OpenDC

Collaborative Datacenter Simulation and Exploration for Everybody

Prof. dr. ir. Alexandru Iosup
Project Lead

Leon Overweel
Product Lead and Software Engineer responsible for the web server, database, and API specification

Georgios Andreadis
Software Engineer responsible for the frontend web application and splash page

Matthijs Bijman
Software Engineer responsible for the datacenter simulator

Sacheendra Talluri
M.Sc. student, TU Delft

Vincent van Beek
Ph.D. student, TU Delft

Tim Hegeman
M.Sc. student, TU Delft

Jesse Donkervliet
M.Sc. student, TU Delft

Laurens Versluis
Ph.D. student, Vrije Universiteit Amsterdam

Mihai Neacsu
M.Sc. student, Vrije Universiteit Amsterdam

Why do we need **OpenDC**?

The **datacenter** industry…

- Is worth over **$15 bn** & growing
  - “Produces” **cloud services**

NETFLIX   

- Has many hard-to-grasp **concepts**
  - Scheduling, workloads, devops, ...
- Must be **accessible to many**

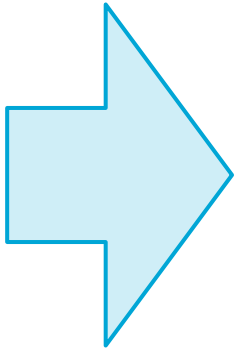

OpenDC provides...

What does **OpenDC** bring to the table?

1. Datacenter Technology & Methods

Risk Analysis + Management

Heterogeneity

Efficiency → SME Availability

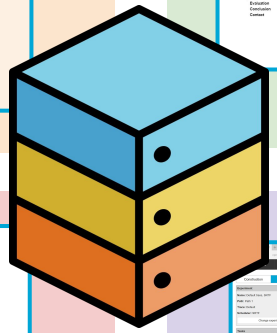

3. Education Practices

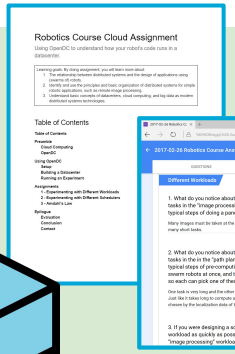

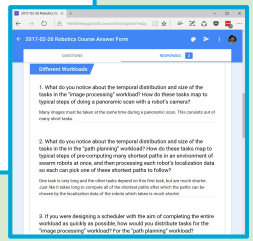

2. Scientific Methods

Mnemos: Self-Expressive Management of Business-Critical Workloads in Virtualized Datacenters

Vincent van Beek[1,2] Jesse Donkervliet Tim Hegeman
Stefan Hugtenburg
Alexandru Iosup
[1] Bitbrains IT Services Inc., Amstelveen, the Netherlands
[2] Delft University of Technology, Delft, the Netherlands
Corresponding author: vincent.vanbeek@bitbrains.nl

May 18, 2015

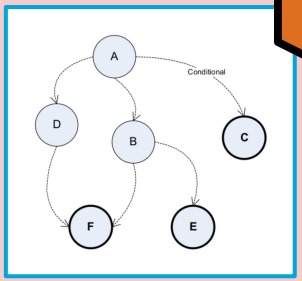

4. Software & Data Artifacts

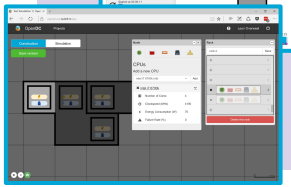

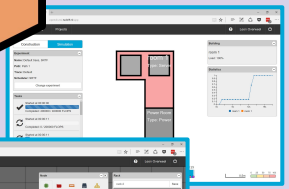

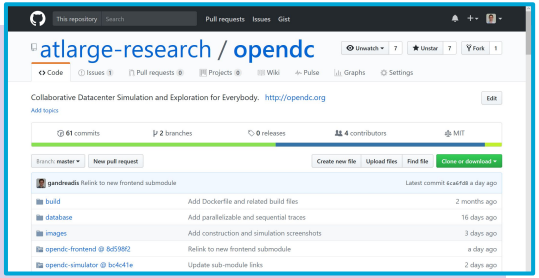

Does **OpenDC** deliver? 1. Datacenter Tech. & Methods

Explore a **variety of concepts**...

