

HPGMG BoF

High Performance Geometric Multigrid Birds of a Feather

Introduction *Mark Adams (LBNL)*

HPGMG-FV on the KNL processor *Samuel Williams (LBNL)*

HPGMG on the Pascal GPU Architecture *Peng Wang (NVIDIA)*

Communication and Vectorization in HPGMG *Vladimir Marjanovic (HLRS)*

Questions and Discussion *all*

Review of HPGMG-FV

Samuel Williams

Lawrence Berkeley National Laboratory

SWWilliams@lbl.gov

Goal was to create a benchmark that is...

- Application- and Mathematically-relevant
- Scale-free and independent of parallelization
- Precise in its definition
- Sufficiently simple that undergraduates can optimize it
- Exercises aspects of system architecture not addressed by HPL or HPCG

HPGMG-FV

- Solves Variable Coefficient Poisson... $Lu = \nabla \cdot \beta \nabla u = f$
- Discretized with the 4th order Finite Volume Method
- Solved on a Cubical Cartesian grid with Dirichlet Boundary Conditions
- Uses the asymptotically exact Full Multigrid (FMG)
- Fully specified stencils and smoothers


Finite Volume Method

- In FVM, we must calculate a flux term on each of the 6 faces on every cell in the entire domain....

$$Lu = \langle \nabla \cdot \beta \nabla u \rangle = \text{[Diagram showing a cube with six faces highlighted in purple and arrows indicating flux directions: left, right, front, back, bottom, and top.]}$$


4th Order Operator $\langle \nabla \cdot \beta \nabla u \rangle$

- In 2nd order, we can approximate each of these flux terms as a 2-point weighted stencil...



- 6 such terms form a 7-point variable-coefficient stencil.
 - Low arithmetic intensity
 - 6 MPI messages/smooth


- For 4th order, additional terms are required...



- 25-point stencil...
 - 9 (unique) 4-point stencils
 - ~no extra DRAM data movement
 - 4x the floating-point operations
 - 3x the MPI messages/smooth
 - 2x the MPI message size

HPGMG-FV uses Full Multigrid (FMG)

- FMG is a single pass, direct solver that provides a solution to the discretization error (4th order)
- The FMG multigrid F-Cycle is a series of progressively deeper geometric multigrid V-cycles



HPGMG-FV Uses Many Different Stencils

- Several dozen stencil sweeps per step
- Stencils vary in shape and size
- Grid sizes vary exponentially.


- **Challenges compiler technology**

$$u^h = \text{interpolate_v4}(u^{2h})$$

$$\begin{aligned} \text{smooth} \\ r^h &= f^h - Lu^h \\ f^{2h} &= \text{restrict}(f^h) \\ u^{2h} &= 0 \end{aligned}$$


$$\begin{aligned} u^h &+= \text{interpolate}(u^{2h}) \\ \text{smooth } Lu^h &= f^h \end{aligned}$$

$$\begin{aligned} &1 \text{ global MPI collective} \\ &\|f^h - Lu^h\|_\infty \end{aligned}$$



HPGMG Has Multiple Communication Patterns

- Work is redistributed onto fewer cores (agglomeration)
- Coarse grid solves can occur on a single core of a single node
- Coarse grid solution is propagated to every thread in the system





BERKELEY LAB

LAWRENCE BERKELEY NATIONAL LABORATORY



U.S. DEPARTMENT OF
ENERGY

HPGMG Performance November 2016

November 2016 Ranking

HPGMG Rank	System Site	System Name	10 ⁹ DOF/s	MPI	OMP	Acc	DOF per Process	Top500 Rank	Notes
1	ALCF	Mira	500	49152	64	0	36M	9	
2	HLRS	Hazel Hen	495	15408	12	0	192M	14	
3	OLCF	Titan	440	16384	4	1	32M	3	K20x GPU
4	KAUST	Shaheen II	326	12288	16	0	144M	15	
5	NERSC	Edison	296	10648	12	0	128M	60	
6	CSCS	Piz Daint	153	4096	8	1	32M	8	K20x GPU
7	Tohoku University	SX-ACE	73.8	4096	1	0	128M	-	vector
8	LRZ	SuperMUC	72.5	4096	8	0	54M	36	
9	NREL	Peregrine	10.0	1024	12	0	16M	-	
10	NREL	Peregrine	5.29	512	12	0	16M	-	
11	HLRS	SX-ACE	3.24	256	1	0	32M	-	vector
12	NERSC	Babbage	0.762	256	45	0	8M	-	KNC