Leveraging One-Sided Communication for Sparse Triangular Solvers

Nan Ding, Samuel Williams, Sherry Li, Yang Liu

We implement and evaluate ***one-sided communication-based, distributed-memory Sparse Triangular Solvers (SpTRSV)*** and construct ***a critical path performance model*** in order to assess our observed performance relative to machine capabilities.

Motivation

- SpTRSV is used in conjunction with Sparse LU to affect preconditioning in linear solvers
- Communication dominates the SpTRSV time…
  - 16 KNL nodes using production SuperLU code
  - Six representative matrices:
    - 1 from FES code m3d-c1
    - 5 from SuiteSparse Matrix Collection

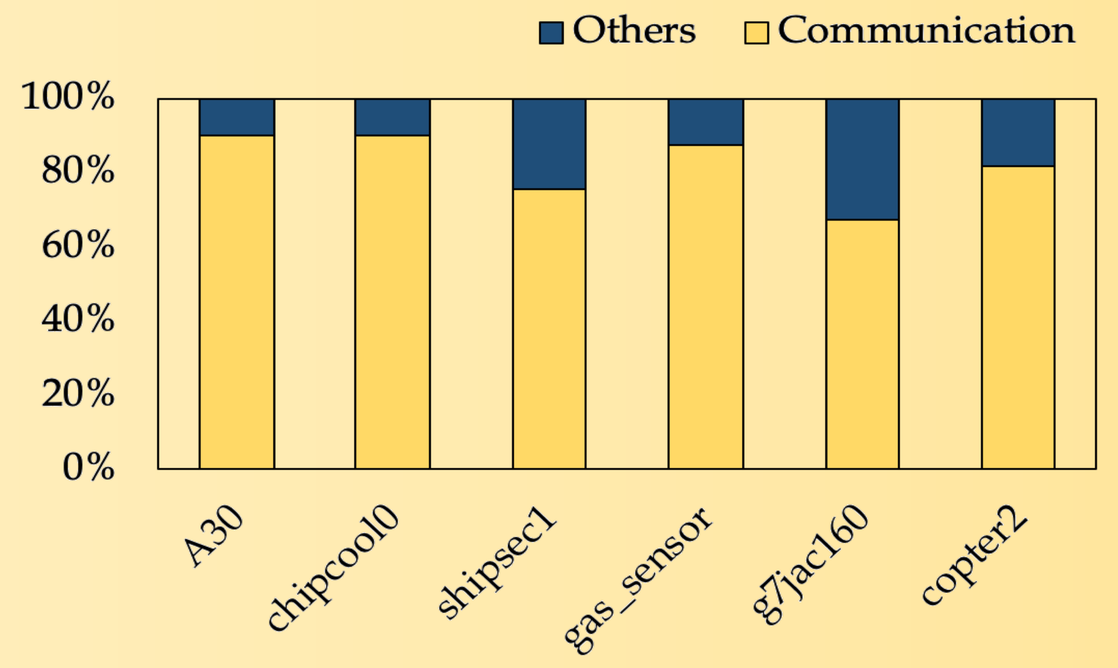

One-sided Communication-based Distributed-memory SpTRSV

- Created a round-trip ping-pong benchmark and ran on 16 KNL nodes
- foMPI one-sided outperforms Cray two-sided (highly optimized) MPI and Cray one-sided MPI by 3x and 8x respectively.
- Created a one-sided MPI version of SpTRSV for Cray Aries systems
- Attain up to a 2.2x speedup on 256 to 4,096 processes on Cori KNL (NERSC) over the existing two-sided version in SuperLU_DIST

