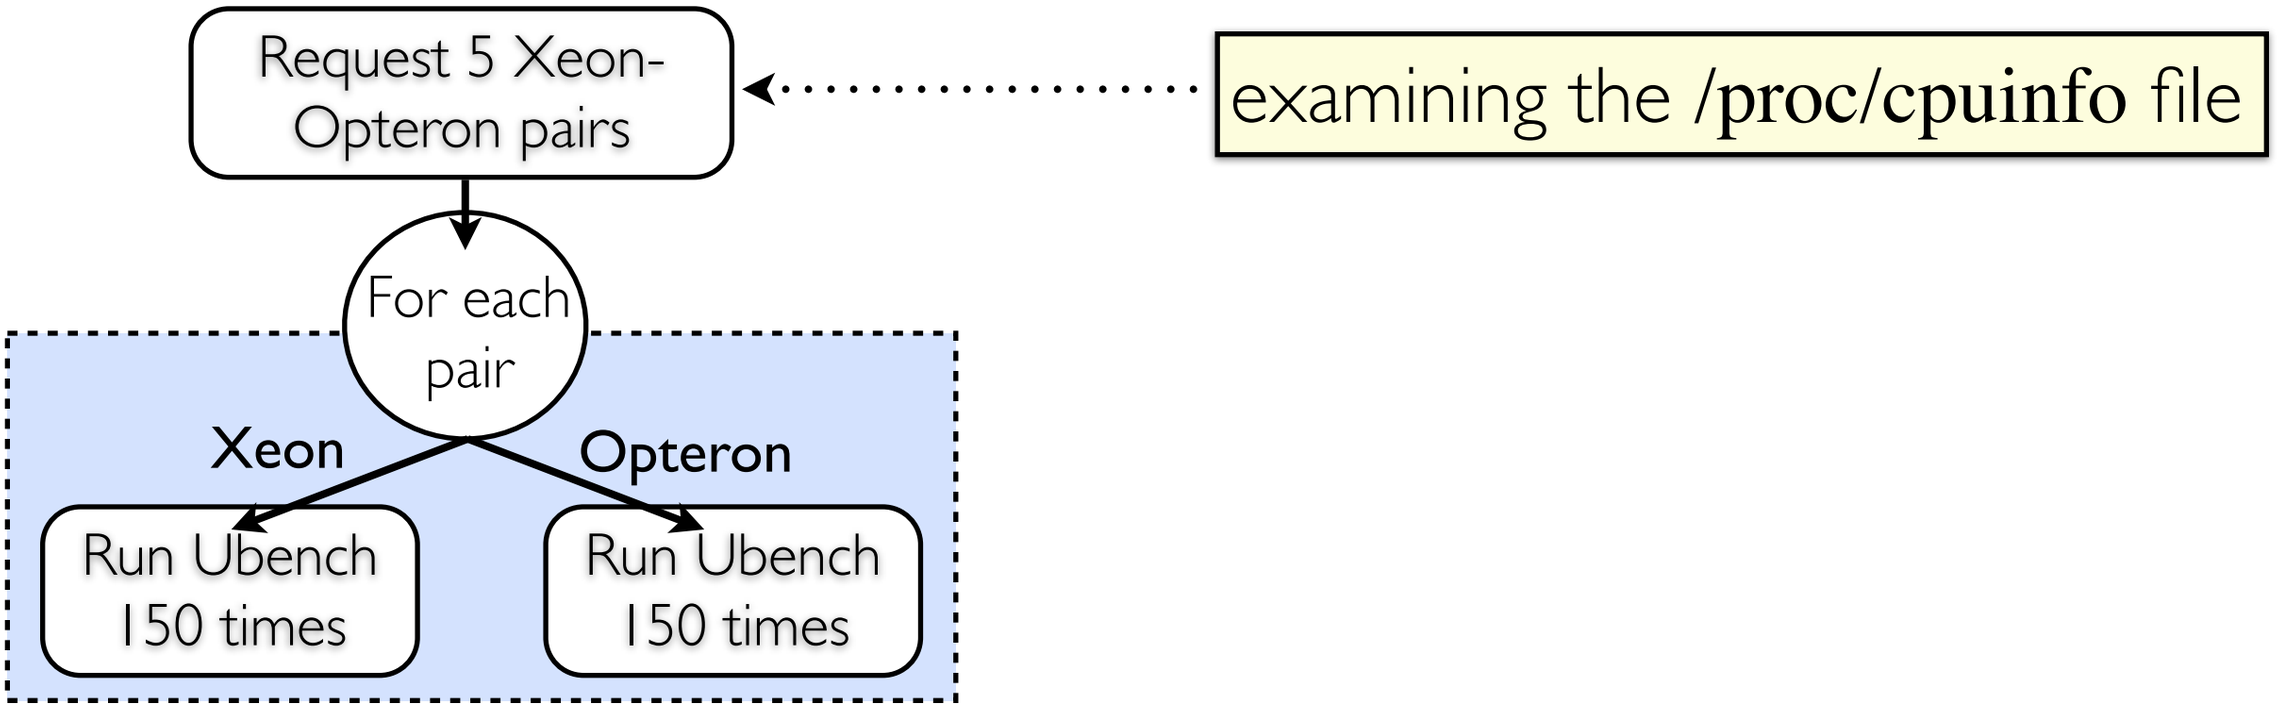

CPU Distribution (2/2)

SAARLAND UNIVERSITY
Information Systems Group

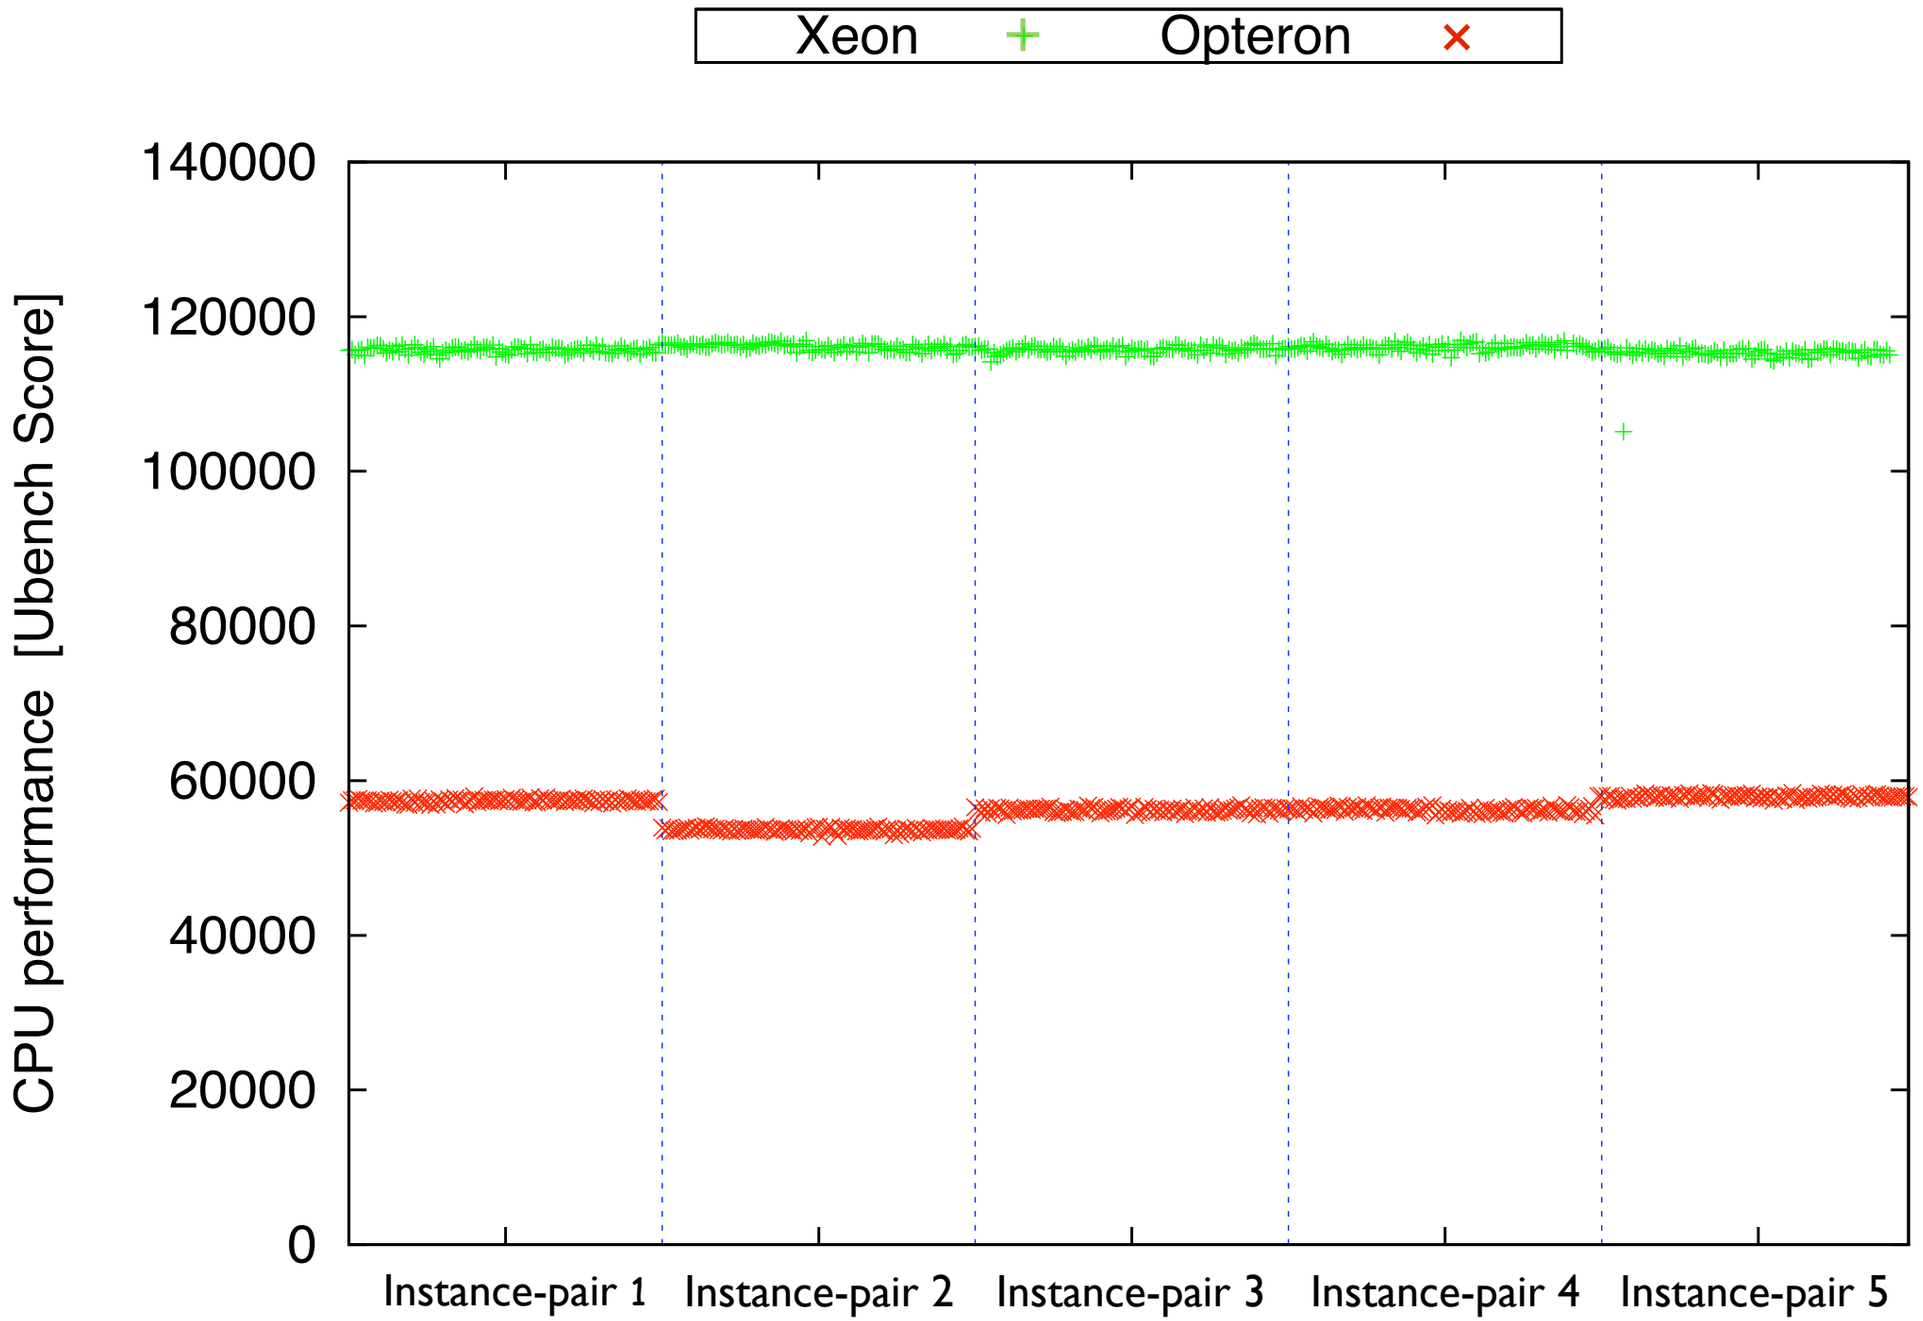

CPU Distribution (2/2)

