

Exasim^{*}: Generating Discontinuous Galerkin Codes for Numerical Solutions of Partial Differential Equations on Graphics Processors

Jordi Vila-Pérez^{*}, R. Loek Van Heyningen, Ngoc-Cuong Nguyen, Jaume Peraire
*Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge,
02139, Massachusetts, USA*

Abstract

This paper presents an overview of the functionalities and applications of Exasim, an open-source code for generating high-order discontinuous Galerkin codes to numerically solve parametrized partial differential equations (PDEs). The software combines high-level and low-level languages to construct parametrized PDE models via Julia, Python or Matlab scripts and produce high-performance C++ codes for solving the PDE models on CPU and Nvidia GPU processors with distributed memory. Exasim provides matrix-free discontinuous Galerkin discretization schemes together with scalable reduced basis preconditioners and Newton-GMRES solvers, making it suitable for accurate and efficient approximation of wide-ranging classes of PDEs.

Keywords: Parametrized PDE models, discontinuous Galerkin, GPU, CPU, automatic code generation, exascale computing

1. Motivation and significance

The use of high-order methods to solve partial differential equations (PDEs) has experienced a growing interest among practitioners [1–3] in different areas of engineering and science. Their increased accuracy at a reduced computational cost and their low diffusion and dispersion errors confer them a major advantage when compared to low-order schemes [4, 5]. In particular, DG formulations have become one of the most adopted high-order approaches in many different areas [6–12]. DG methods rely on a locally conservative formulation that ensures

^{*}<https://github.com/exapde/Exasim>

^{*}Corresponding author.

Email address: jvilap@mit.edu (Jordi Vila-Pérez)

high-order accuracy on unstructured meshes. In addition, they provide a stable definition of the convection operator and allow suitable *hp*-adaptivity strategies [13–16] and an efficient exploitation of parallel computing architectures [17].

Because of these reasons, different classes of DG discretizations have been proposed over the last years [18–23]. In particular, hybridized DG (HDG) methods have gained popularity for the numerical solution of all classes of PDEs [24–34], based on their superiority with respect to other DG alternatives both in terms of accuracy and computational complexity [35, 36]. However, their conjunction with nonlinear solvers imposes a high memory footprint, either in terms of building the system matrices, or dealing with larger systems of equations in matrix-free approaches [37]. As a result, these methods offer a poorer scalability on GPU platforms, what makes them unsuitable for solving large problems [38].

This paper presents `Exasim`, an open-source software for generating high-order DG codes, based on the local DG (LDG) method [21, 39]. `Exasim` performs an implicit matrix-free approach [38, 40] that makes the code suitable for running on multiple architectures using both CPUs and GPUs, allowing the solution of large-scale problems. Moreover, it exploits a parametrized formulation of PDEs, expressing them as first order systems of equations with the eventual inclusion of ordinary differential equations (ODEs) to form a system of differential-algebraic equations (DAEs). This general framework reduces the mathematical description of the model to the definition of state variables, fluxes, source terms, together with initial and boundary conditions, and allows the user to model multiple PDE systems. Furthermore, the software presents a user interface in a high-level language (Julia, Python or Matlab) where one can specify the aforementioned terms symbolically in a seamless effort. Then, the code performs a preprocessing stage that generates C++ code, which interfaces with C++ and CUDA kernels and allows the software to run on different computing platforms.

With this approach, `Exasim` offers an open-source product that can be easily adopted for users with any kind of expertise in DG discretizations. At the same time, it serves as an advanced research tool, capable to run in different architectures with multiple processors and with suitable scalability properties, thus allowing to tackle complex problems that are beyond the capabilities of existing codes [38, 40]. Indeed, whereas the implicit-in-time matrix-free approach shares some common features with some available DG codes [41–43], it contrasts with most of the available open-source DG approaches, such as [44–51], which are CPU-based, and also with other codes running on GPUs which rely on explicit time-marching schemes, like [52, 53].

The remainder of this paper is structured as follows. Section 2 describes the parametrized PDE models and discretization methods that `Exasim` handles. Then, section 3 details the code architecture and its different functionalities. A number of examples are presented in section 4 to illustrate the capabilities of the soft-

ware. Finally, sections 5 and 6 summarize the impact of this work and the main conclusions of this paper, respectively.

2. Models and Discretization Methods

Exasim produces executable DG code to solve a wide variety of PDE models that can be described under general parametrized formulations and classified under convection, diffusion or wave-type equations (models C, D and W in the software), such as those listed in Table 1, and described as follows.

Convection model	Linear and nonlinear convection, Burgers equation, Euler equations, shallow water equations.
Diffusion model	Poisson equation, convection-diffusion equations, linear and nonlinear elasticity, compressible and incompressible Navier-Stokes.
Wave model	Wave equation, linear and nonlinear elastodynamics, Maxwell's equations.

Table 1: Examples of PDEs belonging to each of the model categories in Exasim.

