Resource and Risk Management in Datacenters

Vincent.vanbeek@solvinity.com

- Vincent van Beek
- Distributed Systems Group TU-Delft
- Supervisor: Prof. Dr. ir. Alexandru Iosup
- Promotor: Prof. Dr. ir. Dick Epema

TU Delft VU VRIJE UNIVERSITEIT AMSTERDAM Solvinity

Why Resource and Risk Management in Datacenters

- Datacenters are at the center of all IT systems
- Hosting Business Critical Applications
- New technology is introduced at a rapid rate
- Consolidation is driving costs down
- Many enterprise customers are risk averse and want guarantees

Business-Critical Workloads

- **Business Critical-Critical Workloads are different from well known scientific workloads and grid workloads.***

- Size of the VMs
- Long running nature

Statistical Characterization of Business-Critical Workloads Hosted in Cloud Datacenters
In the IEEE/ACM CCGRID 2015 conference

*[Chen et al. (MASCOTS 2011, PVLDB 2012), Reiss et al. (SoCC 2012), Mishra et al. (SIGMETRICS 2010), Ren et al. (IISWC 2012), Di et al. (CLUSTER 2012)]

Risks for Business-Critical Workloads

Operational Risk

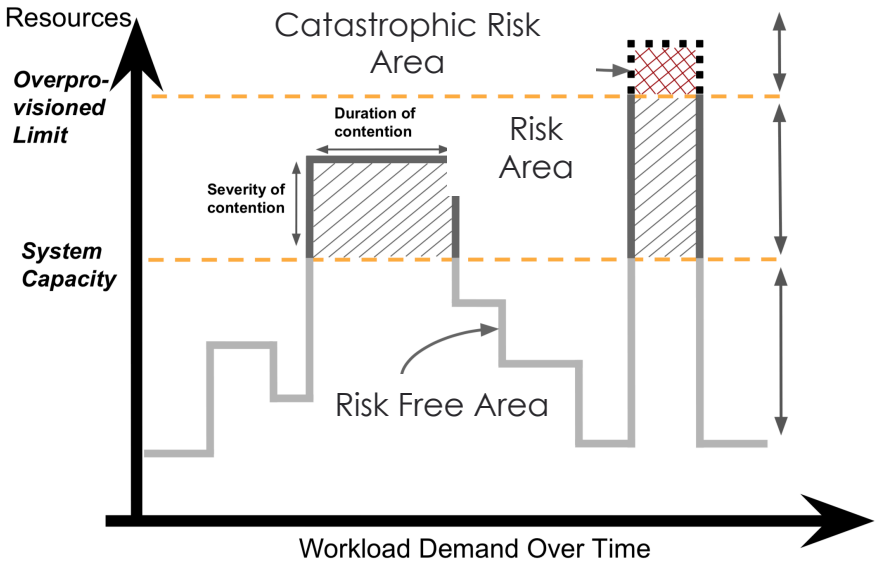

Risks for Business-Critical Workloads

Operational Risk

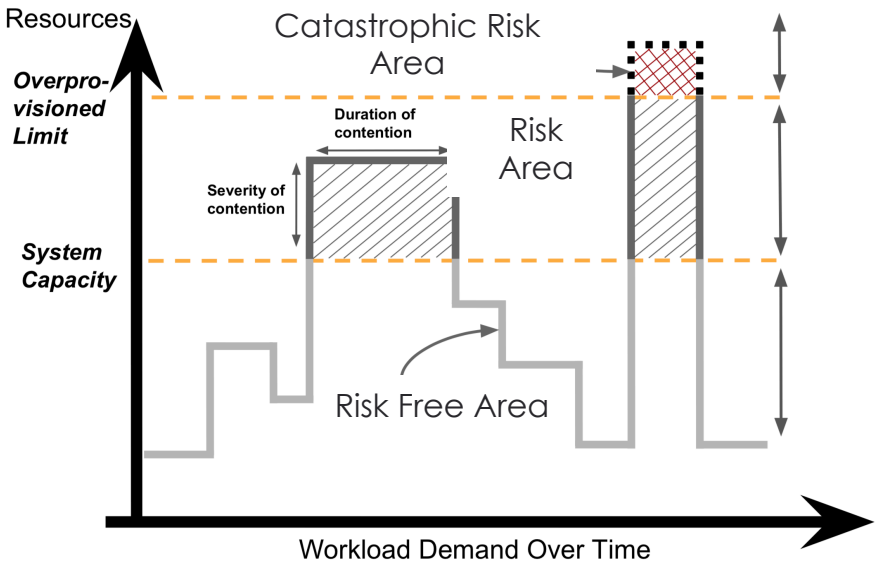

Disaster Recovery Risk

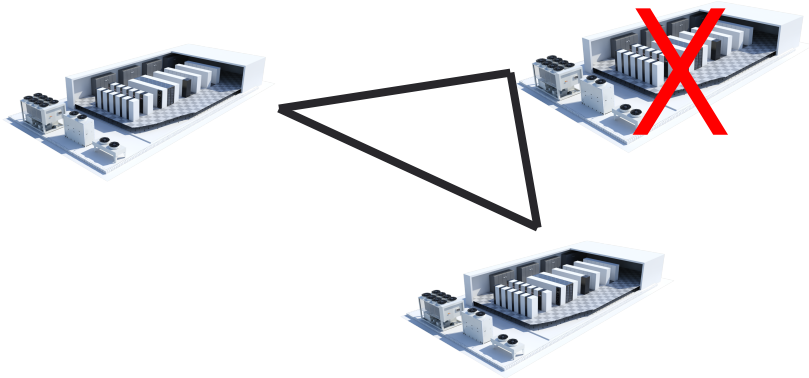

Virtualization in Datacenters

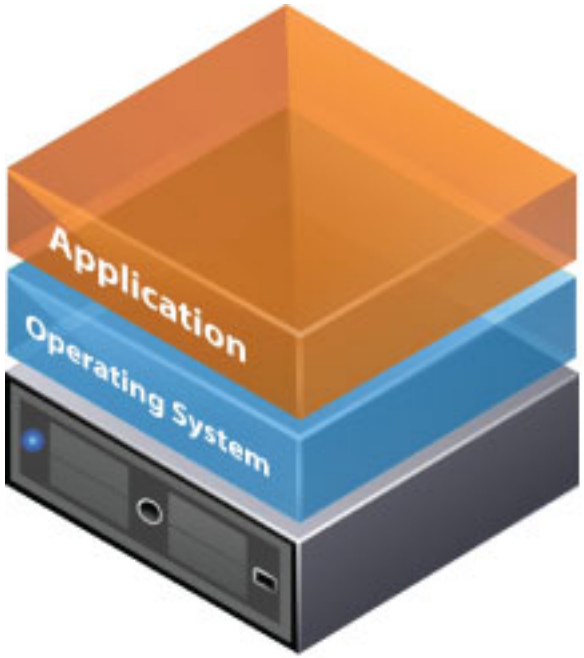

Traditional Architecture

Virtual Architecture

VM Placement

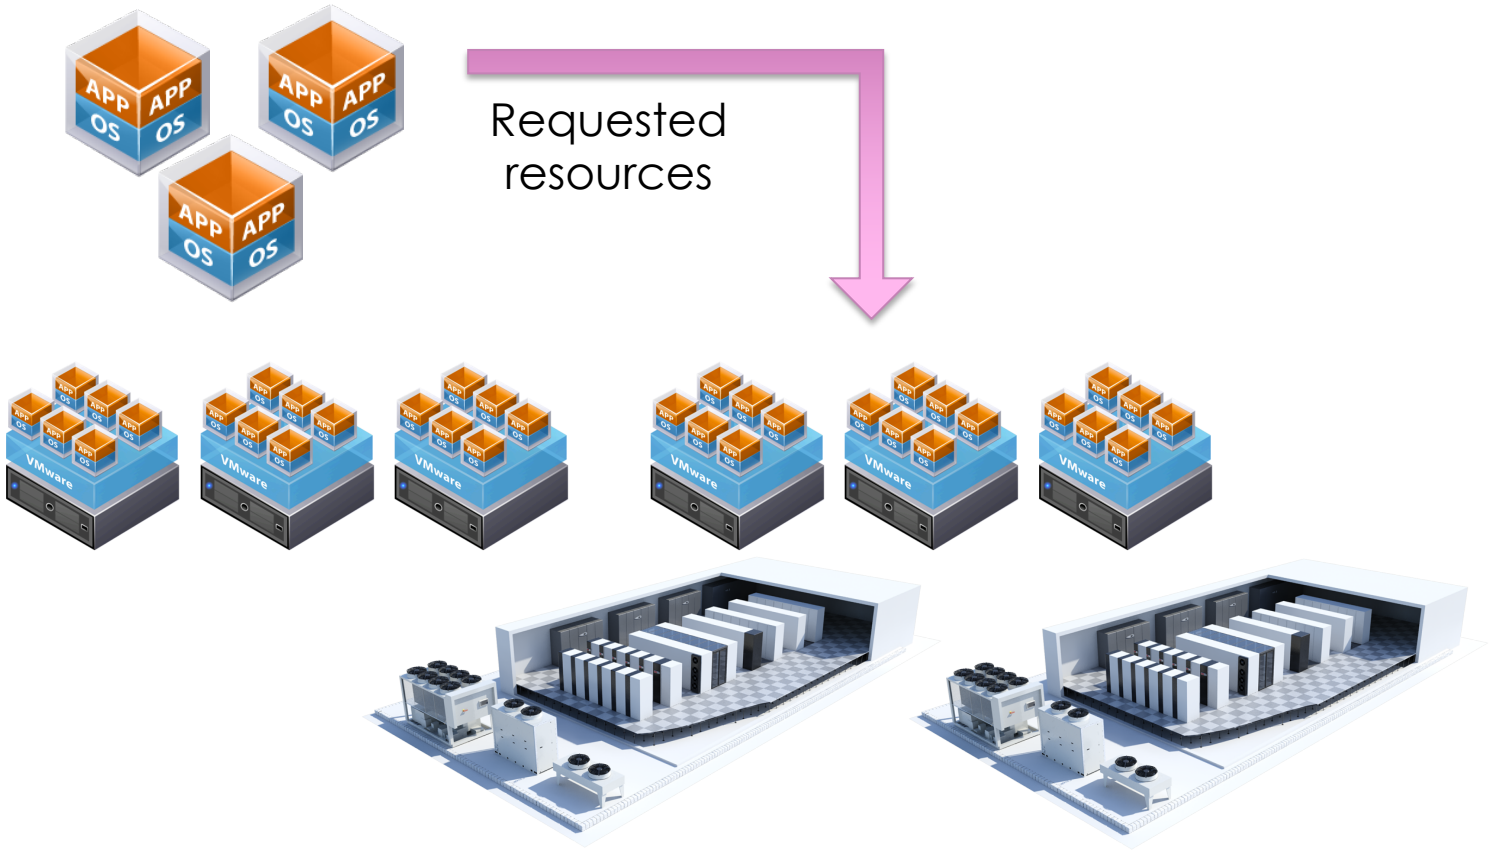

Portfolio Scheduling for Managing Operational and Disaster-Recovery Risks in Virtualized Datacenters Hosting Business-Critical Workloads, ISPDC 2019, *V. van Beek, G.* Oikonomou and A. Iosup

