This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

SAND2018-7003C

Unsorted Sparse Matrix Addition using Kokkos

Sandia National Laboratories
B.M. Kelley
Sandia National Laboratories, New Mexico 87185

Problem

Tpetra (part of Trilinos) implements distributed (row-partitioned) sparse matrices in CRS format:

Proc 0 {
Proc 1 {
Proc 2 {
Proc 3 {

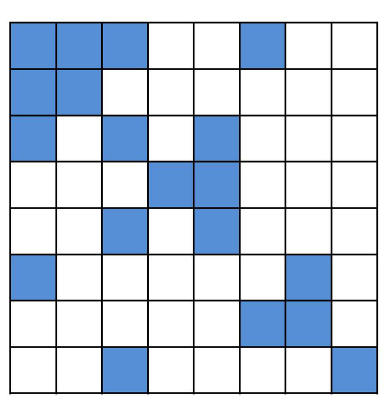

Proc 3 owns entries in columns 2, 5, 6 and 7, so a possible column map is:

Local	0	1	2	3
Global	6	5	2	7

- Memory requirements reduced by storing only local column indices on each processor
- "column map" records the local to global column mappings on each processor
- Adding matrices requires temporary conversion to global indices, which **may not preserve ascending column order**

Approach

All rows computed independently in CRS format. Example for one row of C = A+B:

A values:

9	-6	3

A columns:

3	0	4

B values:

5	2	-1	4

B columns:

2	5	3	4

Symbolic Phase: sort and merge columns

Concatenate columns:	3	0	4	2	5	3	4
A/B entry permutation:	0A	1A	2A	0B	1B	2B	3B

Radix sort by column, updating permutation:

0	2	3	3	4	4	5
1A	0B	0A	2B	2A	3B	1B

Merge to get C columns:

0	2	3	4	5
1A	0B	0A, 2B	2A, 3B	1B

Record where original A, B columns belong in C CRS arrays:

A columns:	3	0	4
Apos:	2	0	3

B columns:	2	5	3	4
Bpos:	1	4	2	3

Numeric Phase: scatter values directly into C

```
For i = 0…Alength: Cvals[Apos[i]] += Avals[i];
For i = 0…Blength: Cvals[Bpos[i]] += Bvals[i];
```

C values:

-6	5	8	7	2

Results

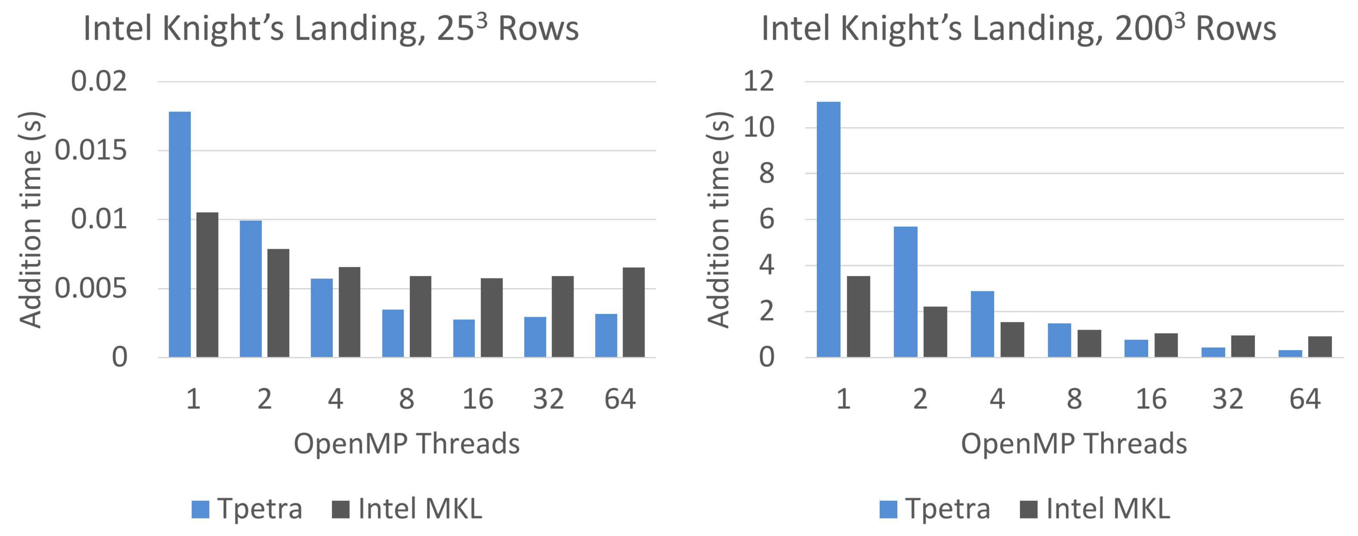

This matrix addition arises in the Jacobi smoother used by algebraic multigrid (MueLu) to interpolate from a coarse grid to a fine grid.

