

LBL Updates

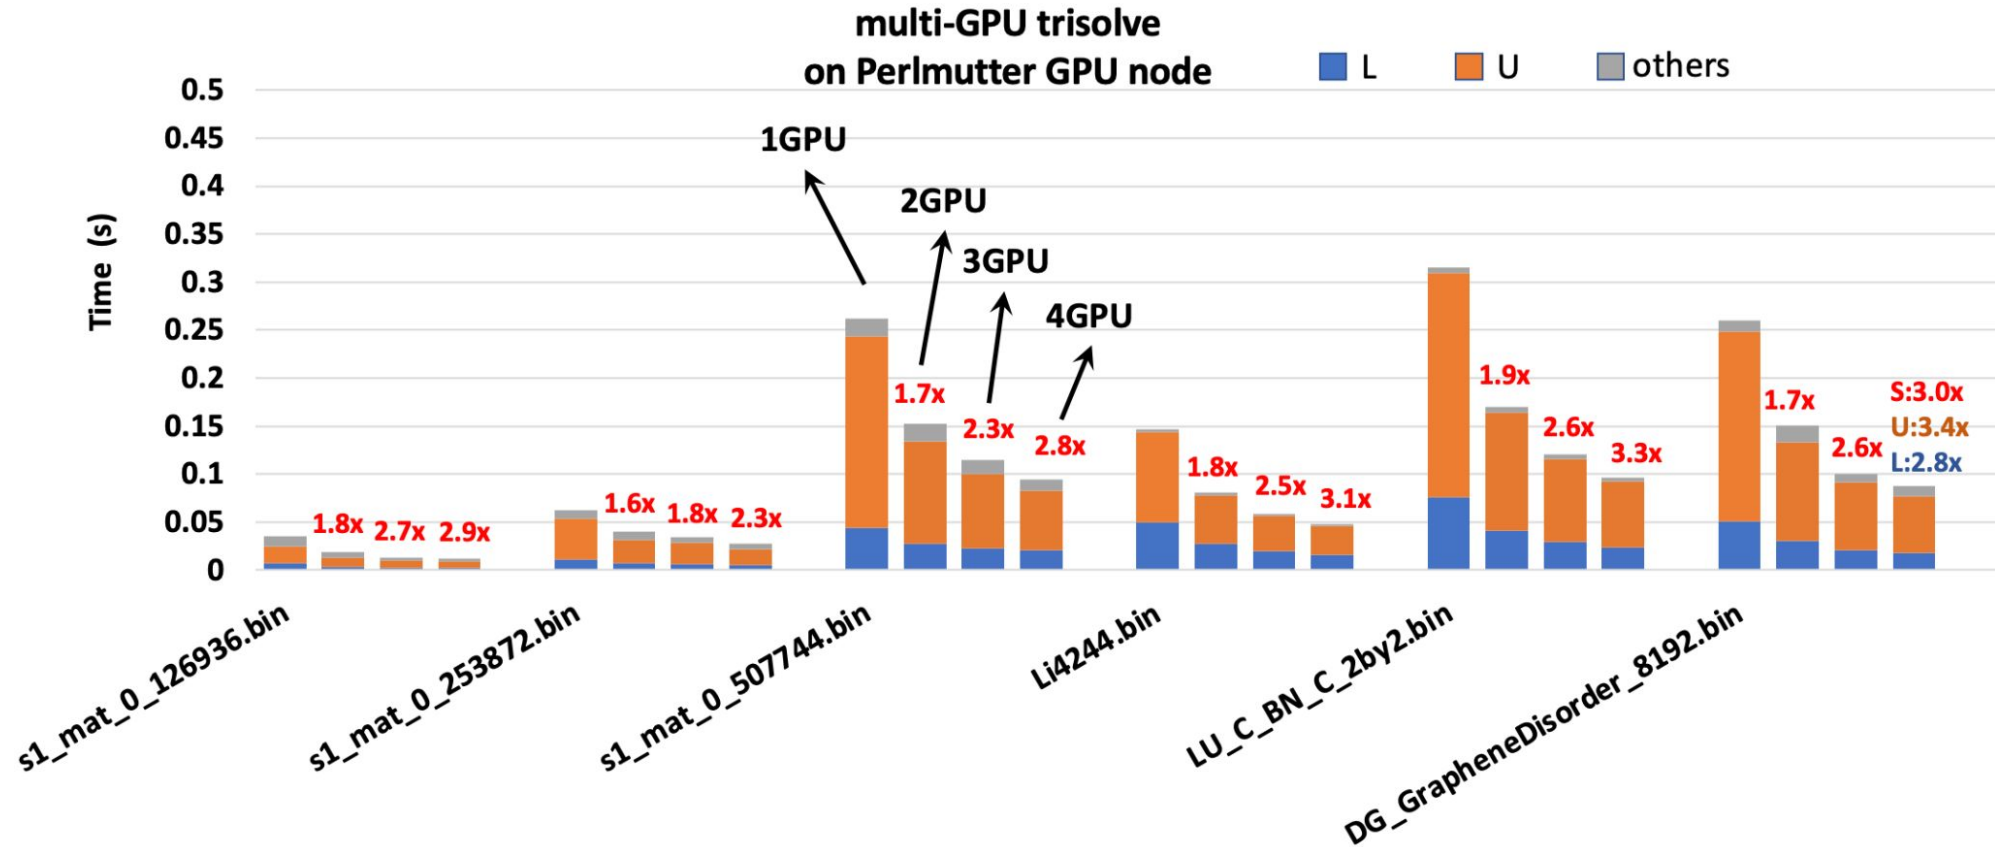
February 2023

Topics

- One-sided Solvers on Perlmutter GPU (Nan)
- Batch 1D Toroidal Solves (Hans)
- Improvements for 3D triangular solves (Yang)
- Memory Usage and Savings (Sherry)
- Q&A

One-sided Solvers on Perlmutter GPU

- Up to 3.3x speedup using 4 GPU vs. 1GPU
- 4 GPUs: Perlmutter is ~5x faster than Summit due to higher GPU-GPU communication performance
 - Perlmutter: NVLINK3: 12links * 25GB/s/direction./link = 300GB/s
 - Summit: NVLINK2: 2links * 25GB/s/direction/link = 50GB/s



Toroidal direction C1 solver

- Got Steve's F90 C1 code working on Cori
 - Block-tridiagonal algorithm, targeting $\sim 128 \times 2$ DoF's in toroidal direction
 - Next step: parallel decomp in each poloidal plane
 - Each "line solve" needs 2 DoF's from each toroidal coor
- Clarifying questions answered:
 - Velocities and diffusion coefficients will vary in toroidal direction
 - This means "batch" solves, 1 system, 1 rhs
 - Need to broadcast 2 DoF's for "rhs" of solve + coefs (vel, etc.) for system
- In parallel, ideally keep all processes busy with batches of solves
 - All-to-all required to gather all the inputs to each proc for its line solves
 - Each process will use batch Parallel Cyclic Reduction + Schur Comp solves
