LBL Updates

December, 2021

LBL Updates

- Yang: Profiling SuperLU/TriSolve on A100 (Perlmutter)
- Sam: Multigrid thoughts

- Note, Nan will be on parental leave thru January

Profiling on A100 GPU

- Numerical factorization time on 1 CPU core, 1 A100 GPU, and 1 V100 GPU
- Roofline analysis (data locality vs. performance) on CPUs vs GPUs
  - in Roofline, up is better, left is memory-intensive, right is compute-intensive
  - Ideally, all loops should be near the roofline (peak FLOPs or bandwidth-limit)
  - Moreover, bulk of the run time should come from high AI (compute-intensive) loops
- For SuperLU
  - Many DGEMM calls within numerical factorization
  - CPU is right on the roofline
  - GPU performance is all over the place (cuBLAS is not tailored for matrix dimensions/aspect ratios in SuperLU?)

- s1_mat_0_126936: large GEMMs with small middle dimension. A100 is best

CPU 1 Core
Perlmutter
Factor time: 3.95 s

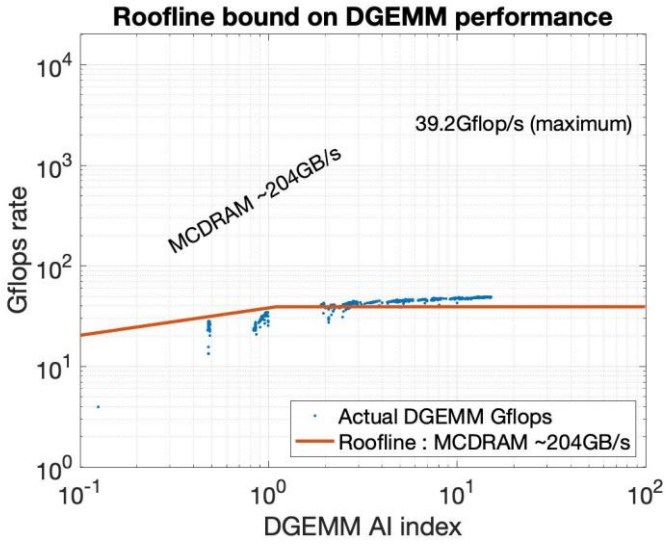

1 GPU (A100)
Perlmutter
Factor time: 3.13 s

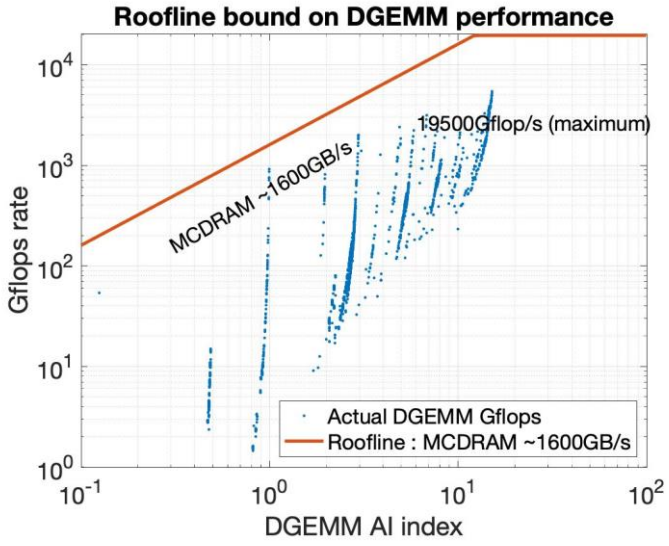

1 GPU (V100)
Cori
Factor time: 3.77 s

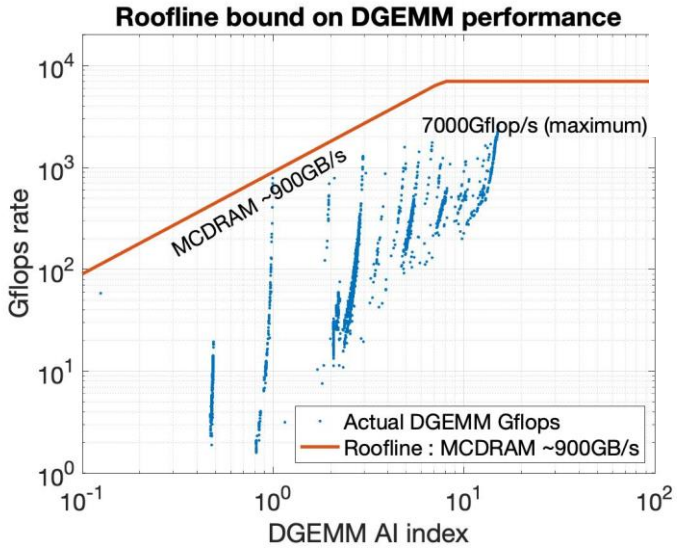

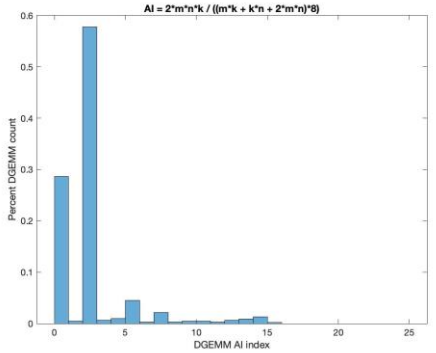

Distribution of AI

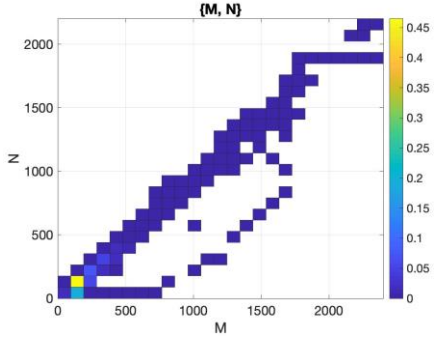

Distribution of M,N

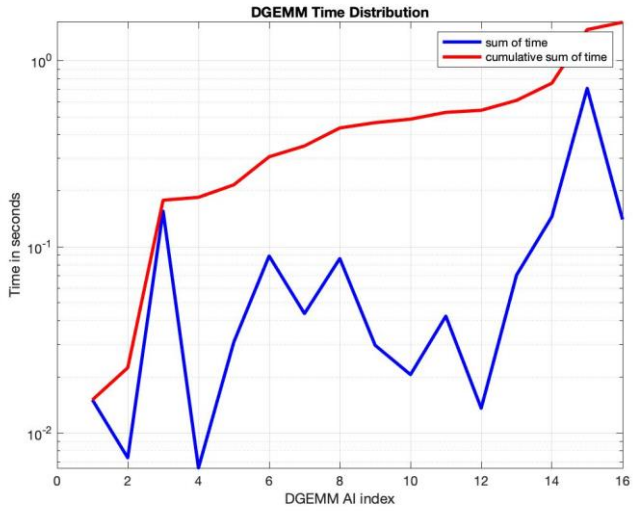

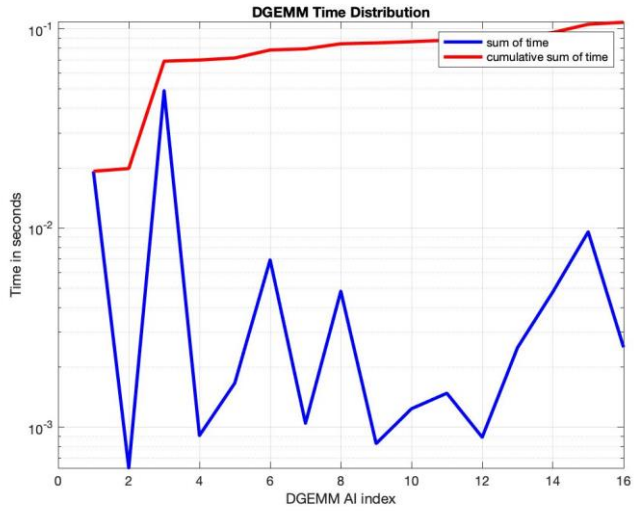

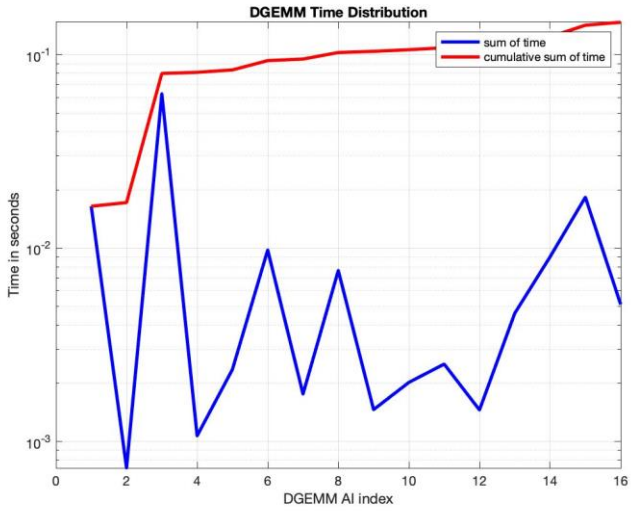

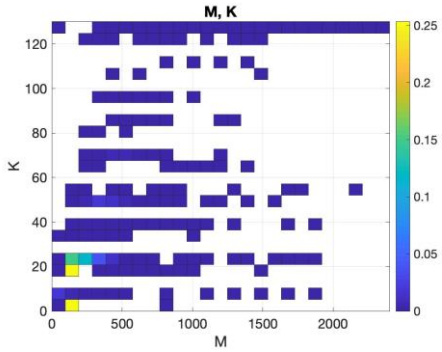

Distribution of M,K

1D (Toroidal) Multigrid

- Consider FGMRES with toroidal (MG) and poloidal (SuperLU) preconditioning
- Assume values (q's) are stored at nodes
- Nodes define a polynomial that is C1 continuous at nodes (can be used to interpolate between nodes)
- In 1D…