Contention of Resources

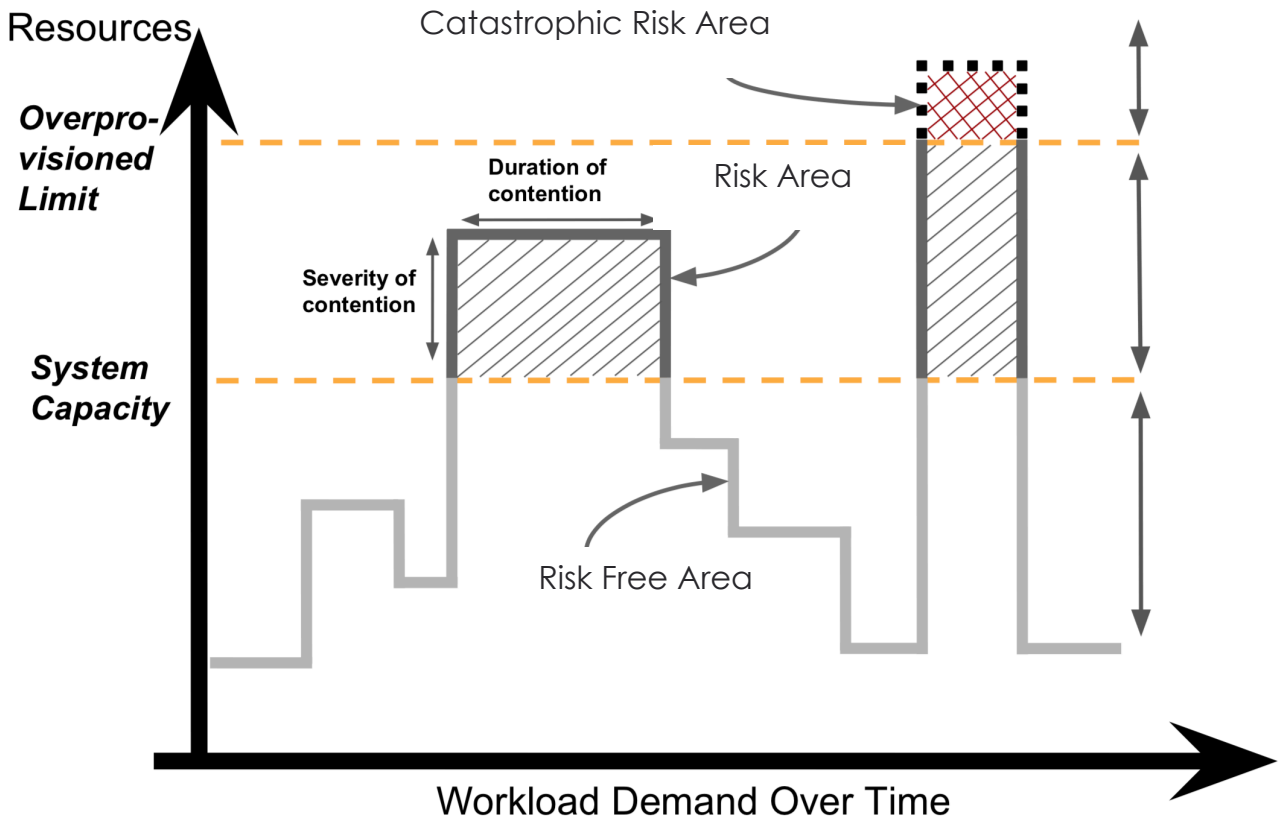

Prediction of CPU contention

3 Stage process

1. Metric selection
2. Model accuracy calculations
3. Model selection

Stage 1: Metric Selection

- **Correlation-based Selection**

  Correlation of metrics with CPU Ready

  Pearson* and *Spearman

- **25 performance metrics of VMs and of clusters**
- **Preserving the concept of *blackbox* VMs**

Stage 2: Model Accuracy Calculations

- ***Regression Models***

  Linear Regression (LR)

  Curvilinear Regression

  k- Nearest Neighbors regression (k-NN)

  Gradient Boosting regression (GB)

- **Time- series cross-validation**
- **R^2 Coefficient of determination**

Stage 3: Model Selection

