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

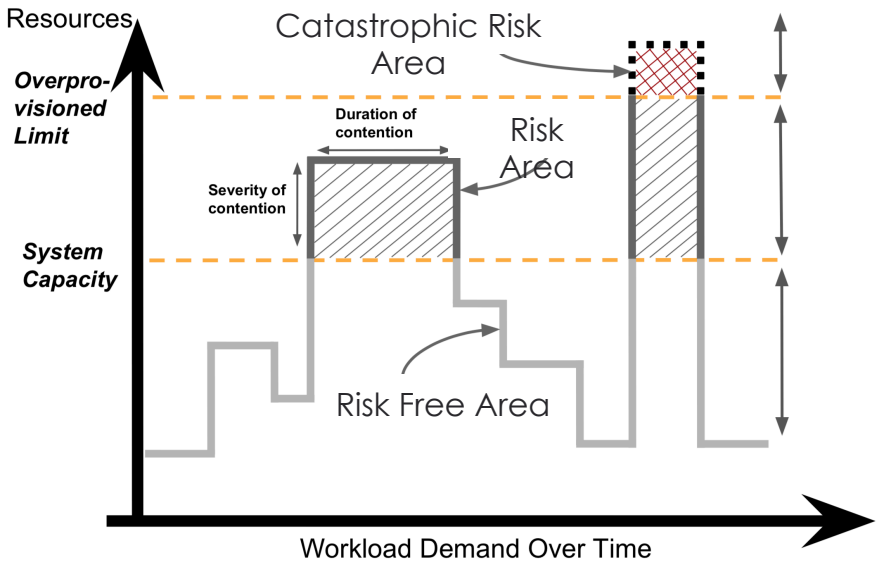

Risks for Business-Critical Workloads

Operational Risk

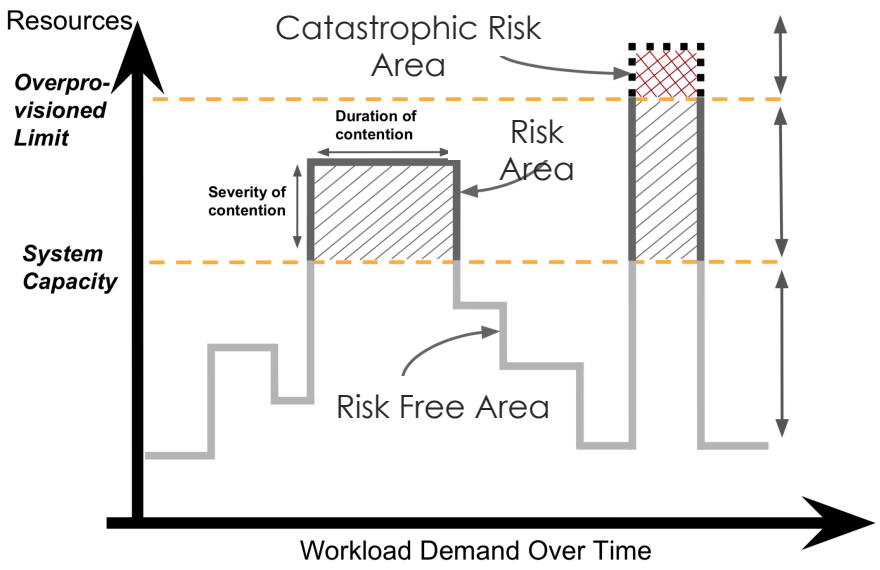

Disaster Recovery Risk

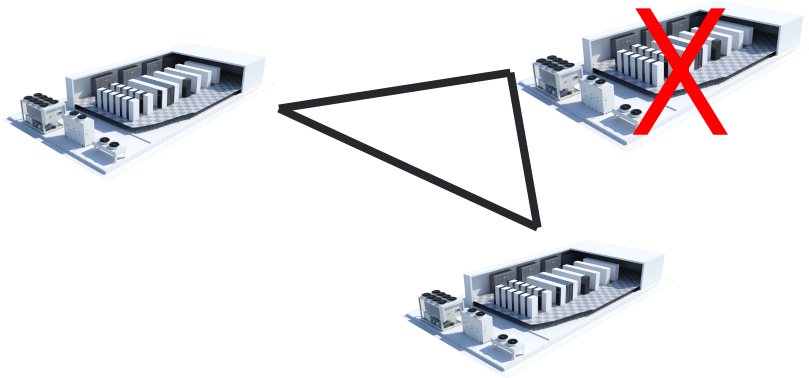

Virtualization in Datacenters

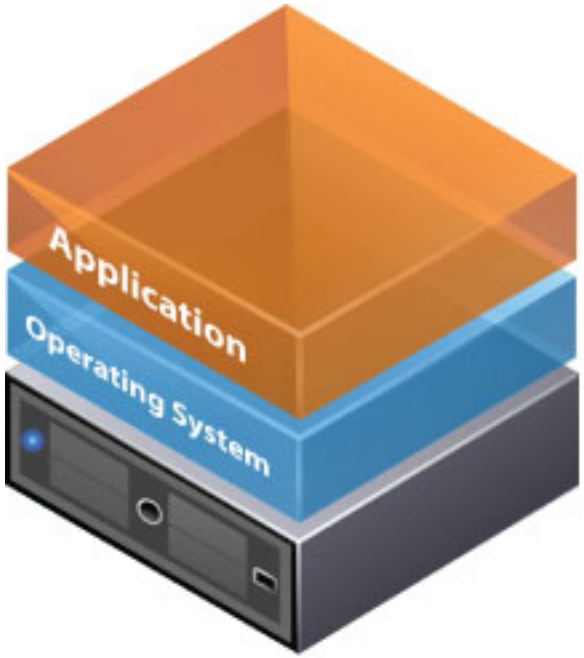

Traditional Architecture

Virtual Architecture

VM Placement

Requested resources

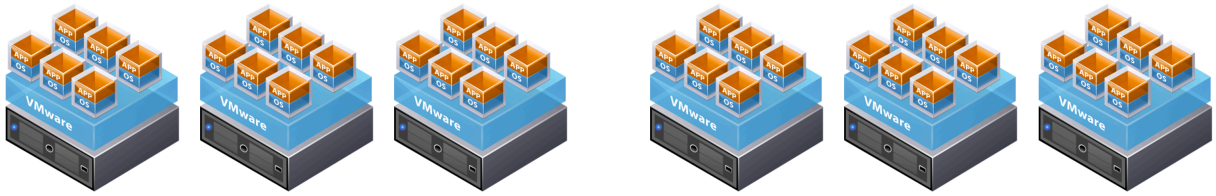

Used resources

