

A Workflow for Fast Evaluation of Mapping Heuristics Targeting Cloud Infrastructures

Roman Ursu¹ Khalid Latif¹ David Novo¹
Manuel Selva¹ Abdoulaye Gamatie¹ Gilles Sassatelli¹
Dmitry Khabi² Alexey Cheptsov²

¹LIRMM (CNRS and University of Montpellier)

²HLRS (High-Performance Computing Center of Stuttgart)

19 January 2016



Proposal

Motivations

- ▶ Cloud infrastructures are more and more used
- ▶ Scientific workloads run on these infrastructures
- ▶ How to assess the quality of the mapping of applications ?

¹<http://www.amalthea-project.org>

Proposal

Motivations

- ▶ Cloud infrastructures are more and more used
- ▶ Scientific workloads run on these infrastructures
- ▶ How to assess the quality of the mapping of applications ?

Proposal

- ▶ Simulation framework
 - ▶ Rapid prototyping
 - ▶ Easily configurable

¹<http://www.amalthea-project.org>

Proposal

Motivations

- ▶ Cloud infrastructures are more and more used
- ▶ Scientific workloads run on these infrastructures
- ▶ How to assess the quality of the mapping of applications ?

Proposal

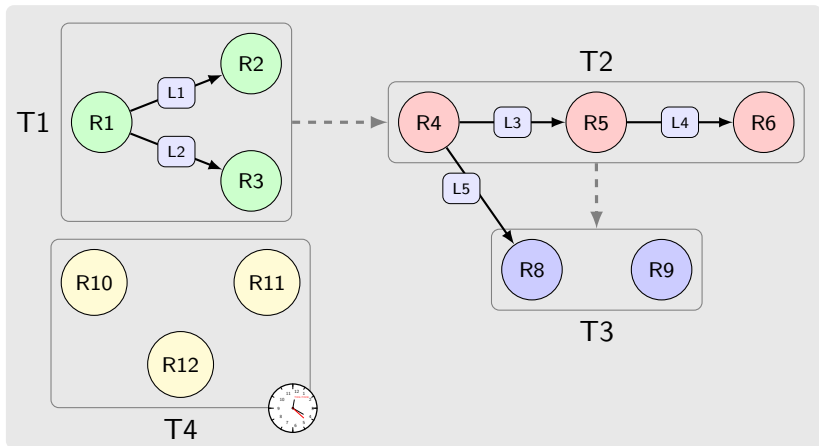
- ▶ Simulation framework
 - ▶ Rapid prototyping
 - ▶ Easily configurable

Hypothesis

- ▶ Application described using AMALTHEA¹

¹<http://www.amalthea-project.org>


AMALTHEA



Application

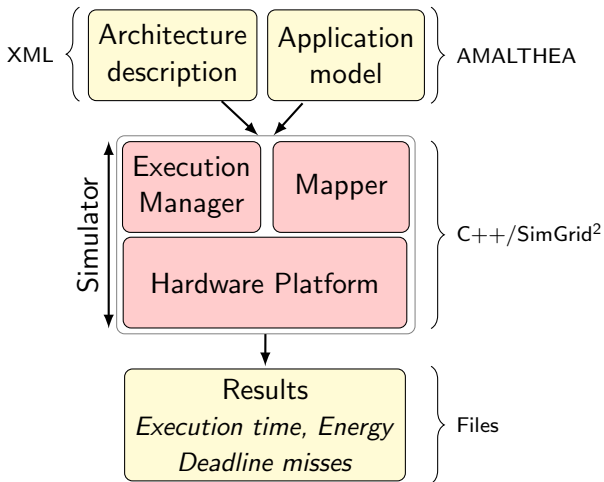


Simulated Hardware



— Mem. bus - - - CPU Interconnect — PCI — Ethernet/IB

Simulation Framework Overview




²Henri Casanova et al. "Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms". In: *Journal of Parallel and Distributed Computing* 74.10 (June 2014), pp. 2899–2917. URL: <http://hal.inria.fr/hal-01017319>.