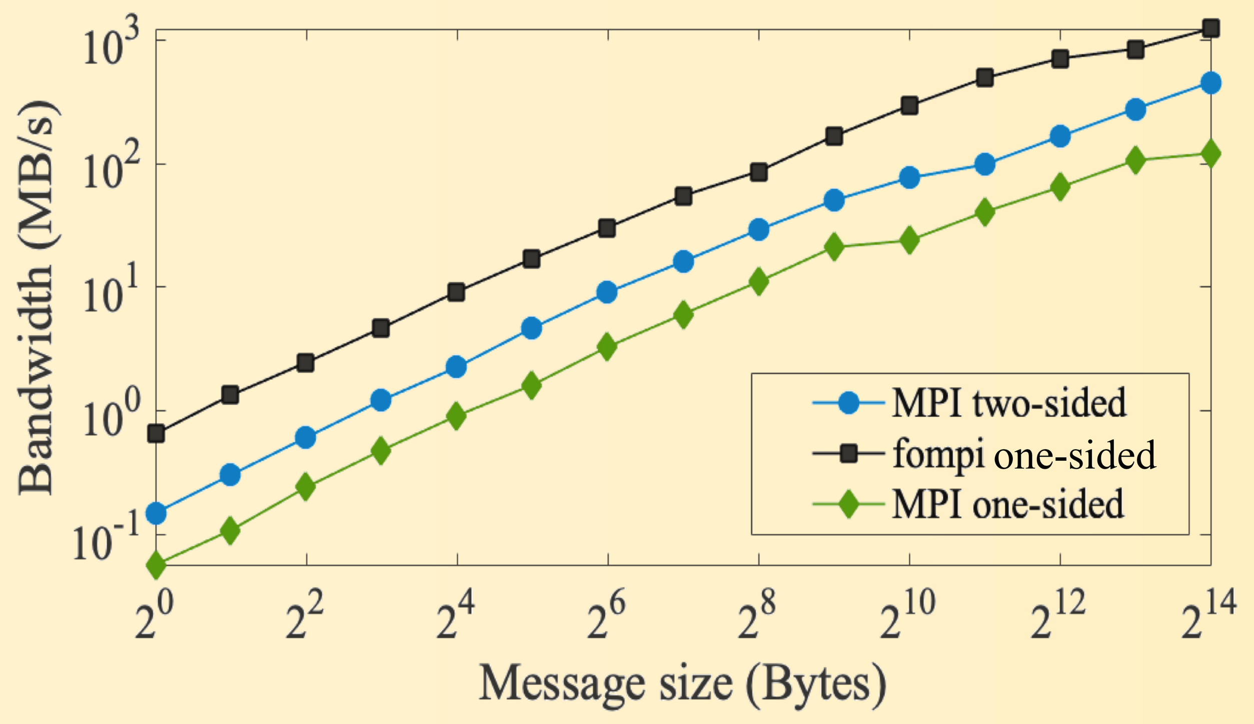

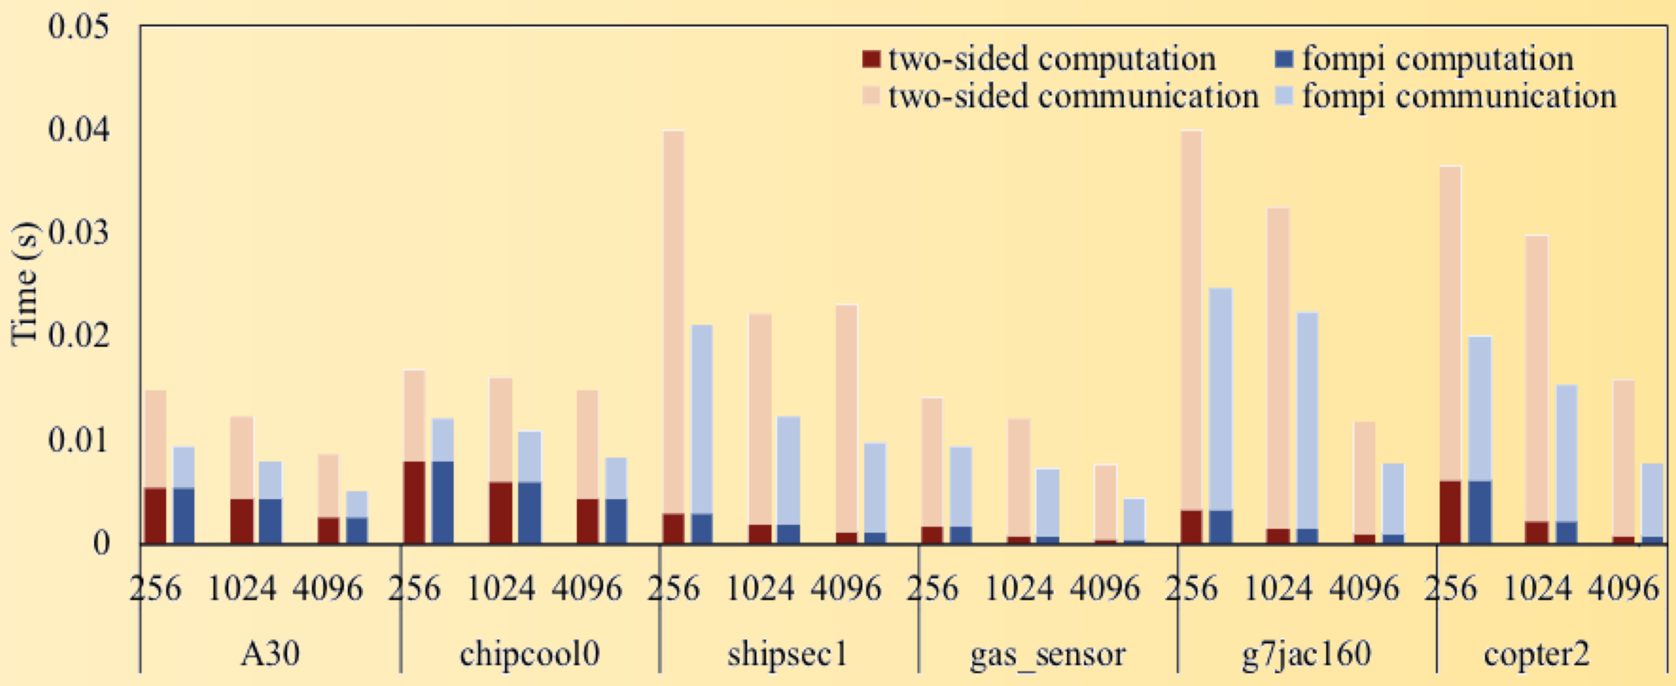

foMPI (ETH) is a fast, MPI-3.0 RDMA library interface on Cray systems. It uses distributed-memory application (DMAPP) and XPMEM for fast 80 inter-node and intra-node one-sided communication.

Critical path performance model for SpTRSV

- Circles represent
  - matrix-vectors in SpTRSV (data movement, memory bandwidth)
- Edges represent
  - data dependencies (diagonal block: in-degree/out-degree, network bandwidth/latency)
- Critical path could become longer due to limited resources

