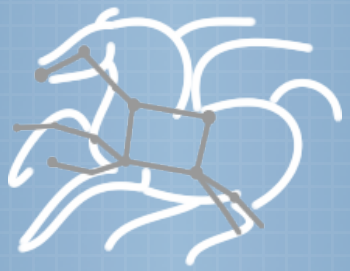

A Unified Approach for Modeling and Optimization of Energy, Makespan and Reliability for Scientific Workflows on Large-Scale Computing Infrastructures

Rafael Ferreira da Silva[1], Thomas Fahringer[2], Juan J. Durillo[2], Ewa Deelman[1]

Workshop on Modeling & Simulation of Systems and Applications
August 13-14, 2014, University of Washington, Seattle, Washington

[1]University of Southern California, Information Sciences Institute, Marina Del Rey, CA, USA
[2]University of Innsbruck, Technikerstrasse 21a, Innsbruck, Austria

USC Viterbi
School of Engineering
Information Sciences Institute

Introduction

- Scientific workflows are often used to manage large-scale computations on HPC and HTC platforms
  - Several studies have been conducted to optimize workflow scheduling
  - However, most existing optimization techniques are limited to single or two objectives

- Research in green computing often address cooling and energy usage reduction in large data-centers
  - There are few studies on how resources are used by applications

- Green computing in scientific workflows
  - Studies are limited to the measurement of energy usage according to resource utilization
  - The energy consumption model is simplistic (e.g., homogeneous execution nodes)

Research Goals

- Development of an energy consumption model to address real large-scale infrastructure conditions
  - e.g., heterogeneity, resource availability, external loads
  - Validation of the model in a fully instrumented platform able to measure the actual temperature and energy consumed by computing, networking, and storage systems
- Development of a multi-objective optimization approach to explore workflow execution tradeoffs

Application Model: Scientific Workflows

- Directed Acyclic Graph (DAG)
  - Nodes denote tasks
  - Edges denote task dependencies
- Tasks
  - Command-line programs that read one or more input files and produce one or more output files
  - Compute-intensive or data-intensive
- Data dependencies
  - Result of output files from one program becoming input files for another program

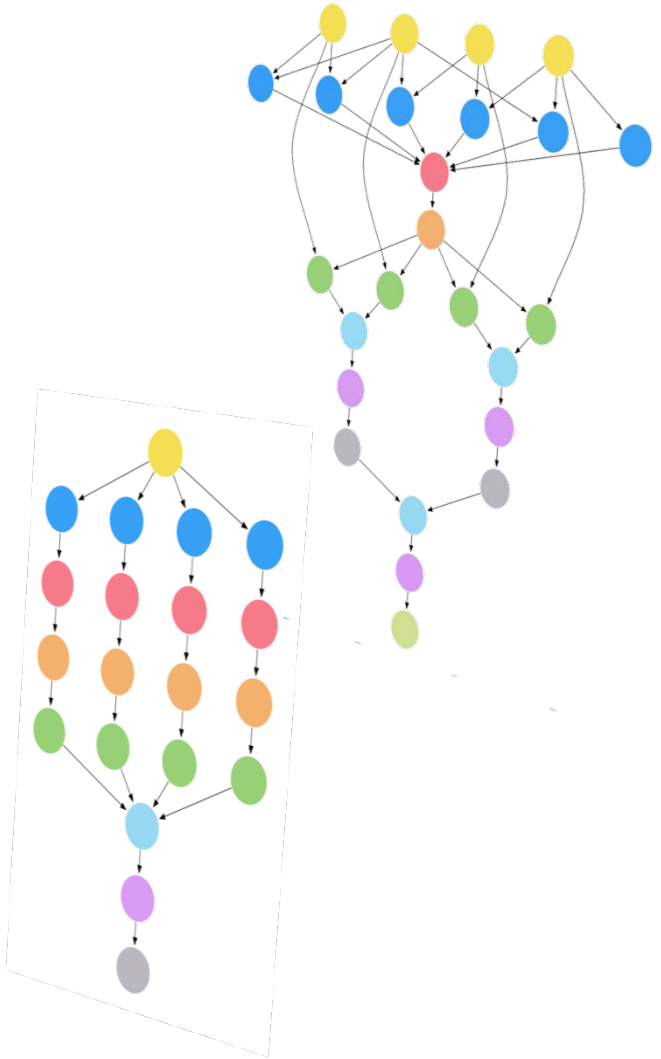

System Model: Distributed Infrastructure

- Infrastructure as a Service (IaaS)
  - Data and task computations are stored/performed in the infrastructure

