

A Low-Communication Method to Solve Poisson's Equation on Locally-Structured Grids

Extended Abstract*

Peter McCorquodale

Lawrence Berkeley National Laboratory
Berkeley, California
PWMcCorquodale@lbl.gov

Brian Van Straalen

Lawrence Berkeley National Laboratory
Berkeley, California
BVStraalen@lbl.gov

Phillip Colella

Lawrence Berkeley National Laboratory
Berkeley, California
PColella@lbl.gov

Christos Kavouklis

Lawrence Livermore National Laboratory
Livermore, California
kavouklis1@llnl.gov

ABSTRACT

This poster describes a new algorithm, Method of Local Corrections (MLC), and a high-performance implementation for solving Poisson's equation with infinite-domain boundary conditions, on locally-refined nested rectangular grids. MLC represents the potential as a linear superposition of small local discrete convolutions, with global coupling represented using a non-iterative form of geometric multigrid. Thus the data motion is comparable to that of only a single V-cycle of multigrid, and hence is an order of magnitude smaller than tradition multigrid iteration. The computational kernels where most of the time is spent are 3D FFTs on small domains. Our results show solution error that is fourth order in mesh spacing, down to a fixed barrier that can be made low using high-order Mehrstellen stencils. Strong scaling tests on 64 to 4096 cores on NERSC Cori I (Haswell) show over 60% efficiency, and weak scaling by replication tests over 64 to 32768 cores show 92% efficiency on the same platform. We find comparable solve times between HPGMG on a uniform grid with one billion grid points, and MLC on the same number of grid points adaptively distributed. Since MLC operates on a nested series of adaptive locally-refined grids, it is able to solve problems with much higher resolution at the finest level than an algorithm on a uniform grid.

CCS CONCEPTS

• **Mathematics of computing** → **Partial differential equations**;
Mathematical software performance;

KEYWORDS

Poisson's equation, local corrections, domain decomposition, adaptive mesh refinement

*Poster has also been submitted.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

ICPP2018, August 2018, Eugene, Oregon, USA

© 2018 Copyright held by the owner/author(s).

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM.

<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

ACM Reference Format:

Peter McCorquodale, Phillip Colella, Brian Van Straalen, and Christos Kavouklis. 2018. A Low-Communication Method to Solve Poisson's Equation on Locally-Structured Grids: Extended Abstract. In *Proceedings of 47th International Conference on Parallel Processing (ICPP2018)*. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Poisson's equation arises in such fields as astrophysics, plasma physics, electrostatics, and fluid dynamics. We are solving Poisson's equation with infinite-domain boundary conditions:

$$\Delta\phi = f, \quad \phi(\mathbf{x}) = -\frac{1}{4\pi|\mathbf{x}|} \int f(\mathbf{y})d\mathbf{y} + o\left(\frac{1}{|\mathbf{x}|}\right) \text{ as } |\mathbf{x}| \rightarrow \infty. \quad (1)$$

The solution of this equation, using the Green's function, G , is

$$\phi(\mathbf{x}) = (G * f)(\mathbf{x}) \equiv \int G(\mathbf{x} - \mathbf{y})f(\mathbf{y})d\mathbf{y}, \quad G(\mathbf{z}) = -\frac{1}{4\pi|\mathbf{z}|}. \quad (2)$$

We solve this problem on a set of locally-refined nested rectangular grids by the Method of Local Corrections (MLC) [2, 3]. We represent the potential ϕ as a linear superposition of small local discrete convolutions of a discretized form of the Green's function with the right-hand side f , with global coupling represented using a non-iterative form of geometric multigrid.

2 LOCAL DISCRETE CONVOLUTIONS

In the discretized version of (2), we perform a discrete convolution

$$(G^h * f^h)[\mathbf{g}] \equiv \sum_{\mathbf{g}' \in \mathbb{Z}^3} h^3 G^h[\mathbf{g} - \mathbf{g}'] f^h[\mathbf{g}'] \quad (3)$$

with mesh spacing h . This convolution can be computed efficiently using 3D FFTs and the discrete Green's function G^h satisfying

$$(\Delta^h G^h)[\mathbf{g}] = \begin{cases} 1, & \text{if } \mathbf{g} = \mathbf{0}; \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

We can use high-order discrete Laplacian stencils of the form

$$(\Delta^h \phi^h)_{\mathbf{g}} = \sum_{\mathbf{s} \in [-s, s]^3} a_{\mathbf{s}} \phi^h_{\mathbf{g}+\mathbf{s}} \quad (5)$$

with truncation error

$$\Delta^h \phi^h - \Delta \phi = C_2 h^2 \Delta(\Delta \phi) + \sum_{Q'=2}^{Q/2-1} h^{2Q'} \mathcal{L}^{2Q'}(\Delta \phi) + h^Q L^{Q+2}(\phi) + O(h^{Q+2}) \quad (6)$$

where C_2 is a constant, and $\mathcal{L}^{2Q'}$ and L^{Q+2} are linear differential operators of order $2Q'$ and $Q+2$, respectively. Modifying the right-hand side of $\Delta \phi = f$ by adding appropriate derivatives of f gives a high-order approximation with compact stencil Δ^h .

