CGW Workshop 14, Krakow, Poland, October 27-29, 2014

Adaptive Scheduling of Workflows on Clouds with Deadlines

Michael Gerhards[1,2], Volker Sander[1], Adam Belloum[2], Marian Bubak[2]

[1] FH Aachen - University of Applied Sciences, Heinrich-Mussmann-Str. 1, 52428 Juelich, Germany
http://www.fh-aachen.de/forschung/projekt-hixforagws/veroeffentlichungen/
[2] University of Amsterdam, Science Park 107, 1090 GH Amsterdam, The Netherlands
http://ivi.uva.nl/ http://www.commit-nl.nl/people
m.gerhards@fh-aachen.de, v.sander@fh-aachen.de, a.s.z.belloum@uva.nl, bubak@agh.edu.pl

FH AACHEN UNIVERSITY OF APPLIED SCIENCES

Problem Statement

Cloud computing offers new scheduling opportunities. Any cloud-aware scheduling algorithm has to decide on the number of actually used resources and for each resource about its instance size, and, to address the occurred monetary cost adequately, start time and end time. Basically, it has to perform a trade-off between deadline and cost constraints in an elastic environment. Neither static nor dynamic scheduling strategies alone address all characteristics of a cloud ecosystem. While dynamic scheduling can react on unexpected behaviors and runtime variations, an initial setup of resources is required. The proposed adaptive scheduling strategy combines the advantages of both mechanisms to form a novel approach for deadline constraint dynamic scheduling of scientific workflows on cloud infrastructures.

Solution

Application Model

The targeted application models are large scale computational workflows having thousands of tasks and following well-structured data dependencies such as the templates of the Pegasus workflows Epigenomics, Cybershake, Inspiral, Montage, and SIPHT. The maximal makespan is limited by a deadline.

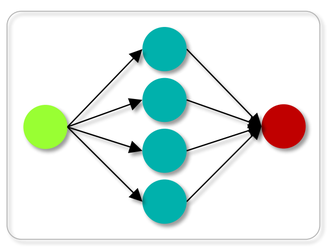

Cloud Ecosystem

Cloud infrastructures provide horizontal and vertical scalable on-demand resources following typically an hourly billed pay-as-you-go cost model. Any cloud-aware scheduling algorithm has to decide on the number of actually used resources and for each resource about its instance size and occupation time to address the occurred monetary cost.

Runtime Profiles

All thousands of tasks of a workflow template only follow few task types. A runtime profile for each task type is created based on gathered monitoring/provenance data of previous runs. The runtime profile of a task type also depends on its multithreading capabilities that are applicable on a particular cloud resource.

Type	Ex Time Mean	Ex Time Variation	Parallel Degree
map	10 min	2 min	0.10
compute	20 min	4 min	0.80
reduce	8 min	1 min	0.30

Resource Allocation Plan

The initial resource allocation plan for a particular workflow run is created using a static scheduling strategy considering the workflow dependency graph, the runtime profiles of the occurring task types, and the cloud billing model as parameters. The workflow engine instantiates on demand all required cloud resources given by the resource allocation plan. Whenever a resource finished executing its currently assigned task, the workflow engine schedules the next task for execution selected from the resource allocation plan assignments.