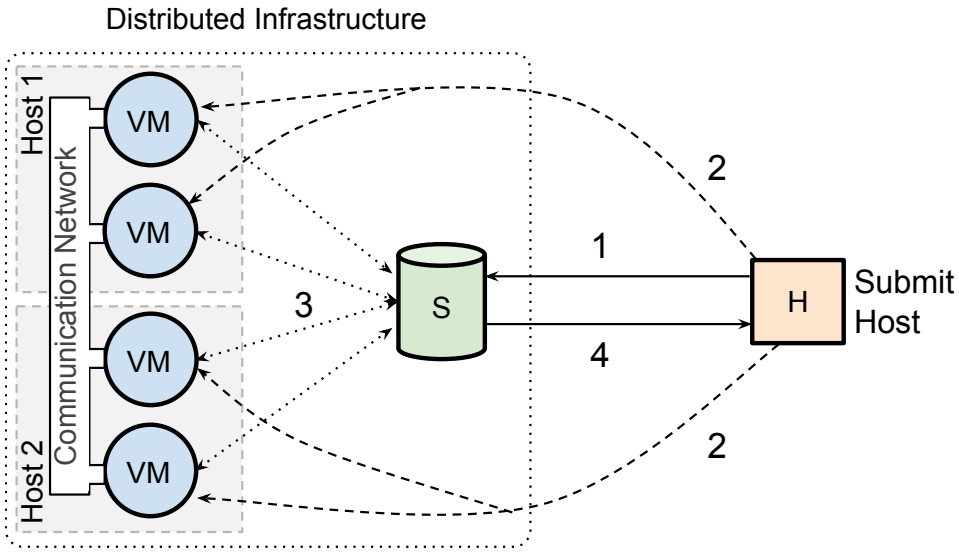

1: Application setup: provision of a set of parameters and input files uploading

2: Workflow task scheduling

3: Output data is stored on the storage server

4*:* Output data required by the user is downloaded from the storage server

Runtime and Reliability Models

- At Workflow Level (Our Expertise)
  - Collect and summarize performance metrics for workflow applications
  - e.g., *process I/O, runtime, memory usage, CPU utilization*
  - Profile data is used to build distributions of workflow applications

- At Infrastructure Level (Looking for a Partner)
  - Collect temperature and energy consumption from execution nodes, storage servers, and network systems
  - Requires a fully instrumented platform

Research Dimensions

- **Goal**: Multi-objective optimization of energy consumption, makespan, and reliability for scientific workflows
- Monitoring
  - Workflow profile data has been collected as part of the DOE dV/dt project (ER26110)
  - Temperature and energy consumption monitoring requires access to a fully instrumented infrastructure

