

Ultra Scalable Simulation with SimGrid

Program Systèmes Embarqués et Grandes Infrastructures 2008-2011

Coordinator: Martin Quinson (Nancy University)



Outcomes

Faster

- ▶ Simulation kernel orders of magnitude faster in 3.7 than in 3.3
- ▶ 10x slowdown only to simulate 1M Chord hosts on a single machine

Stronger

- ▶ Nightly builds on PIPOL (200+ tests, 10+ archs., 75% code coverage)
- ▶ Less corner cases in TCP model validation

Larger

- ▶ Simulate millions of small P2P processes on a single machine
- ▶ Simulate hundreds of large MPI processes using memory folding



Wider

- ▶ Specific models and concepts for P2P simulations
- ▶ Various communication actions to enlarge the user base

SimGrid is now ready to be the LEADER in simulation of Large-Scale Distributed Computing Systems

WP1 Models Highlights


- ▶ More validated network models




- ▶ Vivaldi + Last Mile model

WP5 Parallelization Highlights

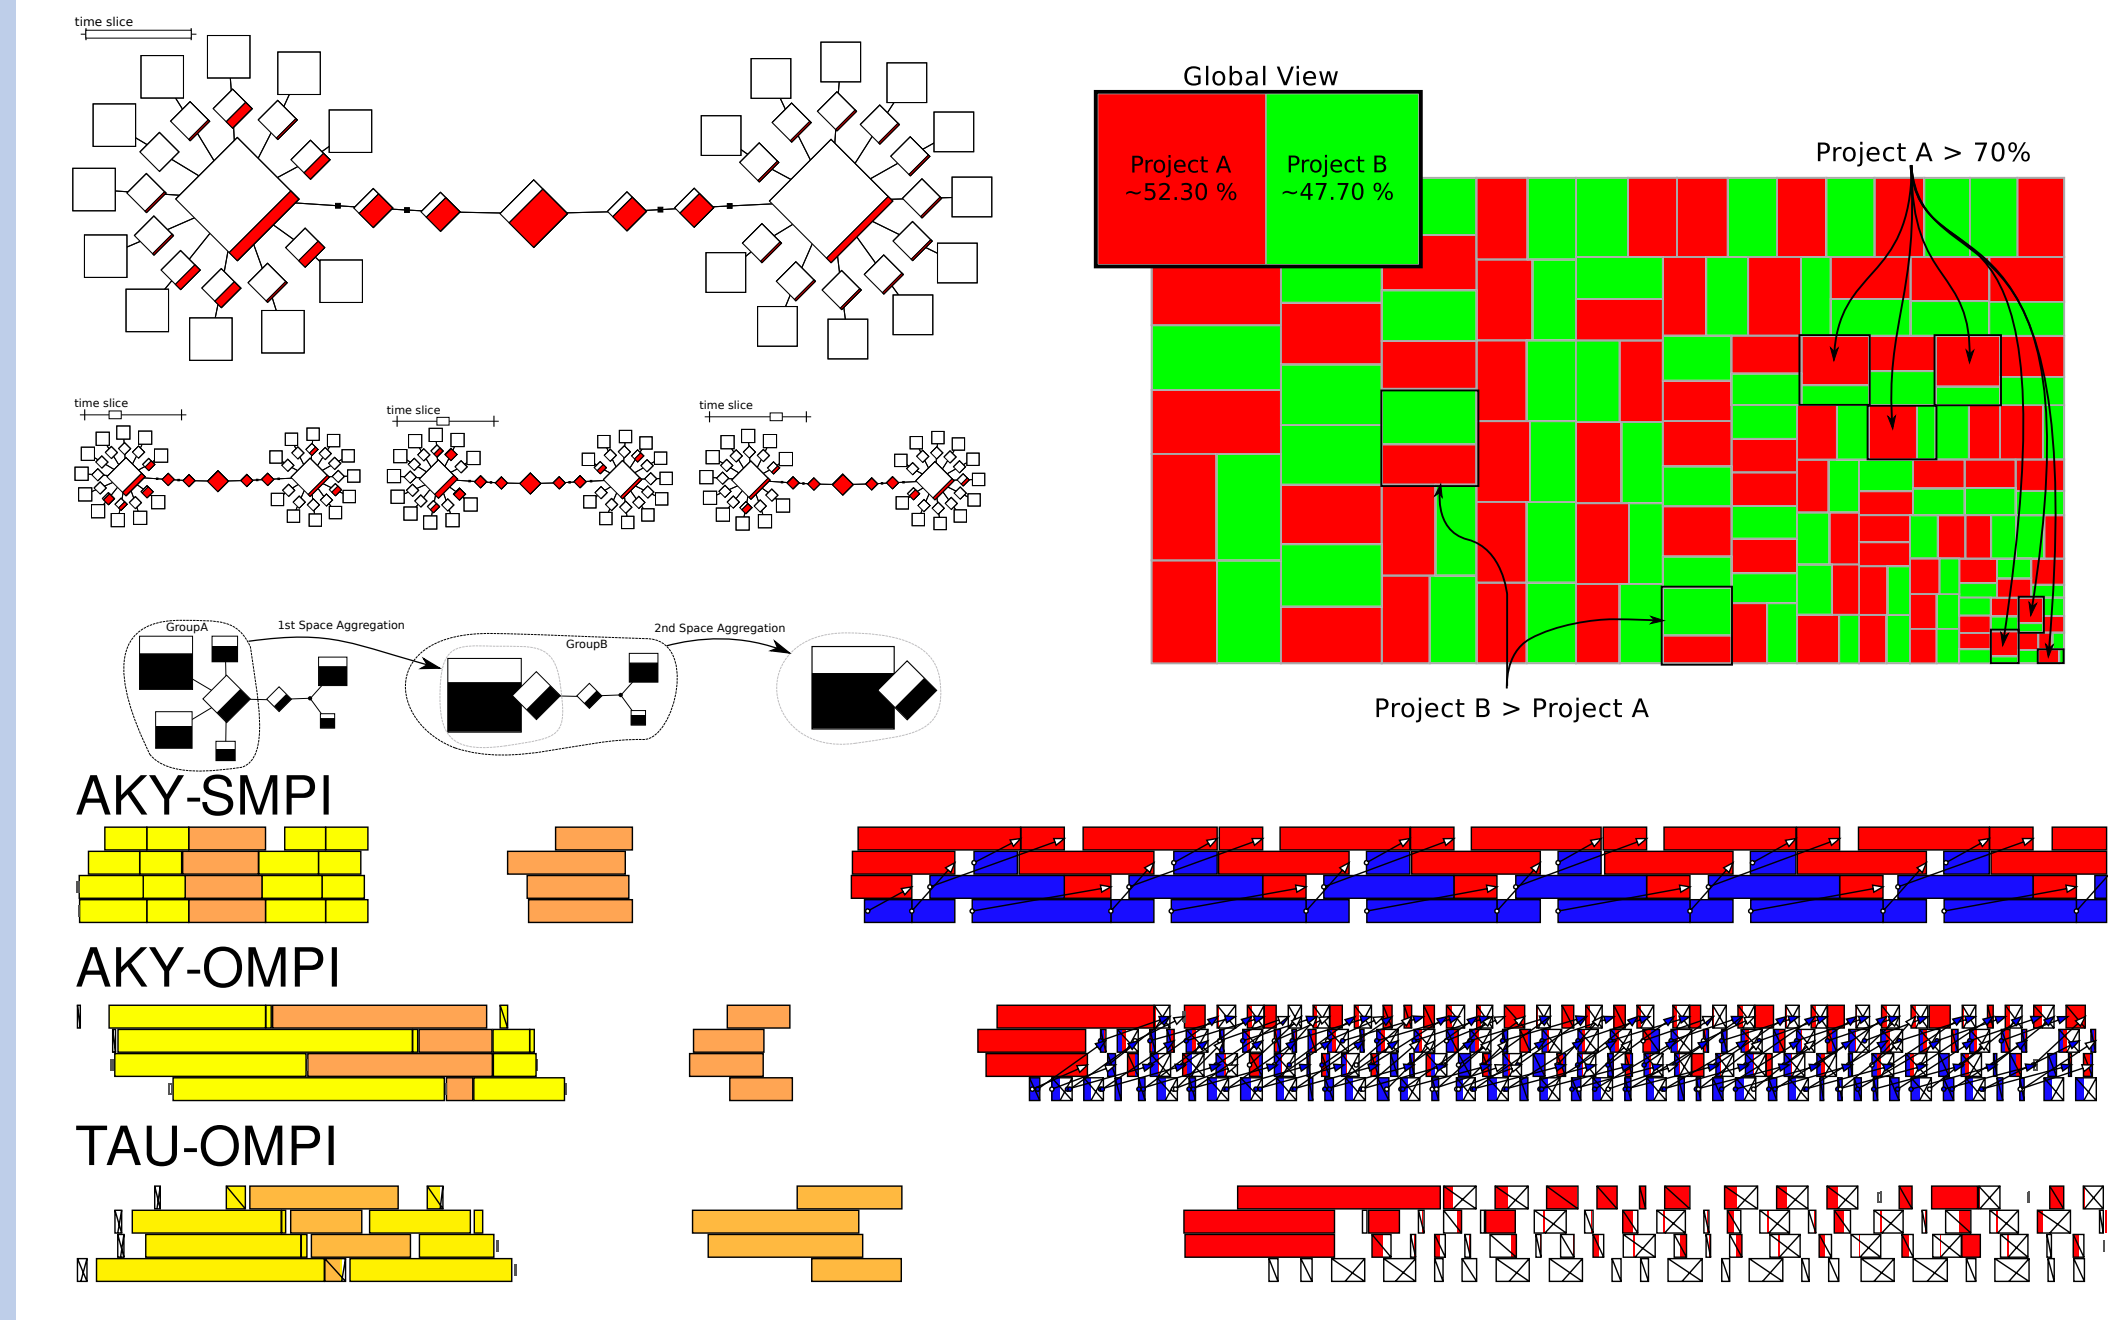
- ▶ Simulate Chord faster than P2P simulators