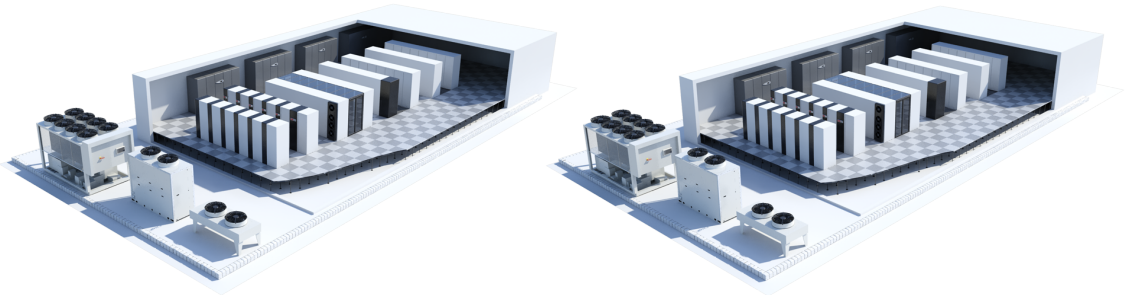

Problem statement

Affinity

SLAs

Dynamic Characteristics

Multiple Resources

Where to put which VM?

Time Component

Many Single Resource Solutions

Scheduling

Workloads

Many Point Solutions

Anti-Affinity

BlackBox Problem

Why Portfolio Scheduling

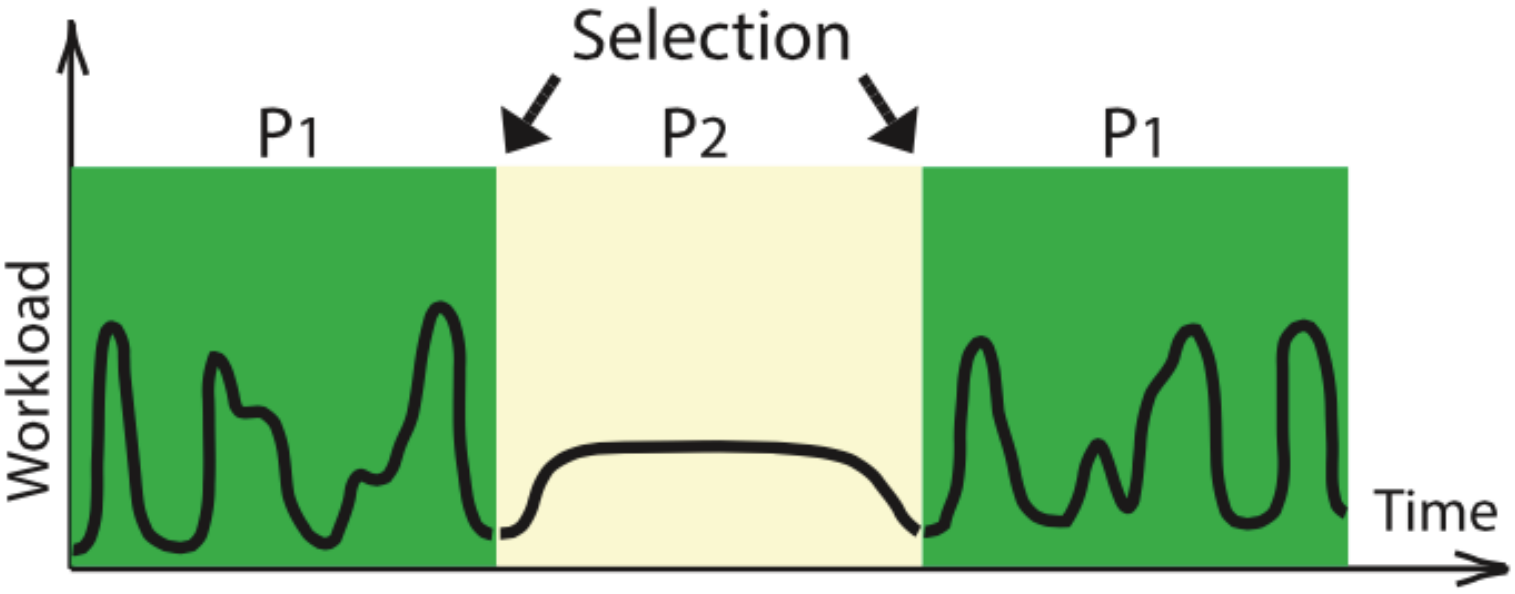

Source: Deng et al. *A periodic portfolio scheduler for scientific computing in the data center*

What is Portfolio Scheduling?