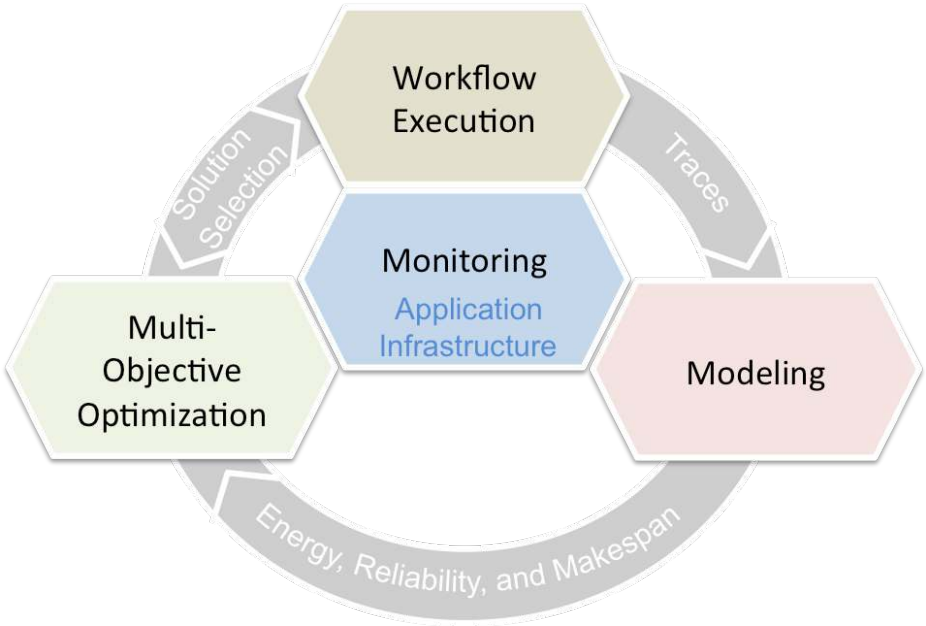

Multi-objective optimization process

Research Dimensions

- Multi-Objective Optimization
  - The improvement of one optimization criteria may imply in the deterioration of another criteria
  - Development of heuristics to reduce the large-search space of workflow executions
- Modeling (Dynamic Optimization)
  - Models will be constantly updated based on the profiling data collected during the workflow execution
- Workflow Execution
  - Conducted with the Pegasus WMS (OCI SI2-SSI #1148515)

Discussions

- Major Contribution
  - Multi-objective optimization of energy consumption, makespan, and reliability for scientific workflows on large-scale computing infrastructures
- Gaps in Current Research
  - There is no energy-aware profiling of scientific workflow applications
  - Research is focused on the optimization of a single or two objectives
  - Strong assumptions are made (e.g., homogeneous environments)
- Synergistic Projects
  - dV/dT: Accelerating the Rate of Progress Towards Extreme Scale Collaborative Science (DOE ER26110)
  - Pegasus WMS (OCI SI2-SSI #1148515)
  - DOE Sustained Performance, Energy and Resilience (SUPER) project

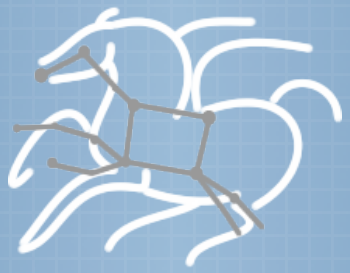

A Unified Approach for Modeling and Optimization of Energy, Makespan and Reliability for Scientific Workflows on Large-Scale Computing Infrastructures

Thank you.

rafsilva@isi.edu

http://pegasus.isi.edu

USC Viterbi
School of Engineering
Information Sciences Institute

Mono-Objective Optimization

- Single optimization goal to improve workflow executions
  - e.g., makespan, cost, etc.

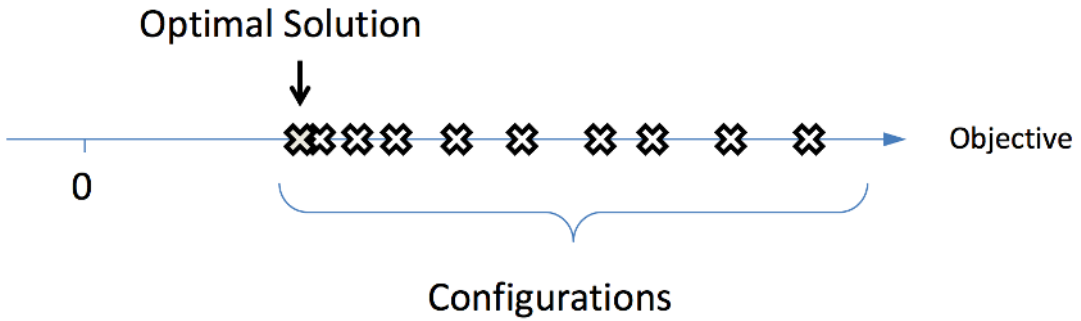

Multi-Objective Optimization

- The improvement of one optimization criteria may imply in the deterioration of another criteria
  - There is no single solution that is optimal with respect to all objectives