A Roofline Performance Analysis of an Algebraic Multigrid Solver

Alex Druinsky, Brian Austin, Sherry Li, Osni Marques, Eric Roman and Samuel Williams

Lawrence Berkeley National Laboratory

Goal: *Develop a performance model of an optimal-order resilient solver for PDEs, using high-order discretization elements and AMG, combined with low-rank sparse factorizations.*

Two-grid Adaptive Smoothed Aggregation Spectral AMG Method[1,2]

1. Pre-smoothing: $y = x_i + M^{-1}(b - Ax_i)$
2. Coarse grid correction
   a. restrict the residual: $r_c = P^T(b - Ay)$
   b. coarse grid equation: $A_c\, x_c = r_c$
   c. update fine grid iterate: $z = y + Px_c$
3. Post-smoothing: $x_{i+1} = z + M^{-T}(b - Az)$

- Agglomerates of elements, aggregates of vertices
- P (interpolation operator): Chebyshev-based polynomial; sparse; requires eigenvectors of agglomerates and SVD
- M (smoother): Chebyshev-based polynomial; sparse
- Coarse grid solver:
  - PCG (Gauss-Seidel or Jacobi)
  - HSS (FGMRES)

$$P = S\begin{bmatrix} \bar{P}_1 & & & \\ & \bar{P}_2 & & \\ & & \ddots & \\ & & & \bar{P}_{na} \end{bmatrix}$$

$S = smoother$
$\bar{P}_i = set\ of\ eigenvectors$
$(depends\ on\ tolerance\ \theta)$
$na = number\ of\ agglomerates$

Hierarchically Semi-Separable (HSS) Low-rank Sparse Solver and Preconditioner[3,4]

- HSS structure for dense, but data-sparse structured matrices (examples: BEM, integral equations, PDEs with smooth kernels)
- Matrix off-diagonal blocks are rank deficient
- Recursion leads to hierarchical partitioning, nested bases lead to nearly linear complexity

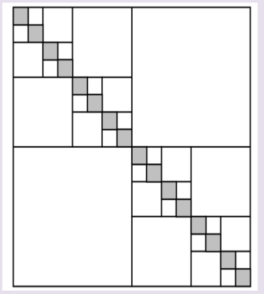

hierarchical partition

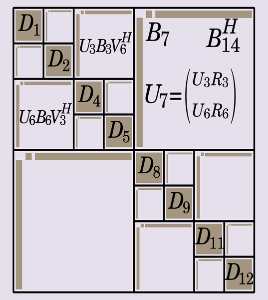

SVD with nested bases

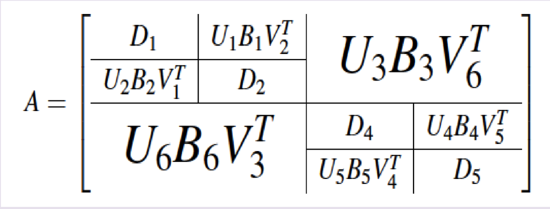

matrix block representation

Case Study: SPE10 Benchmark (model 2)[5]

- Model of a formation in the Brent oil field; 1200'x2200'x170' (fine scale cell size is 20'x10'x2'). The top 70 ft (35 layers) represents the Tarbert formation; the bottom 100 ft (50 layers) represents Upper Ness (fluvial).
- Modelled described by Brinkman equations[6]

$$\begin{cases} -\nu\Delta u + k(x)u + \nabla p = f(x) & \forall x \in \Omega \\ \text{div } u = 0 & \forall x \in \Omega \\ \mathcal{B}u = g & \text{on } \partial\Omega \end{cases}$$

fine grid
n ≈ 1.16 M
nnz ≈ 30.63 M
nnz/n ≈ 26

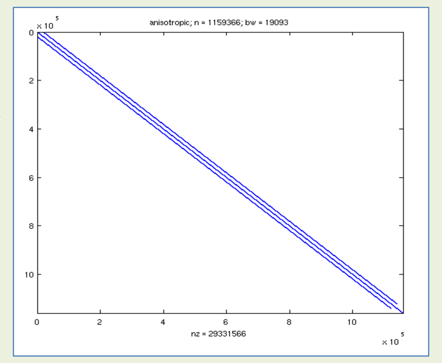

coarse grid (sample)

