Non-invasive Semi-Structured Multigrid on Advanced Architectures

P.I. R. Tuminaro, J. Hu, C. Siefert, L. Berger-Vergiat, M. Mayr
Sandia National Labs, Albuquerque, NM & Livermore, CA.
Univ. of Bundeswehr, Munich, Germany

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, L.L.C., a wholly owned subsidiary of Honeywell International, Inc., for the U. S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Abstract

Structured meshes provide enormous efficiency benefits, but are not convenient for representing complex domains . Partially structured meshes can represent sophisticated geometries, but can be cumbersome to adopt for mature applications.

Consequently, we propose a partially structured mesh methodology that

- ❖ supports flexibility needed by applications,

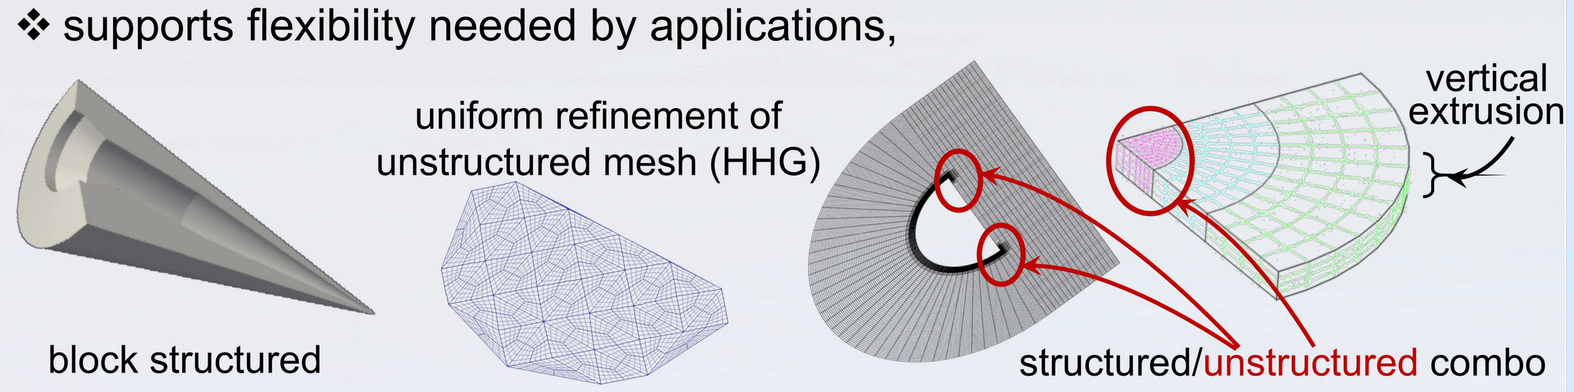

- ❖ is relatively non-invasive for mature applications to adopt,
- ❖ provides fast & efficient multigrid solvers (MG).

A precise mathematical approach is essential in tackling these objectives.

Motivation

Structured mesh advantages for next generation platforms (NGPs):

- ❖ significantly less communication
- ❖ highly efficient kernels
- ❖ lower bandwidth requirements & less indirect addressing
- ❖ low MG fill-in (less coarse matrix nonzeros) & fewer levels
- ❖ special algorithms to improve robustness (e.g., black-box MG, line ideas)
- ❖ vastly simplified MG setup

Approach

1) Efficiency: represent matrices & vectors as union of structured regions
   + develop precise mathematical understanding of relationships between kernels for regionals representations & standard matrices and vectors

2) Design semi-structured MG methods with strong convergence properties
   + leverage math framework to mirror reliable AMG methods when *practical*, but introduce *benign* approximations to improve performance/usability

3) Encourage adoption of semi-structured framework by scientists
   + provide flexibility for complex scenarios
   + facilitate use through a non-invasive methodology

Math challenges:

i) formal understanding of structured regional approximations,
ii) formulate stiffness matrix approximations to facilitate non-invasiveness
iii) coarsening to maintain mesh structure throughout MG hierarchy,
iv) interpolation along interfaces between unstructured & structured regions,
v) non-conformal mesh issues when coarsening block structured meshes,

Results

New Mathematical Framework

Standard vectors & matrices have regional equivalents (where shared interface dofs are replicated).

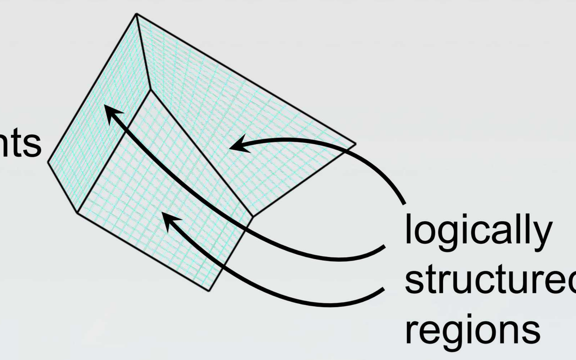

e.g., $u = \psi\,\hat{u}$, $A = \psi\,\check{A}\,\psi^T$ — regional forms of matrices & vectors (precise details omitted)

fine & coarse transformation matrix from regional to standard

Lemma $R\,\psi = \psi_c\,\widehat{R}$ … under some grid
Lemma $\psi^T P = \widehat{P}\,\psi_c^T$ transfer assumptions