SAARLAND UNIVERSITY
Information Systems Group

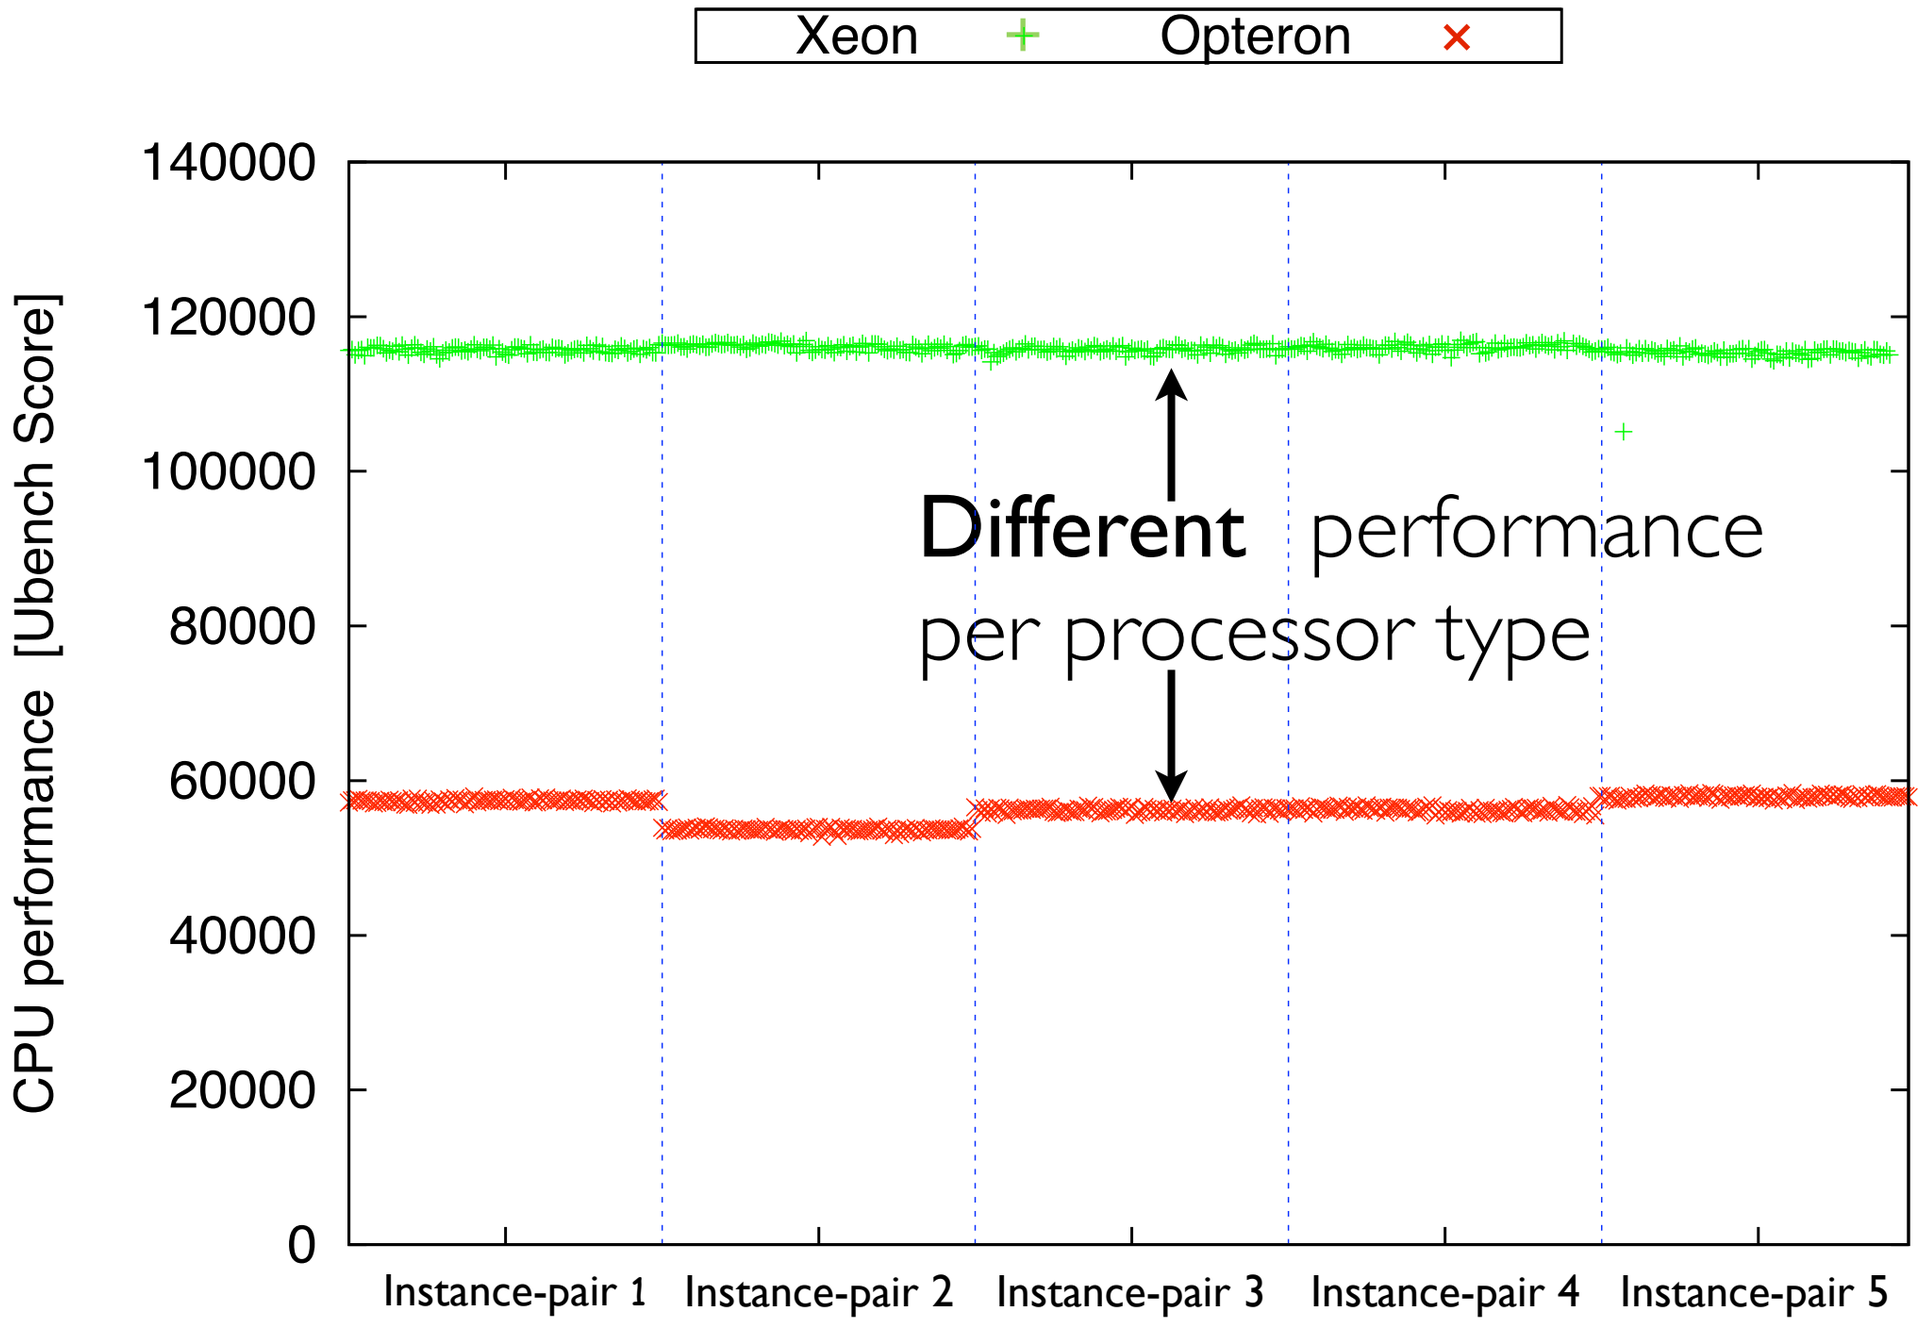

CPU Distribution (2/2)