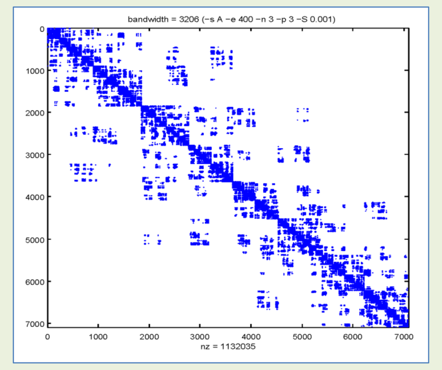

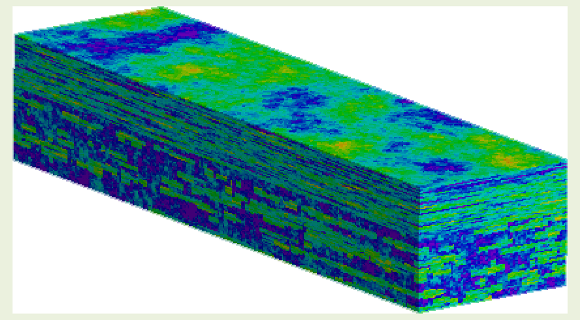

porosity of the whole model

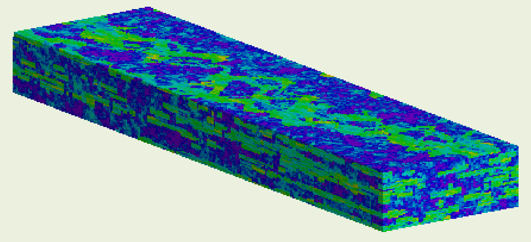

part of the Upper Ness sequence

- Parameters for AMG code: elements per agglomerate (e), degrees of polynomials for M and P (ν), spectral tolerance (θ).
- Target platform: Cray XC30 (Edison @ NERSC), 12-core Intel Ivy Bridge at 2.4 GHz; L1 and L2 caches with 64 KB and 256 KB, respectively; 30-MB L3 cache shared by the 12 cores (only one core used in the experiments below, OpenMP work in progress).
- Increasing the number of elements per agglomerate leads to a smaller but denser coarse grid, reduces the number of AMG cycles, but with a much costlier AMG hierarchy construction.
- For our purposes the important parameter is the spectral tolerance, as P becomes sparser for small tolerances (HSS derives much more benefits from sparse P).

