Modeling, Optimization and Performance Evaluation of Scientific Workflows in Clouds

Kamil Figiela and Maciej Malawski
Department of Computer Science, AGH University of Science and Technology, Kraków, Poland

Motivation

- Scientific workflows are convenient way of expressing complex experiments.
- Clouds can provide on-demand compute resources.
- New billing schemes and heterogeneous infrastructure present in the cloud preclude well-established resource allocations policies, known e.g. from grids.

Real workflows
Montage, FEM solver

Synthetic workflows
from Pegasus Workflow Gallery

Generic benchmarks
from CloudHarmony, official Amazon ECU

Application-specific benchmarks
benchmark application on real infrastructure

Application data
application structure, number of tasks

Infrastructure data
performance of application on given infrastructure

Workflow planner and optimizer

- We rely on MILP (mixed integer linear programming) for optimization.
- Optimization models in AMPL/CMPL languages that are backed by CPLEX and Cbc solvers.

Scheduling model

Used model depends on application class:
- Bag-of-tasks
- Coarse-grained workflows
- Fine-grained workflows

Performance model

- Application model
- Infrastructure model

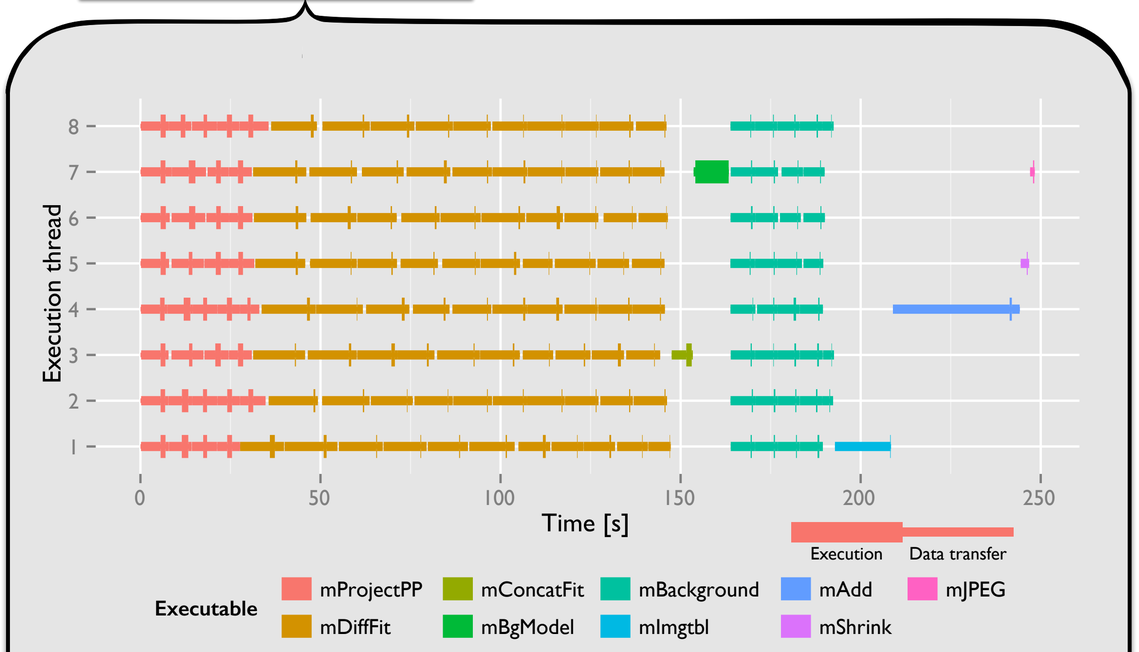

Montage execution trace on 8-core m3.2xlarge instance with a 2x2 degree region. Most of the execution time is spent on I/O.