SAARLAND UNIVERSITY
Information Systems Group

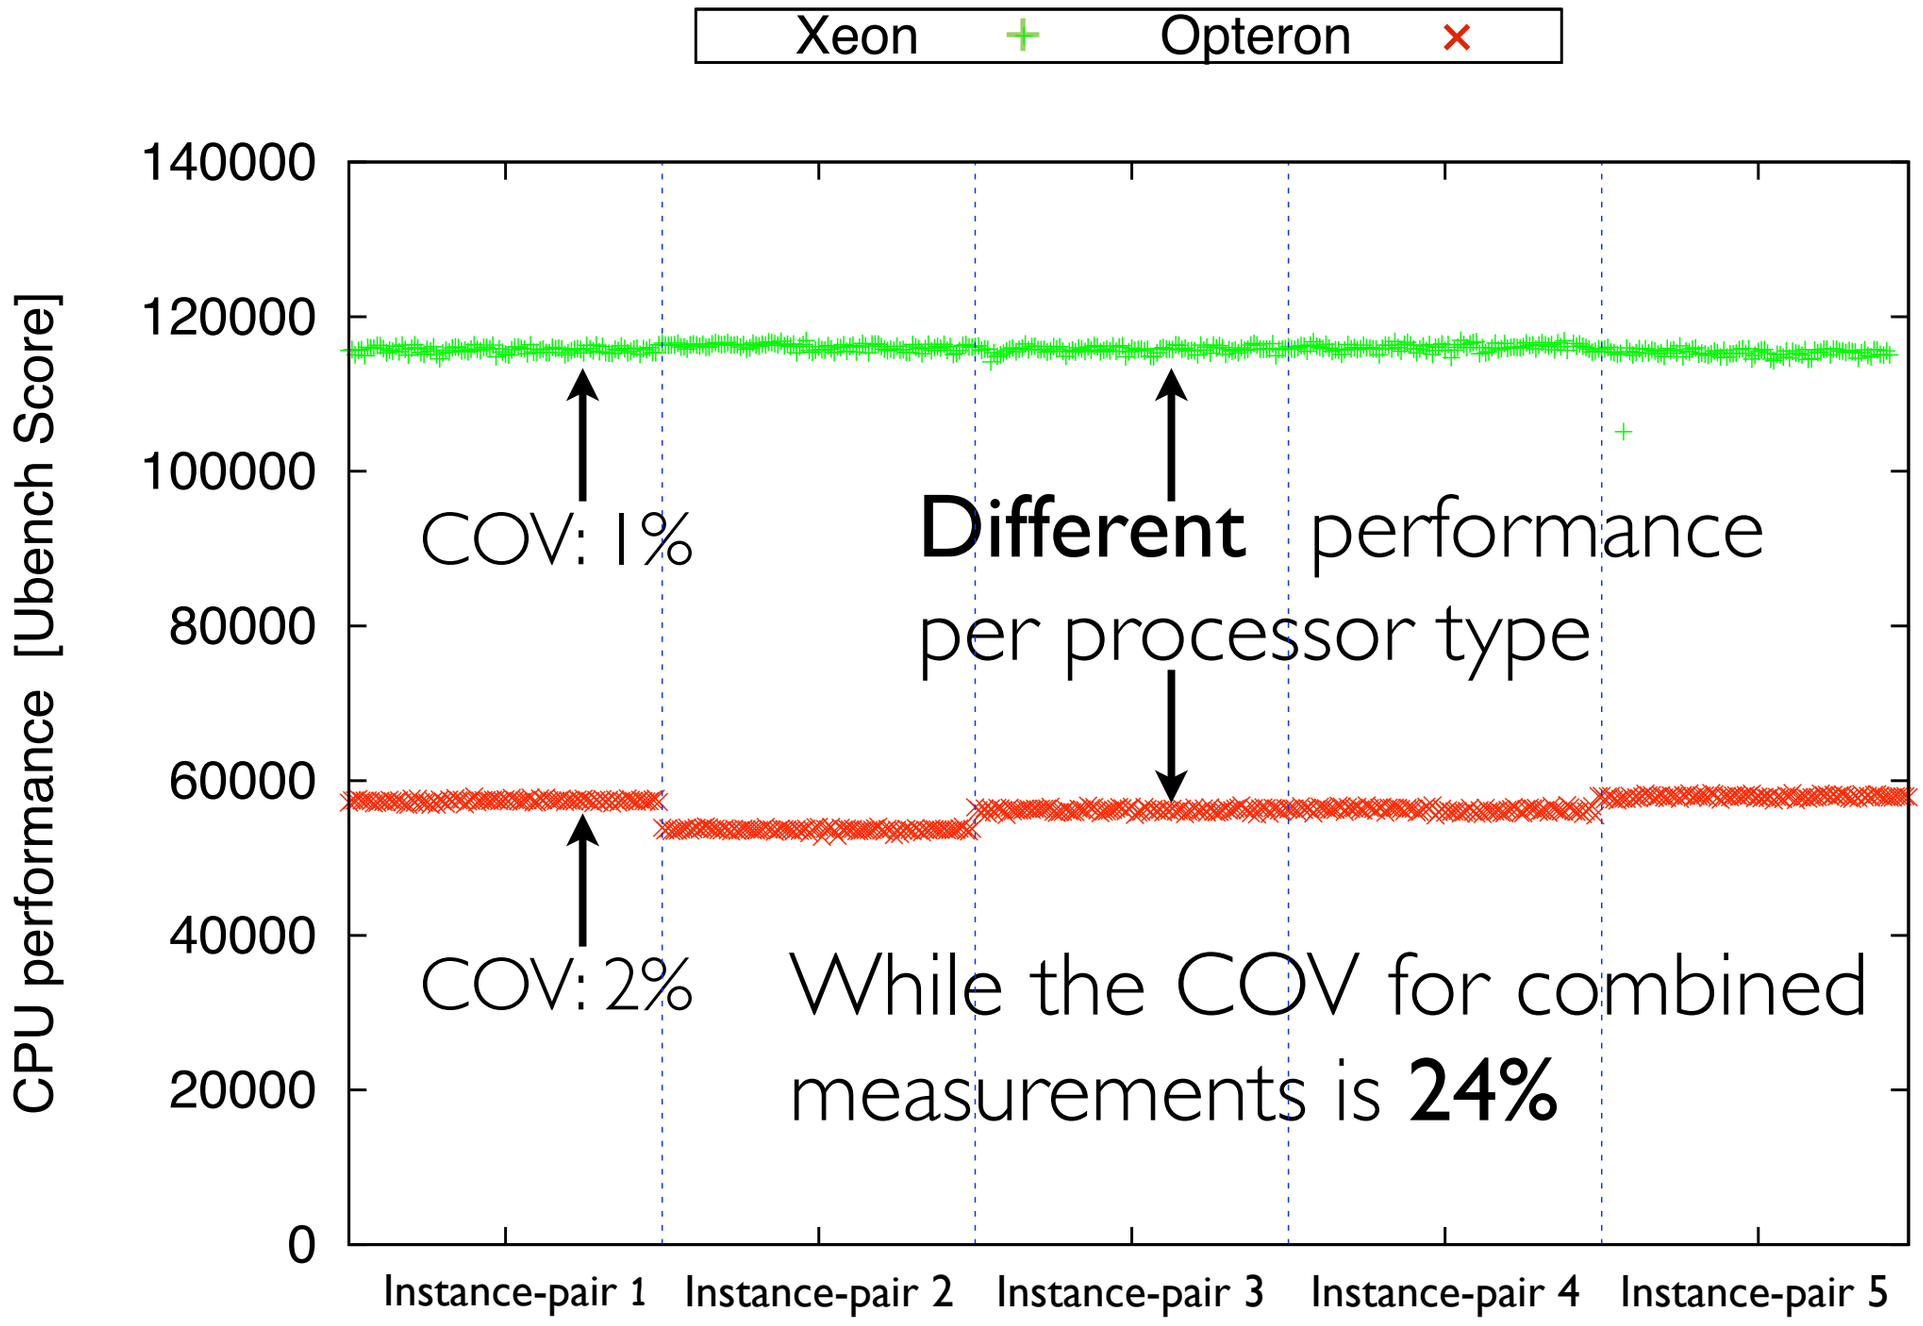

CPU Distribution (2/2)

SAARLAND UNIVERSITY
Information Systems Group

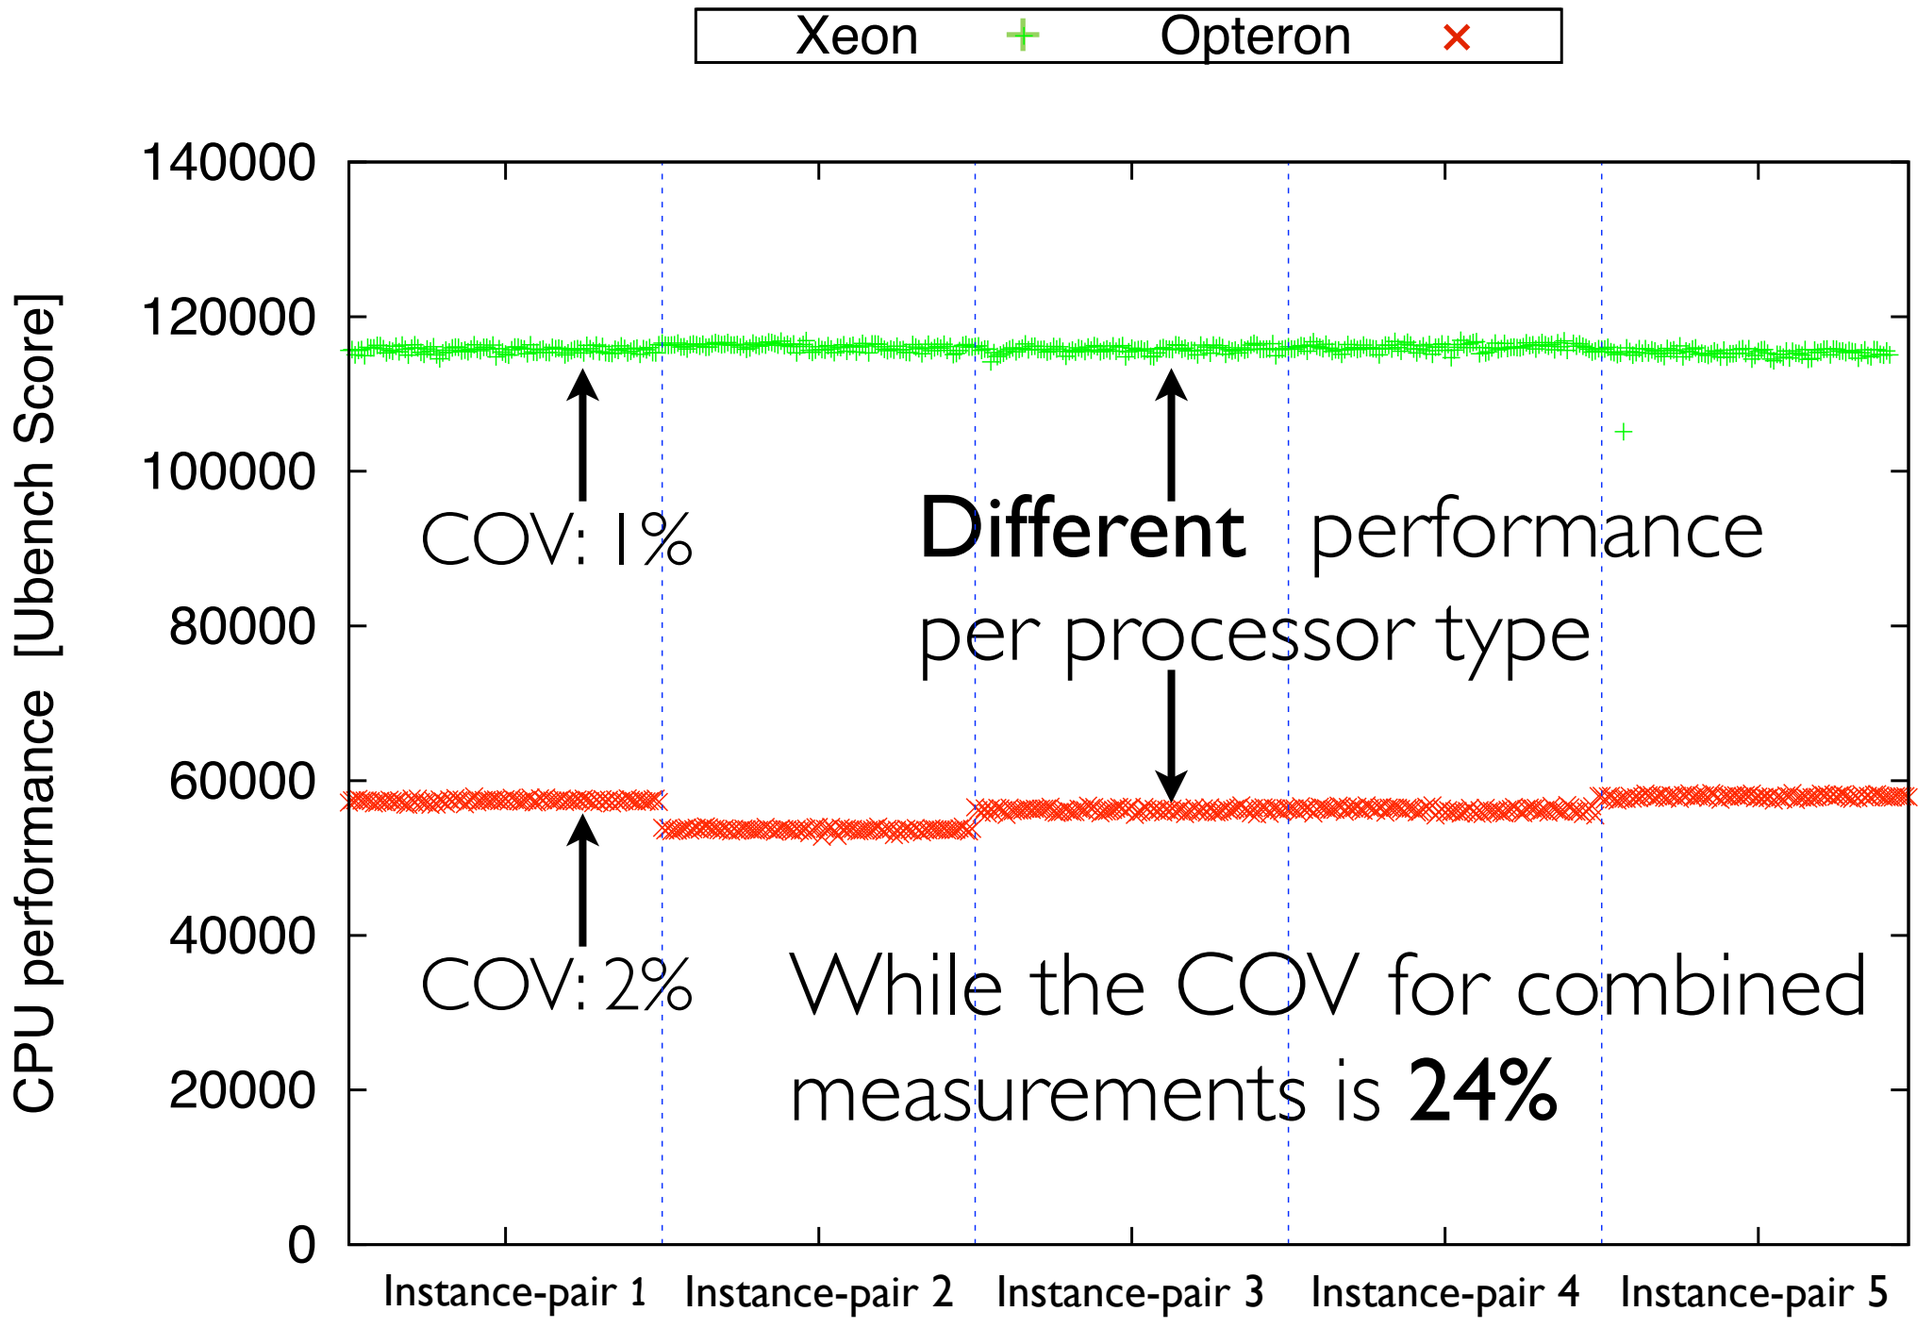

Observation: *1* cpu ⟶ *1* underlying system

[memory and I/O follows this pattern]