Mapper


Responsibilities

- ▶ Map labels on memories
- ▶ Map runnables on computing units

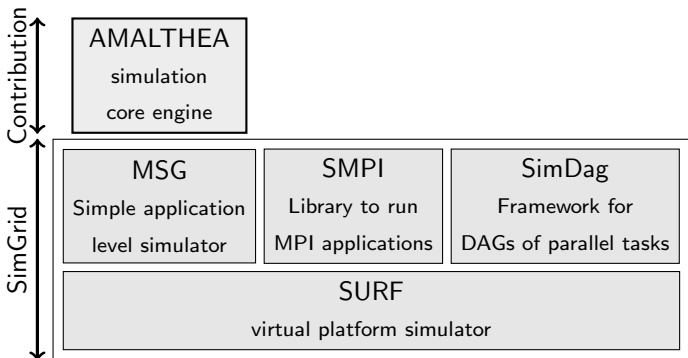


³A.K. Singh et al. "Mapping on multi/many-core systems: Survey of current and emerging trends". In: *50th ACM / EDAC / IEEE Design Automation Conference (DAC)*. 2013.

Execution Manager and Hardware Platform

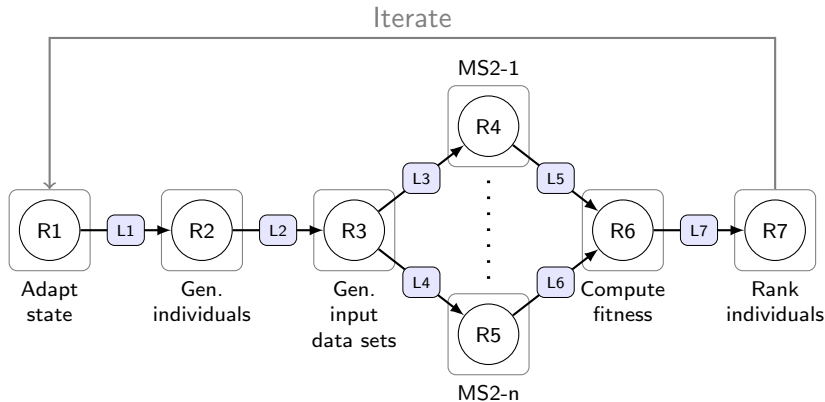


Execution Manager and Hardware Platform



- ▶ Parse and build internal representation of application
- ▶ Release runnables according to dependencies
- ▶ Launch jobs according to mapper answers
- ▶ Manage communication for label accesses

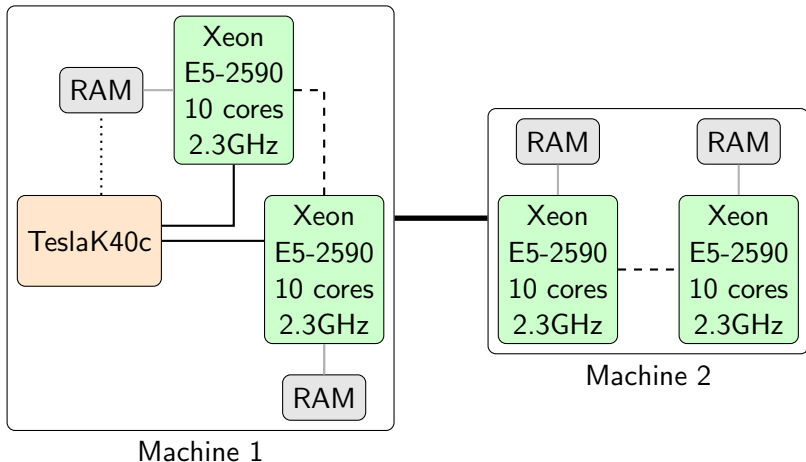
Case Study - Scientific Application



Genetic algorithm that performs molecular dynamic simulation⁴



⁴Colin W. Glass et al. "A molecular simulation tool for thermodynamic properties, new version release". In: *Computer Physics Communications* 185.12 (2014), pp. 3302–3306. ISSN: 0010-4655.

Case Study - Hardware



— 4 channels --- QPI — PCIe3.0 — InfiniBand

Case Study - Results



Conclusion

- ▶ Fully operational simulation flow
 - ▶ From application to results
- ▶ To be released soon⁵

Perspectives

- ▶ Add state of the art mapping strategies
- ▶ Add support for other application modeling formats
- ▶ Digg into the energy model provided by SimGrid

⁵<https://github.com/DreamCloud-Project>

Questions

Thank you for your attention