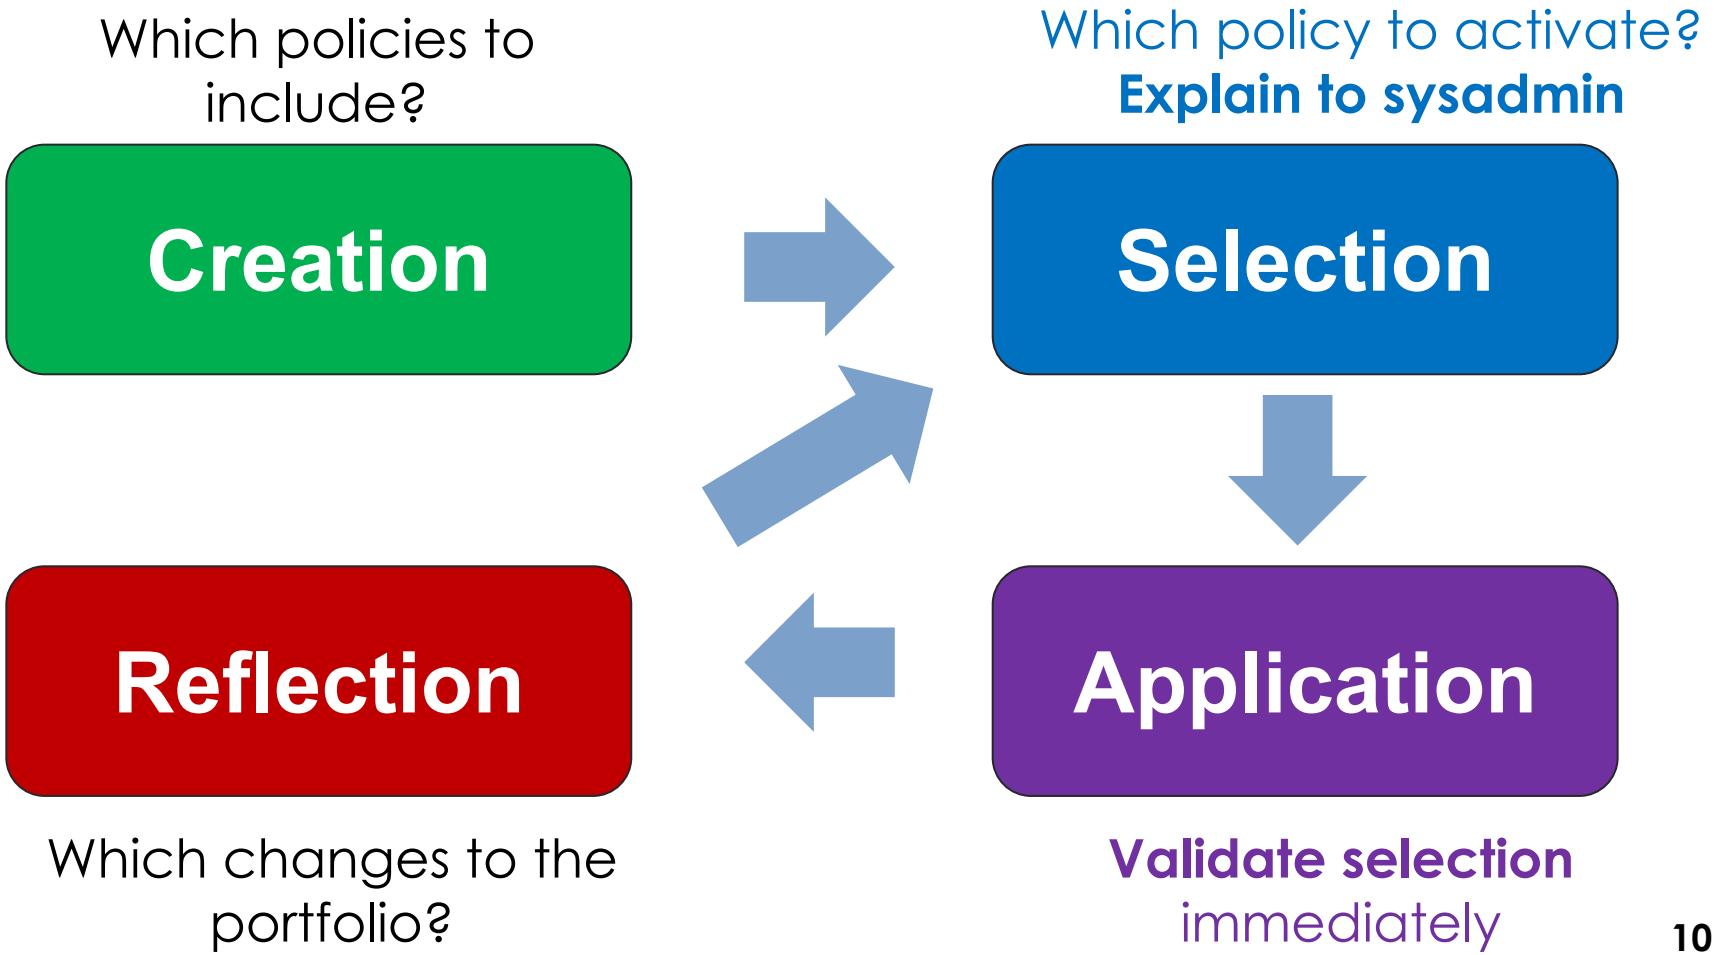

Portfolio Scheduling for Risk Management in Datacenter running Business-Critical Workloads

- Risk predictor for operational risk
- Risk aware utility functions