Larger Clusters

SAARLAND UNIVERSITY
Information Systems Group

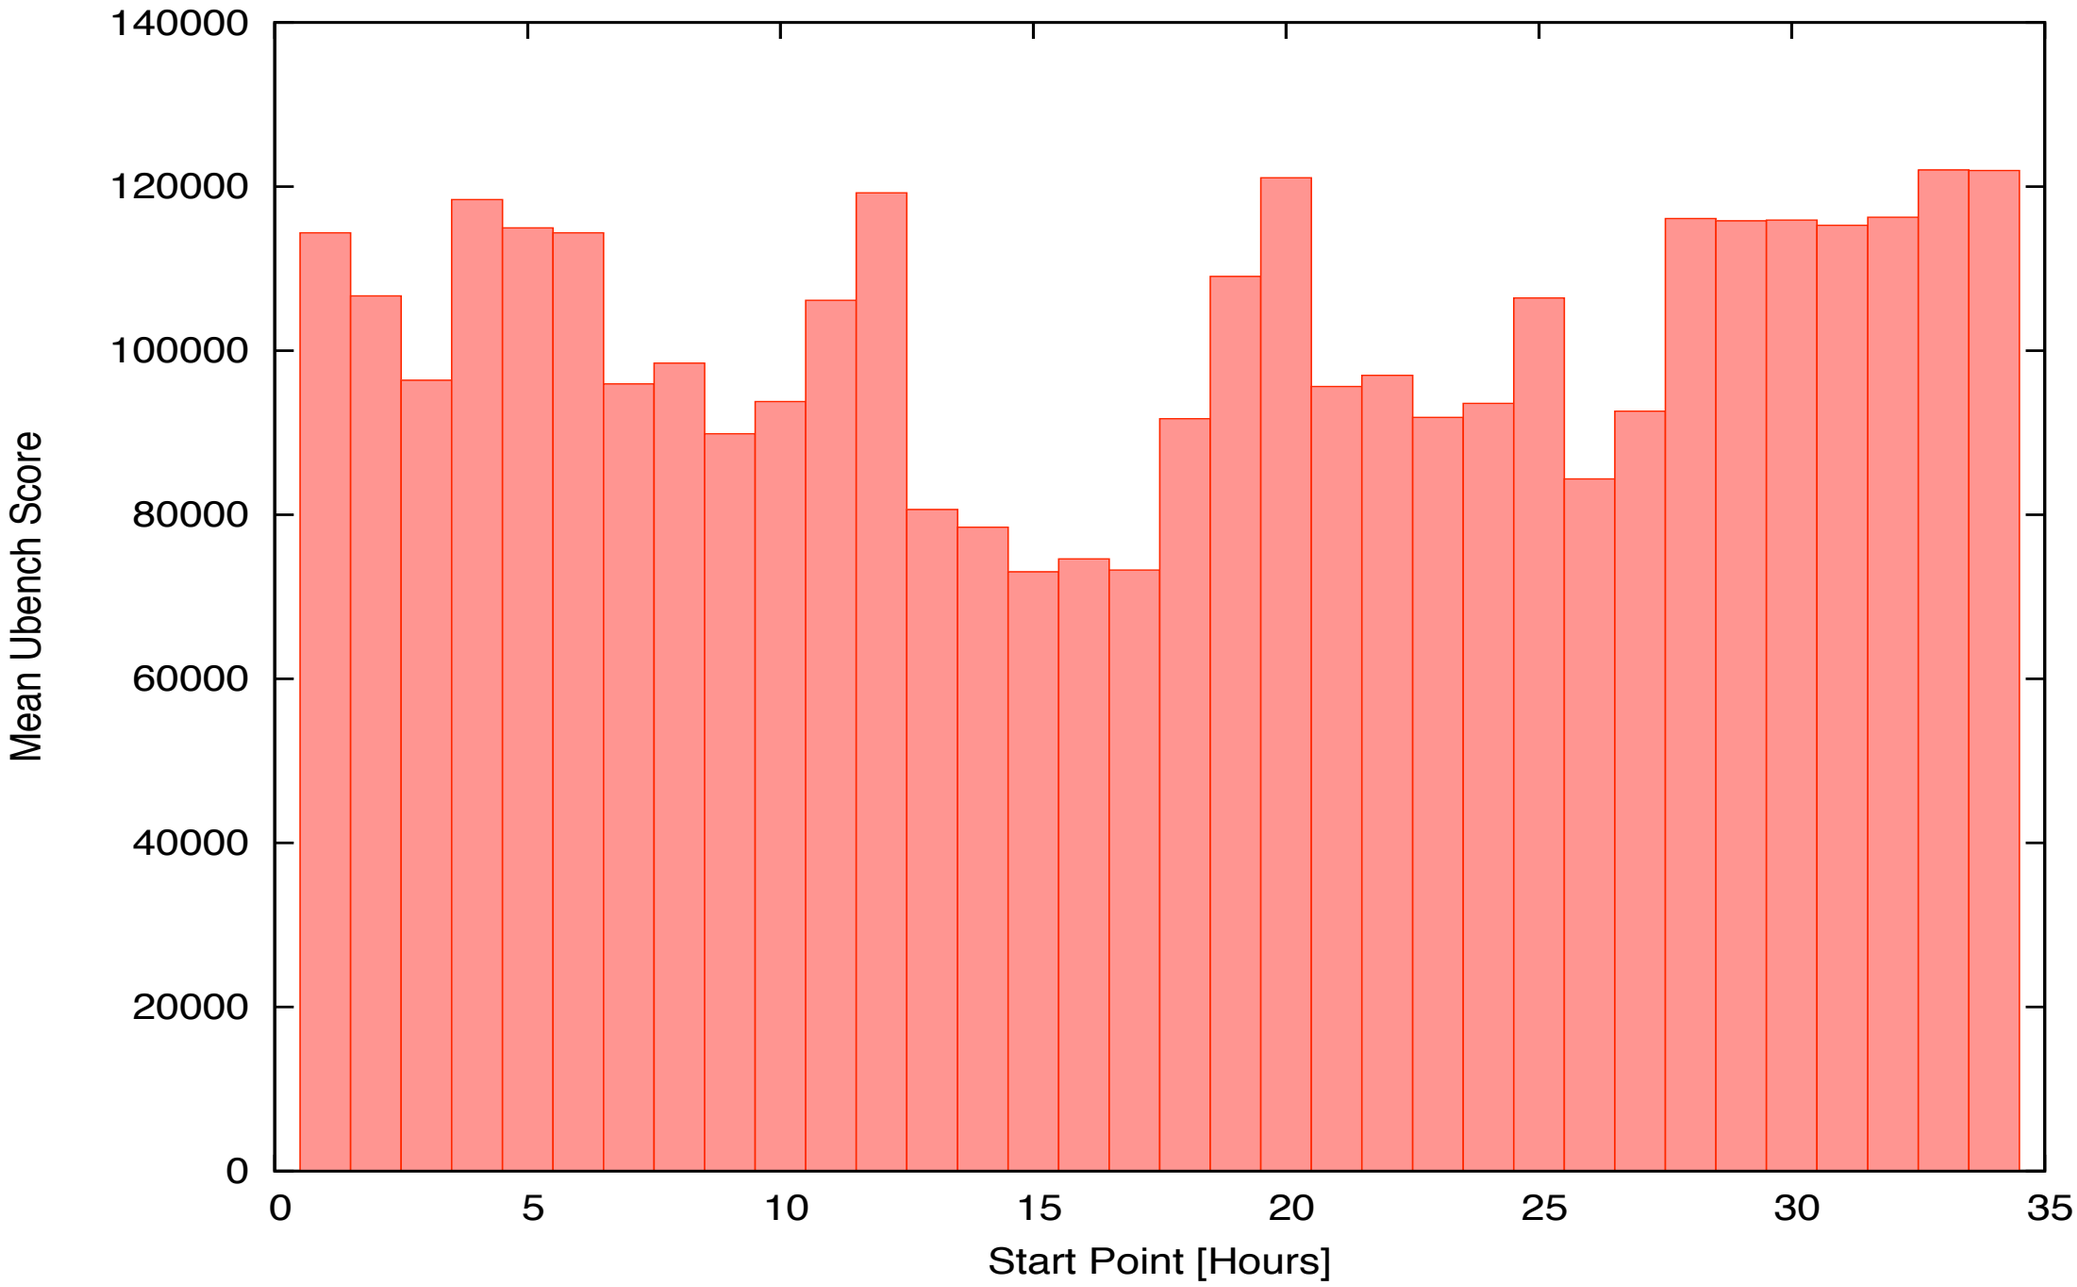

50 Large Instances Cluster

Larger Clusters

SAARLAND UNIVERSITY
Information Systems Group

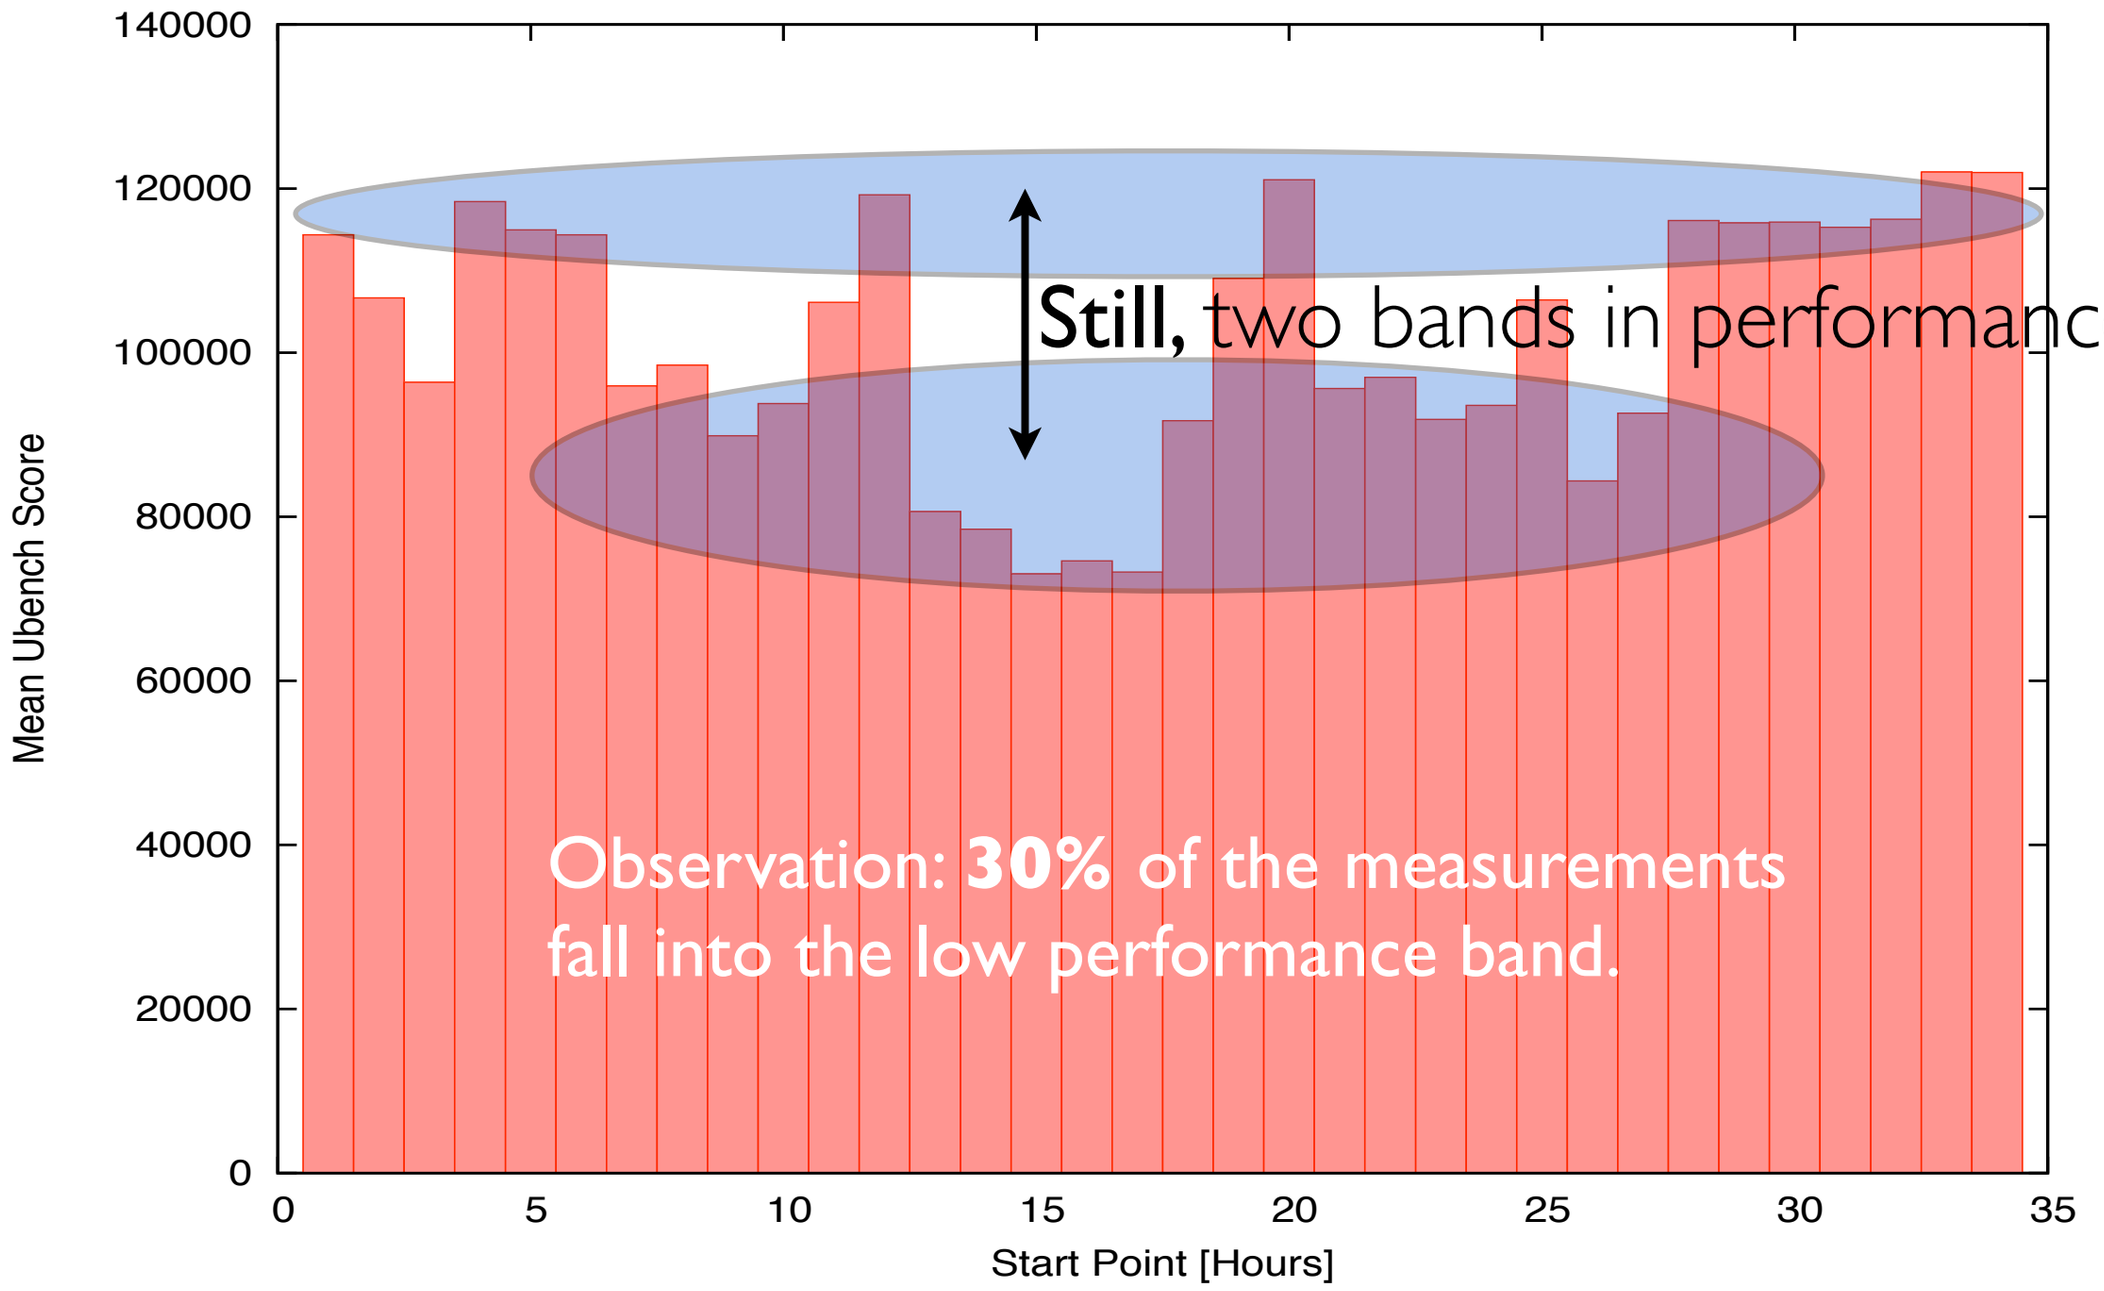

MapReduce Job

SAARLAND UNIVERSITY
Information Systems Group