Thm: $\psi_c\,\widehat{R}\,\check{A}\,\widehat{P}\,\psi_c^T = RAP$

Mathematical framework

- ❖ formalizes relationships between traditional kernels & efficient regional ones
- ❖ clarifies operator approximations for non-invasive simplifications
- ❖ reveals communication-avoiding transfer conditions for region matrix-matrix multiply

facilitates proper implementation & algo. design

function mgCycle$(\check{A}, \hat{u}, \hat{b})$:
$\hat{u} \leftarrow \hat{u} + \omega\,(\hat{b} - \Psi^T\Psi\check{A}\hat{u})\ ./\ \hat{d}$
$\hat{r} \leftarrow \hat{b} - \Psi^T\Psi\check{A}\hat{u}$
$\hat{u}_c \leftarrow 0$
$\hat{u}_c \leftarrow \text{Solve}(\check{A}_c, \hat{u}_c, \Psi_c^T\Psi_c\hat{R}(\widehat{\Psi\Psi^T})^{-1}\hat{r})$
$\hat{u} \leftarrow \hat{u} + \hat{P}\hat{u}_c$

function mgSetup$(\check{A})$
$\hat{d} \leftarrow \Psi^T\Psi\ \text{diag}(\check{A})$
$\hat{P} \leftarrow \text{construct}P(\check{A})$
$\hat{R} \leftarrow \hat{P}^T$
$\check{A}_c \leftarrow \hat{R}\check{A}\hat{P}$
$\Psi_c \leftarrow \text{inject}(\Psi)$

Structured gains

Specialized matrix-vector products achieve over 3.5x speedups relative to standard Trilinos implementations, using standard CSR formats for both.

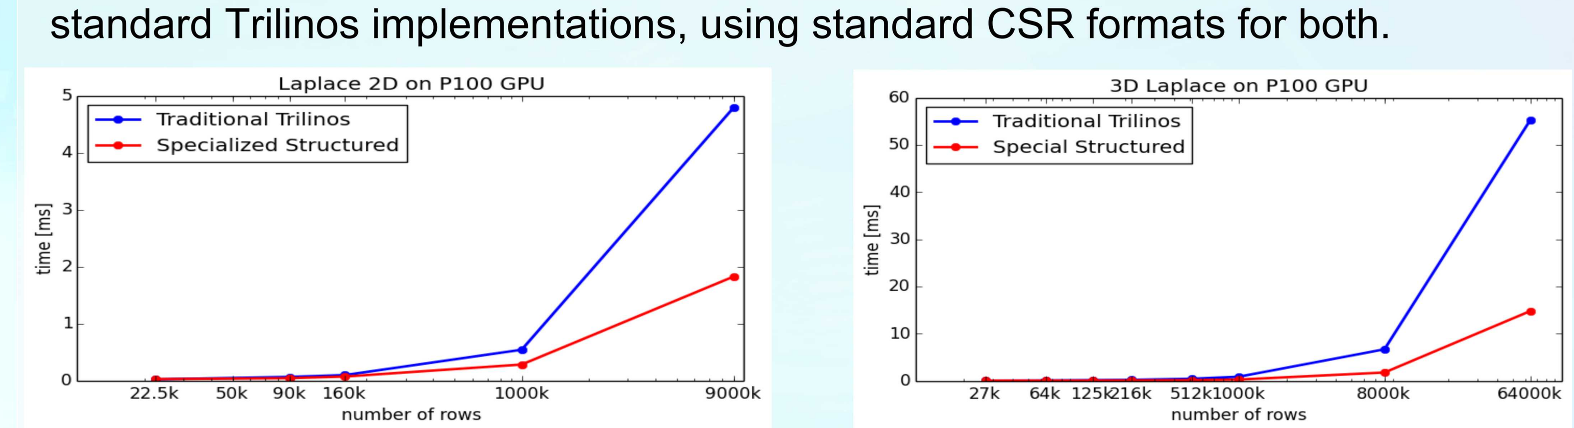

Specialized matrix-matrix multiplies perform RAP triple-matrix products over an order of magnitude faster than standard Trilinos triple-matrix products

Time (sec.) for 2D triple products

MUELU: standard Trilinos
CONST: piecewise-constant basis functions for P
GEO: linear basis functions for P
GENERIC: general coefficients for P

coarse mesh size	9pt stencil for P		25pt stencil for P		
	MueLu	Const	MueLu	Generic	Geo
140 × 36	.0024	.0001	.0109	.0012	.0009
140 × 180	.0124	.0006	.0572	.0061	.0046
700 × 180	.0726	.0070	.2944	.0320	.0240
700 × 900	.3702	.0356	1.4786	.1606	.1208

h to 3h coarsening rate, $R = P^T$

Structured/Unstructured Meshes

Hybrid solution strategy for grids with structured & unstructured regions improves convergence over standard AMG on stretched meshes by allowing for line smoothing algorithms. Properly matching inter-grid transfers at region interfaces is essential for AMG convergence.