2.1. Parametrized PDE models

The underlying PDE system must be written as a set of first-order PDEs and can be coupled as well with a certain ordinary differential equation (ODE) to form a system of differential-algebraic equations. For instance, the diffusion model, defined in the open domain $\Omega \subset \mathbb{R}^{n_d}$ for $t_f > 0$, and expressed in its more general version, reads as

$$\mathbf{q} + \nabla \mathbf{u} = \mathbf{0}, \quad \text{in } \Omega \times (0, t_f], \quad (1a)$$

$$\mathbf{m}(\tilde{\mathbf{u}}, \mathbf{x}, t, \boldsymbol{\mu}) \frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathbf{f}(\tilde{\mathbf{u}}, \mathbf{x}, t, \boldsymbol{\mu}) = \mathbf{s}(\tilde{\mathbf{u}}, \mathbf{x}, t, \boldsymbol{\mu}), \quad \text{in } \Omega \times (0, t_f], \quad (1b)$$

$$\alpha \frac{\partial \mathbf{w}}{\partial t} + \beta \mathbf{w} = \mathbf{s}_w(\tilde{\mathbf{u}}, \mathbf{x}, t, \boldsymbol{\mu}), \quad \text{in } \Omega \times (0, t_f], \quad (1c)$$

with appropriate initial and boundary conditions. Here, the set of state variables $\tilde{\mathbf{u}} = (\mathbf{u}, \mathbf{q}, \mathbf{w}) \in \mathbb{R}^{n_{cu}} \times \mathbb{R}^{n_{cu} \times n_d} \times \mathbb{R}^{n_w}$ is the exact solution of the PDE model, $\mathbf{x} \in \Omega$ is the vector of coordinate variables, t represents time variable in $(0, t_f]$, and $\boldsymbol{\mu} \in \mathbb{R}^{n_{param}}$ is a vector of physical parameters. Additionally, the vector-valued function $\mathbf{m} \in \mathbb{R}^{n_{cu}}$ is a mass function, the matrix-valued function $\mathbf{f} \in \mathbb{R}^{n_{cu} \times n_d}$ is a flux and the vector-valued function $\mathbf{s} \in \mathbb{R}^{n_{cu}}$ is a source term. Similarly, $\alpha \in \mathbb{R}$, $\beta \in \mathbb{R}$ and $\mathbf{s}_w \in \mathbb{R}^{n_w}$ are, respectively, two parameters and a source term for the additional ODE.

The diffusion model (1a–1b) reduces to the convection model when the state variables $\tilde{\mathbf{u}}$ do not contain \mathbf{q} and equation 1a is not included in the model. The wave model derives from the diffusion model when equation 1a is replaced by

$$\frac{\partial \mathbf{q}}{\partial t} + \nabla \mathbf{u} = \mathbf{0}, \quad \text{in } \Omega \times (0, t_f], \quad (1d)$$

and employs the ODE equation 1c to recover the displacement field, \mathbf{w} . Finally, note that besides the diffusion, convection, and wave models, `Exasim` can solve higher-order PDE models by rewriting them as a first-order system of equations.

2.2. Discretization methods

`Exasim` implements the LDG method for the spatial discretization and the diagonally implicit Runge-Kutta (DIRK) method for temporal discretization [54]. The LDG discretization of the PDE model (1) leads to a semi-discrete system of equations involving the source, flux, mass, numerical trace, and numerical flux functions. In `Exasim`, these functions are input as mathematical expressions of symbolic variables in a script file. This allows users to define the LDG discretization of a PDE model by merely writing mathematical functions in a high-level language setting. Note as well that different DG methods can be implemented in `Exasim` for the spatial discretization of PDE systems by providing suitable expressions of the numerical traces and numerical fluxes.

On the other hand, `Exasim` employs several DIRK schemes for the time integration, from first-order implicit Euler to higher-order schemes, such as the three-stage four-order DIRK scheme. The number of stages and order of accuracy can be specified by users.

3. Software description

`Exasim` combines a high-level interface for preprocessing and code generation with C++ language to obtain high-performance codes that can run on both CPU and GPU architectures. In particular, the functions and parameters that define the PDE model (1) are specified via scripts in a high-level language (Julia, Python or Matlab). An automatic code generation module is then responsible for converting them into C++ codes that handle fluxes, source terms, boundary conditions and initial conditions. Finally, the kernel code, written in C++ with MPI-based parallelization and CUDA for GPUs, implements the corresponding discretization and solution methods. `Exasim` leverages a number of external libraries and software, such as the linear algebra libraries BLAS and LAPACK, Gmsh for mesh generation, METIS for mesh partitioning, GPU-aware MPI libraries, CUDA Toolkit for GPU architectures and Paraview for visualization.

A summary of this code architecture is depicted in Figure 1, which illustrates the different tasks and dependencies together with some of the software main functionalities. Some of the main capabilities of the code related to the solution methods, meshing and visualization, besides its GPU parallel performance are described as follows.

