

Exceptional service in the national interest



HTS: A Multithreaded Direct Sparse Triangular Solver Combining Level Scheduling and Recursive Blocking

Andrew M. Bradley

Thanks: E. Boman, J. Booth, C. Dohrmann, S. Hammond, W. Held, M. Heroux, R. Hoekstra, K. Kim, S. Olivier, A. Prokopenko, S. Rajamanickam


SAND2015-XXXX



Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP


Problem Statement

- Solve $R P T Q x = b$
 - upper or lower sparse triangular matrix T
 - row scaling R
 - permutations P, Q
 - solution and RHS x, b
- (Everything that is needed for LDL, LU, incomplete factorizations, etc.)
- For many sequential RHS with
 - the same T
 - the same nonzero pattern $\text{pat}(T)$
- On a multi/many-core node





Solution Approach



- Direct
 - (Not iterative)
- Symbolic phase
 - Find parallelism in $\text{pat}(T)$, the graph of T
- Numerical phase
 - Load data structures with numbers
- Solve phase





Algorithms: Level Scheduling



Algorithms: Level Scheduling






Algorithms: Pruned Point-to-Point







Park, J., M. Smelyanskiy, N. Sundaram, and P. Dubey., "Sparsifying synchronization for high-performance shared-memory sparse triangular solver." In *Supercomputing*, pp. 124-140. Springer International Publishing, 2014.



Algorithms: Recursive Blocking





Algorithms: Hybrid




Algorithms: Hybrid




Algorithms: Hybrid




- Level schedule solve: Park et al.'s method
- Recursive blocking
 - Also with point-to-point
- Switching method
 - Currently: smoothed level size with threshold
- Currently: OpenMP and C++98




Speedup w.r.t. MKL trisolver




Numerical phase w.r.t. parallel solve




Symbolic phase w.r.t. serial solve




Speedup w.r.t. MKL trisolver




Numerical phase w.r.t. parallel solve








Symbolic phase w.r.t. serial solve



UMFPACK LU, Ivy Bridge




UMFPACK LU, Knights Corner, AFFINITY=compact








Ordering

Level-0 incomplete Cholesky, Knights Corner (Compton)


Speedup w.r.t. MKL trisolver




Level-0 Cholesky, NodeND, Ivy Bridge

Level-0 Cholesky, NodeND, Knights Corner, AFFINITY=compact

Numeric phase w.r.t. parallel solve

Numeric phase w.r.t. parallel solve

Symbolic phase w.r.t. serial solve

Symbolic phase w.r.t. serial solve


UMFPACK LU on UF matrices: Ivy Bridge, 819 matrices



UMFPACK LU on UF matrices: Knights Corner, 814 matrices, all cores



Future Work

- Point-to-point level scheduling
 - Group rows into tasks to minimize #dependencies
 - Size tasks to reflect level of synchronization
- Hybrid
 - Switching method(s)
 - Does not have to be 3 blocks; could alternate
- HTS
 - Kokkos interface
 - Support a variety of input formats
 - Pure Kokkos implementation?
- Direct sparse methods on GPU?