- **Application:** a MapReduce aggregation job
- **Number of virtual nodes:** 50
- **Repetitions:** once every hour

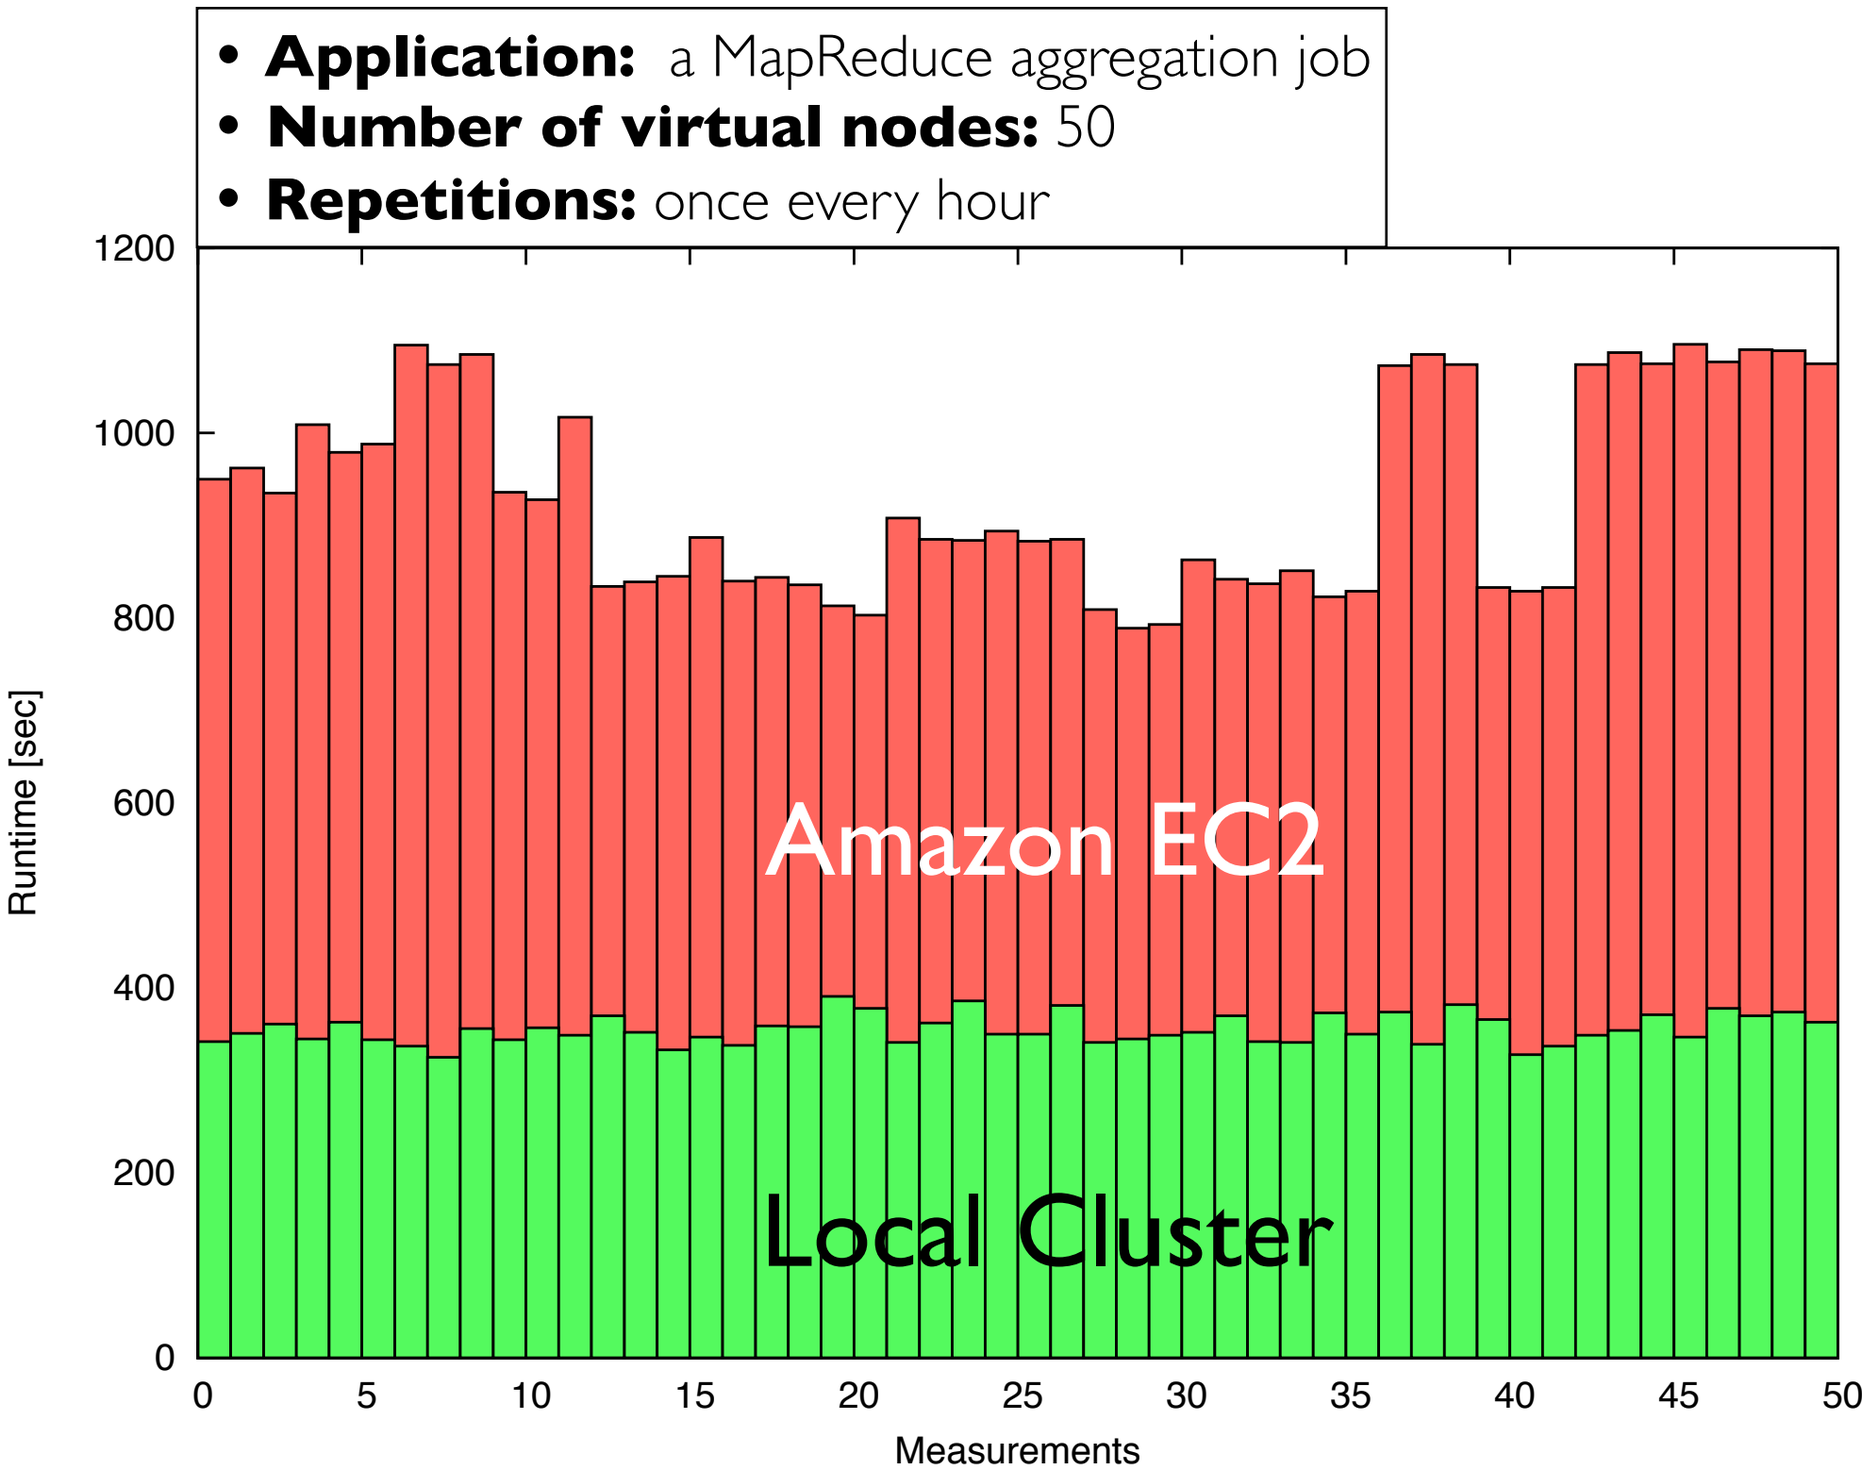

MapReduce Job

SAARLAND UNIVERSITY
Information Systems Group

- **Application:** a MapReduce aggregation job
- **Number of virtual nodes:** 50
- **Repetitions:** once every hour