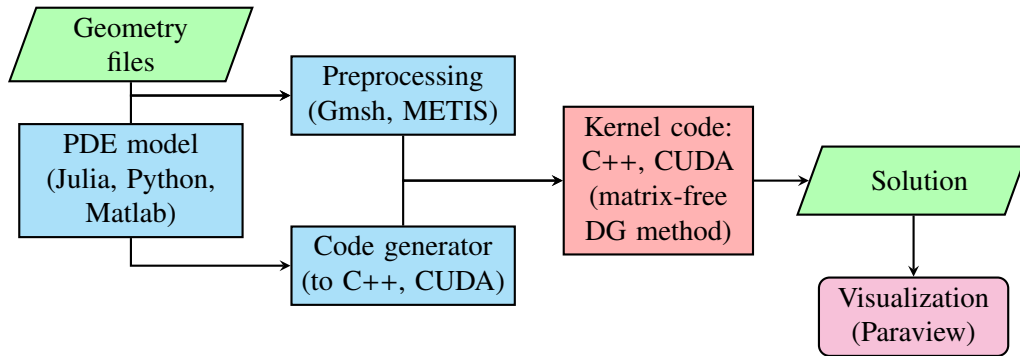


Figure 1: Flowchart summarizing the code architecture of Exasim and the dependencies among the different processes. The color code is the following: processes implemented in high/low-level languages (blue/red, respectively), input and output files (green), external functionalities (magenta).

3.1. Automatic code generation

Exasim generates both standard C++ code and CUDA code from the mathematical functions written in Julia, Python, or Matlab. To this end, Exasim employs symbolic libraries such as Sympy and Matlab’s symbolic toolbox to generate C code which is then converted to C++/CUDA code by Exasim’s code generator. The resulting code is automatically optimized, given that common subexpression elimination (CSE) tool is used to eliminate duplicate expressions.

3.2. High-order mesh generation

Exasim provides a Mesh module to generate meshes for simple geometries. Similarly, Exasim uses Gmsh [55] to generate meshes from geometry model files. Nevertheless, any alternative open-source mesh generators such as CUBIT [56], CGAL [57], DistMesh [58], TetGen [59], Mmg [60], MeshLab [61] or SALOME [62], among others, can be used for complex geometries. Because a high-order mesh is needed for the DG discretization of a PDE model, Exasim produces the high-order mesh from a standard finite element mesh by curving its corresponding boundaries and mapping the nodes of the mesh appropriately [63, 64].

3.3. Jacobian-free Newton-Krylov solvers

The LDG method is Exasim’s default discretization, since it enables an efficient implementation of a matrix-free solution method well suited for GPU architectures. The resulting system of equations arising from DG/DIRK discretization

is solved by using a matrix-free Newton-GMRES method, thus avoiding the need to construct Jacobian matrices. The performance of the GMRES solver is accelerated by means of a matrix-free preconditioner which is constructed using the reduced basis method and a low-rank approximation to the Jacobian matrix. The matrix-vector products in GMRES can be computed in two different ways in `Exasim`, either by means of a finite difference approximation or by an automatic differentiation (AD) approach relying on the external package `Enzyme` [65, 66]. In addition, a number of different algorithms, including tensor-product with sum-factorization for residual evaluation or an automatic tuning for customized GPU allocation, are used to optimize the code performance. More details of the different algorithms employed for the discretization and solution method can be found in [38, 40].

3.4. Visualization

`Exasim` uses `Paraview` to visualize and analyze the numerical solutions. To this end, a postprocessing tool is employed to generate the corresponding `VTK/VTU` files from the solution data. The visualization is performed immediately once the simulation is completed.

3.5. GPU scalability

Whereas `Exasim` can run both on CPU and GPU architectures, the code features a set of numerical algorithms suited to optimize the performance in GPU systems. Indeed, a significant performance gain for GPU architectures has been reported in [38], indicating an improvement of more than one order of magnitude in runtime with respect to CPU machines.

The software presents excellent performance in weak and strong scaling tests, as illustrated in Table 2, for the direct numerical simulation (DNS) of the Purdue flared cone [67, 68]. The simulation employed third-order DG and DIRK schemes and up to 768 nodes at the OLCF’s Summit supercomputer, with one MPI rank per GPU. A degradation of about 5% is obtained in the weak scaling test when increasing from 24 to 768 nodes, whereas a slow degradation can be observed as well in the strong scaling results as the number of nodes increases.

Nodes	Weak scaling			Strong scaling		
	DOFs	Time (s)	Time ratio	DOFs	Time (s)	Time ratio
24	0.408B	3.10	1.000	1.632B	12.45	1.000
48	0.816B	3.12	1.006	1.632B	6.25	0.502
96	1.632B	3.14	1.024	1.632B	3.14	0.252
192	3.264B	3.17	1.023	1.632B	1.59	0.128
384	6.858B	3.20	1.032	1.632B	0.81	0.065
768	13.056B	3.25	1.048	1.632B	0.43	0.035

Table 2: Weak scaling and strong scaling tests of Exasim on NVIDIA Tesla V100 GPUs. The time column indicates the physical time needed per time-step.

4. Illustrative Examples

Exasim has a large collection of examples including convection-diffusion, heat transfer, compressible flows, wave propagation and magnetohydrodynamics problems. This section presents four different examples which are representatives of convection, diffusion and wave models.

4.1. Convergence study on the Poisson equation

A 3D analytical solution of the Poisson equation is presented to verify the convergence properties of the proposed numerical method, in particular for the diffusion model. The example corresponds to a case of heat diffusion with a source term, i.e. $-\nabla^2 u = f$ on the unit cube $\Omega = (0, 1)^3$, with analytical solution $u = \sin(\pi x) \sin(\pi y) \sin(\pi z)$. The problem is solved in a set of structured tetrahedral grids and using different polynomial degrees of approximation, p . The approximation errors in the different mesh refinements for both the primal and mixed variables, u_h and \mathbf{q}_h , are detailed in Table 3. The expected optimal convergence rates of $p + 1$ are obtained for the approximation of the primal variable, u_h , whereas convergence of order p is obtained for the mixed variables, \mathbf{q}_h .