- The accuracy of the model depicted by the R^2 score
- The runtime of the regression model
- The number of regressors in the model

Workload used for the evaluation

- Data from a multi datacenter production environment
- 3-month period
- 1800 VMs
- 12 clusters (200 physical hosts)

Results Stage 1: Metric Selection

- Demanded CPU of VM v at time t
- Demanded CPU of cluster c hosting VM v at time t
- Number of virtual CPU cores of VM v at time t
- CPU Ready of VM v at time t
- CPU Ready of VM v at time t - 1
- CPU Ready of VM v at time t - 2
- Mean CPU Ready of VMs in cluster c hosting VM v at time t − 1

Results Stage 2: Accuracy

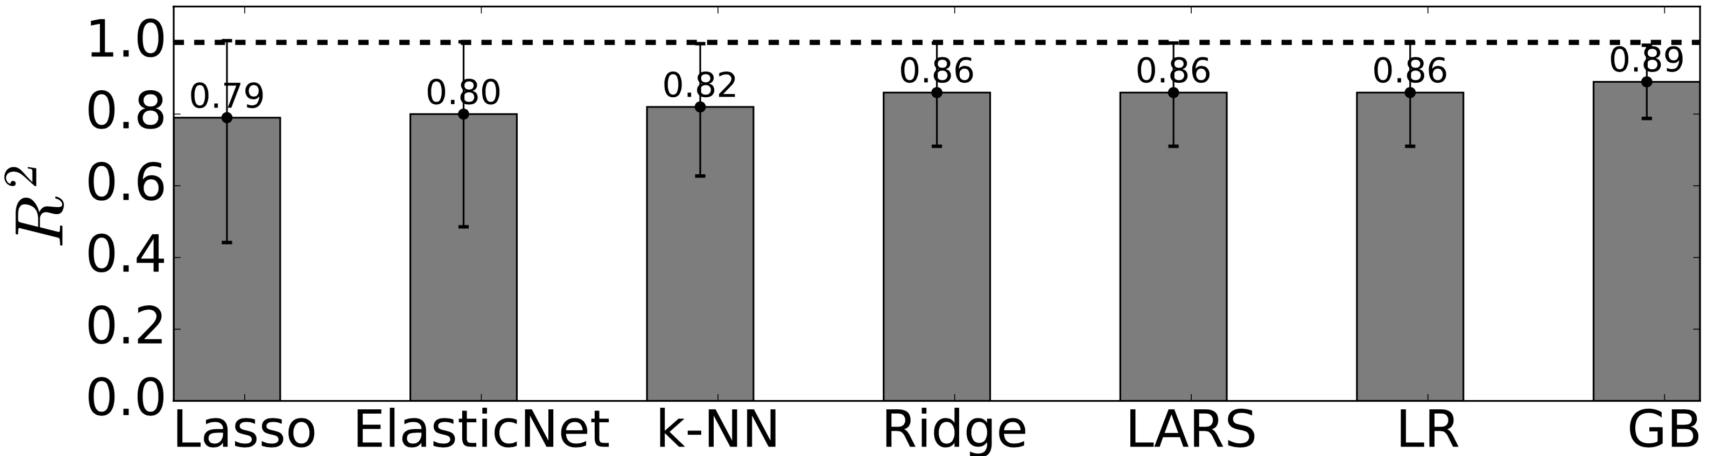

Accuracy of seven regression models.

