

LBL Updates

April, 2022

Topics

- 3D triangular solve on CPU
- Status on multi-GPU trisolve on Perlmutter(NVIDIA) & Spock/Crusher(AMD)
- Multi-GPU trisolve optimization
- Q&A

3D triangular solve on CPU

- Integrating 3D solve into the GPU-3d-opt branch
- CPU 3D solve can achieve 20% speedup comparing to 2D solve for M3DC1 matrices
- Plan: 3D CPU solve -> single GPU solve with new data structure -> 3D GPU solve

		2D	3D	3D	3D	3D	3D	3D
	nprows	16	16	16	8	8	4	4
	npcols	32	32	16	16	8	8	4
	npz	1	1	2	4	8	16	32
s1_mat_0_126936	CPU solve (c code)	0.016	0.02	0.018	0.015	0.014	0.014	0.013
s1_mat_0_253872	CPU solve (c code)	OOM	OOM	0.039	0.04	0.035	0.041	0.037

Performance comparison of 2D and 3D solve on 16 Cori Hawell nodes (flat MPI)

multi-GPU trisolve on Perlmutter(NVIDIA) & Spock/Crusher(AMD)

- Perlmutter: waiting for NVSHMEM v2.5.0 to be available when Perlmutter returns to full service after the system upgrade. NVSHMEM v2.5.0 have several updates in this release to support the Cray environment.
 - TBD: appears as a module, or a library, etc
- Spock/Crusher:
 - ROC_SHMEM will not support slingshot in a short term

multi-GPU trisolve trisolve optimization

- Good speedups using Px1 process layout, but 1xP performance is poor
- Experiments on one summit node
 - Optimized 1x6 GPUs achieves 1.5x comparing to base 1x6
 - Optimized 1x6 GPUs achieves 1.3x comparing to single GPU
 - base 1x6 is slower than single GPU

	base 1x6	opt 1x6	6x1	single GPU
s1_mat_0_507744	8.57E-02	6.06E-02	2.94E-02	7.93E-02

Q&A

1. Does 3D solve scale for M3DC1 matrix blocks regarding block size, nnz, process count?
- 2.