

Scheduling Dynamic, Per-task Non-Functional Requirements in Complex Workflows



Laurens Versluis and Alexandru Iosup

Contact: l.f.d.versluis@vu.nl



The Problem

What is a workflow?

- A set of tasks with precedence constraints among them
- Can be represented as a Directed Acyclic Graph (DAG)

What is a non-functional requirement (NFR)?

- A requirement that specifies criteria to evaluate the operations of a system, rather than its behavior

Previous Research

- State-of-the-Art approaches mostly consider NFRs at the workflow level
 - Wastes resources [1] -> lower resource utilization
 - More costs for (cloud) costumers
- Work that does consider task-based NFRs only consider a subset of NFRs such as deadlines (performance) or cost

Our Approach

A new formalism


- Ability to specify arbitrary, task-based NFRs
- Backwards compatible with existing Resource Management & Scheduling systems
- Support creating groups of tasks having the same NFR

Combine simulation and real-world experimentation

- Simulation allows for testing at scale and many different configurations
- Real-world experimentation produces more meaningful results

Collaborate with third parties that can benefit from this work

- Health sector (e.g. VUmc)
- Industry
- Academia



Preliminary results

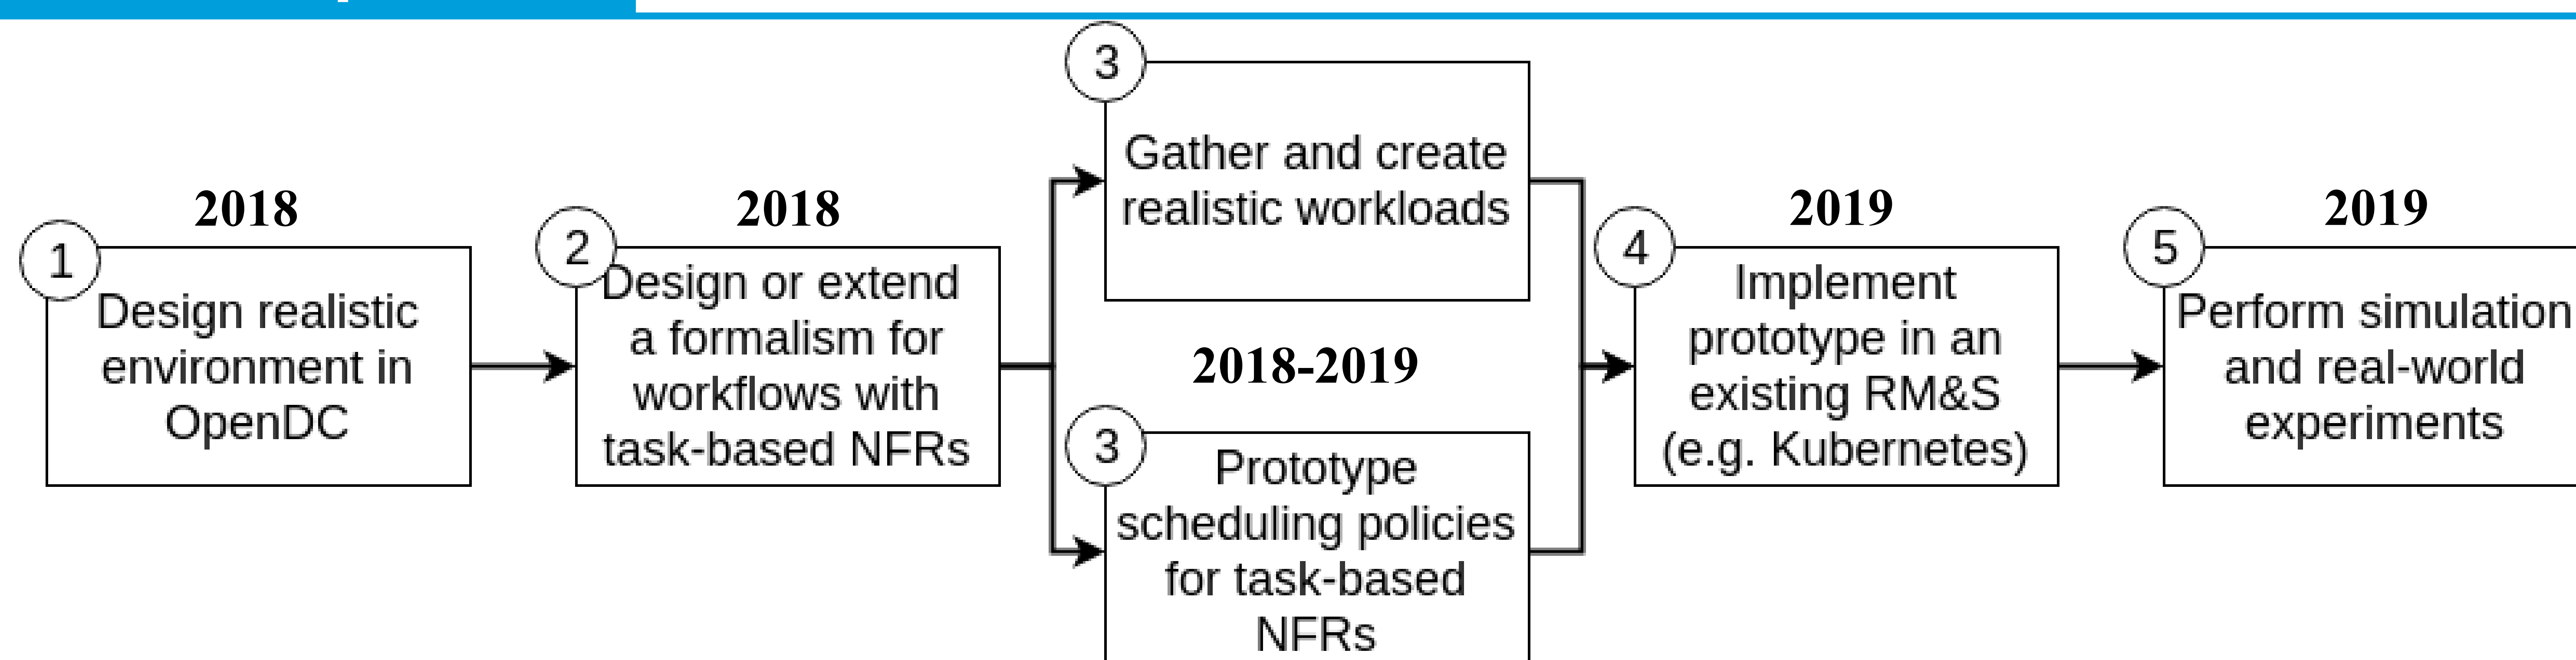
- Investigated formalisms mentioned in scientific work
 - **116 papers across 11 venues inspected [2]**

Conference	Workflow	BPMN	Petri net	DAG
HPDC	2	0	0	0
NSDI	1	0	0	0
SOCC	0	0	0	0
ICPE	1	1	1	0
Cluster	20	0	0	5
OSDI	1	0	0	0
SIGMETRICS	0	0	0	0
CCgrid	50	1	2	15
ICPP	12	0	0	7
IEEECLOUD	21	1	2	12
IPDPS	8	0	0	5
Total	116 (100%)	3 (3%)	5 (4%)	44 (38%)

- Obtained the three most popular formalisms
- Compared these three formalisms on their ability to specify NFRs

Main finding: None of these formalisms are capable of specifying arbitrary NFRs at the task-level

Roadmap



- Do you have workloads that you want to run more efficiently? Contact us, we are looking for collaborators!

References

- [1] An Availability-on-Demand Mechanism for Datacenters, CCGrid 2015
[2] An Analysis of Workflow Formalisms for Workflows with Complex Non-Functional Requirements, HotCloudPerf 2018