Results Stage 2: Accuracy

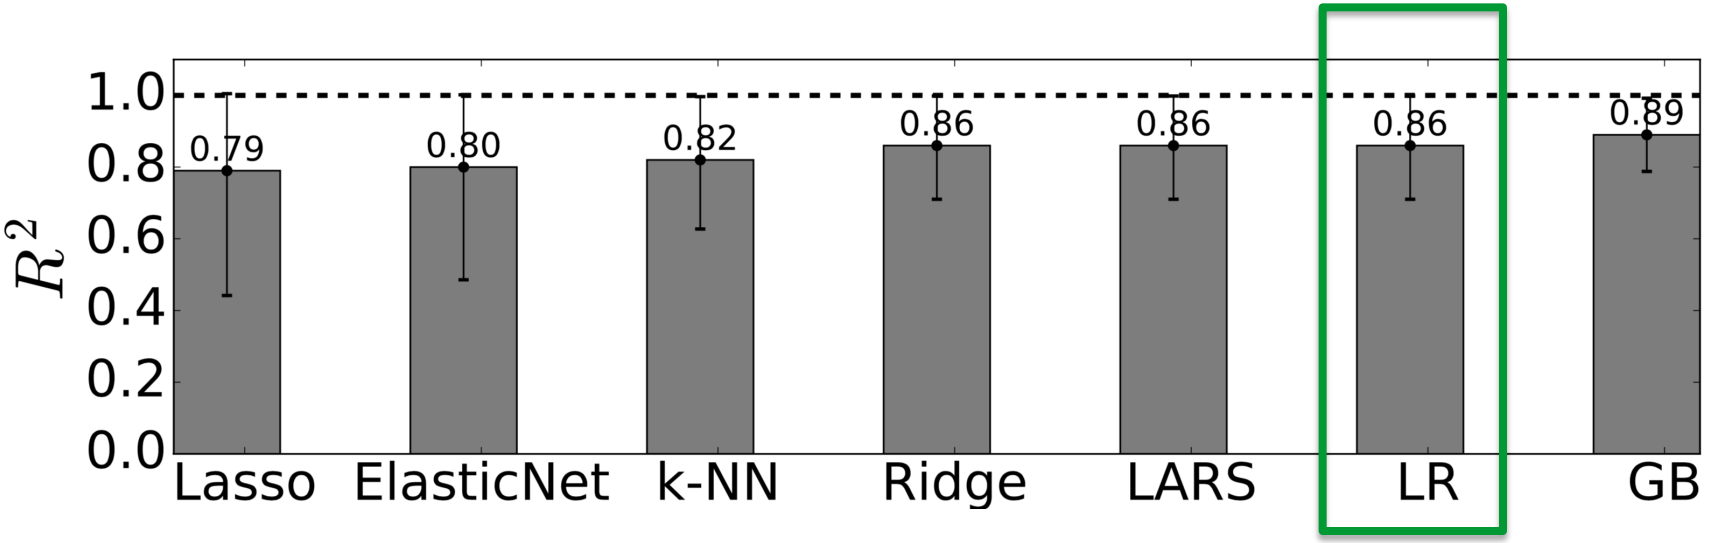

Accuracy of seven regression models.