3 MLC ALGORITHM

Following is our domain decomposition strategy with two levels of grid, with mesh spacings h on the finer level and H on the coarser level, $H/h = 4$. We decompose the fine domain into fixed-size patches of radius $R = (N-1)h/2$, where N is the number of grid points along each dimension.

- (1) For each fine patch, indexed by k , compute local convolutions of equation (3) on patches of radius αR ($\alpha > 1$):

$$\phi_k^{h,\text{init}} = G^h * f_k^h.$$

- (2) Accumulate coarse-grid right-hand side by summing up localized contributions:

$$F^H = \sum_k \Delta^H(\text{Coarsen}(\phi_k^{h,\text{init}})).$$

- (3) Compute global coarse convolution: $\phi^H = G^H * F^H$.
- (4) On each fine patch, solve Dirichlet problem for Poisson, with boundary conditions from a combination of ϕ^h on local patches and interpolation of ϕ^H .

For more than two levels, apply the above algorithm recursively. To improve efficiency at the finest level, we split ϕ_k^h into an inner region of radius βR where $\beta < \alpha$, on which we compute $G^h * f_k^h$, and an outer region between radius βR and radius αR where we use a projection of f_k^h to Legendre polynomials of degree P .

The solution error of this method has the form [2]:

$$\phi^{\text{MLC}} - \phi = O(h^q) + \|f\|_\infty O\left(\left(\frac{H}{\alpha N h}\right)^Q\right). \quad (7)$$

The first term comes from local truncation error and depends on local derivatives of ϕ and number of terms in the right-hand side correction. The second term comes from representation of smooth nonlocal coupling, and has no dependence on derivatives. It can be made arbitrarily low by adjusting α , N , and Q as needed.

4 RESULTS

Figure 1 shows accuracy results for two test problems: a smooth bell-shaped charge distribution, and three separated spherical charges. These are solved on multiple MLC levels, the first on uniformly refined grids and the second on adaptively refined grids. Parameters are $q = 4$, $P = 3$, either $Q = 6$ (27-point operator) or $Q = 10$ (117-point), $\alpha = 3.25$, $\beta = 2.25$, and N either 33 or 65. Both show solution error fourth order in finest mesh spacing, down to a barrier.

Figure 2 shows run times for strong scaling and replication weak scaling tests on NERSC Cori I Haswell cores. We achieve more than 60% strong scaling efficiency over 64 to 4096 cores, and 92% weak scaling efficiency over 64 to 32768 cores.

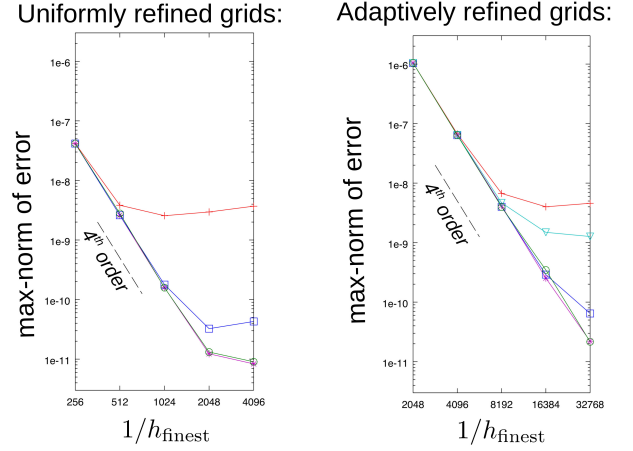


Figure 1: Accuracy results for two test problems.




Figure 2: Run times for (left) strong scaling test from 64 to 4096 cores; (right) weak scaling test from 64 to 32768 cores.

We have run a comparison with 256 cores on 8 nodes on NERSC Cori I. On a uniform grid with 10^9 points, the HPGMG benchmark [1] takes 6.1 seconds with 10 V-cycles [4]. The solve time for our MLC is 10.7 seconds for 10^9 grid points that are adaptively distributed, but with only 0.2% of the domain refined at the finest level. So the run time per grid point of MLC is roughly comparable to that of HPGMG, but the great advantage of MLC is adaptive mesh refinement, concentrating grid points in local regions of interest.

ACKNOWLEDGMENTS

This research was supported at the Lawrence Berkeley National Laboratory by the Office of Advanced Scientific Computing Research of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

REFERENCES

- [1] Mark F. Adams, Jed Brown, John Shalf, Brian Van Straalen, Erich Strohmaier, and Samuel Williams. 2014. HPGMG 1.0: A Benchmark for Ranking High Performance Computing Systems. (05 2014).
- [2] C. Kavouklis and P. Colella. 2016. Computation of Volume Potentials on Structured Grids Using the Method of Local Corrections. *arXiv preprint arXiv:1702.08111* (2016). submitted to CAMCoS.
- [3] P. McCorquodale, P. Colella, G. Balls, and S. Baden. 2007. A local corrections algorithm for solving Poisson's equation in three dimensions. *Communications in Applied Mathematics and Computational Science* 2, 1 (2007), 57–81.
- [4] Samuel Williams. 2018. private communication.