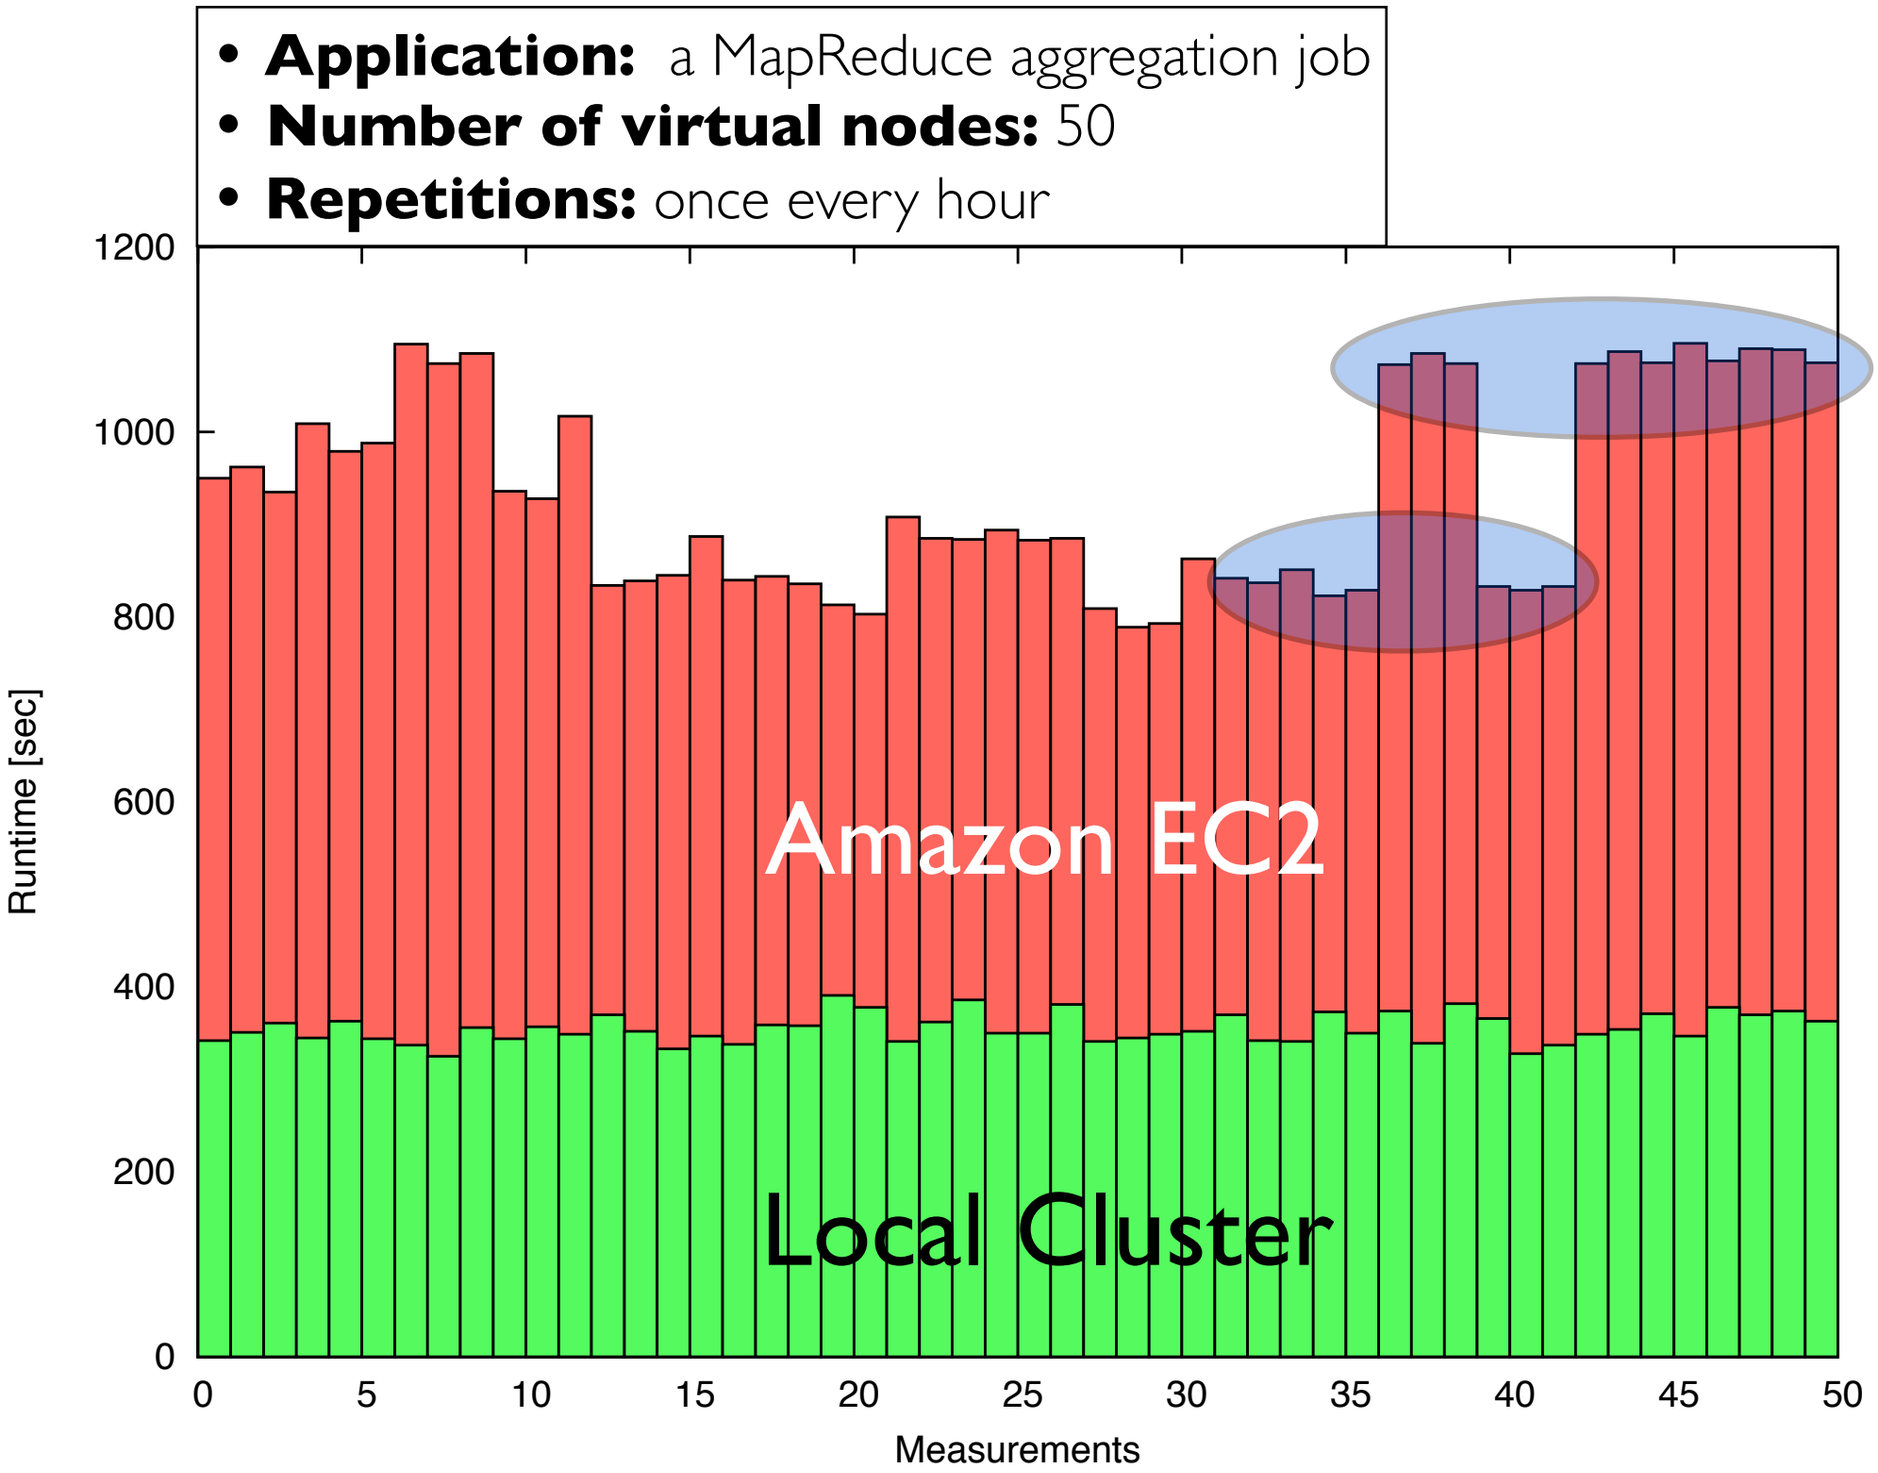

MapReduce Job

SAARLAND UNIVERSITY
Information Systems Group

- **Application:** a MapReduce aggregation job
- **Number of virtual nodes:** 50
- **Repetitions:** once every hour

Observation: upper band results from virtual clusters composed of **80%** of Xeon processors (of **20%** for the lower band).

