

# 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**
- Has many hard-to-grasp **concepts**
  - Scheduling, workloads, devops, ...



**OpenDC** provides...

- Must be **accessible to many**



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

## 1. Datacenter Technology & Methods

Risk Analysis +  
Management

Efficiency →  
SME  
Availability


Heterogeneity

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

Vincent van Beek<sup>1,2</sup> Jesse Donkervliet Tim Hegeman  
Stefan Hugtenburg  
Alexandru Iosup

<sup>1</sup> Bitbrains IT Services Inc., Amstelveen, the Netherlands  
<sup>2</sup> Delft University of Technology, Delft, the Netherlands  
Corresponding author: [vincent.vanbeek@bitbrains.nl](mailto:vincent.vanbeek@bitbrains.nl)

May 18, 2015



## 2. Scientific Methods

@Large Research  
Massivizing Computer Systems



VU  
VRIJE  
UNIVERSITEIT  
AMSTERDAM

TU Delft

## 3. Education Practices

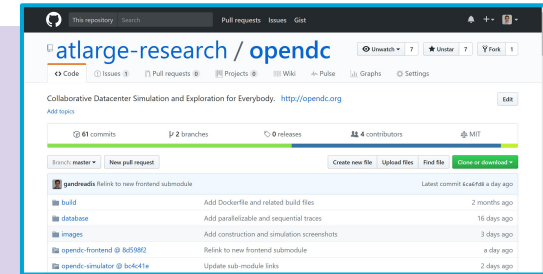
Robotics Course Cloud Assignment  
Using OpenDC to visualize how past robot class runs in a classroom

Table of Contents

1. What do you notice about the temporal distribution and size of the tasks in the "Image processing" workload? How do these tasks map to typical steps of doing a homework case with a robot's camera?  
How do you think to improve the time to do a homework case. This requires a lot of every kind task.

2. What do you notice about the temporal distribution and size of the tasks in the "Path planning" workload? How do these tasks map to typical steps of perceiving a new robot path in an environment of known robot's area, and then processing each robot's localization data so each can pick one of these shortest paths to follow?  
How do you think to improve the time to do a homework case. This requires a lot of every kind task.

3. If you were designing a scheduler with the aim of completing the write workload as quickly as possible, how would you distribute tasks to the "Image processing" workload? To the "Path planning" workload?



