ROMA

Resource Optimization: Models, Algorithms, and Scheduling

Application resilience

Problems and objectives

- Most powerful supercomputers: more than one failure per day
- Resilience: ability to produce correct result in spite of faults
- Extensive experimental campaigns are too expensive
- How to choose a resilience protocol? How to tune it?

Methods

- Probabilistic analyses
  - Exact derivations
  - First-order approximations
- Discrete-event simulations
- Applied to various protocols/techniques:
  - Fault prediction
  - Silent error detection
  - Checkpointing, replication, migration, error correction

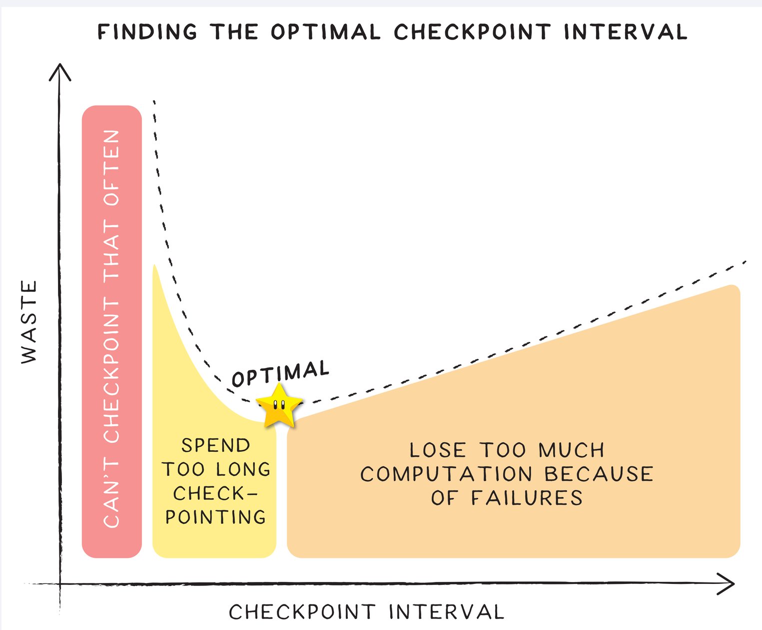

Solvers for sparse linear algebra and related optimization problems

Direct solvers for sparse linear systems

- Target performance and numerical robustness
- Limit resource consumption (e.g., memory)
- Exploit low-rank representations to reduce complexity

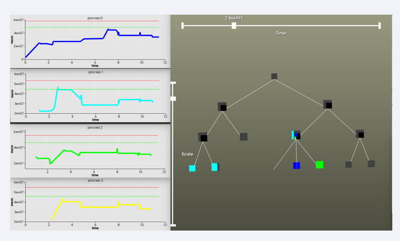

Task scheduling and memory

Combinatorial scientific computing

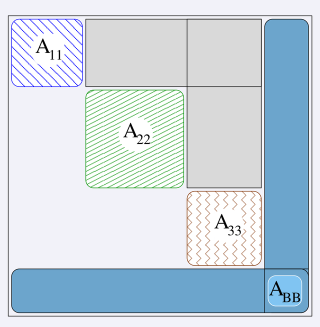

Design, analysis, and implementation of combinatorial algorithms to enable scientific computing

- Matchings and partitioning in graphs and hypergraphs
- High performance computing with matrices and tensors

The MUMPS solver

- Solve systems of linear equations $AX = B$ (A sparse)
- Owners: CERFACS, CNRS, ENS Lyon, INPT, Inria, Univ. Bordeaux
- Software platform to experiment and transfer research
- CeCILL-C license, used worldwide, included within many open-source and commercial packages
- Many features, addresses a wide range of problems: symmetric, unsymmetric, indefinite, multiple (sparse) right-hand sides, Schur, . . .

The MUMPS consortium

- Founded by CERFACS, ENS Lyon, INPT, Inria (manager), Univ. Bordeaux
- Industrial funding in exchange of services
- 11 members

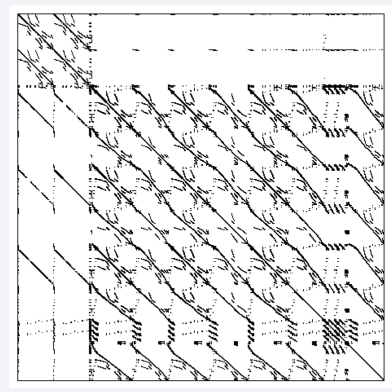

http://mumps-solver.org *http://mumps-consortium.org*

Multi-criteria scheduling strategies

Mix user-oriented objectives (time-to-solution, throughput) with platform-oriented constraints (energy, memory)

Energy-aware algorithms

- Energy consumption of fault-tolerance protocols
- Use of Dynamic Voltage and Frequency Scaling (DVFS)
- Powering cores below nominal voltages

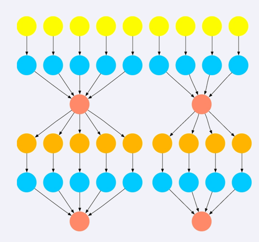

Memory-aware algorithms

- Minimize memory or I/O when processing task graphs
- Memory minimizing traversals for trees and SP-graphs
- Limit memory-peak usage of parallel execution
- Allow for efficient runtime schedulers

Hybrid CPU-GPU scheduling

- Task or task-graph scheduling
- Low-cost scheduling algorithms
- Online scheduling with guarantees

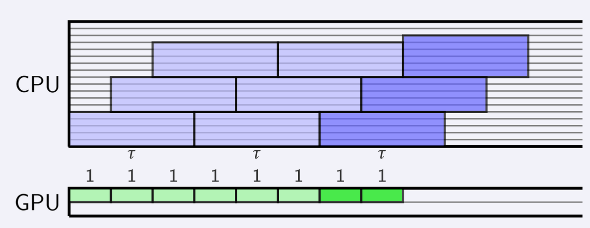

ROMA Members

Permanent members

- Anne Benoit (ENS Lyon assoc. prof.)
- Jean-Yves L'Excellent (INRIA researcher)
- Loris Marchal (CNRS researcher, co-head)
- Yves Robert (ENS Lyon prof. & UT Knoxville)
- Bora Uçar (CNRS researcher)
- Frédéric Vivien (INRIA senior researcher, co-head)

Engineers

- Marie Durand, Chiara Puglisi and Guillaume Joslin (Mumps consortium)

PhD Students

- Yiqin Gao (MILYON Labex)
- Changjiang Gou (CSC, China & ENS Lyon)
- Li Han (CSC, China & ENS Lyon)
- Aurélie Kong-Win-Chang (ENS Lyon)
- Gilles Moreau (MILYON and MUMPS consortium)
- Ioannis Panagiotas (Inria)
- Filip Pawłowski (CIFRE at Huawei)

Administrative assistant

- Evelyne Blesle / Laetitia Lecot (Inria)

Main collaborations

- ANR Solhar
- INPT-IRIT, Toulouse.
- Joint Lab. on Extreme Scale Computing
- Univ. of Tennessee, Knoxville (USA)
- Georgia Tech. (USA)
- Huawei Technologies, Paris

www.ens-lyon.fr/LIP/ROMA/