

Implementing an Adaptive Sparse Grid Discretization (ASGarD) for High Dimensional Advection-Diffusion Problems on Exascale Architectures

M. Graham Lopez

Adam McDaniel

National Center for Computational Sciences

Oak Ridge National Laboratory

Oak Ridge, TN, USA

lopezmg@ornl.gov

Ed F. D'Azevedo

Wael Elwasif

Hao Lu

Lin Mu

Computer Science and Mathematics Division

Oak Ridge National Laboratory

Oak Ridge, TN, USA

David L. Green

Diego Del-Castillo-Negrete

Fusion Energy Division

Oak Ridge National Laboratory

Oak Ridge, TN, USA

greendl1@ornl.gov

B. Tyler McDaniel

Timothy R. Younkin

Electrical Engineering and Computer Science

University of Tennessee

Knoxville, TN, USA

ABSTRACT

Many scientific domains require the solution of high dimensional PDEs. Traditional grid- or mesh-based methods for solving such systems in a noise-free manner quickly become intractable due to the scaling of the degrees of freedom going as $O(N^d)$, sometimes called "the curse of dimensionality." We are developing an arbitrarily high-order discontinuous-Galerkin finite-element solver that leverages an adaptive sparse-grid discretization whose degrees of freedom scale as $O(N * \log 2N^D - 1)$. This method and its subsequent reduction in the required resources is being applied to several PDEs including time-domain Maxwell's equations (3D), the Vlasov equation (in up to 6D) and a Fokker-Planck-like problem in ongoing related efforts. Here we present our implementation which is designed to run on multiple accelerated architectures, including distributed systems. Our implementation takes advantage of a system matrix decomposed as the Kronecker product of many smaller matrices which is implemented as batched operations.

CCS CONCEPTS

• **Applied computing** → **Physical sciences and engineering; Physics;**

This manuscript has been authored by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725 with the U.S. Department of Energy. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a non-exclusive, paid-up, irrevocable, world-wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. The Department of Energy will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

This paper is authored by an employee(s) of the United States Government and is in the public domain. Non-exclusive copying or redistribution is allowed, provided that the article citation is given and the authors and agency are clearly identified as its source.

Conference'17, July 2017, Washington, DC, USA

© 2019

ACM ISBN 978-x-xxxx-xxxx-x/YY/MM...\$15.00

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

KEYWORDS

high dimensional, sparse grids, finite element, portability

ACM Reference Format:

M. Graham Lopez, Adam McDaniel, David L. Green, Diego Del-Castillo-Negrete, Ed F. D'Azevedo, Wael Elwasif, Hao Lu, Lin Mu, B. Tyler McDaniel, and Timothy R. Younkin. 2019. Implementing an Adaptive Sparse Grid Discretization (ASGarD) for High Dimensional Advection-Diffusion Problems on Exascale Architectures. In *Proceedings of ACM Conference (Conference'17)*. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

In many scientific disciplines the underlying processes can be described by high dimensional ($D > 3$) partial differential equations (PDEs). However, our ability to simulate such systems is limited by the $\sim N^D$ exponential growth in the degrees of freedom (unknowns) inherent in such cases. This "curse of dimensionality" has traditionally meant avoiding the use of continuum (mesh- or grid-based) methods for high dimensional problems in favour of Monte-Carlo approaches where the computational resources required are more tractable. While these approaches [e.g., 2] have demonstrated significant successes, they still suffer from Monte-Carlo (MC) noise which is proportional to $\sqrt{N_p}$ (where N_p is the number of MC samples). As such, there is motivation to develop a capability to apply the noise-free continuum methods to high-dimensional problems. This is particularly true for the challenges in simulation presented by magnetically confined plasma in nuclear fusion where the underlying PDE is 6D in nature (cf. Fig. 1).

Here we present the implementation details of the "Adaptive Sparse Grid Discretization" (ASGarD) code which implements a finite-element solver atop an adaptive sparse-grid[3] discretization such that the degrees of freedom required are significantly reduced my dropping grid points associated with high order mixed derivatives not present in the solution.

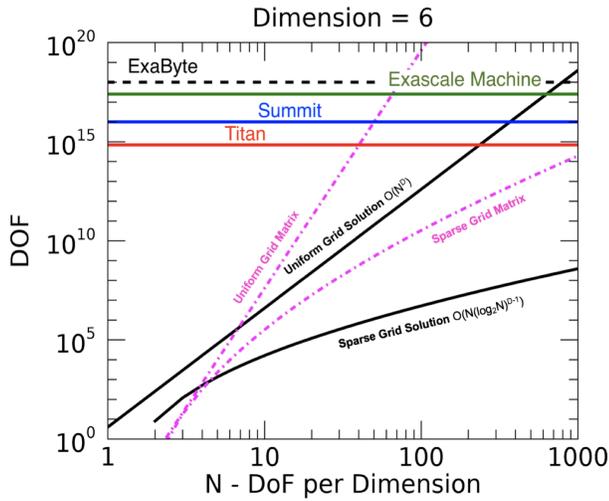


Figure 1: Estimate of degrees-of-freedom vs. number of grid points per dimension for a 6D system assuming a naive "full" grid, and the "sparse" grids implemented in this paper.