$$
\begin{aligned}
v(x) &= v_j\,\nu_{1,j}(x) + v'_j\,\nu_{2,j}(x) + v_{j+1}\,\nu_{1,j+1}(x) + v'_{j+1}\,\nu_{2,j+1}(x), \\
&= v_j + v'_j x + \left(-3v_j - 2hv'_j + 3v_{j+1} - hv'_{j+1}\right)\left(\frac{x}{h}\right)^2 \\
&\quad + \left(2v_j + hv'_j - 2v_{j+1} + hv'_{j+1}\right)\left(\frac{x}{h}\right)^3, \\
&= a_0 + a_1 x + a_2 x^2 + a_3 x^3.
\end{aligned}
$$

$$
\begin{bmatrix} a_0 \\ a_1 \\ a_2 \\ a_3 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ -\frac{3}{h^2} & -\frac{2}{h} & \frac{3}{h^2} & -\frac{1}{h} \\ \frac{2}{h^3} & \frac{1}{h^2} & -\frac{2}{h^3} & \frac{1}{h^2} \end{bmatrix} \cdot \begin{bmatrix} v_j \\ v'_j \\ v_{j+1} \\ v'_{j+1} \end{bmatrix}. \qquad (11.18)
$$

1D (Toroidal) Multigrid

- Consider 1D restriction / prolongation in toroidal direction
- Restriction: Use q_0, q_2 at h to construct(copy) q_0 and q_2 at 2h