- Risk aware scheduling policies
- Tradeoff analysis

Architecture

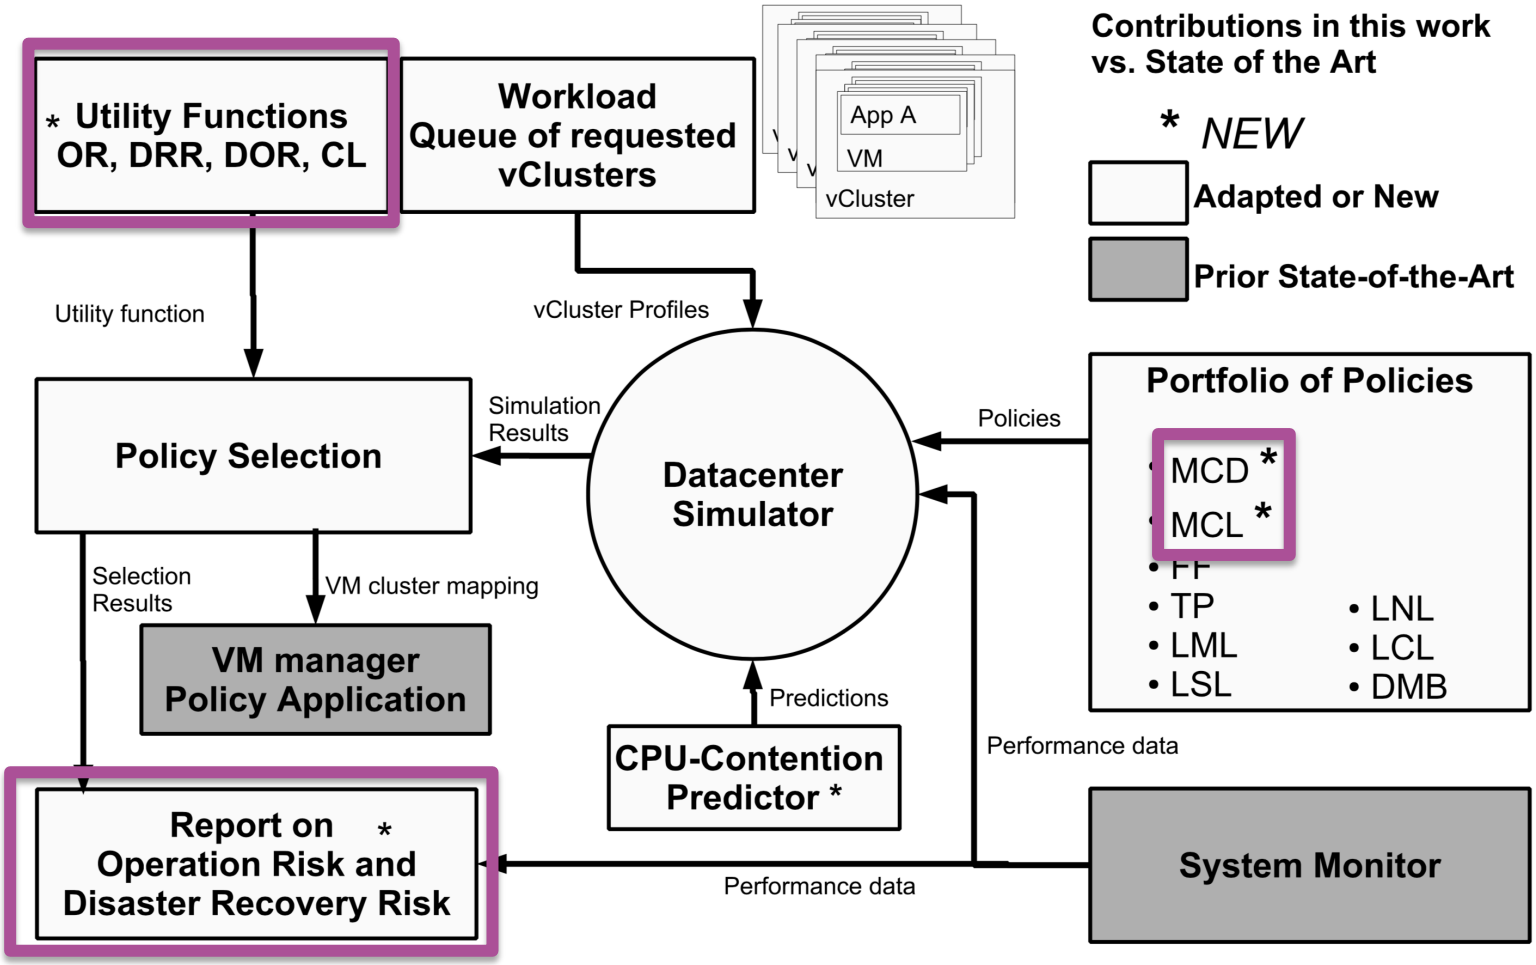

Selection Policies

- **New Utility functions**

  Oprational Risk (OR)