Mesh 1/n	p = 1		p = 2		p = 3		p = 4	
	Error	Rate	Error	Rate	Error	Rate	Error	Rate
$E_{\mathcal{L}_2}(u_h) = \ u_h - u\ _{\mathcal{L}_2(\Omega)} / \ u\ _{\mathcal{L}_2(\Omega)}$								
2	1.51e-01	–	3.65e-02	–	8.95e-03	–	2.31e-03	–
4	5.68e-02	1.41	4.51e-03	3.02	6.99e-04	3.68	8.18e-05	4.82
6	1.68e-02	1.76	5.69e-04	2.99	4.70e-05	3.90	2.65e-06	4.95
16	4.44e-03	1.92	7.28e-05	2.97	3.00e-06	3.97	8.45e-08	4.97
32	1.13e-03	1.97	9.24e-06	2.98	1.91e-07	3.98	5.58e-09	×
$E_{\mathcal{L}_2}(q_h) = \ q_h - q\ _{\mathcal{L}_2(\Omega)} / \ q\ _{\mathcal{L}_2(\Omega)}$								
2	5.71e-01	–	2.19e-01	–	5.67e-02	–	1.31e-02	–
4	3.67e-01	0.64	5.60e-02	1.97	9.47e-03	2.58	8.87e-04	3.89
6	2.05e-01	0.84	1.33e-02	2.07	1.33e-03	2.83	5.41e-05	4.03
16	1.07e-01	0.94	3.26e-03	2.03	1.74e-04	2.94	3.37e-06	4.01
32	5.42e-02	0.98	8.12e-04	2.01	2.20e-05	2.98	7.48e-07	×

Table 3: Poisson example – History of convergence of the primal and mixed variables, u_h and q_h (top and bottom respectively) in the 3D Poisson example, using uniform meshes of tetrahedrons and different polynomial degrees of approximation. × indicates that convergence rates cannot be accurately computed due to finite precision arithmetic issues.

4.2. Scattering of a planar wave by a circular cylinder

The second example presented in this work corresponds to the propagation of an acoustic planar wave of wavenumber $\mathbf{k} = (10, 0)$ in a medium with unit permittivity, $\varepsilon = 1$, and speed of sound, $c = 1$, scattered by a 2D cylinder of unitary radius. The problem is solved in the square domain $\Omega = (-12, 12)^2$, with absorbing boundary conditions on the outer boundaries and Neumann boundary conditions on the cylinder boundary. The numerical simulation is performed using different polynomial orders of approximation, from $p = 3$ to $p = 5$, and a DIRK(3,4) scheme, employing an unstructured mesh composed by 4224 triangles with curved boundaries, shown in Figure 2a. The scattered wave solution is depicted in Figures 2b and 2c for $p = 3$ and $p = 5$, respectively, after 50 periods of time, whereas Figure 2d illustrates the quadratic decay of the displacement field along the line $\theta = \pi/9$. Finally, Figure 2e shows the radar cross-section (RCS) [69] for different orders of approximation, measuring the propagated energy and highlighting the importance of high-order approximations in time and space to ensure low dispersion and diffusion errors on the wave propagation.