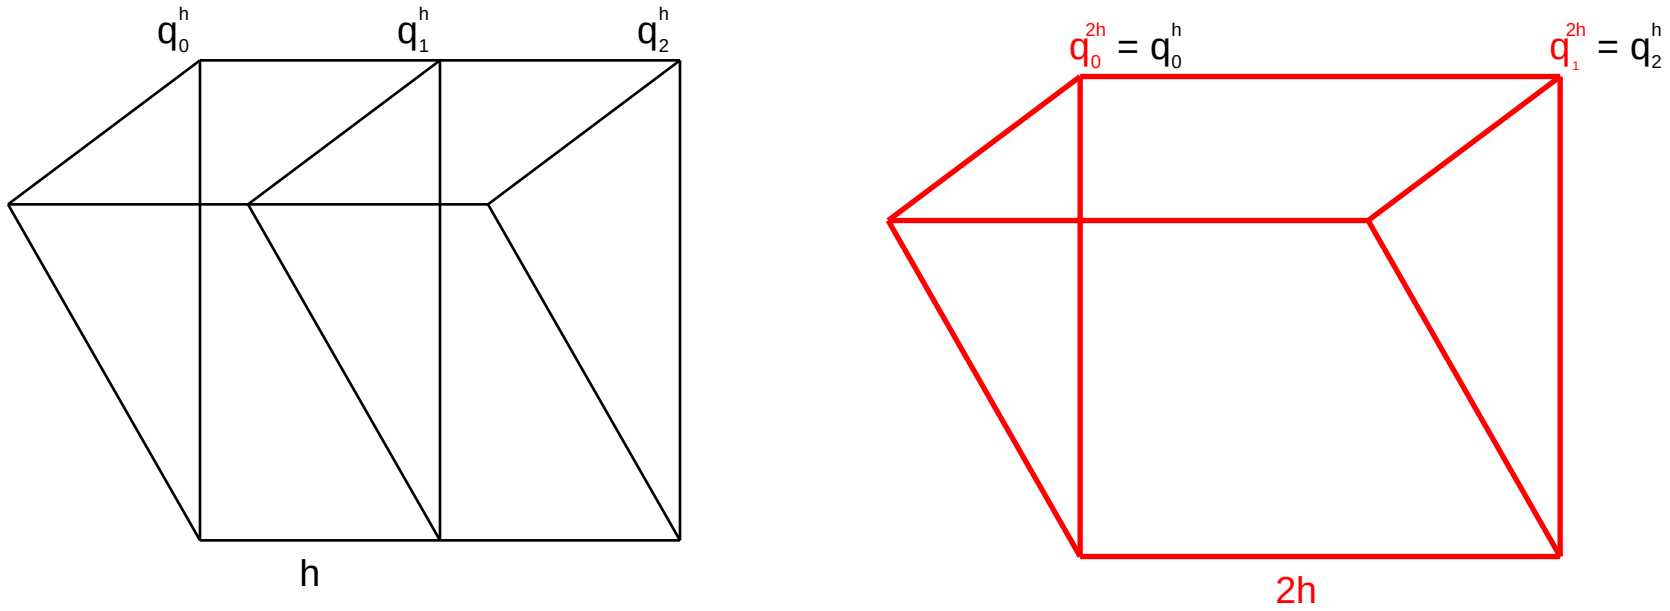

- Prolongation: reconstruct 2h polynomial to interpolate at h to construct q_1 at h
- Construct coarse grid operation A^{2h}

3D Multigrid

- Consider restriction / prolongation within poloidal plane
  - For each coarse grid node, use enclosing fine grid element to interpolate/calculate coarse node values
- Can we do full 3D MG as a preconditioner?
  - Construct coarse grid operator?

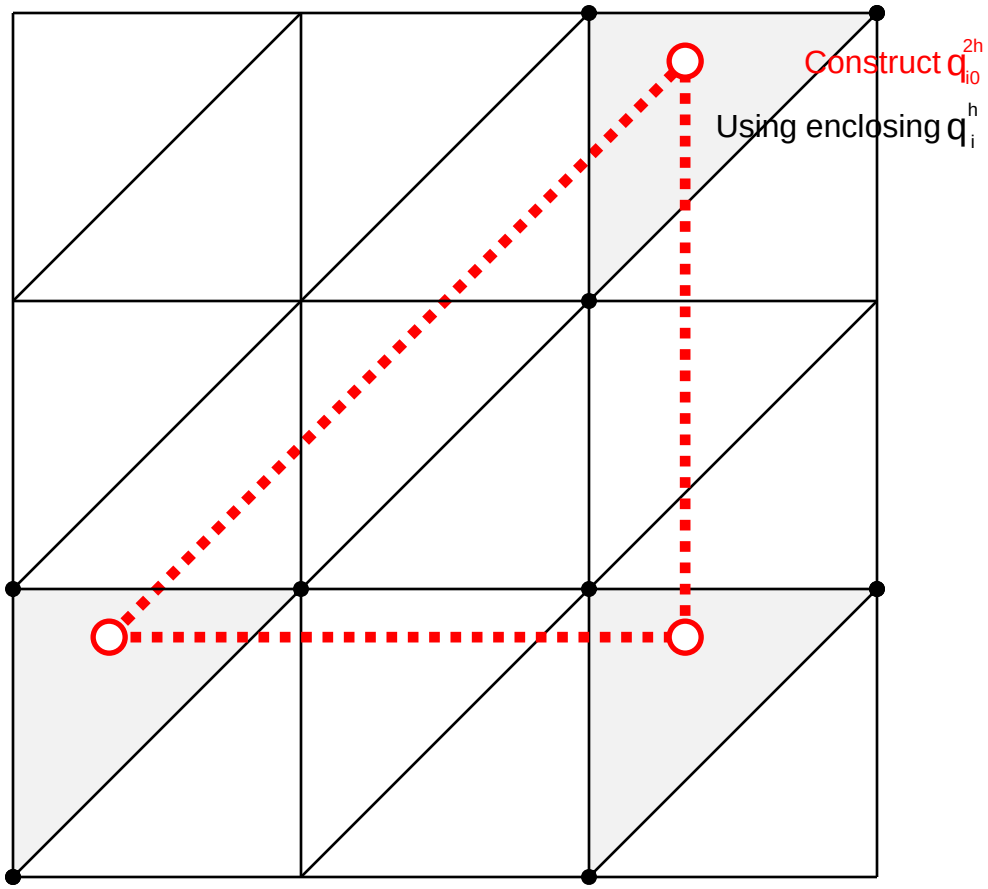

Q&A

1. Restriction’s on using PETSc’s GAMG
2. Steve Jardin has a 1D proxy app for C1 hermite polynomials… can we demo GMG in it?