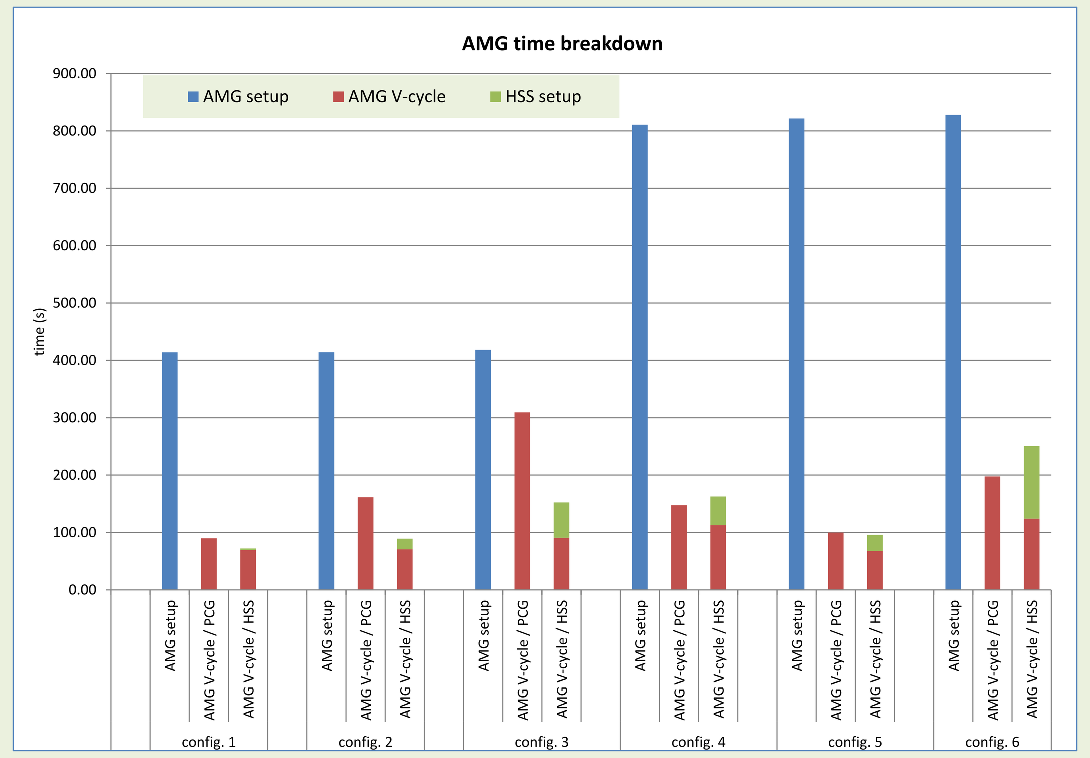

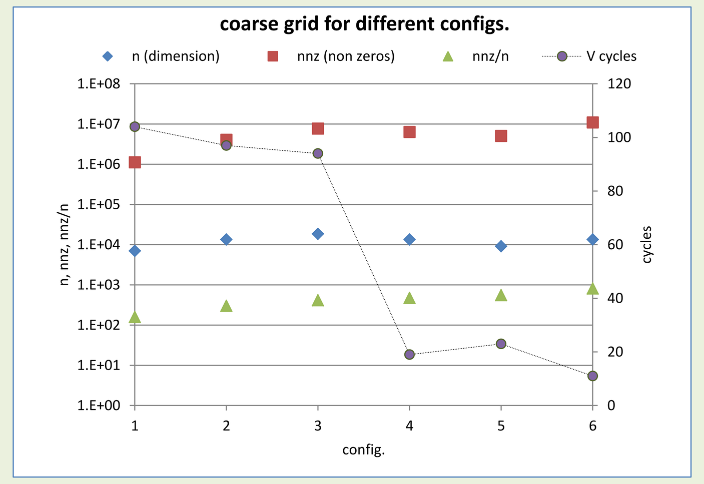

config.	e	ν	θ
1	400	3	0.001
2	400	3	0.003
3	400	3	0.005
4	1000	3	0.003
5	1000	5	0.003
6	1000	5	0.005

Performance Analysis

- A roofline performance model[7] identifies performance bounds and bottlenecks.
- We are looking at the AMG cycle part of the algorithm (ignoring the hierarchy).
- Memory traffic and flops are expressed as simple functions of parameters (number of zeros, matrix dimensions, etc.)
- Byte and flop counts translate into lower bounds on runtimes.

stage	bytes	flops
pre- and post-smooth	$(3\nu+1)(12\,\texttt{nza} + 3\cdot 8n)$	$2(3\nu+1)(\texttt{nza} + 2n)$
restrict	$12\,\texttt{nza} + 12\,\texttt{nzp} + 3\cdot 8n$	$2(\texttt{nza} + \texttt{nzp})$
coarse solve (PCG/J)		
multiply by A_c	$12\,\texttt{nzc}$	$2\,\texttt{nzc}$
preconditioner	$2\cdot 8n_c$	n_c
vector operations	$5\cdot 8n_c$	$2\cdot 5n_c$
interpolate	$12\,\texttt{nzp} + 8n$	$2\,\texttt{nzp}$
termination criterion	$12\,\texttt{nza} + 4\cdot 8n$	$2(\texttt{nza} + n)$

Runtime bounds and actual runtime for config. 1 (e=400, ν=3, θ=0.001) on one socket of the Cray XC30. Effective Gflops ranges from 1.6 for 1 core through 19.2 for 12 cores, obtained by dividing the machine peak rate by 12, based on the assumption that the compiler does not generate optimized code. STREAM benchmark bandwidths: DRAM bandwidths range from 13.2 through 48.5; L3 ranges from 21 through 210.1 GB/s. The estimates are based on traffic between DRAM and L3 for the fine grid matrix and between L3 and L2 for the coarse grid matrix.