## 4. Software & Data Artifacts

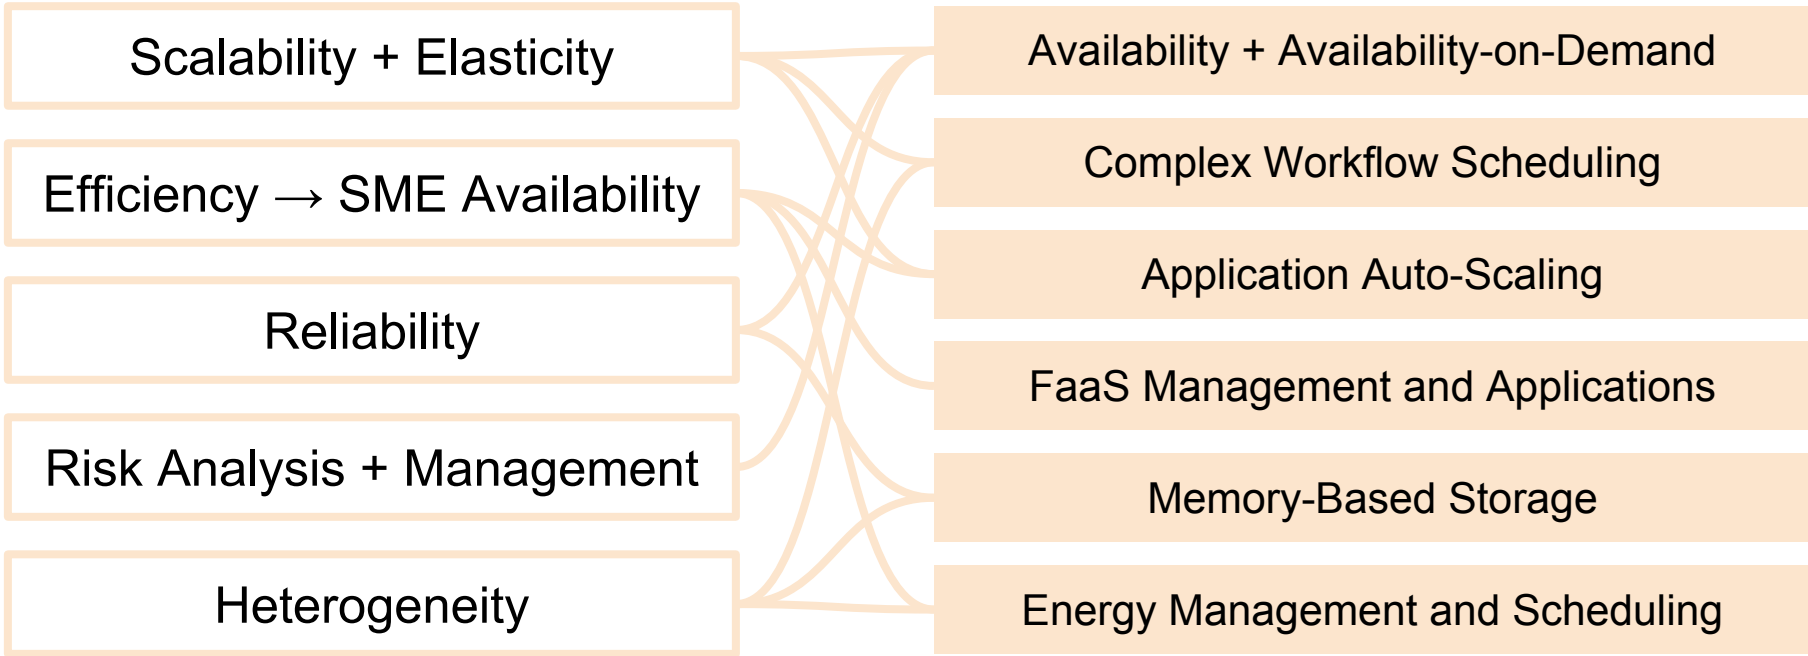
3



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

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

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



# 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**



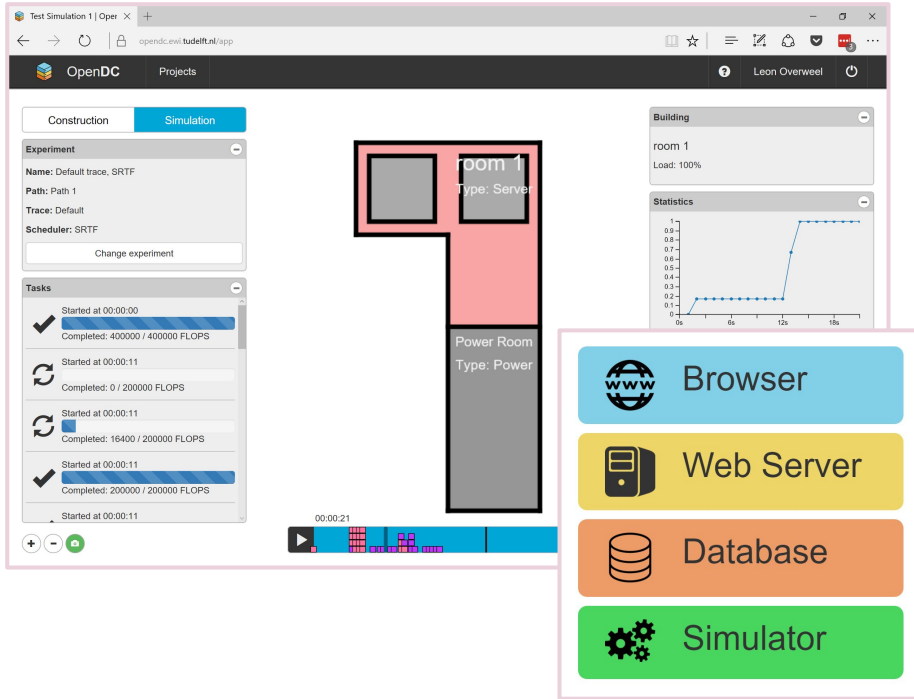
Promoting **science in schools**  
with the **Royal Netherlands**  
**Academy of Arts and Sciences**