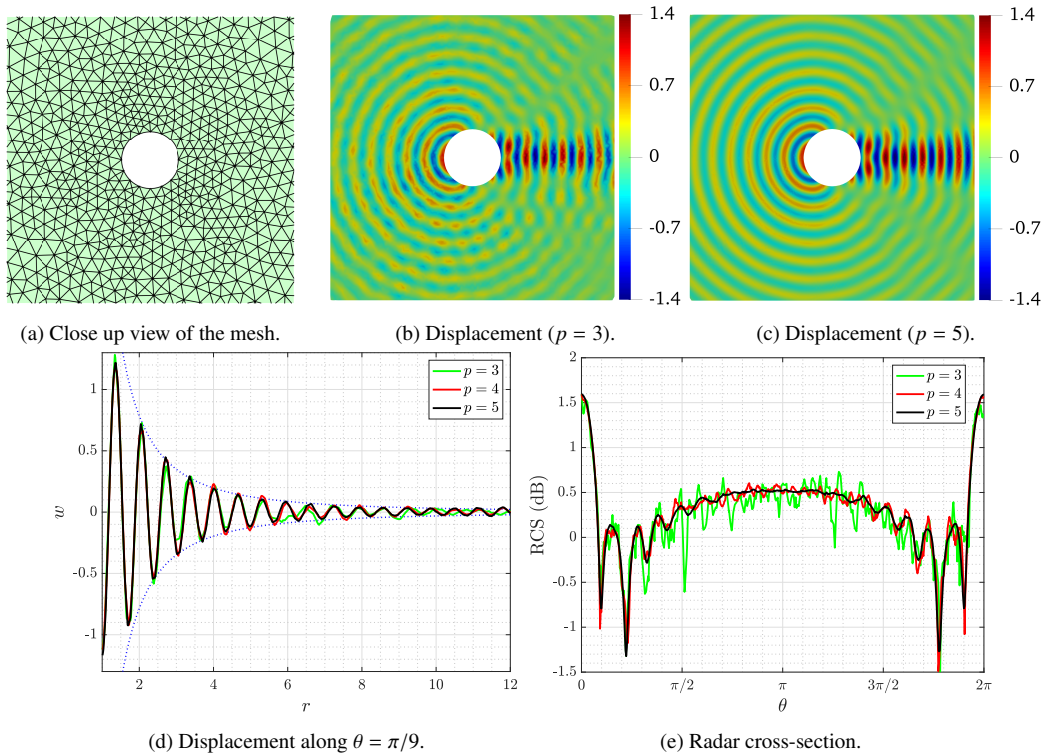


Figure 2: Wave scattering – Detail of (a) the unstructured mesh with curved boundaries, the instantaneous displacement field at $t = 50T$ for (b) $p = 3$, (c) $p = 5$ and (d) along the line $\theta = \pi/9$, and (e) the radar cross-section. The dotted lines in (d) indicate a quadratic decay.

4.3. Bickley jet

A certain configuration of the Bickley jet case (see for instance [70]), using the shallow water equations, illustrates an example of the convection model. The example models the temporal evolution of a jet flow in the square domain $\Omega = (-2\pi, 2\pi)^2$, subject to slight initial perturbations on the velocity field. The problem is solved employing a Cartesian mesh of 128×128 quadrilaterals and periodic boundary conditions, given a non-dimensional gravity value of $g = 10^4$. The simulation employs $p = 4$ polynomials and a DIRK(3,3) temporal integration scheme. Some sketches of the velocity magnitude field for different simulation times are depicted in Figure 3. Exasim solves this case without need of any physical or artificial diffusion, showing a good resolution and propagation of the flow perturbations.

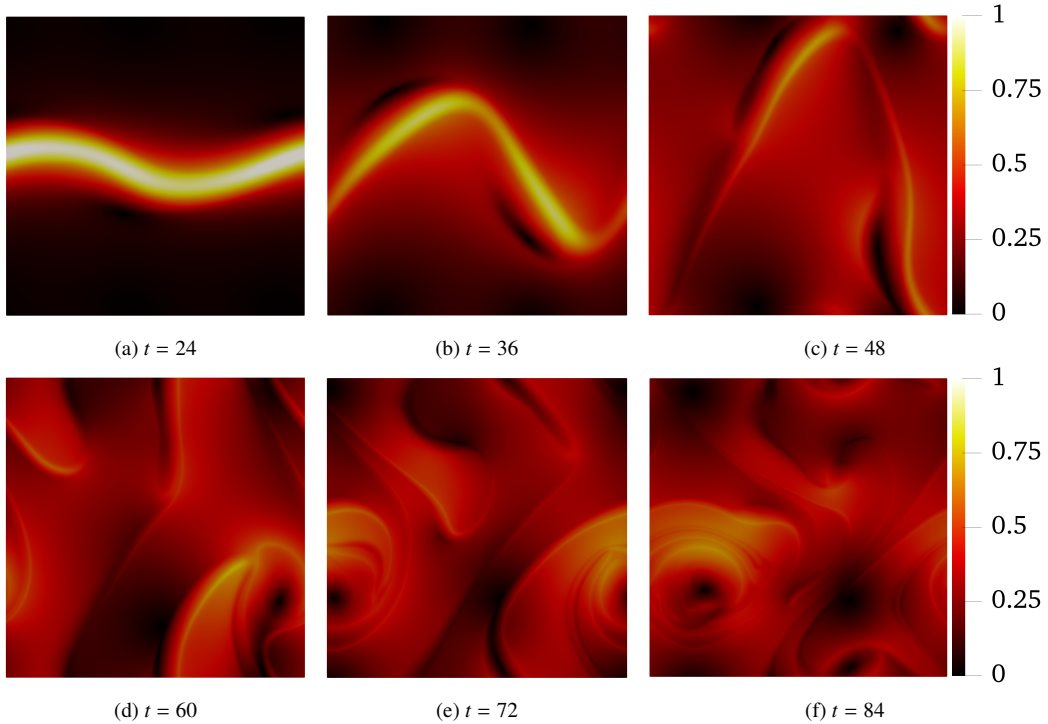


Figure 3: Bickley jet – Magnitude of the velocity field at different simulation times, computed with $p = 4$ polynomials and a DIRK(3,3) temporal scheme.

4.4. Taylor-Green vortex

Finally, the Taylor-Green vortex at $Re = 1600$ is considered to show the potential of Exasim in solving a 3D case featuring several millions of degrees of freedom. The case is solved using an implicit large-eddy simulation (ILES) approach, which relies on the inherent numerical dissipation of the DG discretization, contrary to introducing a subgrid scale (SGS) model to account for the small scales of the flow. The example is solved in $\Omega = (0, 2\pi)^3$ with periodic boundary conditions, employing a structured mesh of 64^3 cubes and fourth-order polynomials.