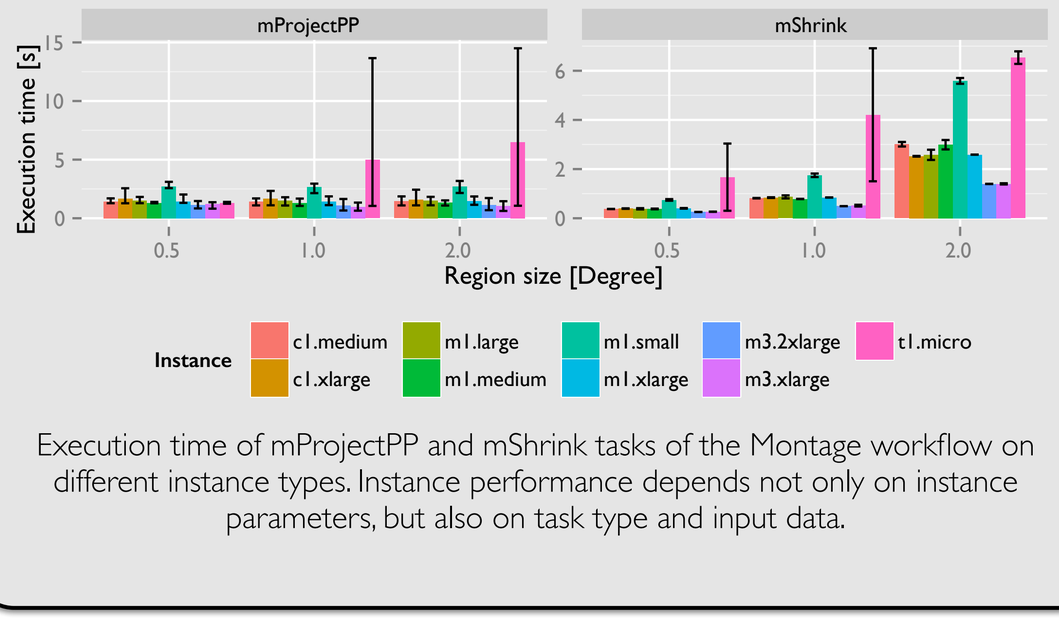

Execution time of mProjectPP and mShrink tasks of the Montage workflow on different instance types. Instance performance depends not only on instance parameters, but also on task type and input data.

Challenges

- Propose a model which describes properties of the application and the underlying infrastructure.
- Take into account different billing schemes e.g. per-hour or per-minute billing, or burstable instances.
- Find balance between model accuracy and acceptable runtime of optimizer and use some approximations.

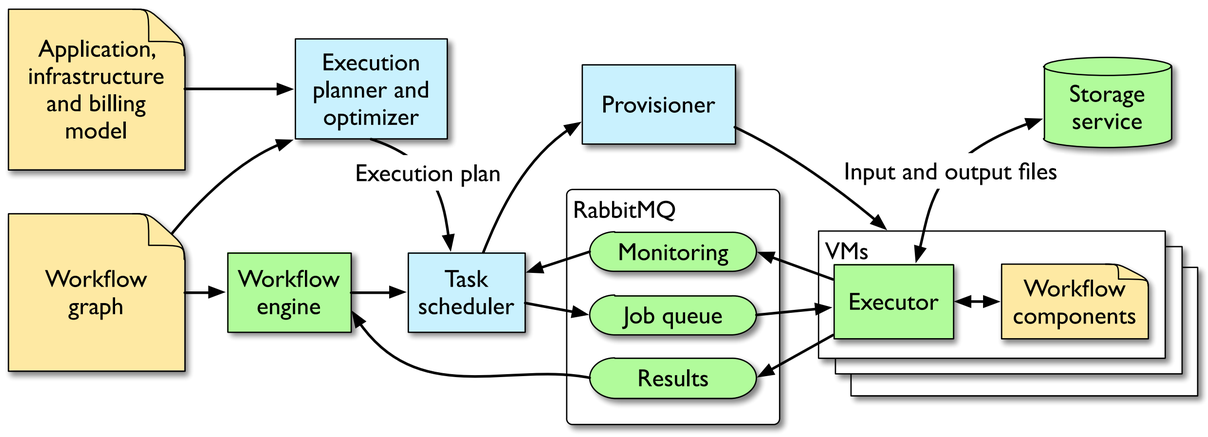

Target execution architecture

Further research

- Integrate optimizer with Hyperflow workflow engine for PL-Grid and PaaSage infrastructures.
- Investigate dynamic workflow scheduling.
- Evaluate the interplay of workflow schedulers with general-purpose cloud autoscaling systems.

References

1. M. Malawski, K. Figiela, M. Bubak, E. Deelman and J. Nabrzyski, *Scheduling multi-level deadline-constrained scientific workflows on clouds based on cost optimization*, Scientific Programming (in print), IOS Press, 2014, DOI: 10.3233/SPR-140406
2. M. Bubak, J. Kitowski, and K. Wiatr, Eds., e*Science on Distributed Computing Infrastructure*, ser. Lecture Notes in Computer Science. Cham: Springer International Publishing, 2014, vol. 8500. DOI: 10.1007/978-3-319-10894-0
3. B. Balis, *Increasing scientific workflow programming productivity with Hyperflow*, in *Proceedings of WORKS'14, 9th Workshop on Workflows in Support of Large-Scale Science*, 2015, (in print).

Acknowledgment K.F. work was supported by AGH Dean's Grant № 15/11/230/131.