

Candidature à un poste de Directeur de Recherche au CNRS

Martin Quinson

Université de Lorraine

March 6, 2013

Computational Science of Computer Systems

Curriculum Vitæ

Martin Quinson, 37 years old, Married, 3 kids.

Professional Experience

- ▶ Since 2011: On leave at Inria, Nancy
- ▶ Since 2005: Assistant Professor of the Université de Lorraine (LORIA)
- ▶ 2004: Temporary teaching assistant at Université de Grenoble (LIG)
- ▶ 2004: Post-doc at University of California, Santa Barbara (USA)

Education

- ▶ Ph.D. ENS-Lyon, december 2003
 - ▶ M.S. ENS-Lyon, 2000
 - ▶ B.S. Université de Saint-Étienne, 1998
- I defend my HDR at Université de Lorraine in two days (on 8. March 2013)

Research Topic: Methodologies of Experimentation

- ▶ Assessing the performance and correctness of Large-Scale Distributed Systems

Research Context

Scientific
Objects

Large Scale Distributed Systems

- ▶ Scientific Computing
- ▶ High Performance Computing
- ▶ Grids
- ▶ Peer-to-peer Systems
- ▶ Volunteer Computing
- ▶ Clouds

Scientific
Questions

Research Context

Scientific Objects

Large Scale Distributed Systems

- ▶ Scientific Computing ▶ High Performance Computing ▶ Grids
- ▶ Peer-to-peer Systems ▶ Volunteer Computing ▶ Clouds




Scientific Questions

- ▶ User/Provider ▶ Time/Energy
- ▶ Throughput/Makespan/#Msg
- ▶ Worst case/Avg/Amortized

Performance

Correction

- ▶ Safety: bad things don't happen
- ▶ Liveness: good things do happen



Research Context

Large Scale Distributed Systems

Scientific Objects

- ▶ Scientific Computing ▶ High Performance Computing ▶ Grids
- ▶ Peer-to-peer Systems ▶ Volunteer Computing ▶ Clouds

Scientific Questions

Performance

- ▶ User/Provider ▶ Time/Energy
- ▶ Throughput/Makespan/#Msg
- ▶ Worst case/Avg/Amortized

Correction

- ▶ Safety: bad things don't happen
- ▶ Liveness: good things do happen

Methodo.

- ▶ Theoretical proofs
- ▶ Direct execution
- ▶ Experimental facilities
- ▶ Simulation
- ▶ Emulation
- ▶ Tests (manual/automated)
- ▶ Theorem proving
- ▶ Model checking
- ▶ Dynamic verification

Research Context

Large Scale Distributed Systems

Scientific Objects

- ▶ Scientific Computing ▶ High Performance Computing ▶ Grids
- ▶ Peer-to-peer Systems ▶ Volunteer Computing ▶ Clouds

Scientific Questions

Performance

- ▶ User/Provider ▶ Time/Energy
- ▶ Throughput/Makespan/#Msg
- ▶ Worst case/Avg/Amortized

Correction

- ▶ Safety: bad things don't happen
- ▶ Liveness: good things do happen

Methodo.

- ▶ Theoretical proofs
- ▶ Direct execution
- ▶ Experimental facilities
- ▶ Simulation
- ▶ Emulation
- ▶ Tests (manual/automated)
- ▶ Theorem proving
- ▶ Model checking
- ▶ Dynamic verification

My Research Interests: Experimental Methodologies

- ▶ Meta-research about how to produce scientifically sound research
- ▶ Strive at developing ready-to-use tools addressing methodological challenges


SimGrid: Simulator of Large-Scale Applications

Scientific Instrument to assess LS Apps through simulation

Idea to test




Experimental setup



Model



Scientific Results



Main Features

- ▶ **Versatile:** Grid, P2P, HPC, Volunteer Computing, Clouds, ...
- ▶ **Valid:** Accuracy limits studied and pushed further for years
- ▶ **Scalable:** 3M chord nodes; 1000× faster than other (despite precise models)
- ▶ **Usable:** Tooling (generators, runner, visu); Open-Source, Portable, ...