The kinetic energy dissipation is evaluated in Figure 4 in comparison to a DNS reference solution [71]. On the one hand, Figure 4a illustrates the evolution of the kinetic energy rate, showing a strong agreement between the present ILES solution and the reference DNS solution. On the other hand, Figure 4b depicts the kinetic energy spectrum at the instant of maximum dissipation, $t = 9$. The ILES solution establishes a close comparison with respect to the reference solution for wavenumbers $k < 80$, whereas it overpredicts the energy in the highest wavenumbers. The wavespectrum also shows a first range of wavenumbers up to $k \simeq 20$ when it follows a power law with exponent close to the theoretical value of $-5/3$.

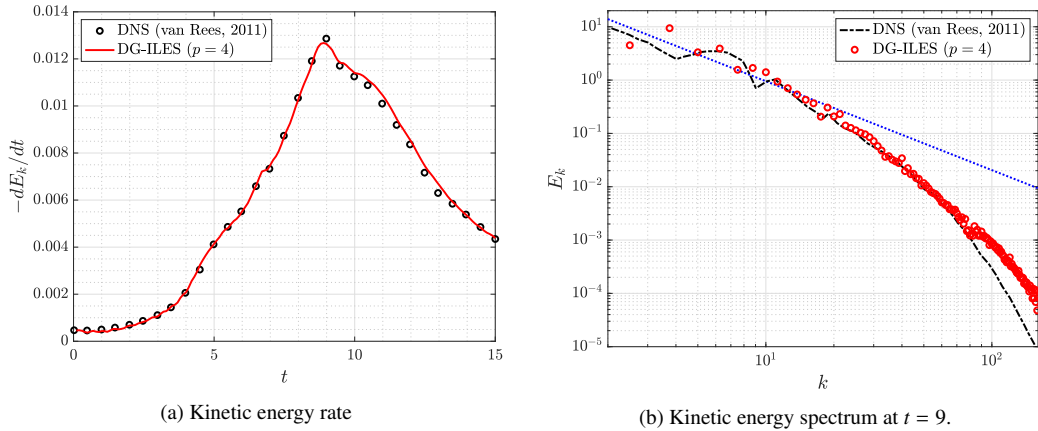


Figure 4: Taylor-Green vortex – Evolution of the kinetic energy rate with time (a) and kinetic energy spectrum at non-dimensional time $t = 9$ (b). The dotted line indicates a slope of $-5/3$.

5. Impact

`Exasim` is aimed at making DG methods accessible to users, offering at the same time a robust solver capable to handle large nonlinear systems of equations by means of innovative numerical algorithms. To do so, a high-level interface allows users to specify the analytical expressions describing fluxes, source terms, boundary conditions and initial conditions of a general parametrized PDE system in simple Python, Julia or Matlab scripts. This simple preprocessing step is then integrated within C++ and CUDA kernels with MPI-based parallelization, giving rise to a high-performance code capable to run on several machines, from laptops to supercomputers, with both CPU and GPU processors.

Furthermore, the software employs a set of GPU-accelerated numerical algorithms suited for this kind of architectures, allowing to take advantage of its computational power and to dramatically increase the scale and size of the numerical simulations `Exasim` is able to tackle. In this manner, the software represents a key tool itself to face practical problems of interest in different physical applications whose computational demands are inaccessible by existing codes.

To this end, `Exasim` has been already used in different large-scale LES computations predicting transitional and turbulent flows in transonic or hypersonic applications [38, 40]. In addition, given the simplicity for formulating PDE models, `Exasim` has the potential to facilitate the introduction of modern high-order DG discretizations in physical systems involving complex descriptions and coupled phenomena.

Finally, the modular structure of `Exasim` permits to exploit the full capabilities of external libraries and packages. For instance, this is the case of `Enzyme`, which has been integrated within the software to perform forward mode automatic

differentiation in the matrix-vector products associated to the GMRES iterations.

6. Conclusions

This paper presents `Exasim`, an open-source software for generating discontinuous Galerkin codes for the numerical solution of PDEs. The code features a high-level user interface in Julia, Python and Matlab for preprocessing and code generation and combines it with C++ and CUDA/MPI kernels, producing a high-performance code able to run on CPU and GPU platforms. The code exploits a matrix-free solution method that provides full GPU functionality and excellent scalability on this kind of architectures. The software has been validated on many different applications and examples, and it constitutes both a basic learning framework and an innovative numerical tool for advanced computational research used to improve our understanding of complex flow physics.

7. Conflict of Interest

The authors declare that there are no known conflicts of interest associated with this publication and that there has been no significant financial support for this work that could have influenced its outcome.

Acknowledgements

The authors would like to thank Sebastien Terrana his contribution in the development of the software, as well as William Moses for his assistance in setting up Enzyme within `Exasim`. We gratefully acknowledge the United States Department of Energy (under contract DE-NA0003965) and the National Science Foundation (under grant number NSF-PHY-2028125) for supporting this work. Finally, the authors thank the Oak Ridge Leadership Computing Facility for providing access to their GPU clusters.

References

- [1] N. Kroll, C. Hirsch, F. Bassi, C. Johnston, K. Hillewaert (Eds.), *IDIHOM: Industrialization of High-Order Methods - A Top-Down Approach*, Springer International Publishing, 2015.
- [2] N. Kroll, *ADIGMA: A European Project on the Development of Adaptive Higher Order Variational Methods for Aerospace Applications*, Aerospace Sciences Meetings, American Institute of Aeronautics and Astronautics, 2009.