Results Stage 3: Runtimes

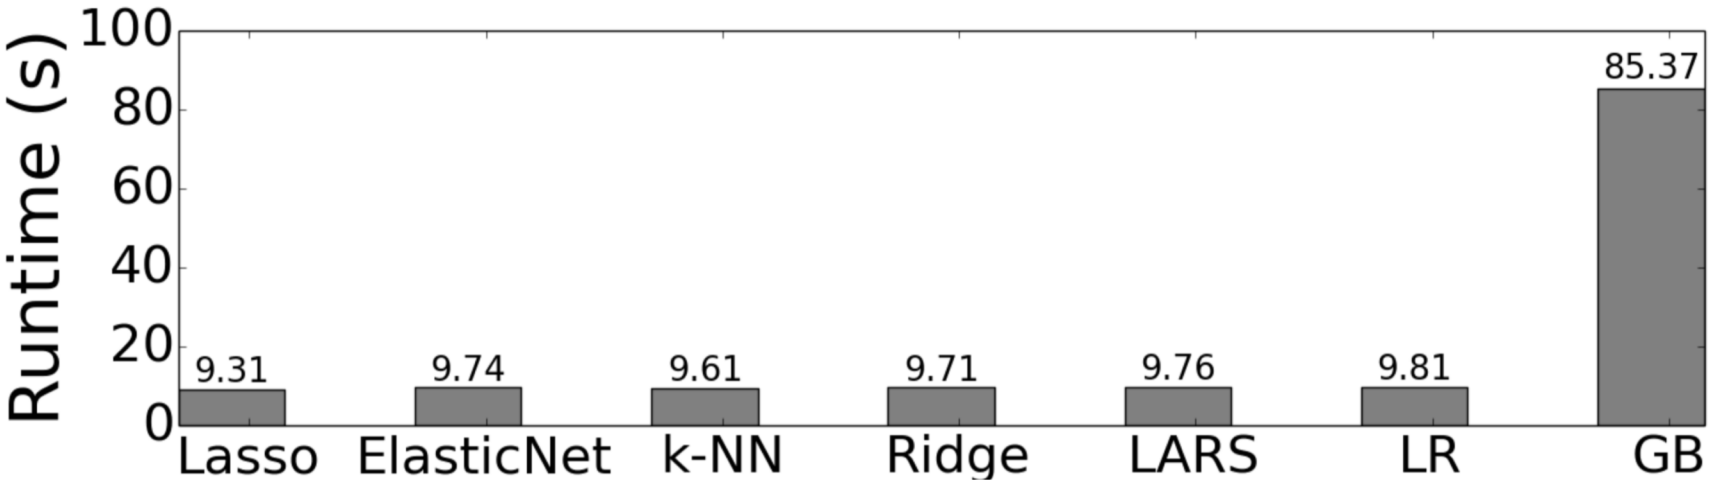

Run time of seven regression models.

Results Stage 3: Runtimes

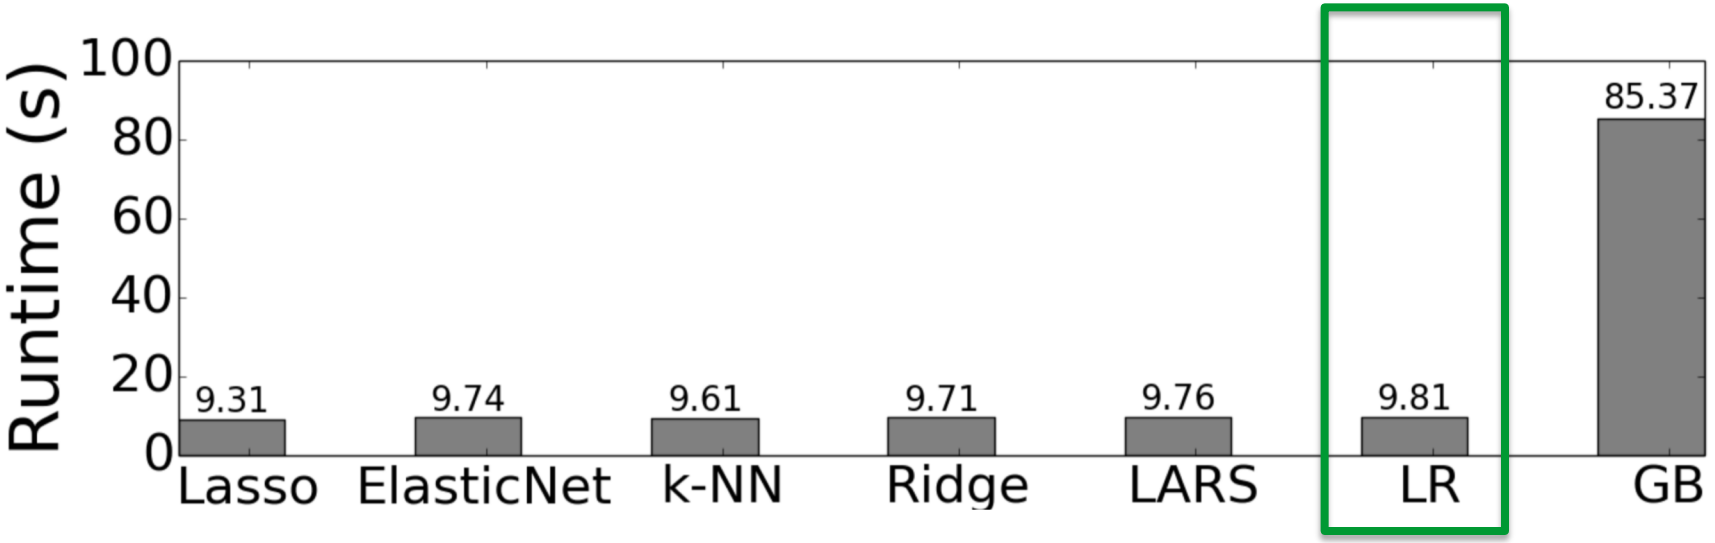

Run time of seven regression models.