- Scalability + Elasticity
- Efficiency → SME Availability
- Reliability
- Risk Analysis + Management
- Heterogeneity

… with **PhD**, **MSc**, and **BSc** projects.

- Availability + Availability-on-Demand
- Complex Workflow Scheduling
- Application Auto-Scaling
- FaaS Management and Applications
- Memory-Based Storage
- Energy Management and Scheduling

Does **OpenDC** deliver? 2. Scientific Methods

How to conduct **scientific surveys** of **resource management** and **scheduling techniques** in datacenters?

How to provide a **useful yet reduced** set of **metrics** for modern datacenter operation?

How to design a **deep yet practical methodological apparatus** for obtaining such metrics?

How to design a **reference architecture** for **cloud schedulers**?

How do we conduct a **global scheduling competition**?

How to efficiently **validate** datacenter simulations?

How to build datacenter-simulation **environments** where **reproducibility** is **ensured by the instrument**?

What is the **performance-validity trade-off** for datacenter simulation?

Does **OpenDC** deliver? 3. Education Practices

OpenDC software **already used** for:

M.Sc. **Project-Based Learning**
@ VUA & TUD

B.Sc. Honours Programme
Classroom-Based Courses

B.Sc. Honours Programme
Project-Based Learning

… and we **plan to use** it for:

<RE/START> Periodic **workshops** for **refugees** in the Netherlands with **Restart Network**

Promoting **science in schools** with the **Royal Netherlands Academy of Arts and Sciences**

KIVI Engineering Society Engaging **high school** students through **workshops** with the **Royal Dutch Engineers Society**

Does **OpenDC** deliver? 4. Software Artifacts

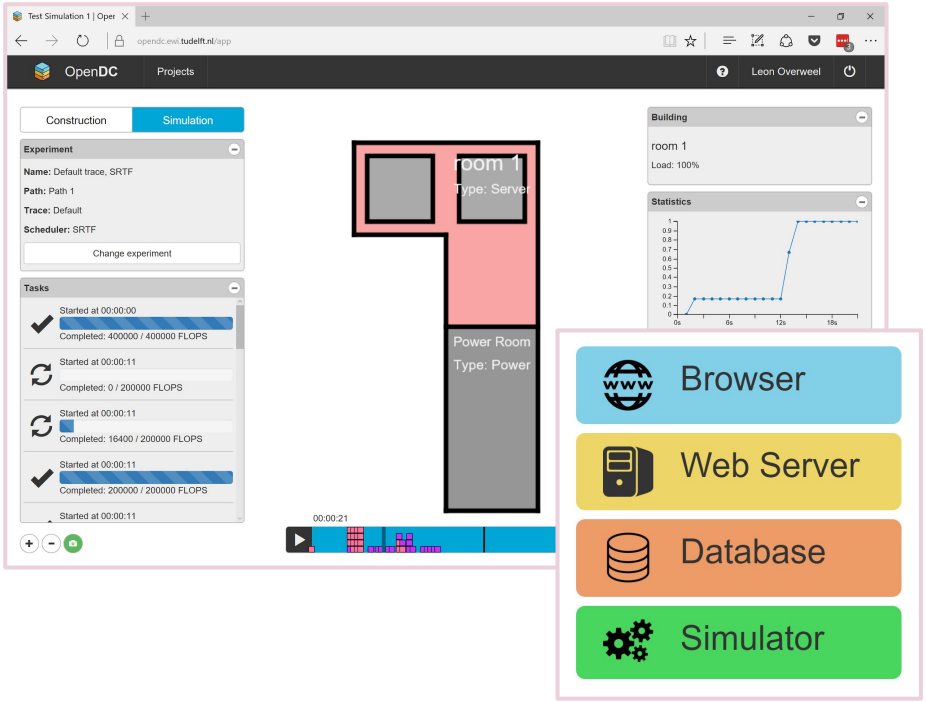

Current capabilities:

- Define dynamic DC **topologies**
- Run experiments on different **schedulers** and **workloads**
- Playback experimental results

Roadmap:

- **UI + API** for workloads + schedulers
- Componentized sim. for research

Availability:

- **Online** → Hosted by TU Delft
- **Locally** → Source on GitHub

What does **OpenDC** bring to the table?