- [3] Z. J. Wang, K. Fidkowski, R. Abgrall, F. Bassi, D. Caraeni, A. Cary, H. Deconinck, R. Hartmann, K. Hillewaert, H. T. Huynh, N. Kroll, G. May, P.-O. Persson, B. van Leer, M. Visbal, High-order CFD methods: current status and perspective, *Int. J. Numer. Methods Fluids* 72 (8) (2013) 811–845.
- [4] J. Slotnick, A. Khodadoust, J. Alonso, D. L. Darmofal, W. Gropp, E. Lurie, D. J. Mavriplis, CFD vision 2030 study : A path to revolutionary computational aerosciences, Tech. Rep. NASA-CR-2014-218178, NASA (2014).
- [5] J. A. Ekaterinaris, High-order accurate, low numerical diffusion methods for aerodynamics, *Prog. Aerosp. Sci.* 41 (3-4) (2005) 192–300.
- [6] B. Cockburn, G. E. Karniadakis, C.-W. Shu, The development of discontinuous Galerkin methods, in: *Discontinuous Galerkin Methods*, Springer-Verlag Berlin Heidelberg, Berlin, Germany, 2000, pp. 3–50.
- [7] T. W. Jan S. Hesthaven, *Nodal Discontinuous Galerkin Methods*, Springer New York, 2010.
- [8] B. Cockburn, S.-Y. Lin, C.-W. Shu, TVB Runge-Kutta local projection discontinuous Galerkin finite element method for conservation laws III: one-dimensional systems, *J. Comput. Phys.* 84 (1) (1989) 90–113.
- [9] B. Cockburn, C.-W. Shu, The Runge–Kutta Discontinuous Galerkin Method for Conservation Laws V, *J. Comput. Phys.* 141 (2) (1998) 199–224.
- [10] F. Bassi, S. Rebay, High-order accurate discontinuous finite element solution of the 2D Euler equations, *J. Comput. Phys.* 138 (2) (1997) 251–285.
- [11] F. Bassi, S. Rebay, A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier-Stokes equations, *J. Comput. Phys.* 131 (2) (1997) 267–279.
- [12] F. Bassi, S. Rebay, Numerical evaluation of two discontinuous Galerkin methods for the compressible Navier-Stokes equations, *Int. J. Numer. Methods Fluids* 40 (1-2) (2002) 197–207.
- [13] R. Hartmann, P. Houston, Adaptive discontinuous Galerkin finite element methods for nonlinear hyperbolic conservation laws, *SIAM Journal on Scientific Computing* 24 (3) (2003) 979–1004.
- [14] A. Balan, M. Woopen, G. May, Hp-adaptivity on anisotropic meshes for hybridized discontinuous Galerkin scheme, AIAA paper (jun 2015).

- [15] G. Giorgiani, S. Fernández-Méndez, A. Huerta, Hybridizable discontinuous Galerkin with degree adaptivity for the incompressible Navier-Stokes equations, *Comput. Fluids* 98 (2014) 196–208.
- [16] A. Cangiani, Z. Dong, E. H. Georgoulis, P. Houston, *hp-Version Discontinuous Galerkin Methods on Polygonal and Polyhedral Meshes*, Springer-Verlag, 2017.
- [17] X. Roca, C. Nguyen, J. Peraire, Scalable parallelization of the hybridized discontinuous Galerkin method for compressible flow, *AIAA Paper* (jun 2013).
- [18] B. Cockburn, J. Gopalakrishnan, A characterization of hybridized mixed methods for second order elliptic problems, *SIAM Journal on Numerical Analysis* 42 (1) (2004) 283–301.
- [19] B. Cockburn, J. Guzmán, S.-C. Soon, H. K. Stolarski, An Analysis of the Embedded Discontinuous Galerkin Method for Second-Order Elliptic Problems, *SIAM Journal on Numerical Analysis* 47 (4) (2009) 2686–2707.
- [20] N. C. Nguyen, J. Peraire, B. Cockburn, A class of embedded discontinuous Galerkin methods for computational fluid dynamics, *Journal of Computational Physics* 302 (2015) 674–692.
- [21] B. Cockburn, C.-W. Shu, The local discontinuous Galerkin method for time-dependent convection-diffusion systems, *SIAM Journal on Numerical Analysis* 35 (6) (1998) 2440–2463.
- [22] B. Cockburn, N. C. Nguyen, J. Peraire, A comparison of HDG methods for Stokes flow, *J. Sci. Comput.* 45 (1-3) (2010) 215–237.
- [23] J. Peraire, P.-O. Persson, The compact discontinuous Galerkin method for elliptic problems, *SIAM Journal on Scientific Computing* 30 (4) (2008) 1806–1824.
- [24] J. Vila-Pérez, M. Giacomini, R. Sevilla, A. Huerta, Hybridisable Discontinuous Galerkin Formulation of Compressible Flows, *Archives of Computational Methods in Engineering* 28 (2) (2020) 753–784.
- [25] R. Sevilla, A. Huerta, Tutorial on Hybridizable Discontinuous Galerkin (HDG) for second-order elliptic problems, in: J. Schröder, P. Wriggers (Eds.), *Advanced Finite Element Technologies*, Vol. 566 of CISM International Centre for Mechanical Sciences, Springer International Publishing, 2016, pp. 105–129.

