Parallel Numerical Solvers on top of Runtime Systems

HiePACS, Runtime, SAGE

C2S@Exa
Computer and Computational Sciences at Exascale

Algorithm - Runtimes - Kernel

The goal of this initiative is to design dense and sparse linear algebra methods that achieve the fastest possible time to an accurate solution on large-scale emerging parallel heterogeneous multicore platforms, using all the processing power that future high end systems can make available. We propose a framework for describing numerical linear algebra algorithms at a high level of abstraction and delegating the actual execution to a runtime system in order to design software whose performance is portable across architectures.

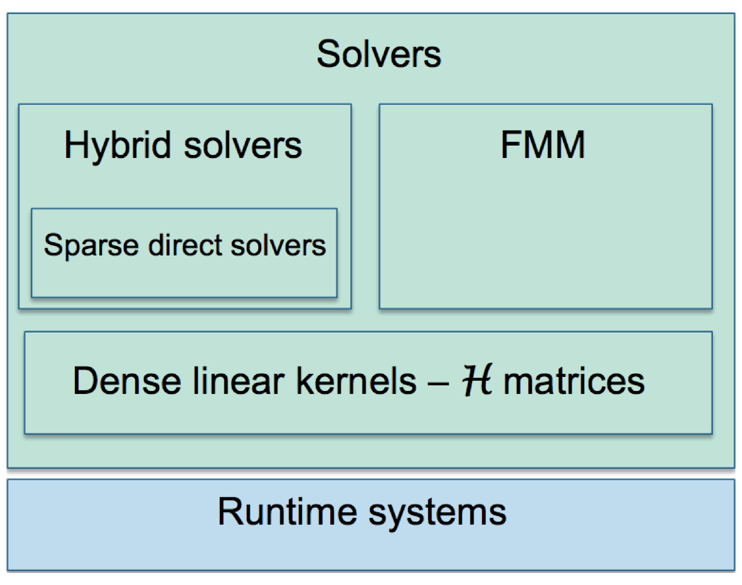

MAGMA MORSE

MATRIX ALGEBRA ON GPU AND MULTICORE ARCHITECTURE

The MAGMA project aims to develop a dense linear algebra library similar to LAPACK but for heterogeneous/hybrid architectures. We present here how MAGMA has been extended using the StarPU / Quark runtime systems in order to handle multicore nodes enhanced with multiple GPUs. The MORSE extension of MAGMA includes one-sided factorizations (LLT, LU, QR) and relative solvers as well as main level-3 BLAS operations.