- ▶ User process concurrent execution

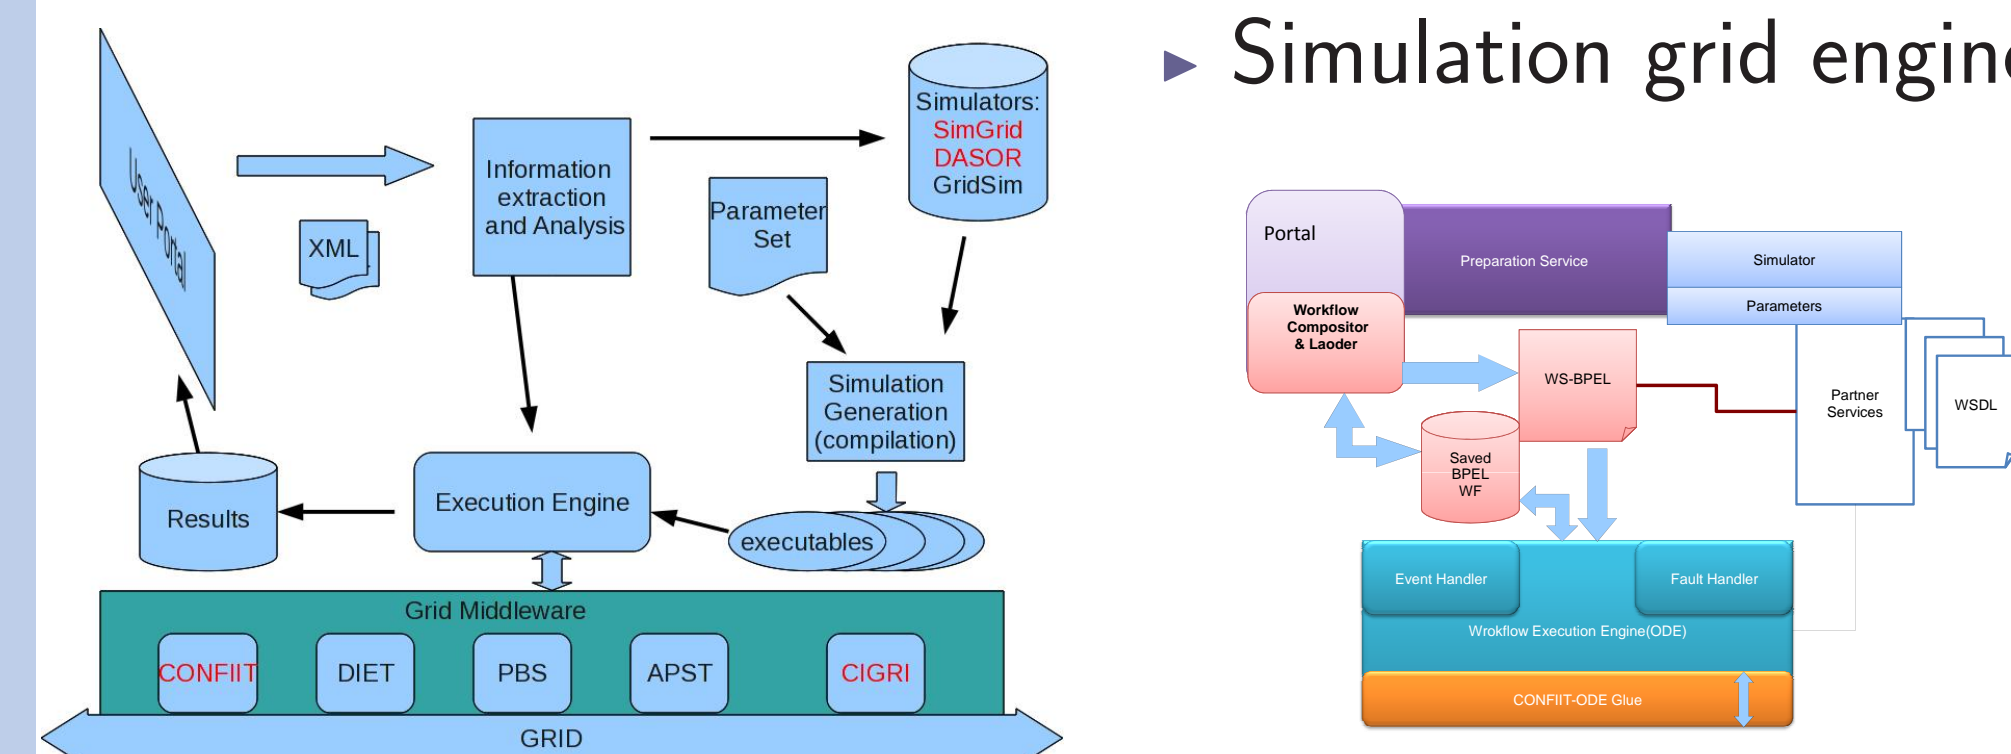


WP3 Simulation Analysis Highlight



WP4 Campaign management Highlight



- ▶ Simulation grid engine



- ▶ Web service based workflow

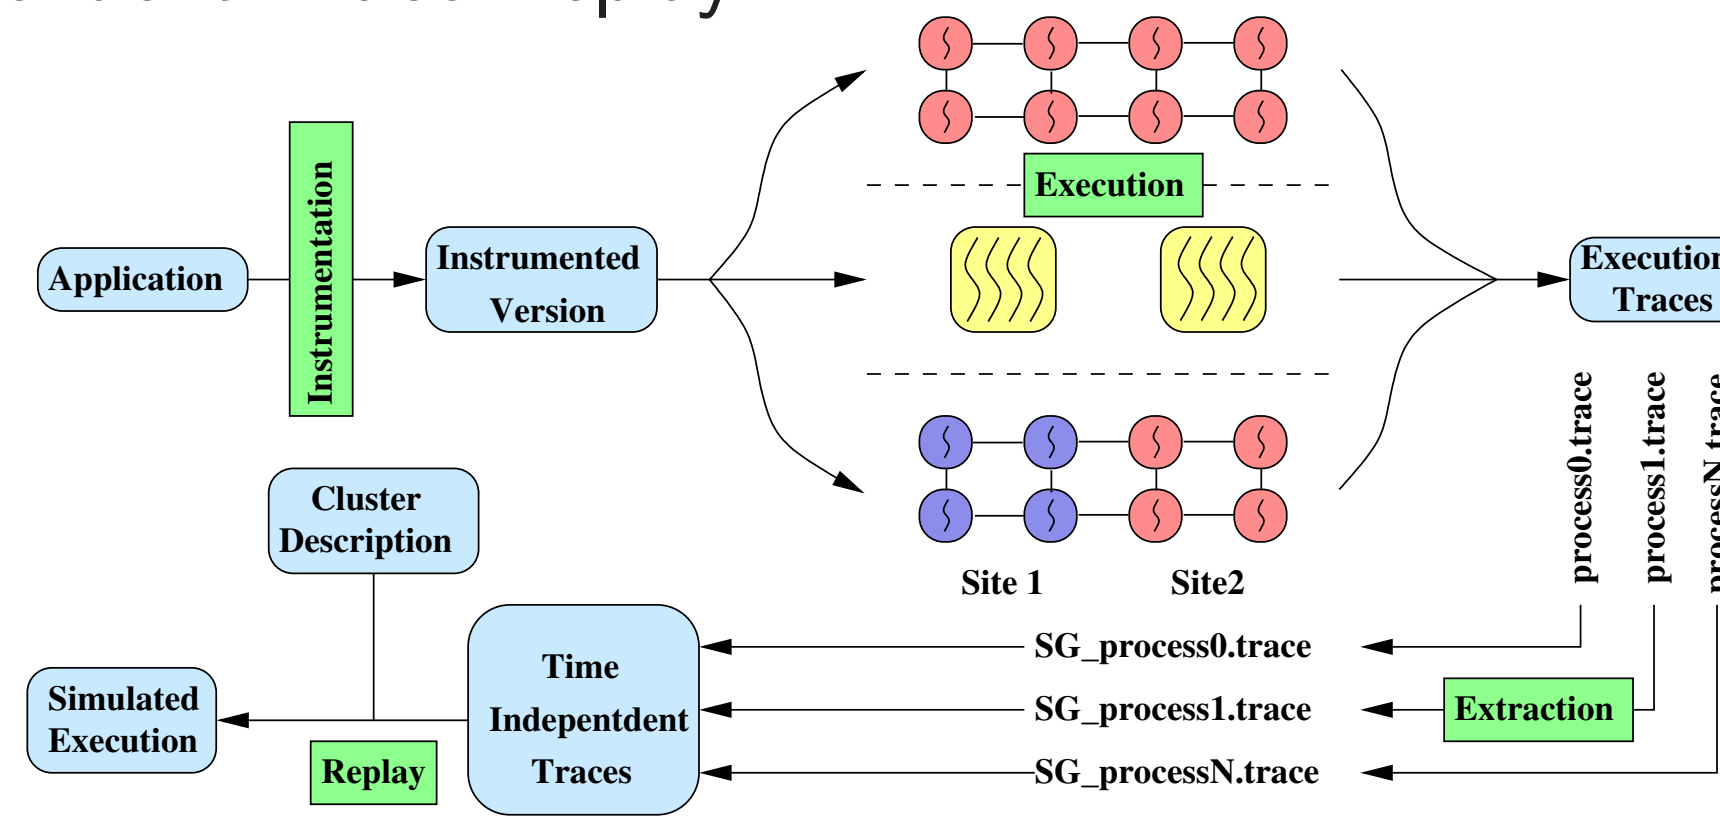
WP6 Applicability Highlight

- ▶ Online and offline MPI simulation



WP2 Model Instantiation Highlight

- ▶ Time Independent Trace Replay



Partners


- ▶ LORIA - Nancy Université
- ▶ INRIA Bordeaux Sud Ouest
- ▶ INRIA Sophia Antipolis Méditerranée
- ▶ Centre de Calcul de l'IN2P3 - CNRS
- ▶ LIG - CNRS
- ▶ INRIA Grenoble Rhône Alpes
- ▶ CRESTIC - Université de Reims Champagne Ardenne
- ▶ University of Hawai'i at Manoa



And the story continues ... with the INFRA SONGS project

Application Domains

- ▶ Data Grids
 - ▶ Distributed Data mgnt for LHC; Hierarchical Storage System
- ▶ Peer-to-Peer and Volunteer Computing
 - ▶ Replica Placement in VOD, Affinities in VC
- ▶ Cloud
 - ▶ Study from client or provider POV, other metrics (energy)
- ▶ High Performance Computing
 - ▶ Exascale, memory and energy models



Simulation Pillars

- ▶ Efficient Simulation Kernel
 - ▶ Optimization and standardization
- ▶ Concepts and Models
 - ▶ Storage, memory, energy, HPN and volatility
- ▶ Analysis and Visualization
 - ▶ Scalable visualization and trace comparison
- ▶ Support to Simulation Methodology
 - ▶ DoE, campaign management, and Open Science