Large Established Project

- ▶ **Impact (2008-2012):** 63 publications (98 distinct authors, 4 continents), 4 PhD
- ▶ **Community-driven:** 26 contributors (5 not affiliated), 5 contributed tools
- ▶ **Co-leader** with A. Legrand (CNRS Grenoble) and F. Suter (CNRS IN2P3)

Major Contributions

SimGrid Scalability

- ▶ **Fast Enough:** Innovative parallelization, efficient algorithms and implem.
- ▶ **Big Enough:** Scalable and versatile platform representation

Correctness Studies in SimGrid

- ▶ Seamless integration of a complete Model Checker (enforces code invariants)
- ▶ Exhaustive reachability analysis, with innovative versatile DPOR technique

Assessing Real Applications

- ▶ **GRAS:** Middleware to run simulation prototypes on real platforms
- ▶ **SMPI:** Study real MPI applications within SimGrid

Scientific Community Management


- ▶ **Project Coordinator:** 2 ANR projects, 1 regional CPER project (total: 4M€)
- ▶ **Methodological convergence:** Board member of Grid'5000 experimental grid
- ▶ **Scientific Animation (SimGrid, Grid'5000):** 4 summer schools, 3 R&D engineers

Also leader of several projects in scientific outreach, pedagogical tools and didactic

Zoom: Parallel P2P Simulation


State of the Art: dPeerSim (best known PDES of P2P)

- ▶ Spreads the load over Logical Processes (LP)
- ▶ 2LP \sim 4h / 16LP \sim 1h. But 47s in sequential PeerSim!!




Our approach: Run user processes in parallel; keep kernel sequential


- ▶ Processes separation through a OS-oriented approach: **simcalls**



Functional View




Temporal View



Implementation

Evaluation: Chord simulations (normalized workload)



Simulator	size	time (seq/par)
PeerSim	100k	4h36
OverSim	300k	10h
SimGrid	300k	32mn
	2M	6h23 / 5h50

Early Track Record

Publications and other Professional Activities

- ▶ 5 journals: PPL '10 on methodologies (cited 22), ParCo'01 (cited 26).
- ▶ 2 book chapters: HDR chapter in book on philosophy of sciences
- ▶ 28 conferences: CCGrid'12 (twice, rate: 27%), Forte'11, IPDPS'11 (rate: 19%, cited 6), UkSim'08 (cited 213), EuroPar'07 (rate: 26%, cited 14), PDCS'06 (best paper, cited 19), IPDPS'02 (cited 61), EuroPar'02 (cited 118).
- ▶ Dissemination: 4 Keynotes, 1 tutorial, 4 invited presentations
- ▶ Supervision of Research: 4 post-docs, 2 Ph.D., 11 masters, 4 R&D engineers
- ▶ Teaching: OS and programming at Telecom Nancy; Coordinator of first year

Service to the Community

- ▶ Leader of research group on Grids performance (10 members) since 2013
- ▶ 7 ACM/IEEE conferences: CCGrid, IPDPS, SIMUTools, PADS
- ▶ 3 international workshops: WATERS'13 (Analysis Tools and Methodologies)
- ▶ Committee member: 7 Ph.D. (France and Belgium), 4 recruitments.
- ▶ Leader of a working group on the didactic of computer science (6 members)

Research Program

Computational Science of Computer Systems convergence of Simulation, Dynamic Verification and Emulation

1/ Modeling of Large-Scale Systems

- ▶ Scalability and Accuracy still not enough for Exascale studies
- ▶ Semantic modeling of MPI 3.0 collectives (implementation-depend)

2/ Simulation of Real Applications

- ▶ OS Virtualization layer for the simulation of legacy code
- ▶ Distributed simulation, and increase parallelism in our simulation

3/ Formal Methods for Large-Scale and HPC Systems

- ▶ Liveness properties on legacy code (OS-level introspection tooling)
- ▶ Domain-specific properties and reduction techniques

4/ Scientific Instrument and Open Science

- ▶ Produce a *de facto* standard tool, with associated tools
- ▶ Foster the emergence of a vivid research community, with best practices

(ERC project currently under evaluation)