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

Efficiency →  
SME  
Availability

Heterogeneity




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

Vincent van Beek<sup>1,2</sup> Jesse Donkervliet Tim Hegeman  
Stefan Hugtenburg  
Alexandru Iosup


<sup>1</sup> Bitbrains IT Services Inc., Amstelveen, the Netherlands  
<sup>2</sup> Delft University of Technology, Delft, the Netherlands  
Corresponding author: [vincent.vanbeek@bitbrains.nl](mailto:vincent.vanbeek@bitbrains.nl)

May 18, 2015



## 2. Scientific Methods

## 3. Education Practices



Robotics Course Cloud Assignment

Using OpenDC to understand how past robot class runs in a classroom

Table of Contents

1. Introduction

2. Getting Started

3. Building a Scheduler

4. Building a Simulator

5. Building a Scheduler

6. Building a Scheduler

7. Building a Scheduler

8. Building a Scheduler

9. Building a Scheduler

10. Building a Scheduler

11. Building a Scheduler

12. Building a Scheduler

13. Building a Scheduler

14. Building a Scheduler

15. Building a Scheduler

16. Building a Scheduler

17. Building a Scheduler

18. Building a Scheduler

19. Building a Scheduler

20. Building a Scheduler

21. Building a Scheduler

22. Building a Scheduler

23. Building a Scheduler

24. Building a Scheduler

25. Building a Scheduler

26. Building a Scheduler

27. Building a Scheduler

28. Building a Scheduler

29. Building a Scheduler

30. Building a Scheduler

31. Building a Scheduler

32. Building a Scheduler

33. Building a Scheduler

34. Building a Scheduler

35. Building a Scheduler

36. Building a Scheduler

37. Building a Scheduler

38. Building a Scheduler

39. Building a Scheduler

40. Building a Scheduler

41. Building a Scheduler

42. Building a Scheduler

43. Building a Scheduler

44. Building a Scheduler

45. Building a Scheduler

46. Building a Scheduler

47. Building a Scheduler

48. Building a Scheduler

49. Building a Scheduler

50. Building a Scheduler

51. Building a Scheduler

52. Building a Scheduler

53. Building a Scheduler

54. Building a Scheduler

55. Building a Scheduler

56. Building a Scheduler

57. Building a Scheduler

58. Building a Scheduler

59. Building a Scheduler

60. Building a Scheduler

61. Building a Scheduler

62. Building a Scheduler

63. Building a Scheduler

64. Building a Scheduler

65. Building a Scheduler

66. Building a Scheduler

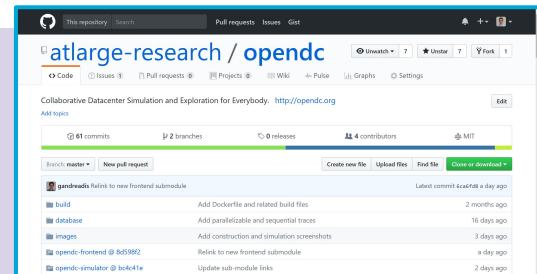
67. Building a Scheduler

68. Building a Scheduler

69. Building a Scheduler

70. Building a Scheduler

## 4. Software & Data Artifacts





# Find **OpenDC** online!



[opendc.org](https://opendc.org)



[github.com/atlarge-research/opendc](https://github.com/atlarge-research/opendc)



[opendc@atlarge-research.com](mailto:opendc@atlarge-research.com)



[atlarge-research.com](https://atlarge-research.com)