Results: Verification

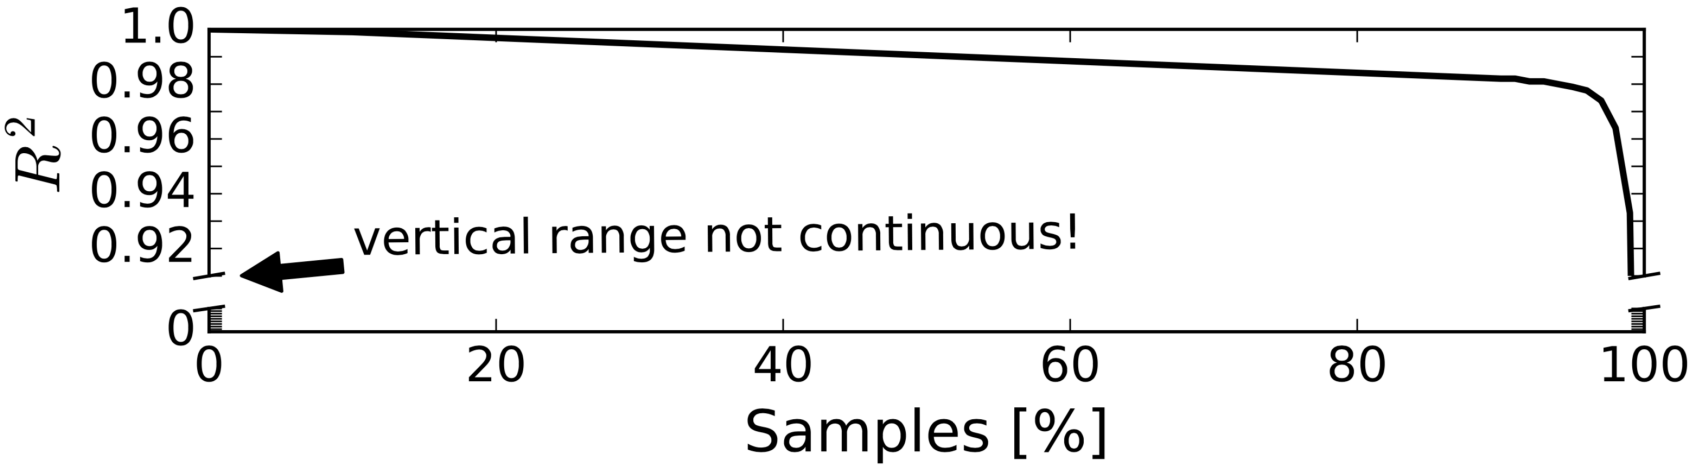

Accuracy of the tuned predictor.

The results indicate an accurate predictor of CPU-Contention for 99% of the values (0.93),

Conclusions

- 3 Stage process for selecting a contention predictor
- We find that it is possible to create an accurate CPU contention predictor in a Real world setting.
- We release workload traces that can be used by others to build predictors and or schedulers.

Workload Traces

- **http://gwa.ewi.tudelft.nl**
- Last year we added new traces from a datacenter in Germany.
- This year we will release new traces used for our research on CPU contention