Static Scheduling

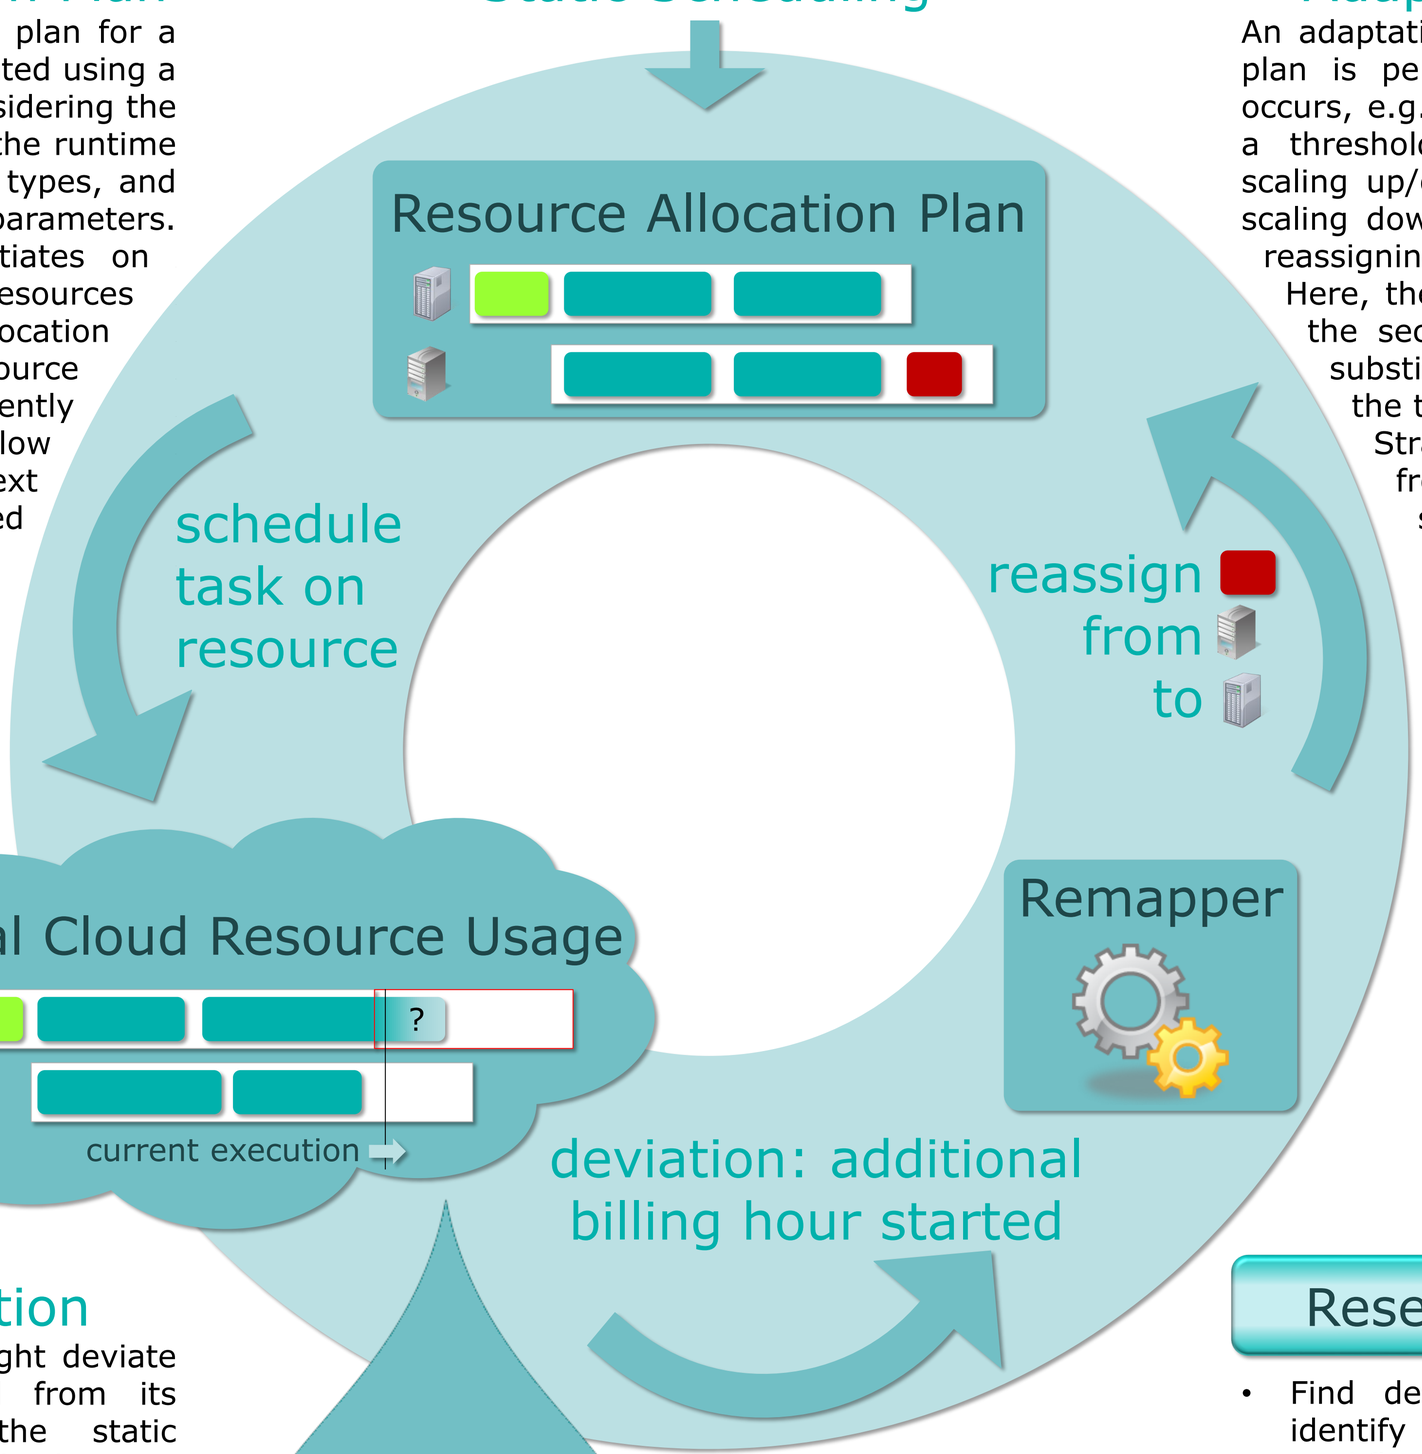

Adaptive Scheduling

An adaptation of the resource allocation plan is performed whenever an event occurs, e.g. a runtime deviation exceeds a threshold. Possible adaptations are scaling up/out to hold the deadline and scaling down/in to reduce cost but also reassigning a task to another resource. Here, the red task is reassigned from the second to the first resource to substitute the unused capacity of the there started new billing hour. Strategic decision points derived from workflow template's flow structure enable a proactive adaptation for future tasks much earlier.

Runtime Deviation

A task's runtime behavior might deviate from its prediction derived from its runtime profile. Hence, the static resource allocation plan might become uneconomic or might even fail in holding the deadline. Here, the third task hosted by the first resource starts an unexpected additional billing hour during its execution.

Research Objectives

- Find deviation thresholds/events to identify adaptation demand for a workflow instance according to its actual execution in a particular cloud economy.
- Derive strategic decision points from the flow structure of a workflow template.

FH AACHEN

COMMIT/

University of Amsterdam