 - Can be optimized for GPU
- This month: stand-alone PoC on Perlmutter CPU (with C interface to C++ libs) to get timings and confirm bottlenecks. GPU next.

Improvements for 3D triangular solves

- Further improve 3D solve by incorporating tree communication for the 2D solves
- Add support for 3D GPU solve, with one GPU per 2D solve

		3D (old)	3D (new, 2Dtree)	3D (old)	3D (new, 2Dtree)	3D (old)	3D (new, 2Dtree)
	nprows	16	16	8	8	4	4
	npcols	32	32	16	16	8	8
	npz	1	1	4	4	16	16
s1_mat_0_126936		0.015	0.007	0.012	0.006	0.01	0.007
s1_mat_0_253872		0.033	0.017	0.026	0.016	0.027	0.018
s2D9pt2048		0.245	0.085	0.094	0.053	0.048	0.044

- 3D solve on Cori Hawell

		1x1x1						1x1x4					
		1CPU core per rank			1GPU per rank			1CPU core per rank			1GPU per rank		
		Flop cnt	time	Gflops	time	Gflops	Flop cnt	time	Gflops	time	Gflops		
s1_mat_0_126936		2.20E+08	9.40E-02	2.34	2.70E-02	8.14	2.66E+08	3.50E-02	7.60	2.30E-02	11.56		
s1_mat_0_253872		8.51E+08	3.10E-01	2.74	7.10E-02	11.98	1.05E+09	1.14E-01	9.18	5.30E-02	19.75		
s2D9pt2048		1.71E+09	1.30E+00	1.31	3.05E-01	5.60	1.84E+09	3.91E-01	4.69	2.33E-01	7.88		

- 3D solve (GPU) on Summit

Feedback for LBL

- C1 Hermite MG (interpolation)
- Jin will send Stellerator test problem (stronger toroidal coupling?)