OpenDC for datacenter simulation

1. Datacenter Technology & Methods

- Risk Analysis + Management
- Heterogeneity
- Efficiency → SME Availability

2. Education Practices

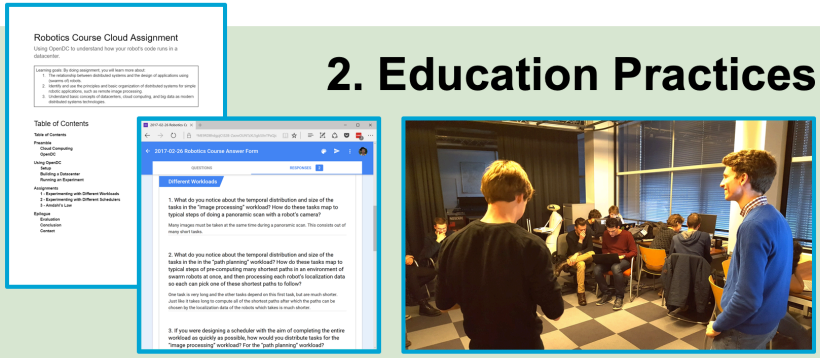

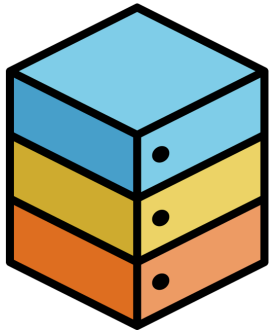

opendc.org

3. Software & Data Artifacts

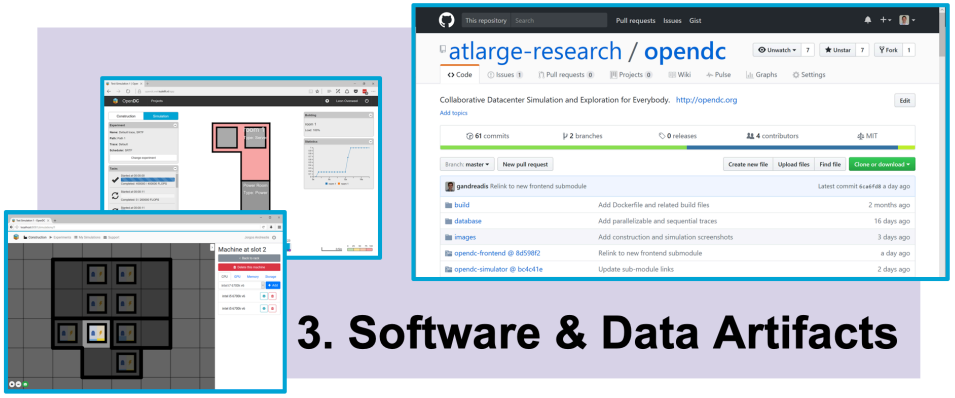

@Large Research
Massivizing Computer Systems

TU Delft VU Vrije Universiteit Amsterdam Solvinity This is IT

Vincent van Beek Georgios Andreadis
Giorgos Oikonomou Alexandru Iosup
OpenDC team

Try **OpenDC** online!

Help us: We need more workload traces to better understand resource usage in datacenters

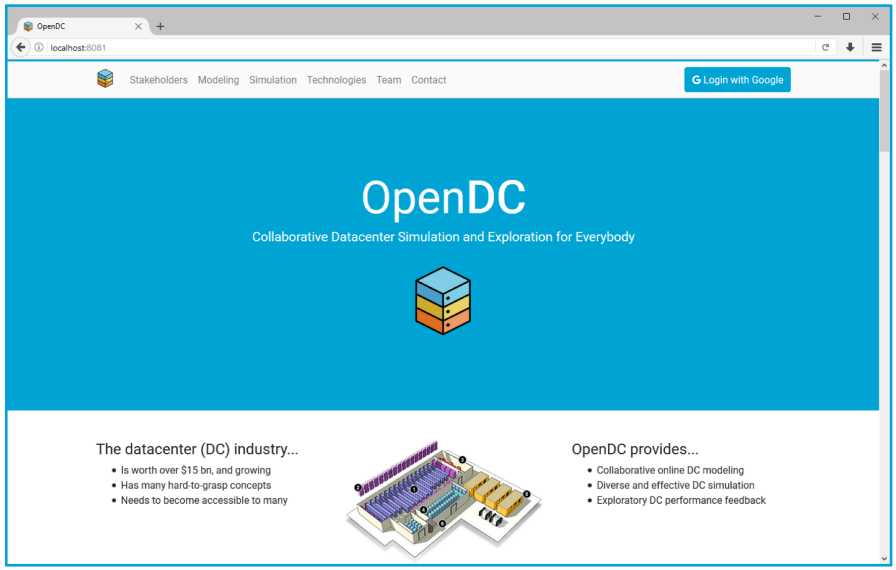

opendc.org

github.com/atlarge-research/opendc

Vincent van Beek Georgios Andreadis
Giorgos Oikonomou Alexandru Iosup
OpenDC team

Resource and Risk Management in Datacenters

- Vincent van Beek

- Vincent.vanbeek@solvinity.com
- linkedin.com/in/vsvanbeek

TU Delft VU VRIJE UNIVERSITEIT AMSTERDAM Solvinity

Recent work

G. Andreadis, L. Versluis, F. Mastenbroek, A. Iosup
"*A Reference Architecture for Datacenter Scheduling: Design, Validation, and Experiments*"
SC18

V. van Beek, J. Donkervliet, T. Hegeman, S. Hugtenburg, and A. Iosup
"*Self-expressive Management of Business-critical Workloads in Virtualized Datacenters*"
IEEE Computer 2015

V. van Beek , G. Oikonomou , A. Iosup
"*Portfolio Scheduling for Managing Operational and Disaster-Recovery Risks in Virtualized Datacenters Hosting Business-Critical Workloads*"
ISPDC 2019

V. van Beek , G. Oikonomou , A. Iosup
"*A CPU Contention Predictor for Business-Critical Workloads in Cloud Datacenters*"
HotCloudPerf 2019