

LBL Updates

December, 2022

Topics

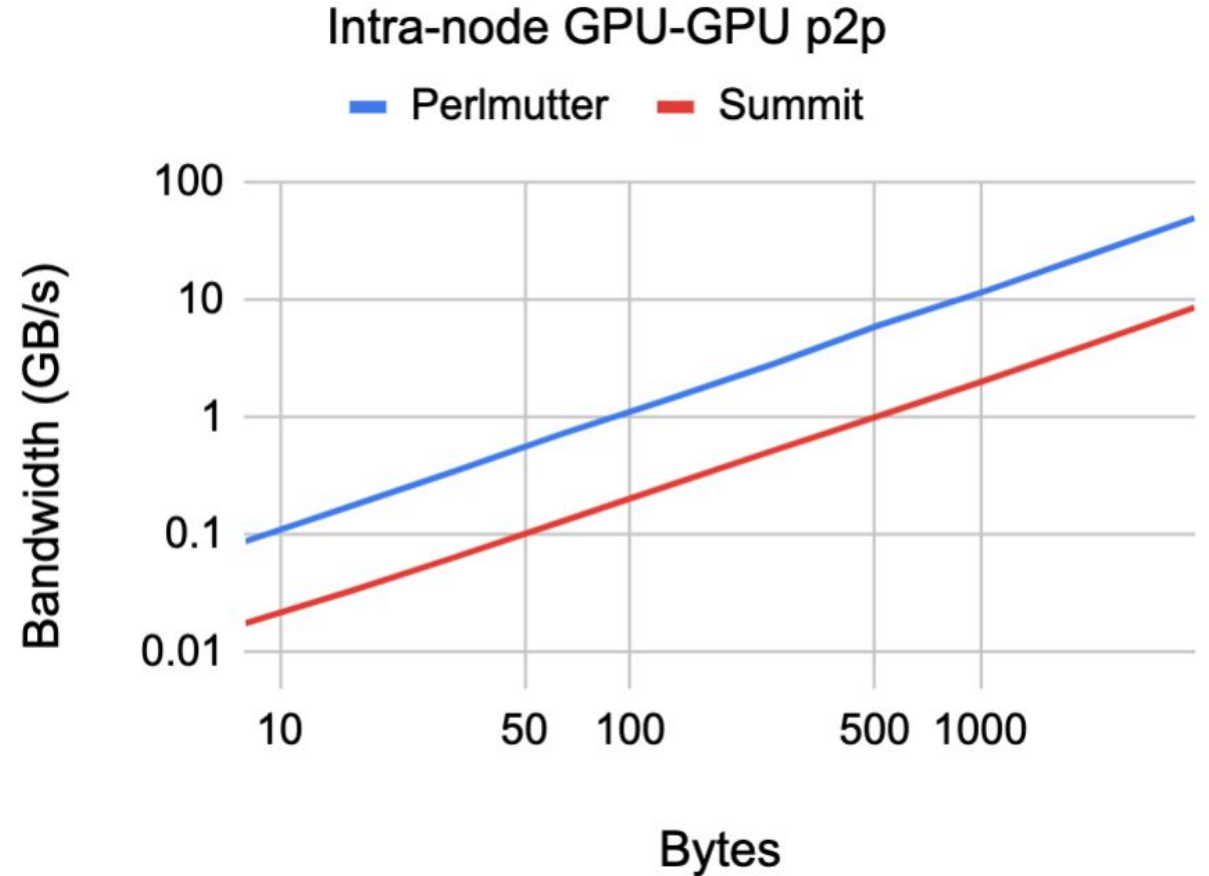
- Scaling Experiments on Perlmutter CPU
- One-sided Communication for Solvers on Perlmutter GPU
- Improvements for 3D triangular solves
- Q&A

For next meeting (Jan 2 is an LBL holiday?):

- Batch 1D Toroidal Solves (Hans)

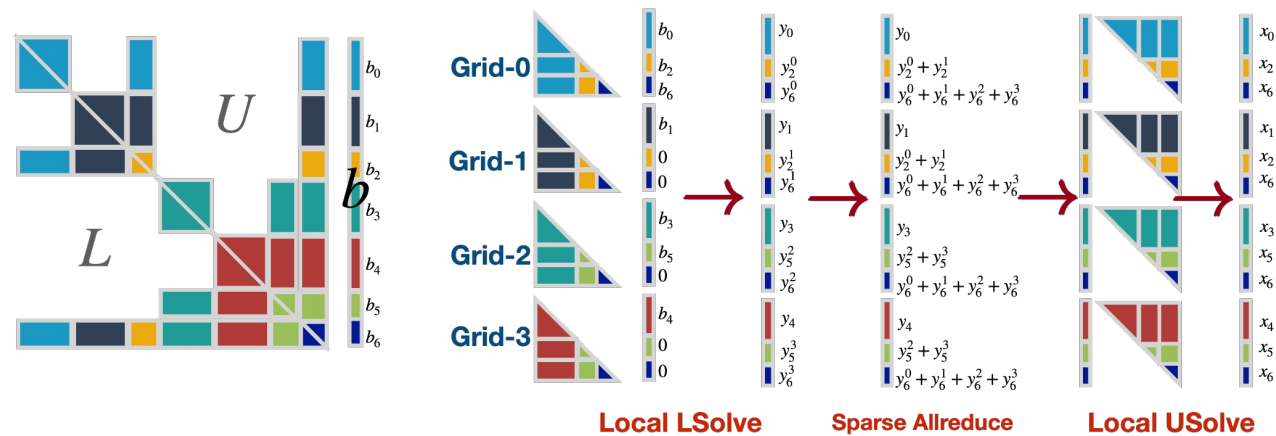
One-sided Communication for Solvers on Perlmutter GPU

- NVSHMEM can run on Perlmutter GPUs
 - 6x higher bandwidth vs. Summit intra-node GPU-GPU
- Working on getting the multi-GPU code run on Perlmutter



Improvements for 3D triangular solves

- Recall that 3D factorization results for s1_mat_0_126936: 0.98s (16x32x1) -> 0.44s (8x8x8)
- Improve 3D solve: trade duplicated computation with inter-grid communication



		2D (pddrive)	3D (old)	3D (new)	3D (old)	3D (new)	3D (old)	3D (new)	3D (old)	3D (new)	3D (old)	3D (new)
	nprows	16	16	16	16	16	8	8	8	8	4	4
	npcols	32	32	32	16	16	16	16	8	8	8	8
	npz	NA	1	1	2	2	4	4	8	8	16	16
s1_mat_0_126936 (16 nodes)		0.01	0.015	0.01	0.014	0.009	0.012	0.008	0.01	0.008	0.01	0.008
s1_mat_0_253872 (32 nodes)		0.022	0.033	0.023	0.025	0.02	0.026	0.019	0.026	0.021	0.027	0.021

Feedback for LBL

- Memory issues... 138MB (symbolic) -> 7GB (numerical... 4.5GB is actual LU values)
- Form matrix in 32b, factor in 32b, use 32b preconditioner
- Jin's MG experiments