  Disaster recovery Risk (DRR)

  OR and DRR combined (DOR)

- **Existing Utility functions**

  Cluster Load (CL)

Scheduling Algorithms

- **Two new algorithms**
  - **Mean Contention Duration (MCD)**
  - **Maximum Consolidation Load (MCL)**
- **Existing algorithms**
  - **First-Fit (FF)**
  - **Type priority (TP)**
  - **Lowest Memory Load (LML)**
  - **Lowest Storage Load (LSL)**
  - **Lowest Network Load (LNL)**
  - **Lowest Cluster Load (LCM)**
  - **Datacenter Memory Balance (DMB)**

Datacenter Simulation Experiments

- **Why simulation**

  ~~**Real world**~~

  ~~**Modeling and Emulation**~~

  Simulation

- **Workload traces**

  - Business Critical Workloads

  - 3 month real world traces

  - 1000 VMs

  - Utilization is low compared to scientific workloads

  - Peak loads can be up to 10.000 times higher than mean the mean of the workload

Distribution of selected scheduling policies

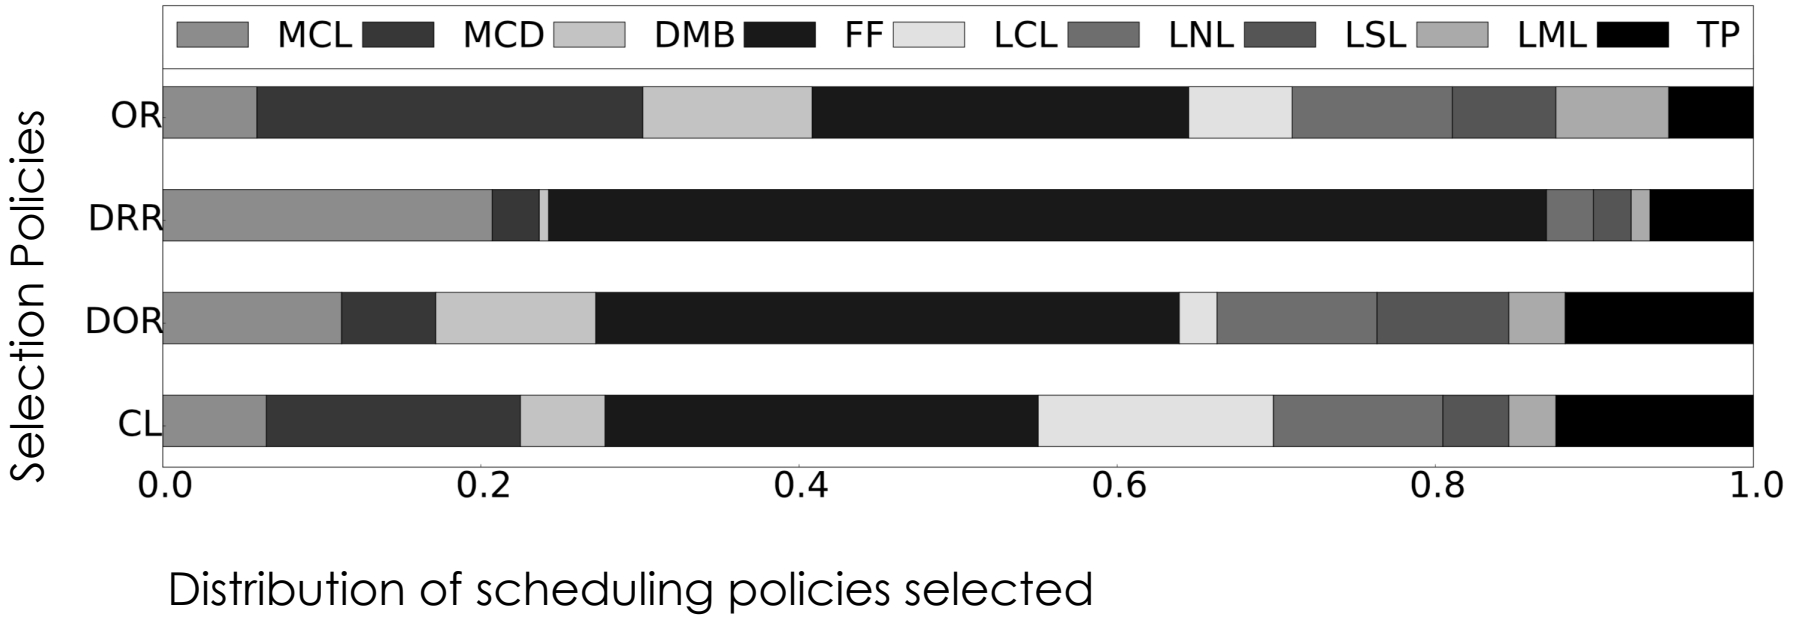

DRR

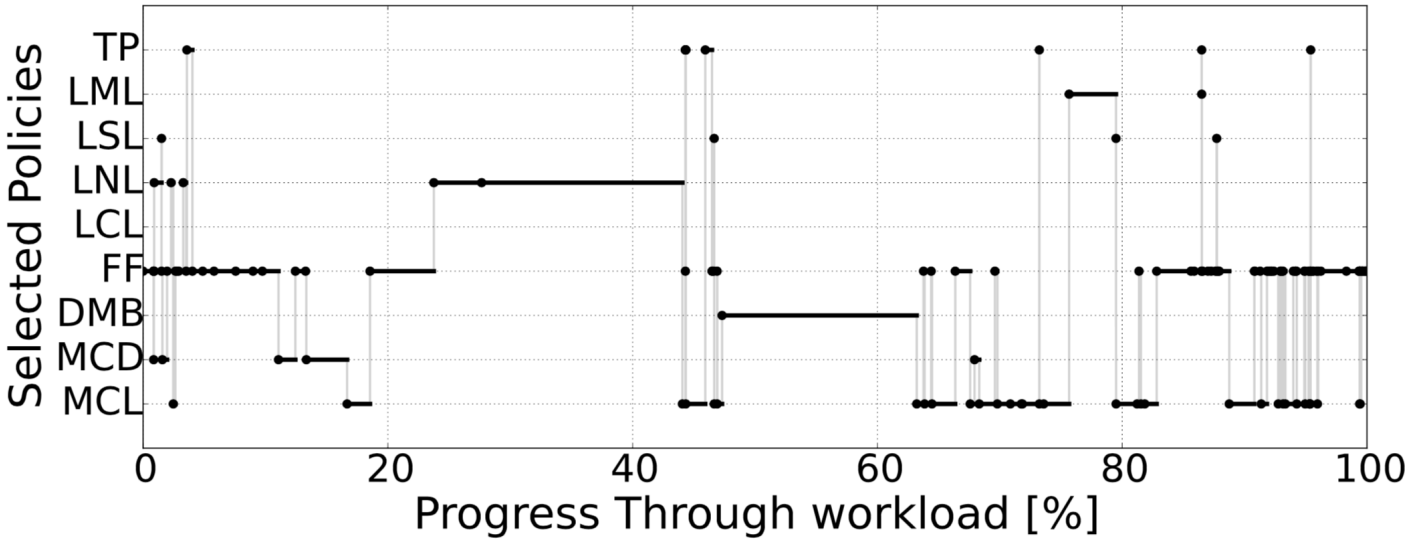

(a) Selected scheduling policies for the DRR utility function.

