

Writing a Reproducible Article

Luka Stanisic and Arnaud Legrand

MESCAL team, LIG, Univ. of Grenoble

COMPAS, 22.04.2014

Context

- HPC applications nowadays use both multi-core CPUs and GPUs
- Managing efficiently computation execution and data transfer is extremely complex
- Need for portable performance \leadsto **Runtime** system

Context

- HPC applications nowadays use both multi-core CPUs and GPUs
- Managing efficiently computation execution and data transfer is extremely complex
- Need for portable performance \leadsto Runtime system

Many configuration parameters:

- ① Task granularity
- ② Scheduling strategies
- ③ Application structure

Emerging challenges:

- ① Finding optimal combination of parameters for a given machine
- ② Evaluate configurations on a wide variety of platforms
- ③ Quickly identify performance issues (e.g., bottlenecks)

Possible solution: Simulation

Our proposal

StarPU

Dynamic runtime for hybrid architectures. StarPU execution consists in scheduling a graph of tasks with data dependencies on the different computing resources


Simgrid

Versatile simulator for distributed systems

Implementation:

- StarPU applications and runtime are **emulated**
- All operations related to thread synchronization, actual computations and data transfer are **simulated**
- Control part of StarPU is modified to dynamically inject computation and communication tasks into the simulator
- StarPU calibration and platform description is used by Simgrid

Results




Results

Results

- Works fine **now**, but coming to this point was not easy
- We had to do many iterations of:
 - ① Running **complex beta code** on **several** not always **dedicated** machines
 - ② Comparing with simulations, **debugging**, **understanding**, **remodeling** and going back to step 1 until not satisfied
- With good results, we decided to make a reproducible article
 - ① From outside it looks like any other pdf paper
 - ② From inside ...

I will try to convince you that our article is not only **reproducible** but also **readable** and **understandable**!

<http://dx.doi.org/10.6084/m9.figshare.928338>

A large, stylized blue iceberg is shown floating in a body of water. The visible portion above the surface is labeled with white text: "article" at the top, "analysis" in the middle, and "data" below it. The submerged portion of the iceberg is labeled with red text: "experimentation" on the left side and "source code" at the bottom. The background consists of a light blue gradient.

article

analysis

data

experimentation

source code

article

analysis

data

Requires a **daily usage of a labbook** (org, git/svn, ...)

experimentation

source code