

LBL Updates

June 2023

Topics

- Batch 1D Solves for Toroidal Preconditioning (Hans)
- One-sided Solvers (Nan)
- 3D Solvers (Yang)
- Reduced Precision (Sherry)
- Q&A

Batch 1D Solves for Toroidal Preconditioning

- Working on Jin's pressure matrix solve w/ 128 ranks
 - → 4 planes, 32 procs/plane.
 - # DoF's seems unbalanced in plane? (ranges from ~700 - ~1200)
 - Is this a good test size? (too few planes means very small toroidal systems)

Questions:

- In PETSC config, FGMRES settings indicates *right* preconditioning?
 - <https://petsc.org/release/manualpages/KSP/KSPFGMRES/>
"Only right preconditioning is supported."
 - Only requires matrix-vector solve

$$Ax = b$$



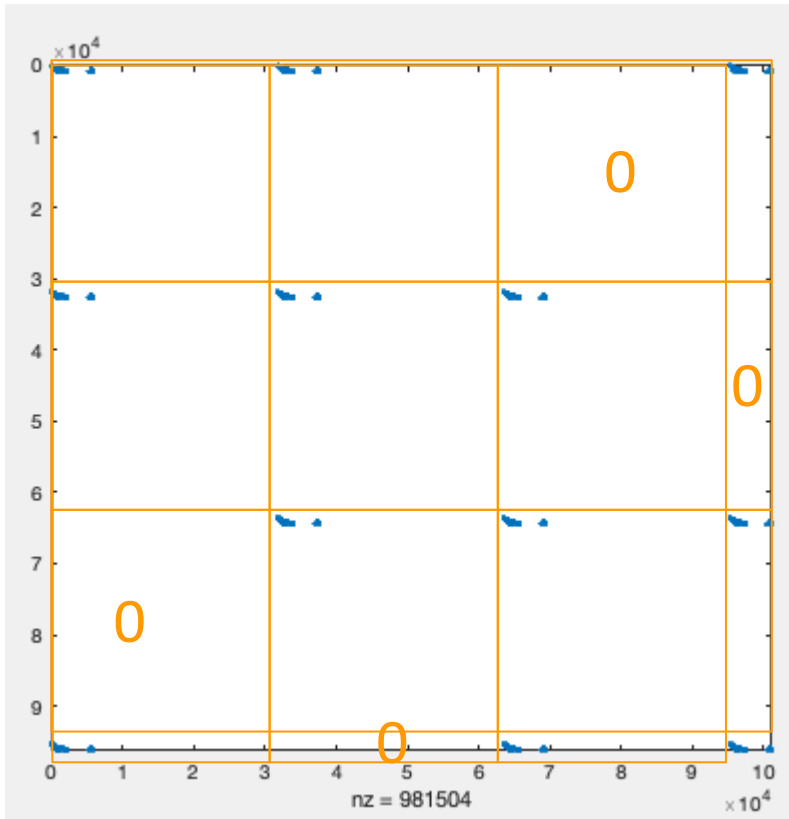
- Permutation of matrix layout seems to be:
 - 12 (toroidal) independent diagonal DoF's?

$$Mx = u \rightarrow x = M^{-1}u$$

- Next steps: finish stand-alone preconditioner tests

Batch 1D Solves for Toroidal Preconditioning

- Is this right? (block cyclic outermost, 12x12 dense blocks innermost)



One-sided Solvers

- Adding 2D support for whole multi-GPU trisolve
- Currently multi-GPU trisolve work on Perlmutter and Summit via NVSHMEM, will work on Frontier once ROC_SHMEM release the new version

3D Factorization

- Yang is implementing a new redistribution routine to reduce memory usage

Using single-prec SuperLU_DIST in double-prec PETSc

- Barry created this branch, and is merged to PETSc main branch https://gitlab.com/petsc/petsc/-/merge_requests/6402
- Runtime option: `-mat_superlu_dist_single_precision`
- Sherry fixed a memory leak reported by Barry

- Jin tested it on Perlmutter: (with `mat_superlu_dist_statprint`)
 - similar numerical results
 - but, no memory reduction
- Sherry checked the output: double and single reports identical mem
 - but, standalone tests `pddrive.c` and `psdrive.c` do show memory difference