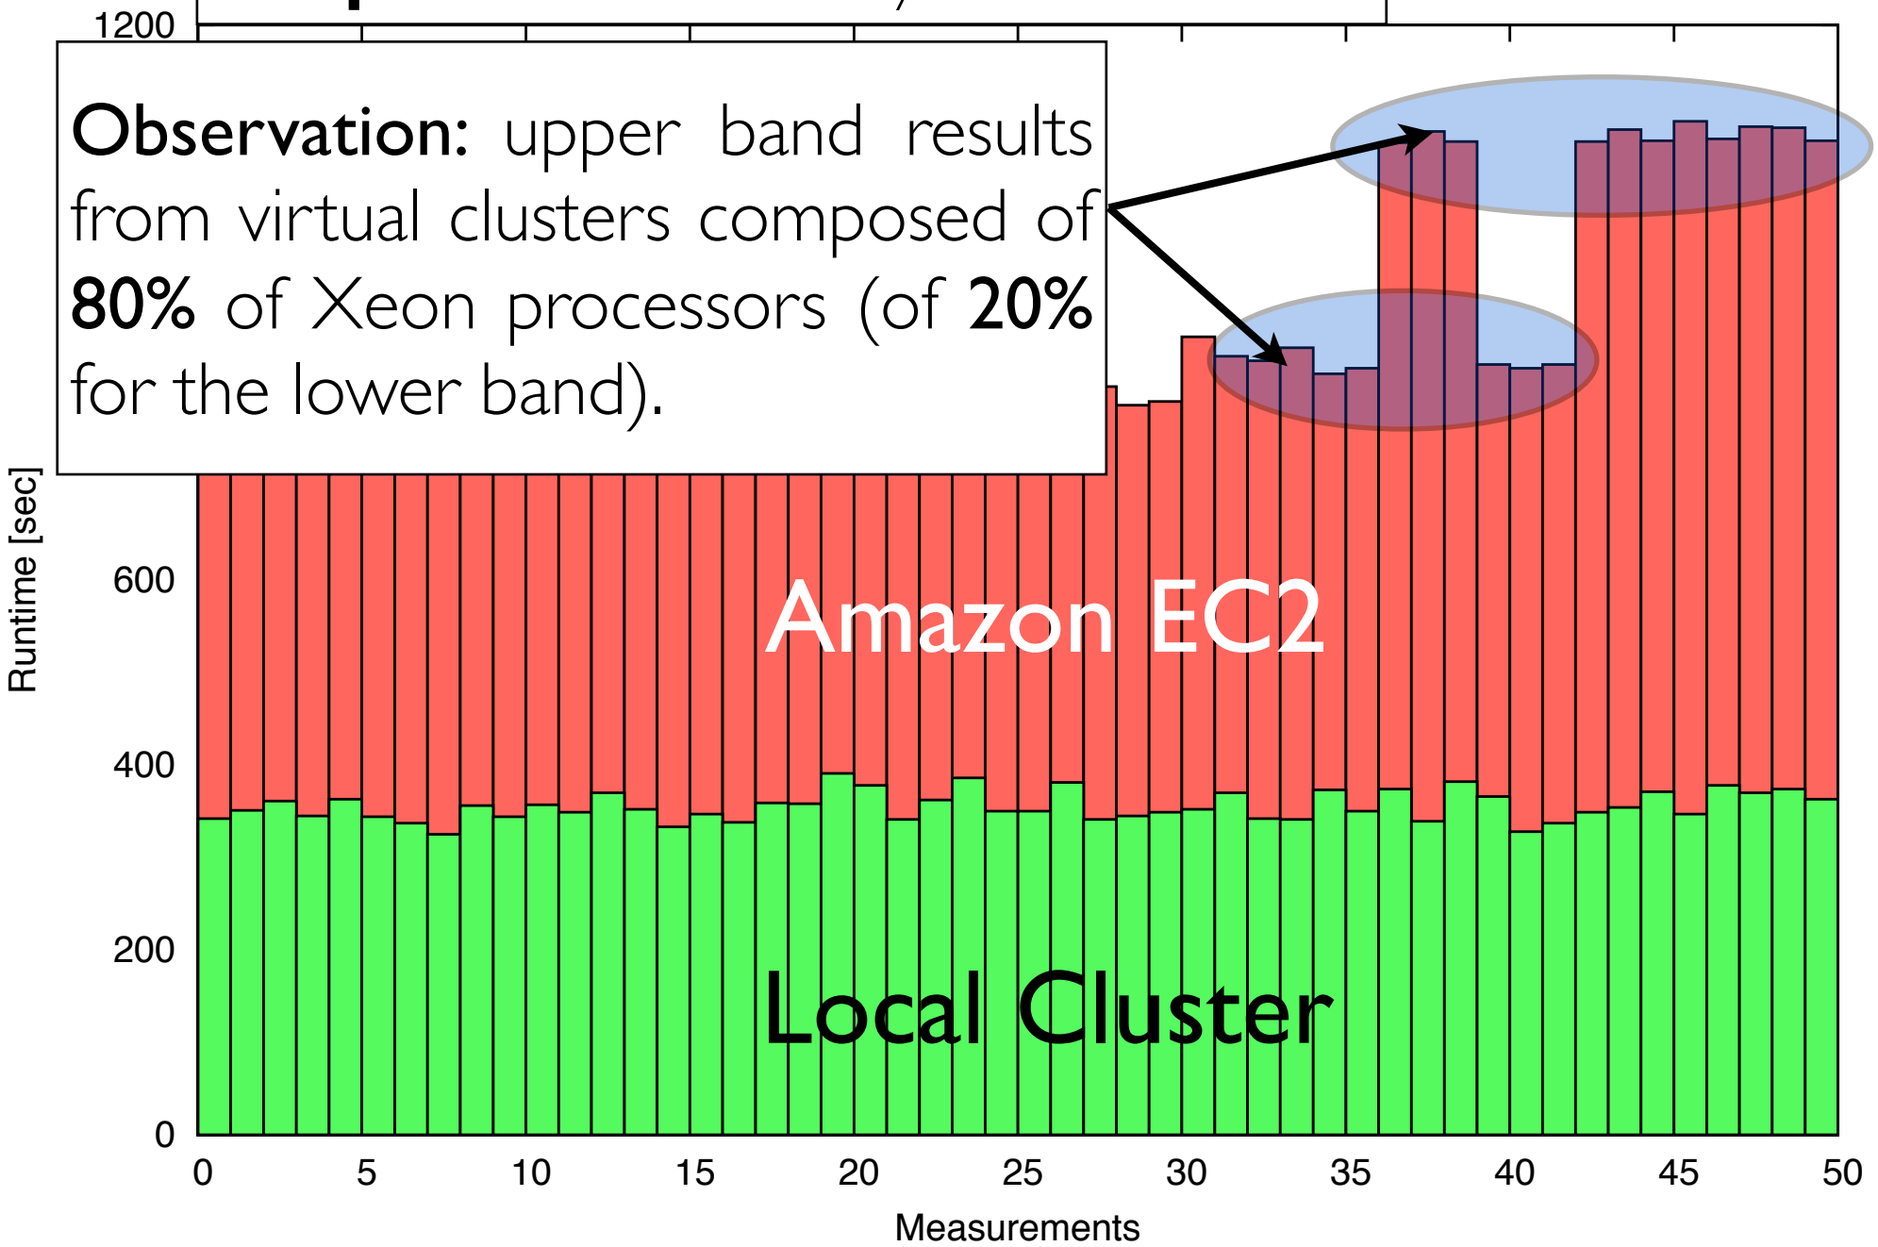

Lessons Learned

SAARLAND UNIVERSITY
Information Systems Group

Lessons Learned

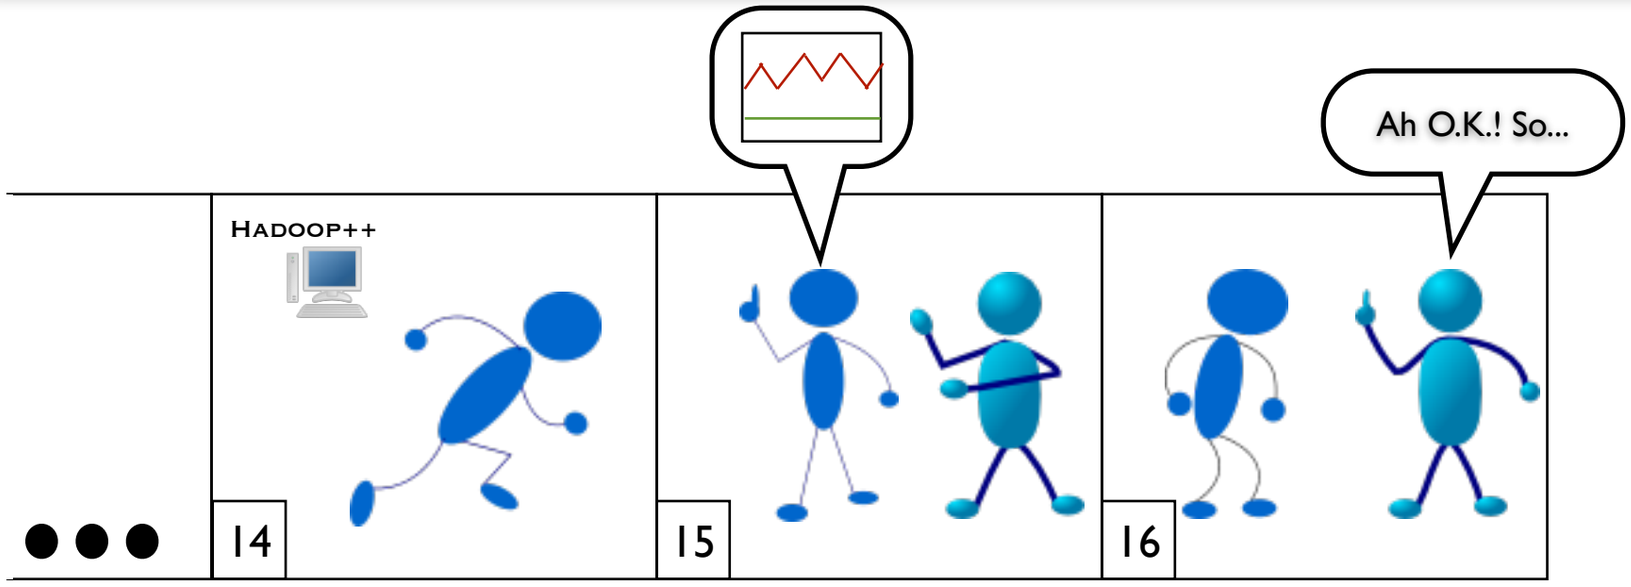

Lessons Learned

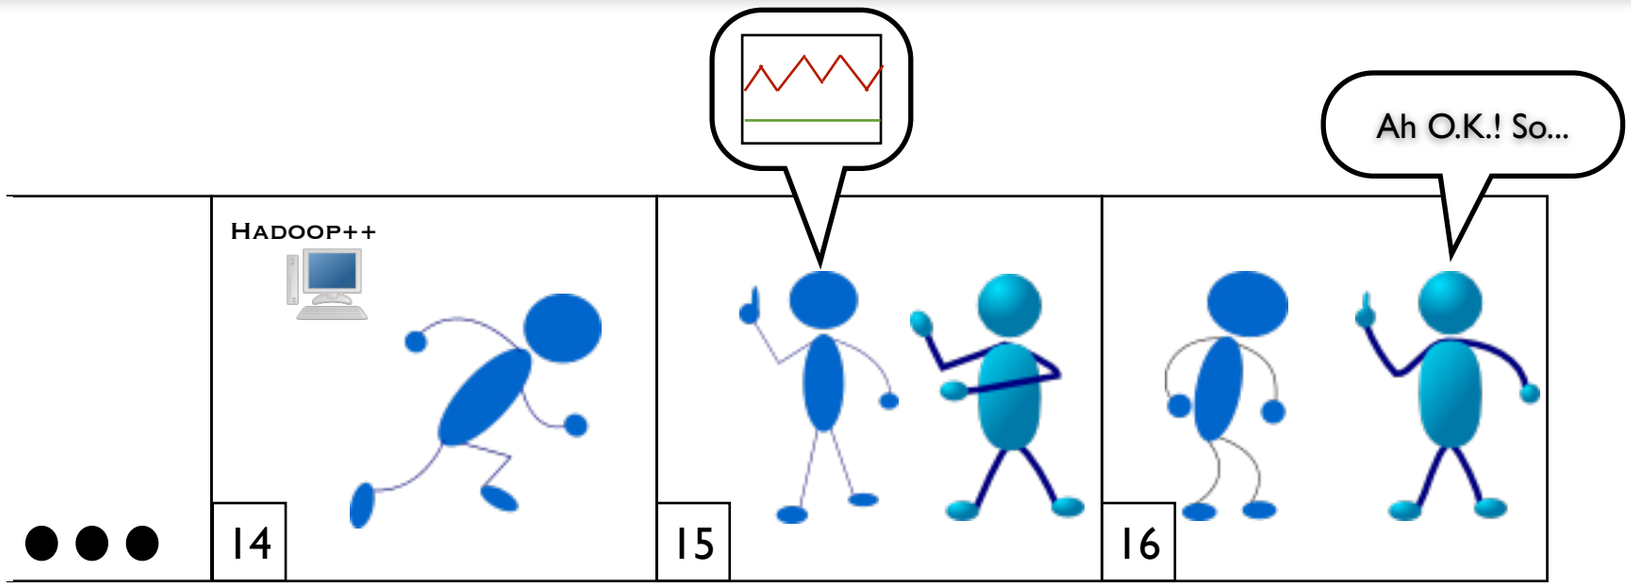

- Be careful!

Lessons Learned

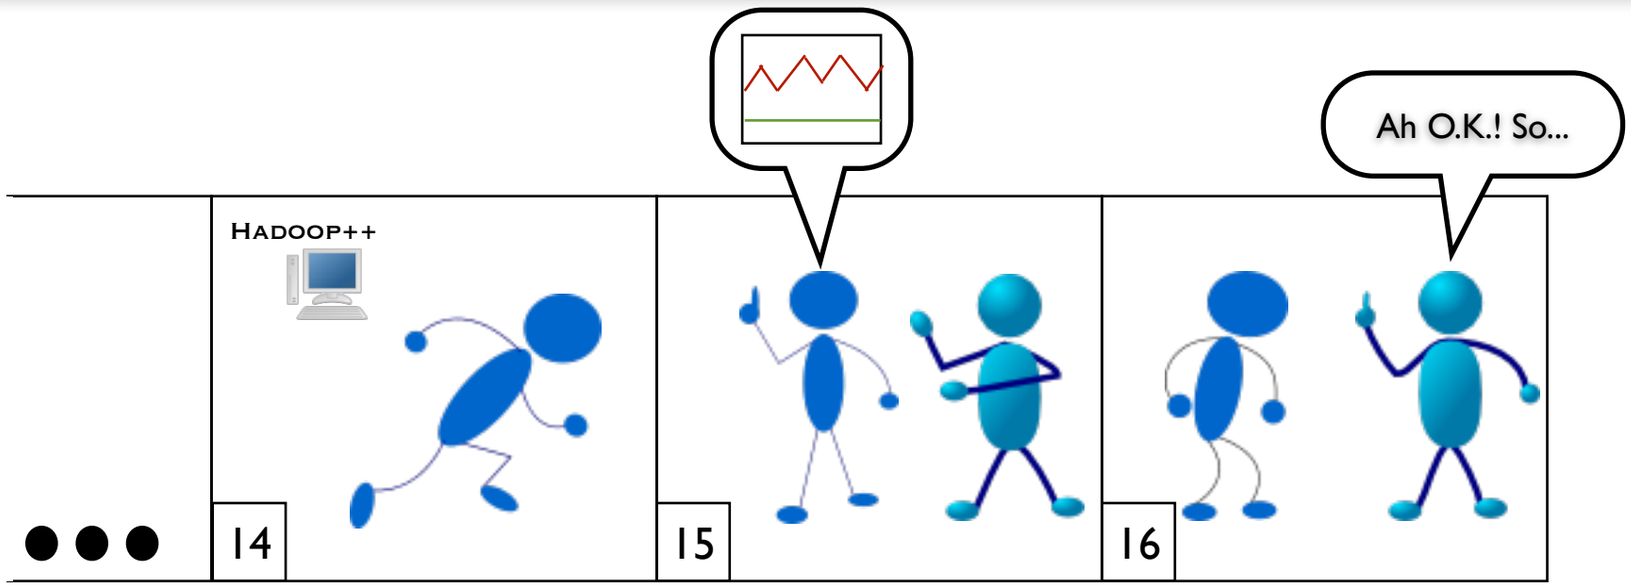

- Be careful!
- High variance in performance: COV up to **24**%
- Hard to interpret results
- Repeatability to limited extent