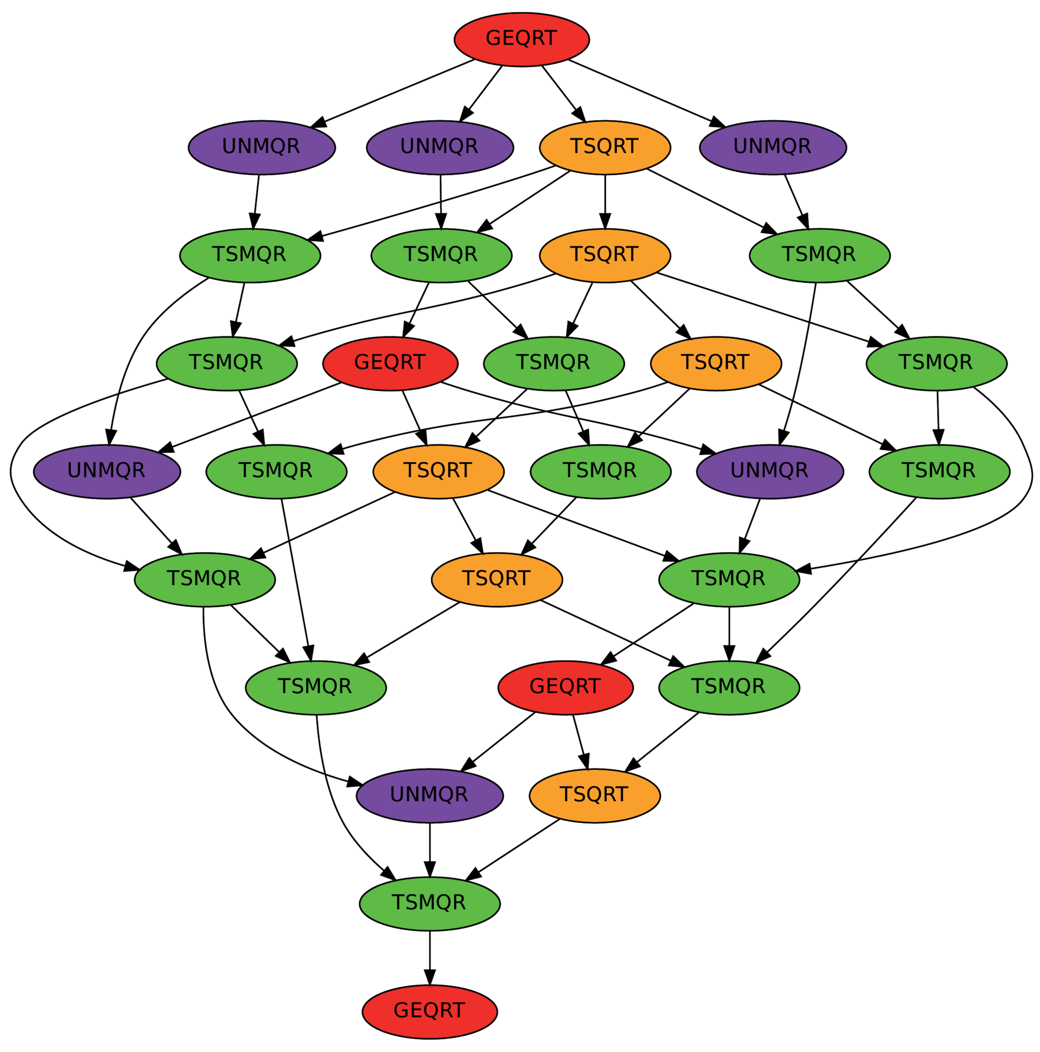

PaStiX

PARALLEL SPARSE MATRIX PACKAGE

The Parallel Sparse matrix package (PaStiX) is a scientific library that provides a high performance parallel solver for very large sparse linear systems based on direct methods.
Numerical algorithms are implemented in single or double precision (real or complex) using LLT, LDLT and LU factorizations. It also provides an adaptive blockwise parallel iLU(k) factorization. Using a generic runtime (such as StarPU or PARSEC), the solver fully exploits heterogeneous nodes (multicore and multiGPU).

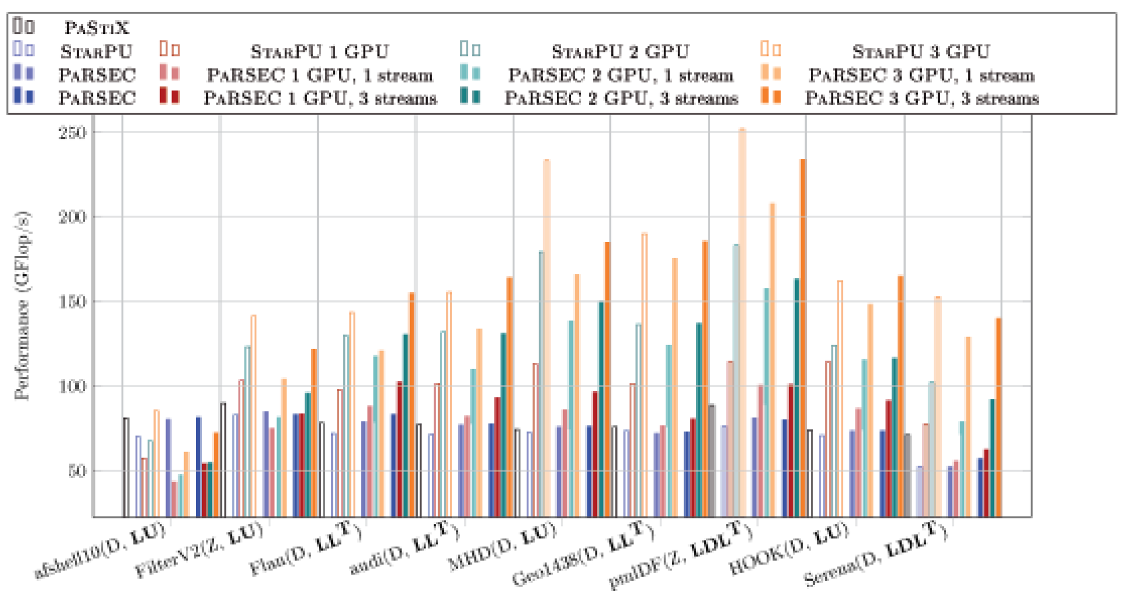

ScalFMM

PARALLEL FAST MULTIPOLE LIBRARY FOR LARGE SIMULATIONS

The Parallel Fast Multipole Library for Large Scale Simulations (ScalFMM) library aims at simulating N-body interactions using the Fast Multipole Method (FMM). This software intends to offer all the functionalities needed to perform large parallel simulations while enabling an easy customization of the simulation components: kernels, particles and cells.