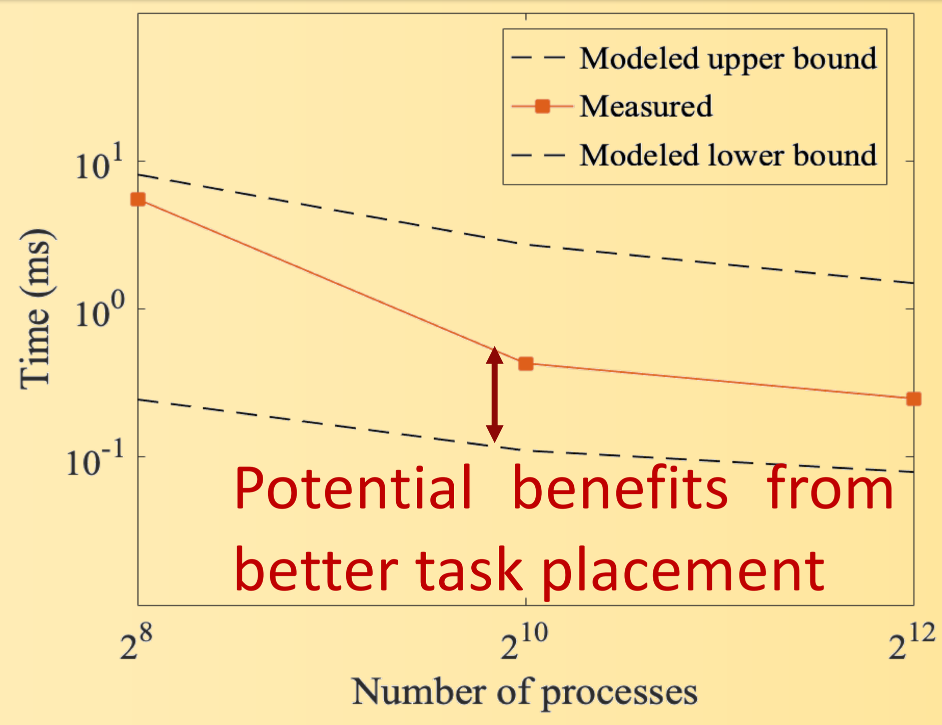

2 processes path length: 11

4 processes path length: 10

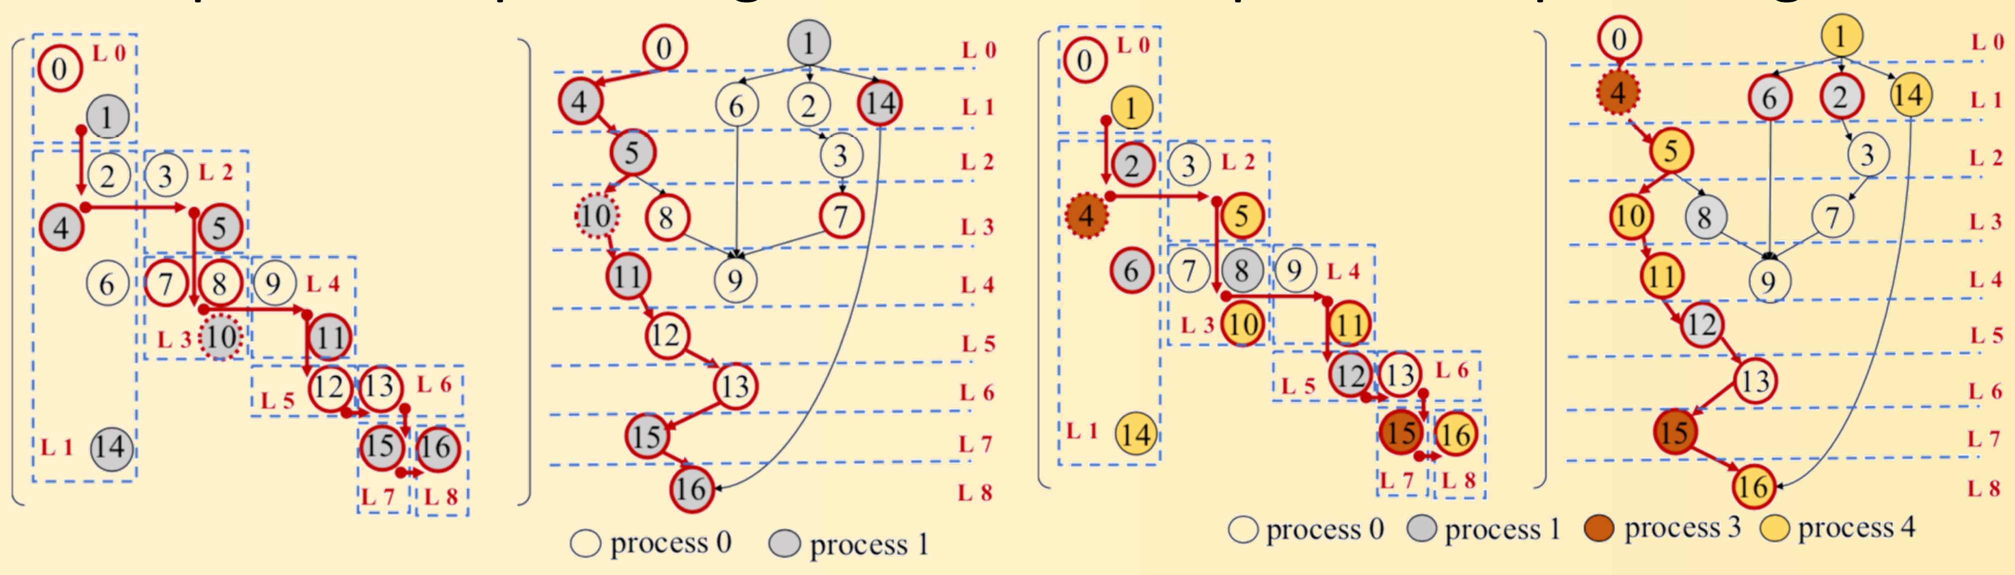

- Assess observed performance relative to machine capabilities
- Drive future code optimization
- Predict performance on future machines/architectures

Future Plans

- Evaluation on Cray Slingshot/NERSC9
- Port to NVIDIA GPUs (single GPU, then multiple GPUs on a node)
- One-sided communication on GPU-accelerated Perlmutter
- Extend the critical path model to other DAG-like computations

More Information:
Nan Ding, Samuel Williams, Sherry Li, Yang Liu
{nanding, SWWilliams, xsli, liuyangzhuan}@lbl.gov
Lawrence Berkeley National Lab

U.S. DEPARTMENT OF ENERGY Office of Science