

# A Short Introduction to Cloud Computing

ir. Laurens Versluis

[l.f.d.versluis@vu.nl](mailto:l.f.d.versluis@vu.nl)

<https://atlarge.science>





VRIJE  
UNIVERSITEIT  
AMSTERDAM



# Cloud popularity and usage at all-time high


- Surveys: **86% of companies use >1 cloud service**, **>\$200B market by 2020**
- Cloud computing increasingly important
  1. Improve competitive position for a company
  2. Reduce costs for company & customer

Public IT Cloud Spending (\$Billions)

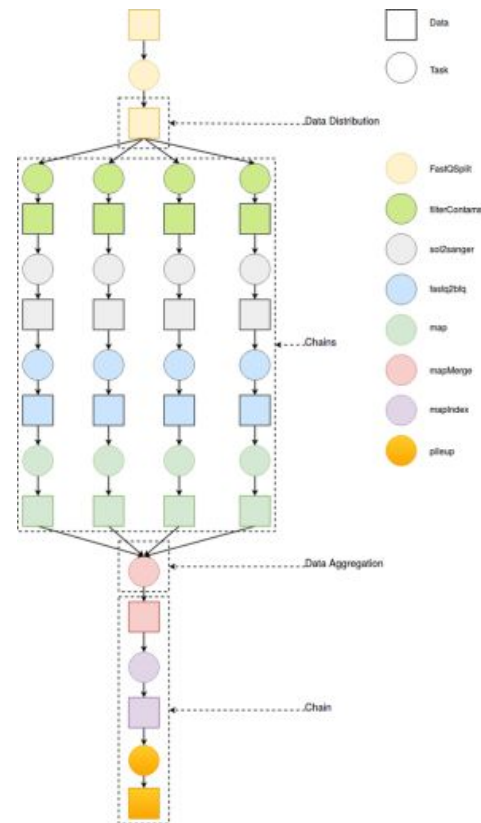



# Workflow execution is common

- Workflows = set of tasks with precedence constraints
  - Usually represented as a Directed Acyclic Graph (DAG)
  - Used to model applications in many domains
- Today: thousands of applications in use




finding cancer early  
epigenomics



# Executing workflows in the cloud

1. Workflows in the cloud



**On which resource do we execute the workflow?**

**How many resources to acquire and when?**

Workload of workflows



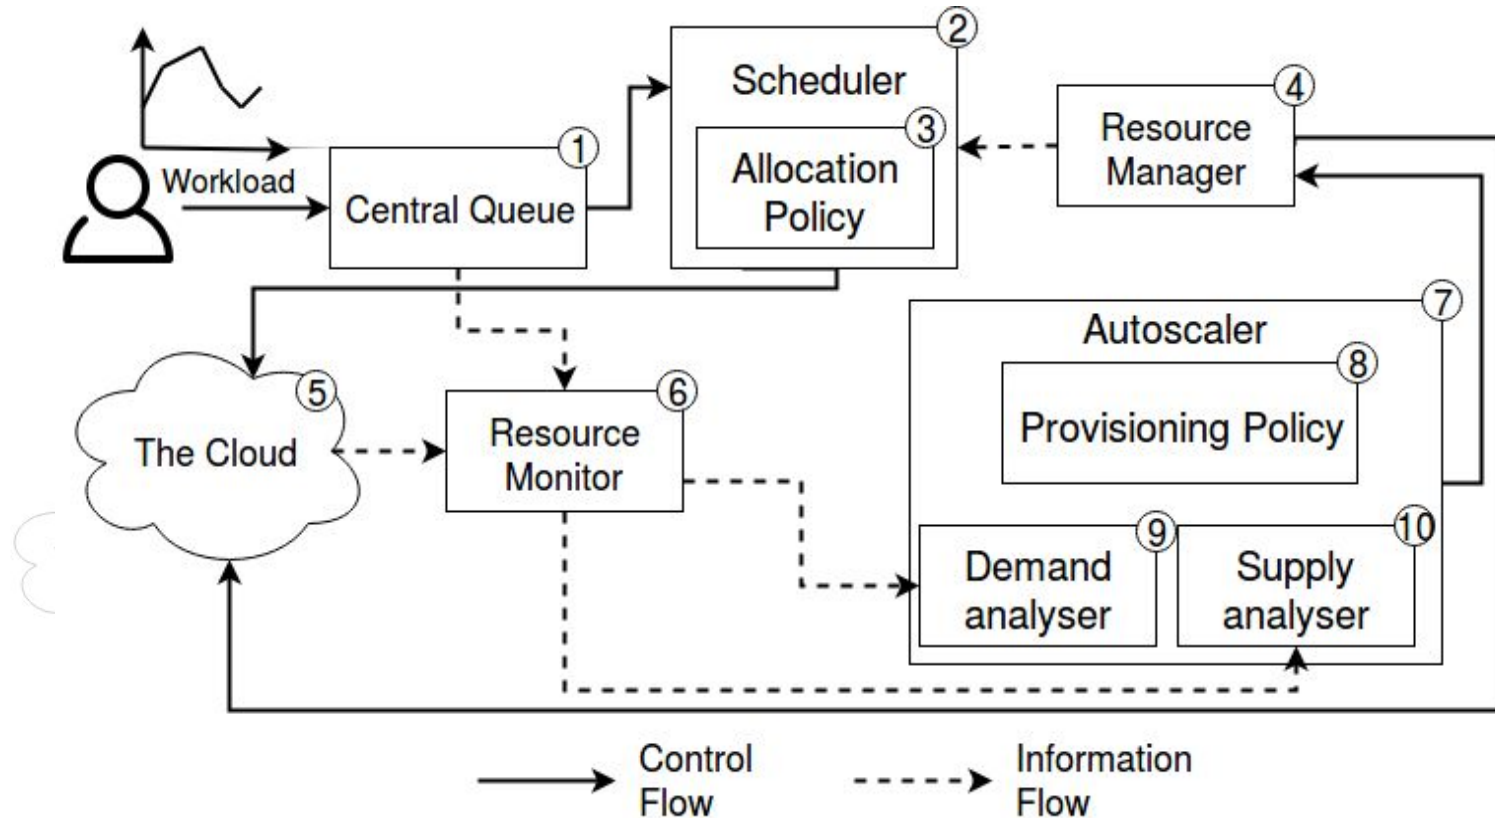
# A problem for cloud providers

Cloud providers preferably want:

1. (Close to) optimum workflow allocation
  - Efficient use of resources
2. Minimize overprovisioning (allocating too many resources)
  - Reduces costs
3. Adhere to the Quality-of-Service (QoS) requirements of clients
4. Automate this process
  - Poor user estimates of resource requirements



# Overview of a workflow management system



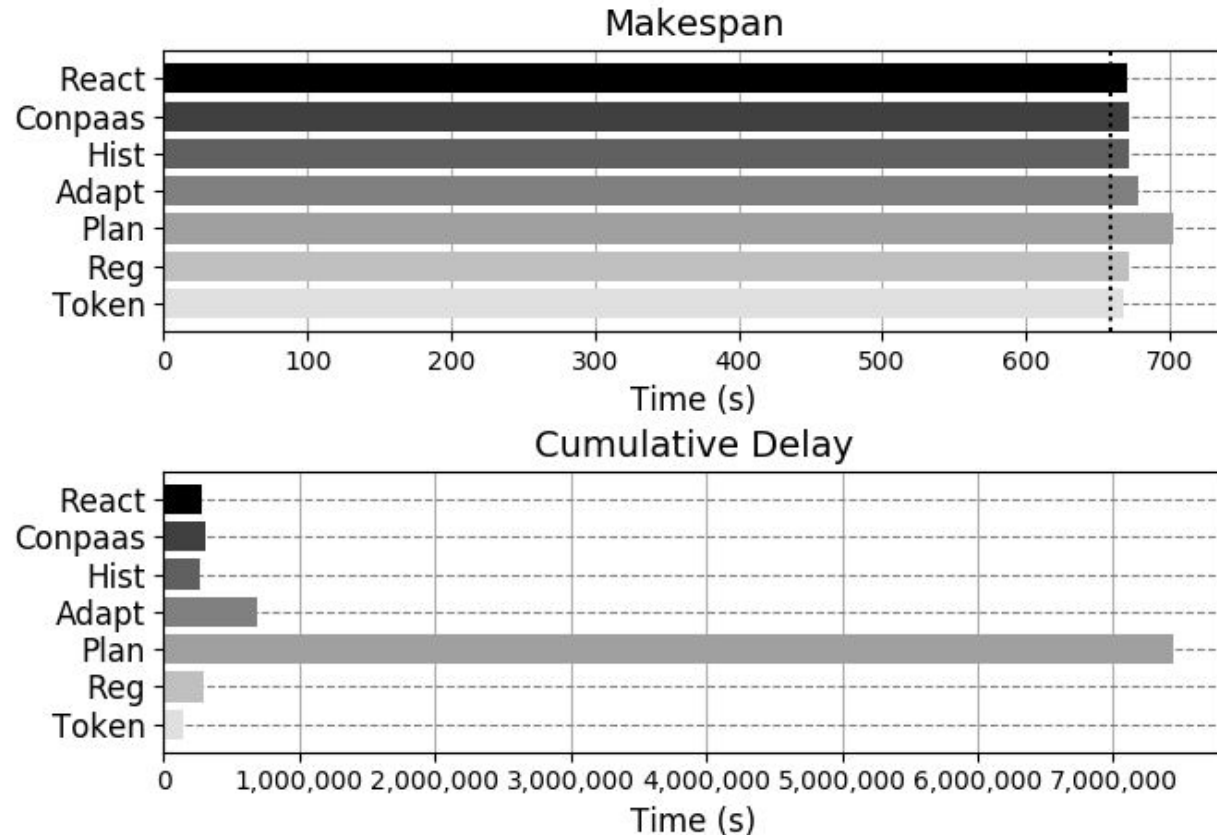
# Example of our work: CCGrid 2018

- Four distinct workload traces
  - Remember: a **workload** is a set of workflows (applications)
- Use a rich set of metrics
  - 10+ forms of **elasticity**
- Four distinct experiments; varied {workload, environment, metrics}

Medium scale

ID	Source	Domain	Workflows	# Tasks
T1	SPEC Cloud Group	Scientific	200	13,876
T2	Chronos	Industrial	1,024	3,072
T3	Askalon EE	Engineering	757	45,786
T4	Askalon EE2	Engineering	3,551	122,105

# An experimental result





# Wrapping up

Scheduling in clouds, a complex problem:

1. Allocation (where to put tasks)
2. Provisioning (when to change resources)
3. Difficult to predict performance given {workload, environment, metric}

@Large is actively working on these topics.

Interested in this work? Feel free to contact me/us!

# A Short Introduction to Cloud Computing

ir. Laurens Versluis

[l.f.d.versluis@vu.nl](mailto:l.f.d.versluis@vu.nl)

<https://atlarge.science>



VRIJE  
UNIVERSITEIT  
AMSTERDAM



# We are actively working on these topics