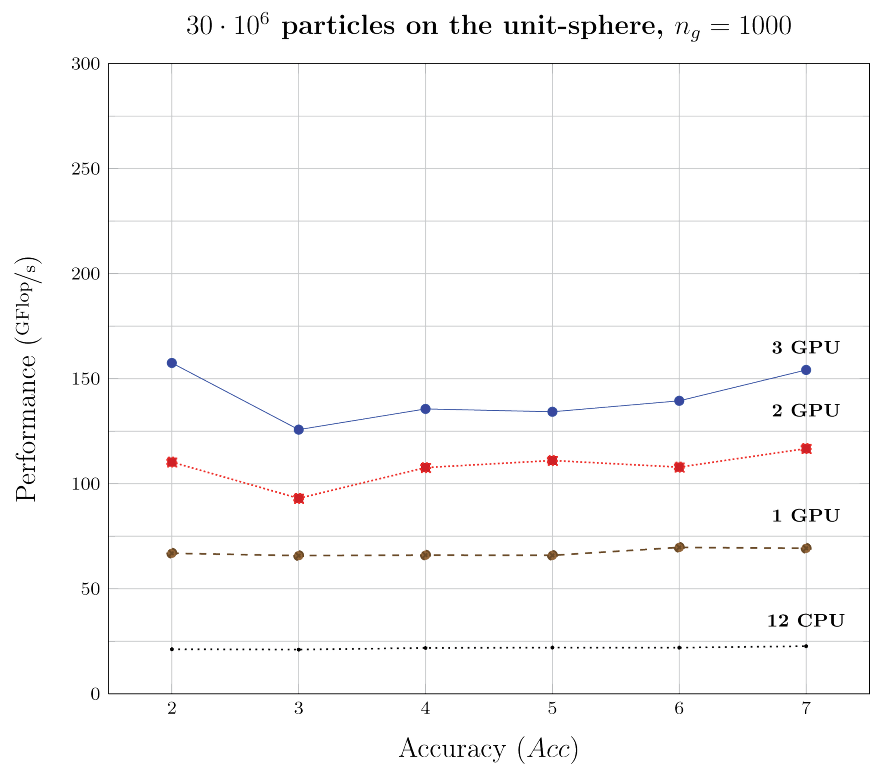

Hybrid

One route to the parallel scalable solution of large sparse linear systems in parallel scientific computing is the use of hybrid methods that hierarchically combine direct and iterative methods.
These techniques inherit the advantages of each approach, namely the limited amount of memory and natural parallelization for the iterative component and the numerical robustness of the direct part.
Variants of Krylov suspace methods suited for parallel computing such as Newton-GMRES with adaptative déflation or augmentation, are developed.
Part of this activity is conduced in the FP7 project Exa2CT (www.exac2ct.eu)

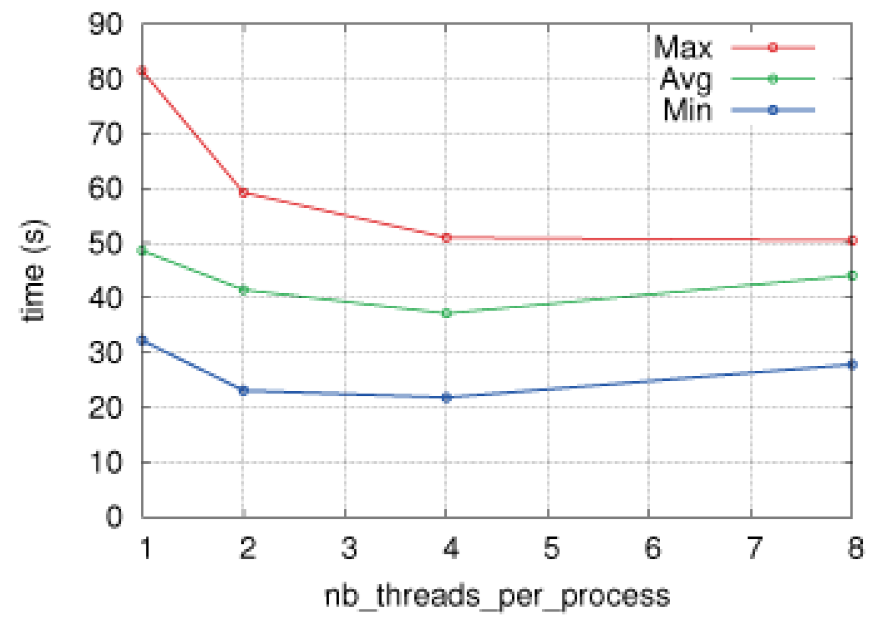

MORSE Inria associate team FastLA Inria Associate team