In this case, multigrid is being run on a 3D Laplacian grid with 25^3 vertices (left) and 200^3 vertices (right).

Although Intel's MKL library is faster than Tpetra for low thread counts, Tpetra has much better scaling and is faster for high thread counts (the intended use case for the many-core KNL architecture).

Significance

The new, Kokkos-based matrix addition in Tpetra saves a significant amount of time in real-world finite element problems (compared to the previous implementation).

Drekar (a multiphysics application) saw an overall speedup of almost 20% in a small (20^3 element) plasma fluid dynamics simulation. The proportion of time spent in matrix addition was reduced from about 20% to under 5%. This test was run on 1 KNL node.

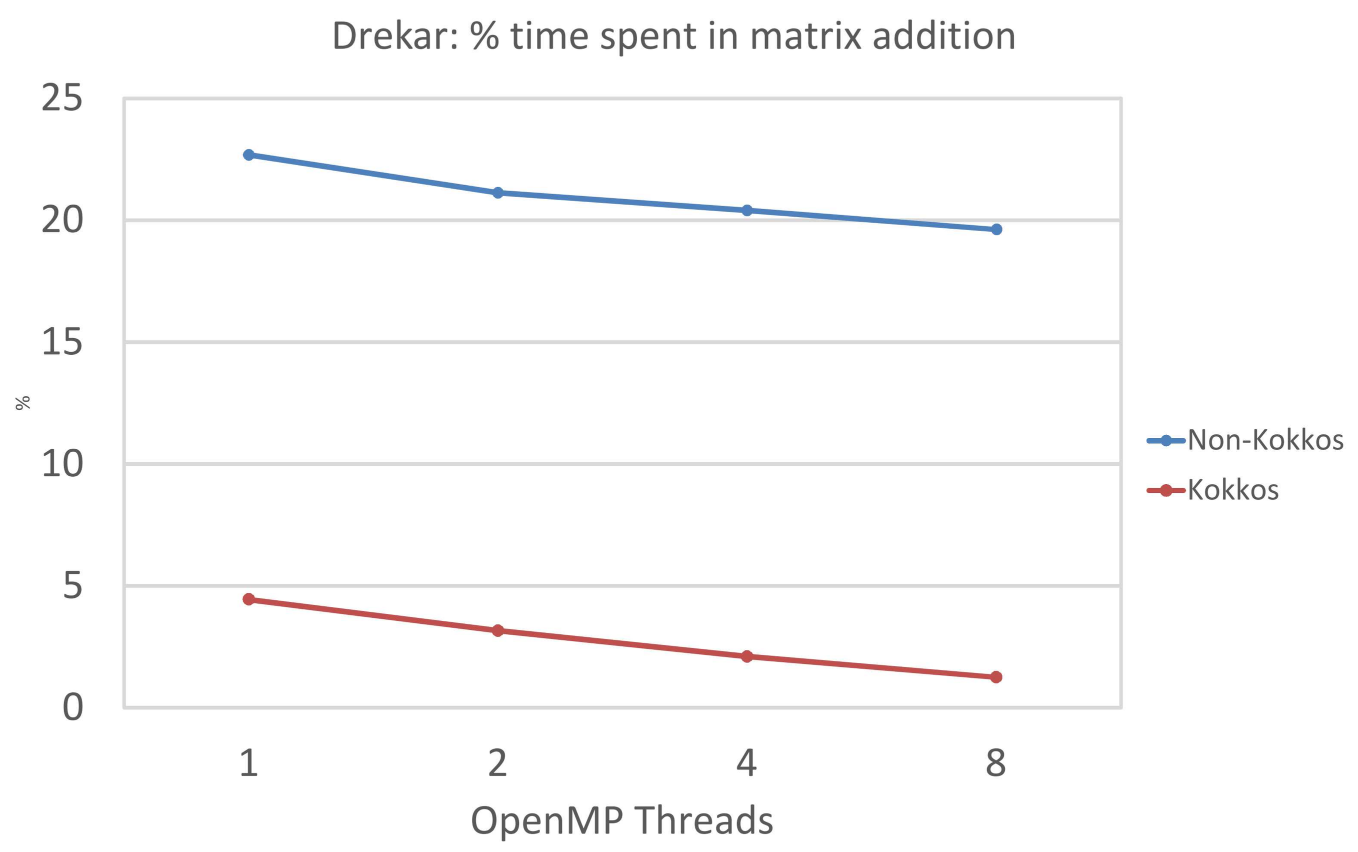

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under Contract DE-NA0003525.

SAND 2018-xxxP