1. Datacenter Technology & Methods

Risk Analysis + Management

Heterogeneity

Efficiency → SME Availability

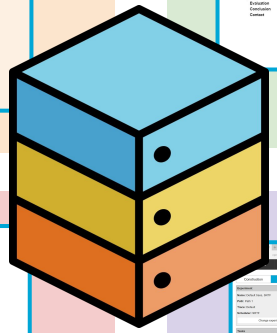

3. Education Practices

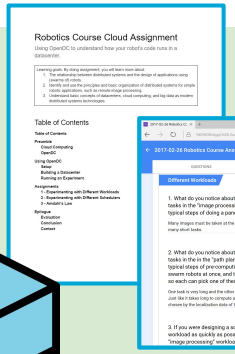

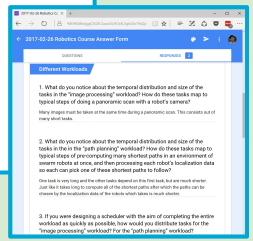

2. Scientific Methods

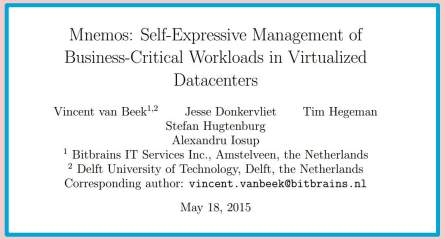

Mnemos: Self-Expressive Management of Business-Critical Workloads in Virtualized Datacenters

Vincent van Beek[1,2] Jesse Donkervliet Tim Hegeman
Stefan Hugtenburg
Alexandru Iosup
[1] Bitbrains IT Services Inc., Amstelveen, the Netherlands
[2] Delft University of Technology, Delft, the Netherlands
Corresponding author: vincent.vanbeek@bitbrains.nl

May 18, 2015

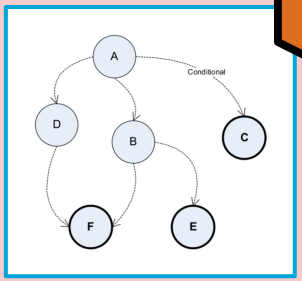

4. Software & Data Artifacts

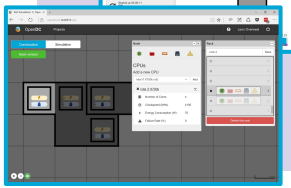

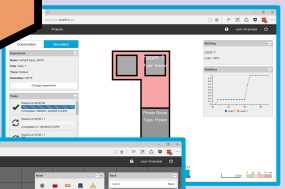

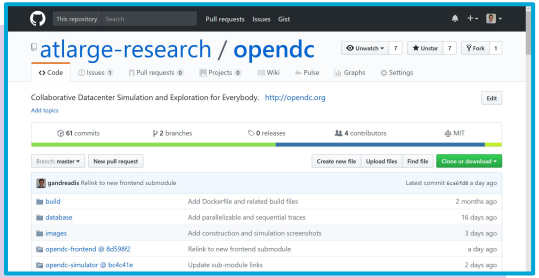

Find **Open**DC online!

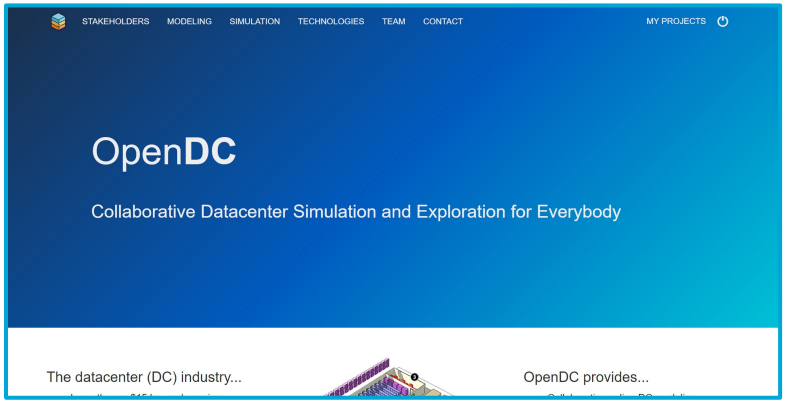

opendc.org

github.com/atlarge-research/opendc

opendc@atlarge-research.com

atlarge-research.com