Lessons Learned

SAARLAND UNIVERSITY
Information Systems Group

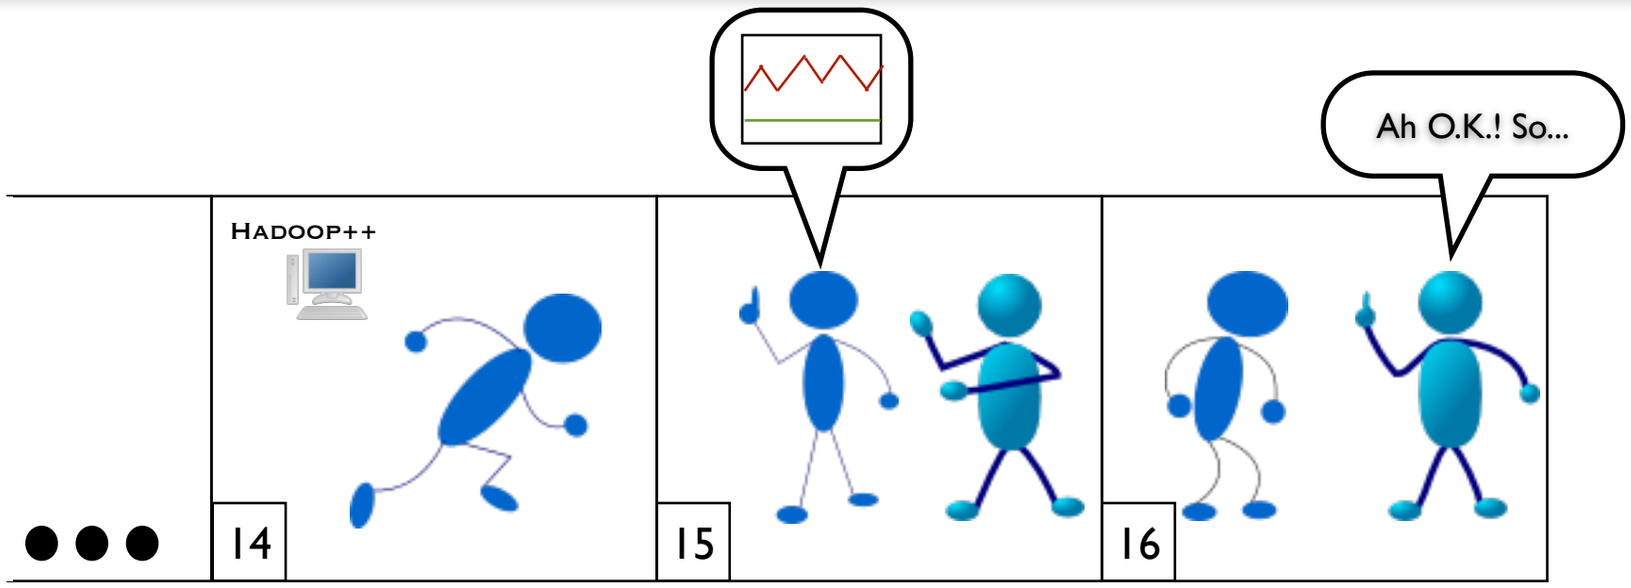

- Be careful!
- High variance in performance: COV up to **24**%
- Hard to interpret results
- Repeatability to limited extent

- **Two bands** in performance

Lessons Learned

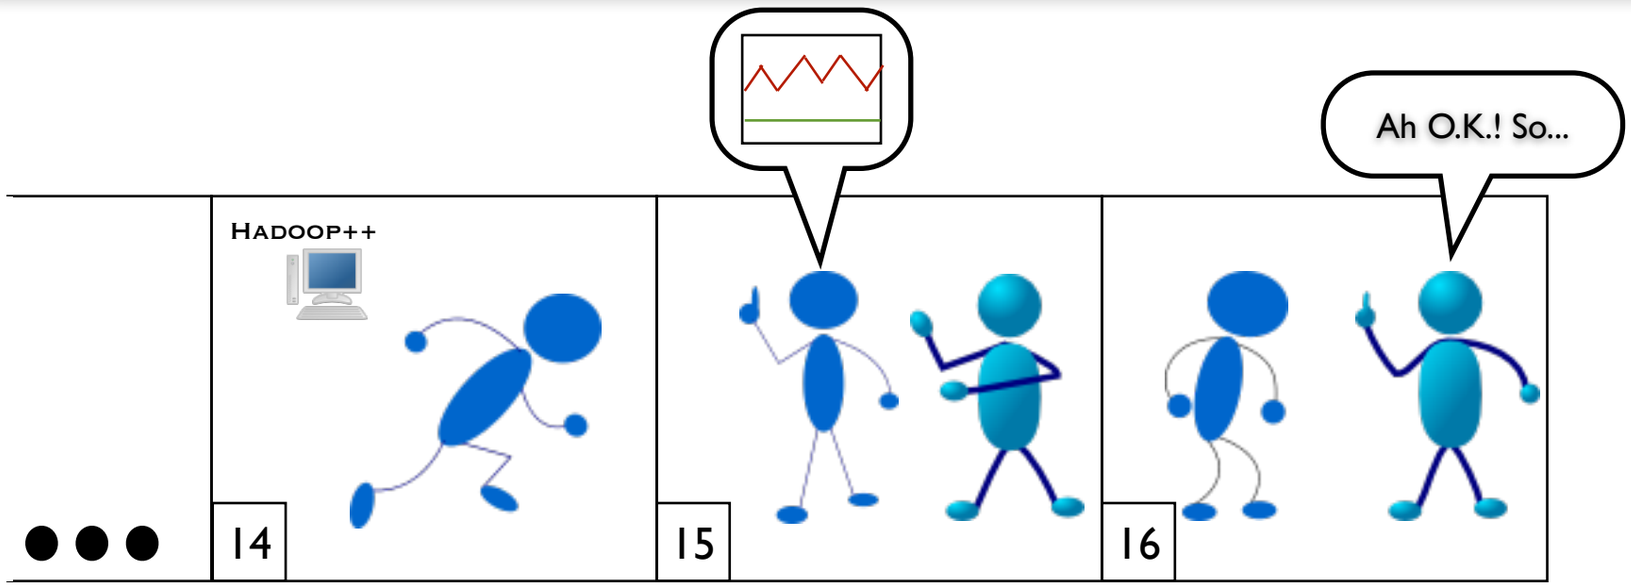

- Be careful!
- High variance in performance: COV up to **24**%
- Hard to interpret results
- Repeatability to limited extent

- **Two bands** in performance
- Partially due to different physical CPU types

Conclusion

SAARLAND UNIVERSITY
Information Systems Group

Conclusion

SAARLAND UNIVERSITY
Information Systems Group

- Amazon should:
  - reveal the **physical** details
  - **allow** users to specify physical characteristics

Conclusion

SAARLAND UNIVERSITY
Information Systems Group

- Amazon should:
  - reveal the **physical** details
  - **allow** users to specify physical characteristics
- Researchers should
  - use **equivalent** virtual clusters to **compare** systems
  - report **underlying** system type with the results

Conclusion

SAARLAND UNIVERSITY
Information Systems Group

- Amazon should:
  - reveal the **physical** details
  - **allow** users to specify physical characteristics
- Researchers should
  - use **equivalent** virtual clusters to **compare** systems
  - report **underlying** system type with the results

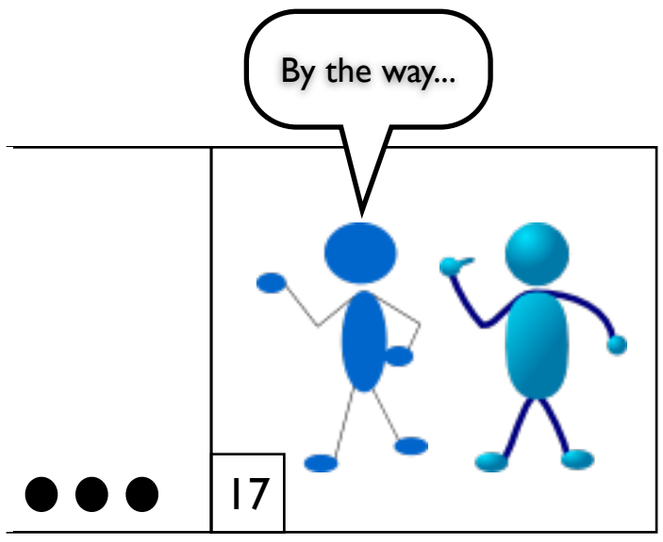

Amazon **recently** introduced the **cluster-compute** Instances

[after VLDB'10 deadline]

Conclusion

SAARLAND UNIVERSITY
Information Systems Group

- Amazon should:
  - reveal the **physical** details
  - **allow** users to specify physical characteristics
- Researchers should
  - use **equivalent** virtual clusters to **compare** systems
  - report **underlying** system type with the results

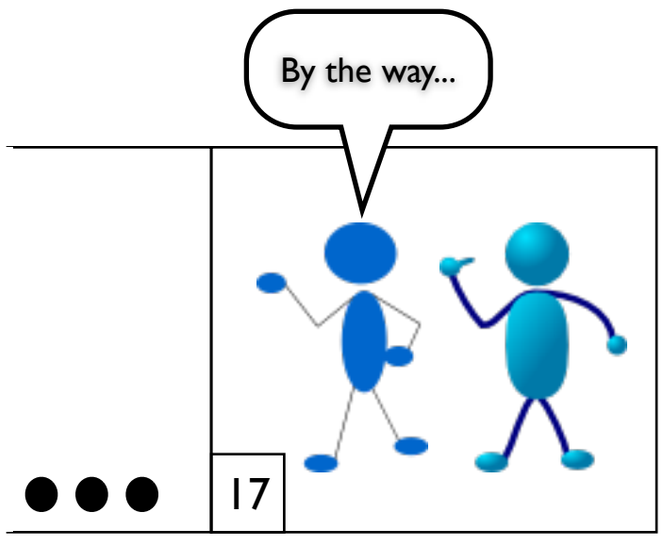

Amazon **recently** introduced the **cluster-compute** Instances

[after VLDB'10 deadline]

Still significantly higher than in a local cluster

Conclusion

SAARLAND UNIVERSITY
Information Systems Group

- Amazon should:
  - reveal the **physical** details
  - **allow** users to specify physical characteristics
- Researchers should
  - use **equivalent** virtual clusters to **compare** systems
  - report **underlying** system type with the results

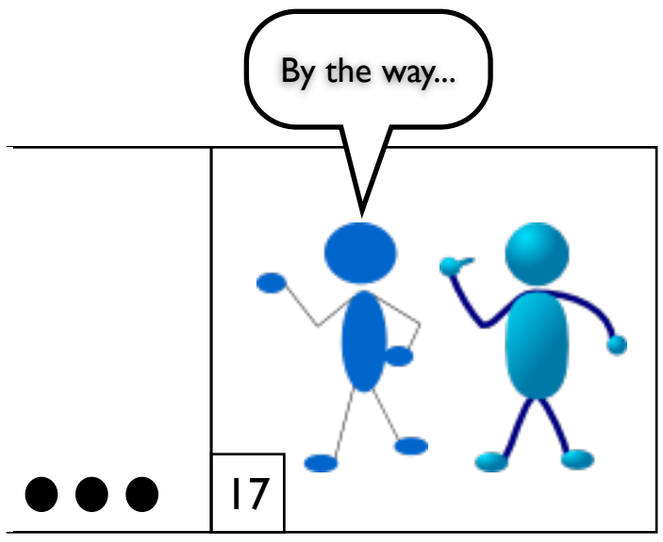

Amazon **recently** introduced the **cluster-compute** Instances

[after VLDB'10 deadline]

Still significantly higher than in a local cluster