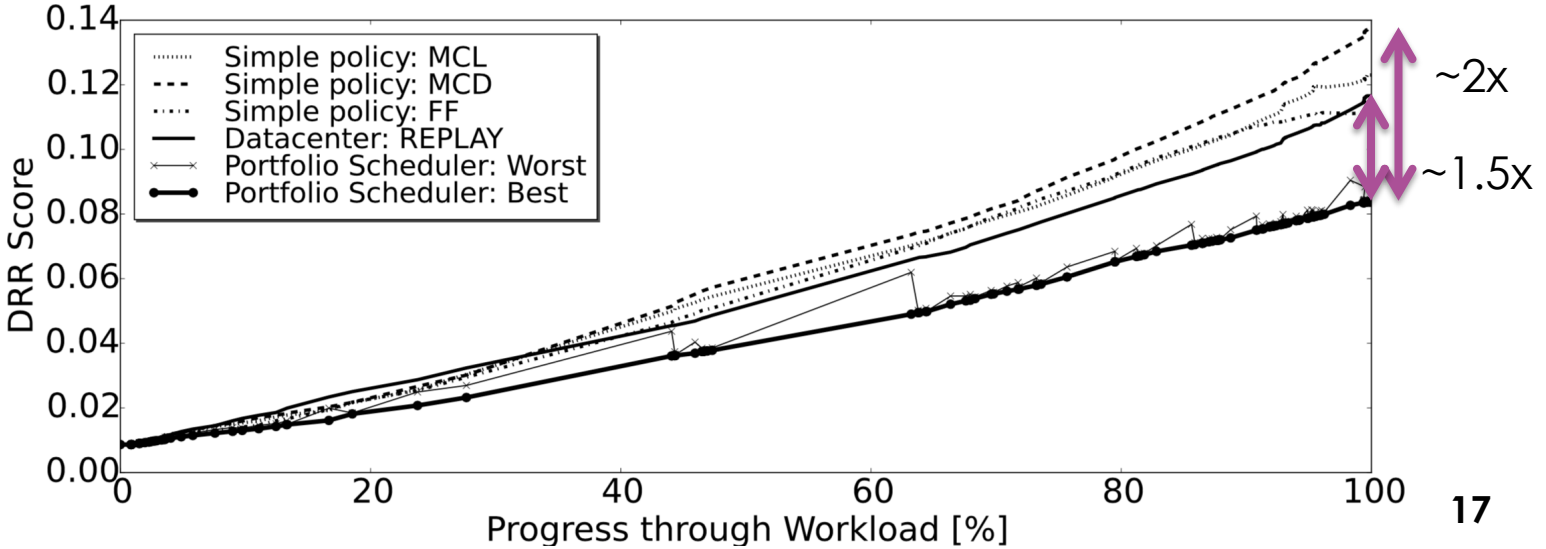

OR

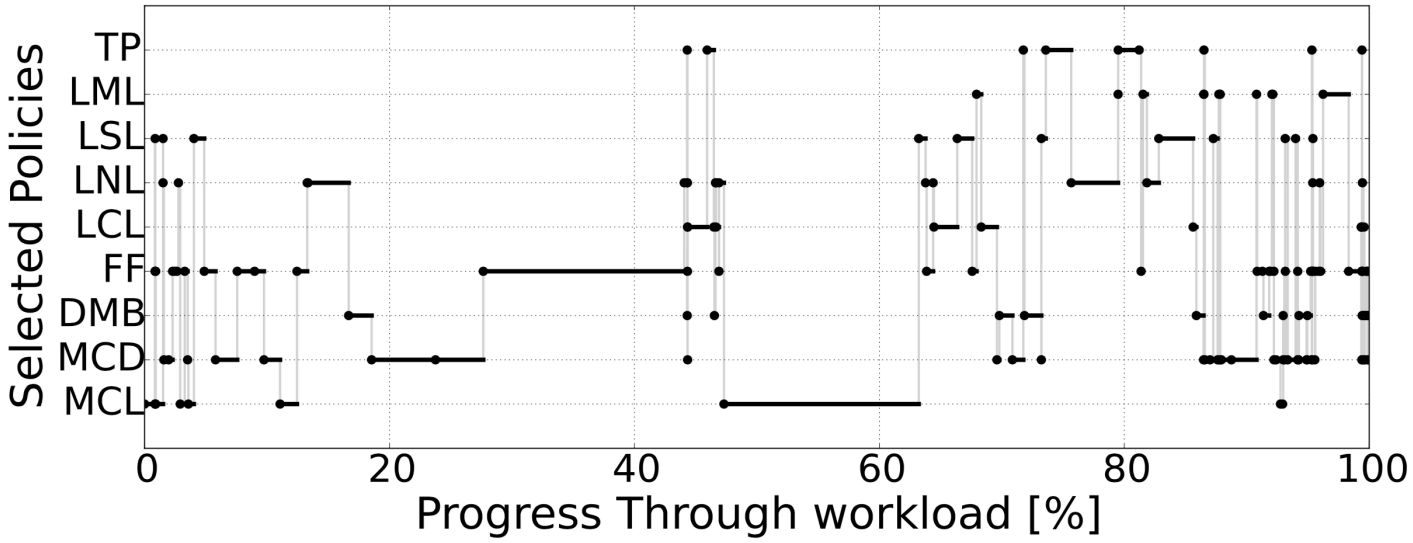

(a) Selected scheduling policies for the OR utility function.

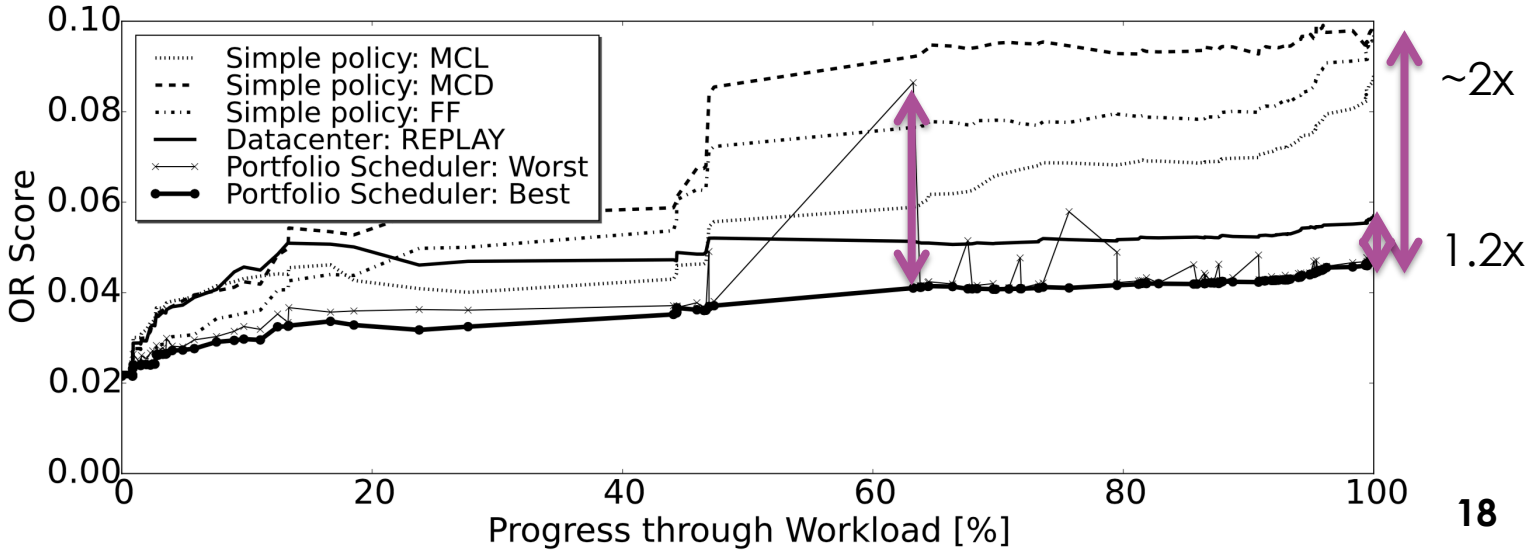

Conclusions

- **We address two risks**

  Operational Risk

  Disaster Recovery Risk

- **We extend the state-of-the-art in portfolio scheduling**

  risk-aware scheduling policies

  risk-aware utility functions

- **We introduce graphs that show the behavior of a portfolio scheduler**

Future work

- Further develop OpenDC
- More risks
- More resource types
- Capacity planning
- Microservice platforms
- Workload Modeling
- Multi-Cloud

