CharmVSS: Very Sparse Direct Linear Solver in Charm++

Presented by: Eric Bohm

CharmVSS

- Collaboration between CASI Corp and Charmworks Inc to produce a distributed memory very sparse direct linear solver.
  - Baseline starting point: shared memory parallel very sparse solver
- Goals:
  - Distributed Memory to solve problems too large for single host
  - Greater parallelization to improve time to solution
  - Retain factor in memory to run many RHS

CharmVSS decomposition

- Blocknodes represent contiguous rows with shared non-zero pattern
- Provide natural decomposition approach
- Grainsize too small for very sparse case as many blocknodes have small number (perhaps 1) row with a small number of nonzeros.
- Blocksets represent contiguous rows of Blocknodes
  - Controllable granularity by selecting max set size in rows, or in total number of nonzero values.

Message handling

- Shared messages:
  - Exposed parallelism at any stage is based on having data dependencies satisfied
  - Most blocknodes have multiple upward and downward dependencies
  - Therefore the data (factor or solve) can be reused by many blocksets within a node.
- Message throttling:
  - Factor data can be a significant fraction of host memory
  - Per host credit scheme prevents the creation of a message until there is sufficient credit on the destination to hold it.
  - Message deletion (when all dependents are done) restores credit

Charm++ Features in CharmVSS

- Zero copy protocol
- NodeGroup
  - Host level shared index
  - Host level shared RHS
  - Host level shared solution
- OmpCharm
  - Trigger MKL DGEMM for large blocknodes in factorization
  - Loop parallelization of beyond extent computation in backward solve
  - Evaluating MKL DDOTI vs DGTHR DDOT in solve

Performance Data

- Platform is NCSA's iForge commercial cluster normal queue
  - 64 GB per node
- Target Problem:
  - Nodes: 231580972
  - OffDiagonals: 687797827
  - Number of Separating Levels: 8
  - Indistinguishable Row Blocks: 173598430
  - Nonzeros in Upper Matrix : 9.19379e+08
  - Nonzeros in Factor: 2.31862e+10
  - Fill Ratio: 25.2194
  - Factor Index Size (MB): 11179.6
  - Factor Value Size (MB): 185490
  - Factor Size (MB): 202227

Factorization Scaling on iForge

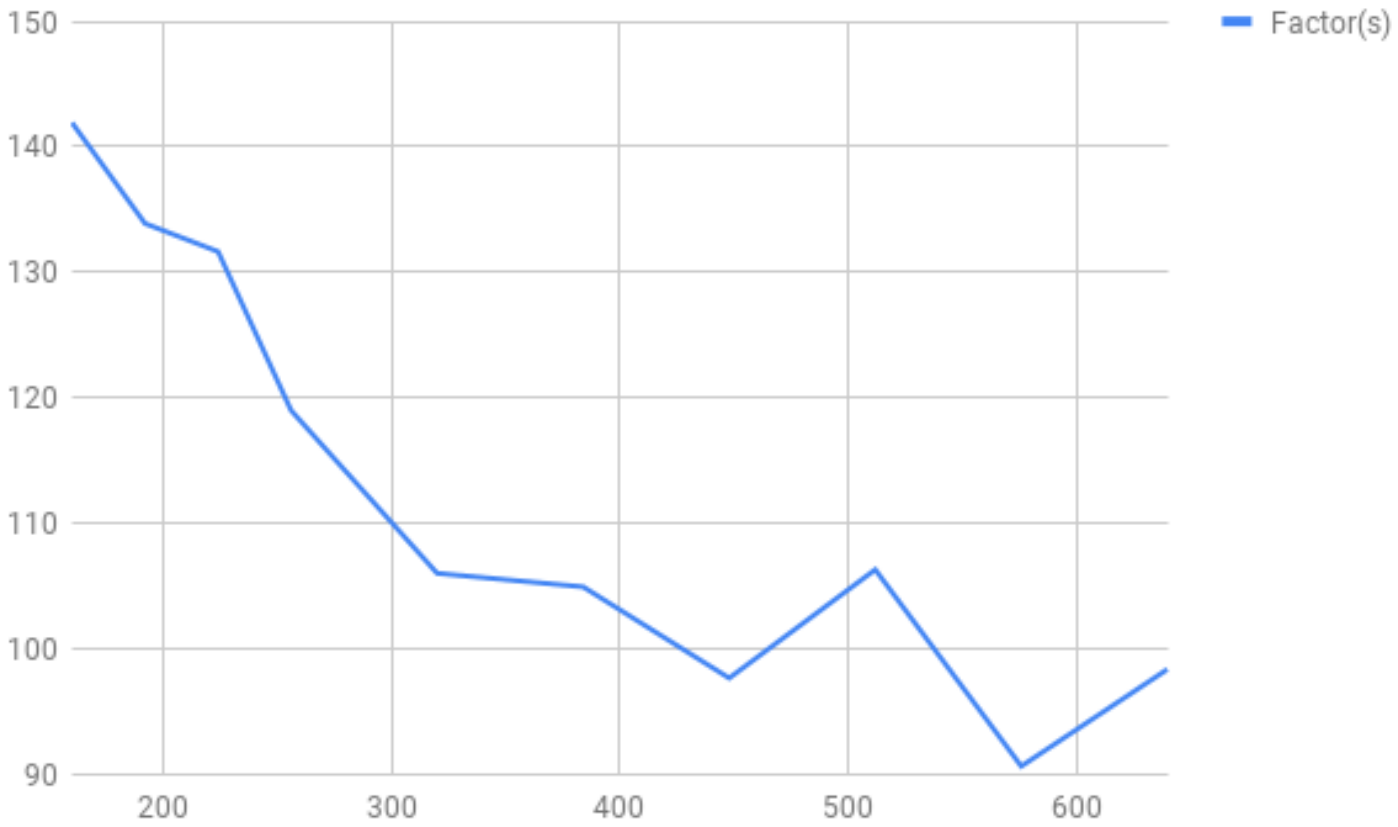

Future work

- Improve scaling of solve
  - OpenMP + MKL
- Improve parallelization of launch
- Integrate with shared memory framework
- Multiple RHS
  - Pipeline and/or aggregate solve
- Compare to other solver codes