

Sparse direct solver on top of large-scale multicore systems with GPU accelerators

Solving large sparse linear systems is a crucial and time-consuming step, arising in many numerical simulations. Consequently, many parallel sparse matrix factorization techniques have been studied.

Since sparse direct solvers are built with dense linear algebra kernels, we are writing prototype solvers on top of PLASMA and MAGMA libraries. Modern architectures induce strongly hierarchical topologies, and we have developed dynamic scheduling designed for those architectures in the PASTIX solver. We are evaluating a way to replace this scheduler by a generic framework (DAGUE or STARPU) to execute the sparse factorization task graph. This work aims at designing algorithms and parallel programming models to implement direct methods on emerging GPU-equipped computers.

Keywords: Sparse linear algebra, Direct solver, High performance computing, GPU accelerators.

References

- [1] George, T., Saxena, V., Gupta, A., Singh, A., Choudhury, A.R.: Multifrontal factorization of sparse spd matrices on gpus. In: IPDPS. (2011) 372–383
- [2] Yu, C.D., Wang, W., Pierce, D.: A cpu-gpu hybrid approach for the unsymmetric multifrontal method. *Parallel Computing* **37**(12) (2011) 759 – 770
- [3] Faverge, M., Ramet, P.: Dynamic scheduling for sparse direct solver on numa architectures. In: Proceedings of PARA’2008, Trondheim, Norway (2008)
- [4] Hénon, P., Ramet, P., Roman, J.: PaStiX: A High-Performance Parallel Direct Solver for Sparse Symmetric Definite Systems. *Parallel Computing* **28**(2) (2002) 301–321

Xavier Lacoste

INRIA Bordeaux, 351 cours de la Liberation, 33405 Talence Cedex

Xavier.Lacoste@inria.fr

Pierre Ramet

University Bordeaux, 351 cours de la Liberation, 33405 Talence Cedex

ramet@labri.fr