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

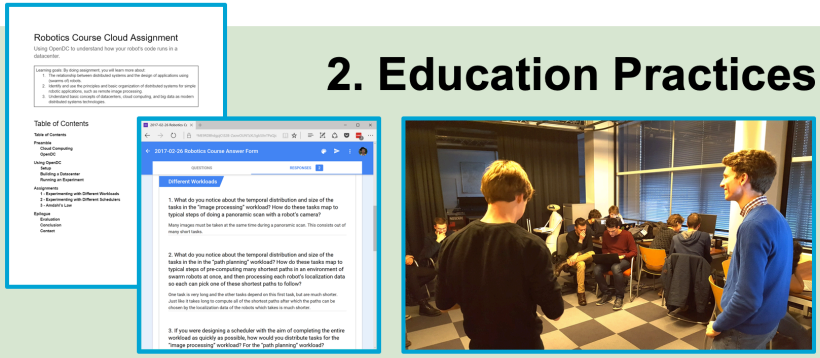

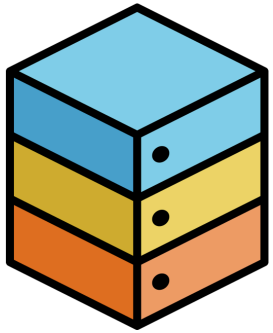

opendc.org

3. Software & Data Artifacts

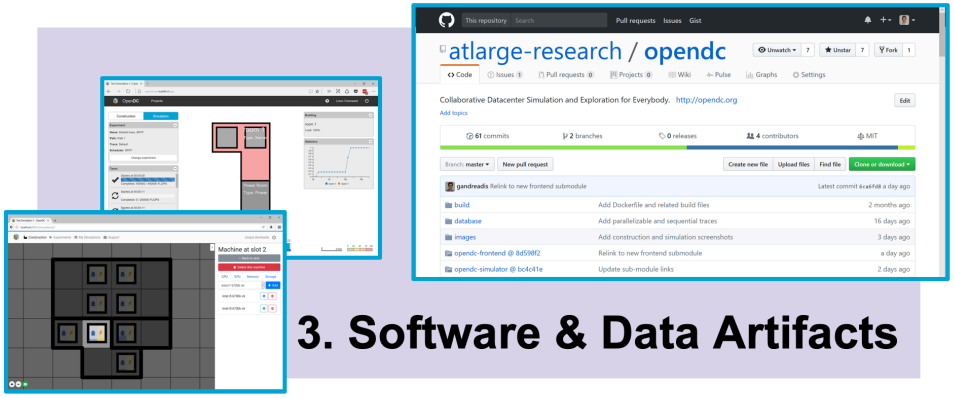

@Large Research
Massivizing Computer Systems

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

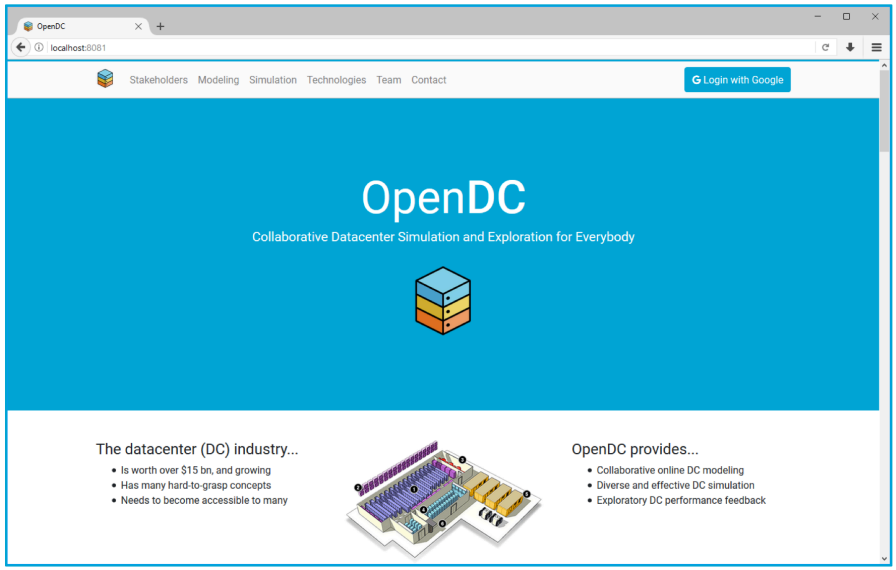

opendc.org

github.com/atlarge-research/opendc

Vincent van Beek Georgios Andreadis
Giorgos Oikonomou Alexandru Iosup
OpenDC team

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