We implement the discontinuous-Galerkin [1] method using a truncated tensor product D -many 1D hierarchical bases. Our implementation allows arbitrary dimension and order and employs either explicit or implicit time stepping. The arbitrary dimensionality and order aspects of the formulation manifests as the Kronecker product of D -many small matrices of size order by order (where the order is typically less than 10). Fig. 2 illustrates a 2D case.

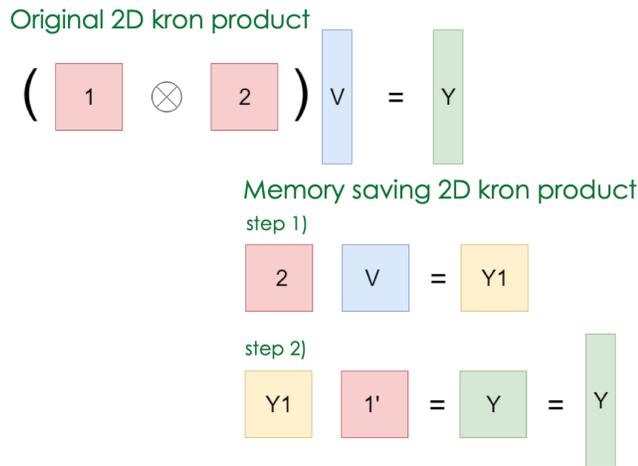


Figure 2: Upper: The Kronecker product which forms the core computational kernel for a 2D problem, i.e., advance the system from a state represented by vector V to state Y . Lower: Splitting of the dimensions to avoid explicit construction of the tensor product expanded system matrix.

The core time advance kernel is the application of the tensor product encoded system matrix to the solution vector at every

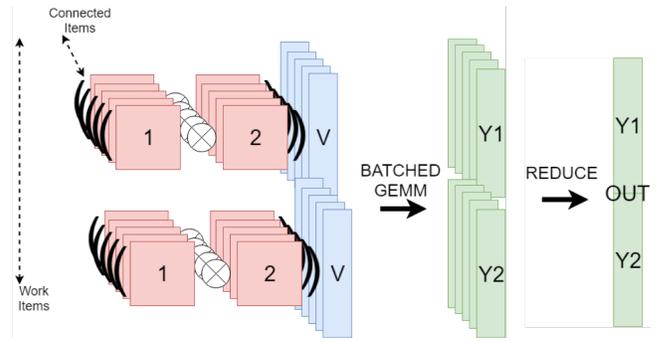


Figure 3: Illustration of the resulting batched-GEMM + reduction operations where the batching is over each finite element (vertical) and connected elements (horizontal).

timestep. Formation of this matrix is infeasible for large problems and is instead done in a matrix-free manner according to Figs. 2 and 3 where we decompose the problem and calculate the contribution of each element-to-element relationship (a connection) in parallel. Each connection requires the tensor product of small (e.g., 8×8) sections of D -many operator matrices, then multiply the result with a segment of the solution vector, which are then batched across connections. The GEMMs within each batch can be performed in parallel, with outputs from one dimension becoming inputs to the next higher dimension. Finally, we reduce the connections for each element to form a segment of the solution vector.

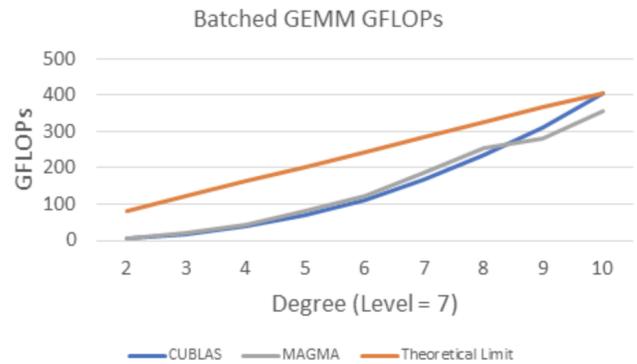


Figure 4: Computational efficiency of GPU enabled batched-GEMM for a 2D problem of various basis function order.

The observed performance of this approach is shown in Fig. 4. While the batching reaches the theoretical peak for the provided matrix sizes, it is far from the device peak. To address this, we will discuss options for fusing pieces of the Kronecker product, thereby improving performance.

ACKNOWLEDGMENTS

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

REFERENCES

- [1] P. F. Antonietti. 2016. Review of Discontinuous Galerkin Finite Element Methods for Partial Differential Equations on Complicated Domains. In *Building Bridges: Connections and Challenges in Modern Approaches to Numerical Partial Differential Equations. Lecture Notes in Computational Science and Engineering*, Cangiani A. Georgoulis E. Barrenechea G., Brezzi F. (Ed.). Vol. 114. Springer, The address of the publisher.
- [2] A. Bottino and E. Sonnendrüjcker. 2015. Monte Carlo particle-in-cell methods for the simulation of the Vlasov-Maxwell gyrokinetic equations. *Journal of Plasma Physics* 81, 5 (2015), 435810501. <https://doi.org/10.1017/S0022377815000574>
- [3] Hans-Joachim Bungartz and Michael Griebel. 2004. Sparse grids. *Acta Numerica* 13 (2004), 147-269. <https://doi.org/10.1017/S0962492904000182>