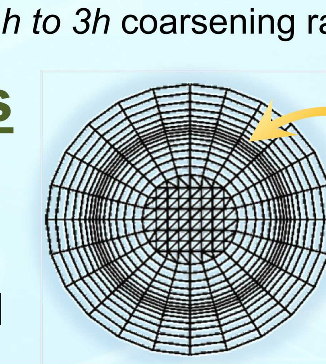

# dofs	max aspect ratio	iterations	
		AMG	Hybrid AMG
285	132.7	31	12
2417	124.4	89	14
21501	121.6	nc	19
192773	120.6	nc	26

Impact

Solver algorithms that leverage mesh structure are essential on NGPs that include GPUs. Our new solvers leverage structure & demonstrate:

- major performance benefits within primary MG kernels
- significant convergence benefits on stretched meshes
- scalability on large systems for extruded mesh solver

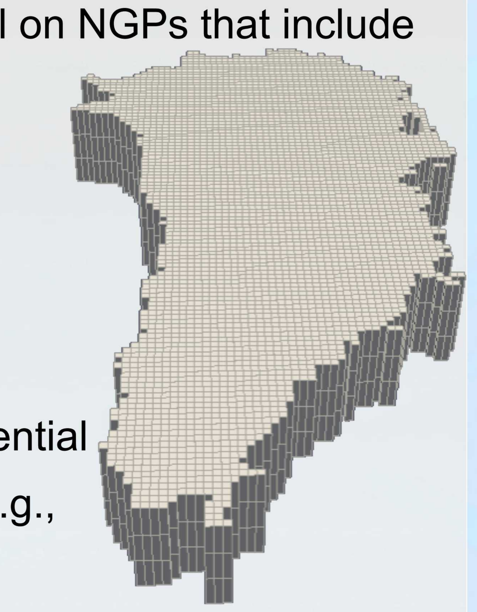

5 yr DOE impact: extruded mesh MG algorithms will be essential for scalability of future extreme scale climate applications (e.g., ice sheet modeling). New algorithms imperative for rapid convergence and insure that some key kernels are communication-free. We expect that block structured MG algorithms will significantly boost scalability for hypersonic re-entry vehicle modeling. Line solvers already play a significant role in this area suggesting promise for structured MG. HHG algorithms have already been shown to deliver enormous scalability benefits [Ruede et. al] . We project that our HHG generalizations will deliver similar gains for finite element incompressible magneto-hydrodynamics (MHD) capabilities.

10 yr impact we believe that leveraging partial structure will be key toward maximizing NGP performance for a wide range of DOE applications. We are starting with climate, hypersonics, & MHD, but envision several other areas (e.g., wind energy).

Synergy

We partner with application teams (2 ASCR & 1 ASC funded) who anticipate significant solver performance benefits on NGPs. The partially structured paradigm leads to new math/computer science questions (e.g., alternative solver algorithms, MG smoother approximations, communication avoiding, software design, kernel optimization) that spur new research. We partner with SciDAC (PISCEES, FASTMATH), the Exascale Computing Project, an ASCR base math program, and ASC projects on NGP kernels & specific applications. The switch to partially structured algorithms is a major shift requiring significant effort. Further partnering will be needed to tackle the transformation toward partially structured meshes.

References

Mayr, Berger-Vergait, Ohm, Tuminaro, "An Algebraic Multigrid Method Tailored to Semi-Structured Grids", in preparation, 2019

Chen, Cambier, Boman, Rajamanickam, Tuminaro, Darve, "A Robust Hierarchical Solver for Ill-conditioned Systems with Applications to Ice Sheet Modeling", submitted, 2018.

Hoffman, Perego, Price, Lipscomb, Jacobsen, Tezaur, Salinger, Tuminaro, Zhang, "MPAS-Albany land ice (MALI): A variable resolution ice sheet model for earth system modeling using Voronoi grids, Geosci. Model Develop., 2018, pp. 1-47.

Tuminaro, Perego, Tezaur, Salinger, Price, "A Matrix Dependent/Algebraic Multigrid Approach for Extruded Meshes with Applications to Ice Sheet Modeling", SIAM J. Sci. Comput.,38(5), pp. C5041-C532, 2016.

P.B. Ohm and R.S. Tuminaro, Hybrid Multigrid Methods for Nearly Structured Meshes, in Center for Computing Research Summer Proceedings 2016, J.B. Carleton and M.L. Parks, eds., Technical Report SAND2017-1294R, Sandia National Laboratories, 2016, pp. 53--62, https://cfwebprod.sandia.gov/cfdocs/CompResearch/docs/proceedings/ccr16.pdf .