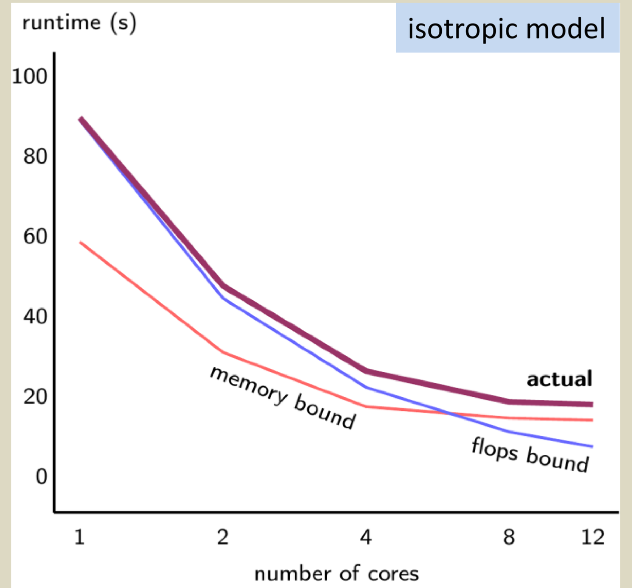

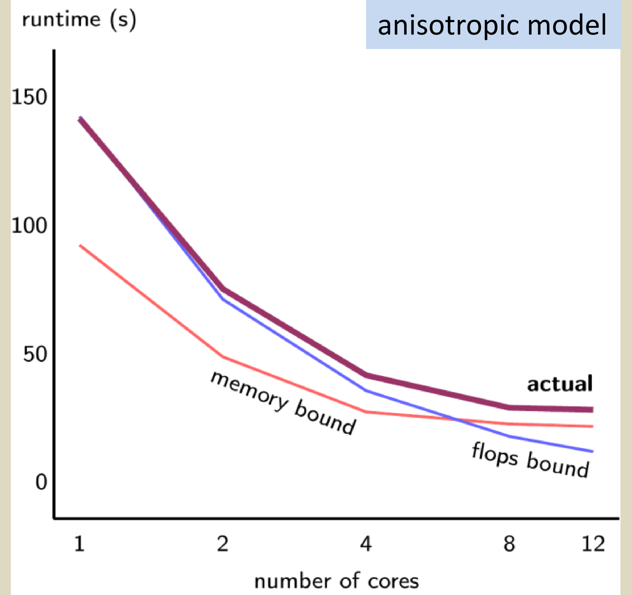

Conclusions

- For SPE10, convergence is not impaired by a small spectral tolerance and performance is improved.
- Roofline bound analysis is quite accurate (within 20%) for the PCG/GS solver, but remains an open problem for HSS (FGMRES)

References

1. D. Kalchev, C. Ketelsen and P. Vassilevski, Two-Level Adaptive Algebraic Multigrid for a Sequence of Problems with Slowly Varying Random Coefficients, SIAM J. Sci. Comp., 35:1215-1234, 2013.
2. M. Brezina and P. Vassilevski, Smoothed Aggregation Spectral Element Agglomeration AMG: SA-pAMGe, LLNL-PROC-490083, June 29, 2011.
3. J. Xia, S. Chandrasekaran, M. Gu and X. S. Li, Fast Algorithms for Hierarchically Semiseparable Matrices Numer. Linear Algebra Appl., 17:953-976, 2010.
4. S. Wang, X.S. Li, J. Xia and M. de Hoop, Efficient Scalable Algorithms for Solving Dense Linear Systems with Hierarchically Semiseparable Structures, SIAM J. Sci. Comp., 35:519-544, 2013.
5. Society of Petroleum Engineers. http://www.spe.org/web/csp//datasets/set02.htm
6. U. Villa, Scalable Efficient Methods for Incompressible Fluid-dynamics in Engineering Problems, PhD Thesis, Emory University, 2012.
7. S. Williams, A. Waterman, and D. Patterson, Roofline: An Insightful Visual Performance Model for Multicore Architectures, Commun. ACM, 52:65–76, 2009.

BERKELEY LAB

U.S. DEPARTMENT OF ENERGY | Office of Science

This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR). Project: Towards Optimal Order Resilient Solvers at Extreme Scale (TOORSES). For more information, contact Sherry Li: xsli@lbl.gov

Acknowledgments: The authors thank P. Vassilevski, D. Kalchev, A. Barker and U. Villa for providing their AMG code and helping us understand it.