- [26] C. Ciucă, P. Fernandez, A. Christophe, N. Nguyen, J. Peraire, Implicit hybridized discontinuous Galerkin methods for compressible magnetohydrodynamics, *Journal of Computational Physics: X* 5 (2020) 100042.
- [27] P. Fernandez, A. Christophe, S. Terrana, N. C. Nguyen, J. Peraire, Hybridized Discontinuous Galerkin Methods for Wave Propagation, *Journal of Scientific Computing* 77 (3) (2018) 1566–1604.
- [28] P. Fernandez, N. C. Nguyen, J. Peraire, The hybridized discontinuous Galerkin method for implicit large-eddy simulation of transitional turbulent flows, *Journal of Computational Physics* 336 (2017) 308–329.
- [29] N. C. Nguyen, J. Peraire, Hybridizable discontinuous Galerkin methods for partial differential equations in continuum mechanics, *Journal of Computational Physics* 231 (18) (2012) 5955–5988.
- [30] N. Nguyen, J. Peraire, B. Cockburn, High-order implicit hybridizable discontinuous galerkin methods for acoustics and elastodynamics, *Journal of Computational Physics* 230 (10) (2011) 3695–3718.
- [31] M. Giacomini, A. Karkoulas, R. Sevilla, A. Huerta, A superconvergent HDG method for Stokes flow with strongly enforced symmetry of the stress tensor, *J. Sci. Comput.* 77 (3) (2018) 1679–1702.
- [32] R. Sevilla, M. Giacomini, A. Karkoulas, A. Huerta, A superconvergent hybridizable discontinuous Galerkin method for linear elasticity, *Int. J. Numer. Methods Eng.* 116 (2) (2018) 91–116.
- [33] F. Vidal-Codina, N. Nguyen, S.-H. Oh, J. Peraire, A hybridizable discontinuous galerkin method for computing nonlocal electromagnetic effects in three-dimensional metallic nanostructures, *Journal of Computational Physics* 355 (2018) 548–565.
- [34] F. Vidal-Codina, N.-C. Nguyen, C. Ciraci, S.-H. Oh, J. Peraire, A nested hybridizable discontinuous galerkin method for computing second-harmonic generation in three-dimensional metallic nanostructures, *Journal of Computational Physics* 429 (2021) 110000.
- [35] A. Huerta, A. Angeloski, X. Roca, J. Peraire, Efficiency of high-order elements for continuous and discontinuous Galerkin methods, *Int. J. Numer. Methods Eng.* 96 (9) (2013) 529–560.
- [36] M. Wopen, A. Balan, G. May, J. Schütz, A comparison of hybridized and standard DG methods for target-based hp-adaptive simulation of compressible flow, *Comput. Fluids* 98 (2014) 3 – 16.

- [37] M. Kronbichler, K. Kormann, W. A. Wall, Fast matrix-free evaluation of hybridizable discontinuous Galerkin operators, in: *Lecture Notes in Computational Science and Engineering*, Springer International Publishing, 2019, pp. 581–589.
- [38] S. Terrana, N. C. Nguyen, J. Peraire, GPU-accelerated Large Eddy Simulation of Hypersonic Flows, in: *AIAA Scitech 2020 Forum*, 2020, pp. AIAA–2020–1062.
- [39] B. Cockburn, G. Kanschat, D. Schötzau, A locally conservative {LDG} method for the incompressible {N}avier-{S}tokes equations, *Math. Comp.* 74 (2005) 1067–1095.
- [40] N. C. Nguyen, S. Terrana, J. Peraire, Large-eddy simulation of transonic buffet using matrix-free discontinuous galerkin method, *AIAA Journal* (2022) 1–18.
- [41] D. Arndt, N. Fehn, G. Kanschat, K. Kormann, M. Kronbichler, P. Munch, W. A. Wall, J. Witte, ExaDG: High-order discontinuous Galerkin for the exascale, in: H.-J. Bungartz, S. Reiz, B. Uekermann, P. Neumann, W. E. Nagel (Eds.), *Software for Exascale Computing - SPPEXA 2016-2019*, Springer International Publishing, Cham, 2020, pp. 189–224.
- [42] D. Arndt, W. Bangerth, D. Davydov, T. Heister, L. Heltai, M. Kronbichler, M. Maier, J.-P. Pelteret, B. Turcksin, D. Wells, The deal.II finite element library: Design, features, and insights, *Computers & Mathematics with Applications* 81 (2021) 407–422.
- [43] R. Anderson, J. Andrej, A. Barker, J. Bramwell, J.-S. Camier, J. Cervený, V. Dobrev, Y. Dudouit, A. Fisher, T. Kolev, W. Pazner, M. Stowell, V. Tomov, I. Akkerman, J. Dahm, D. Medina, S. Zampini, MFEM: A modular finite element methods library, *Computers & Mathematics with Applications* 81 (2021) 42–74.
- [44] F. Hindenlang, G. J. Gassner, C. Altmann, A. Beck, M. Staudenmaier, C.-D. Munz, Explicit discontinuous galerkin methods for unsteady problems, *Computers & Fluids* 61 (2012) 86–93.
- [45] C. Cantwell, D. Moxey, A. Comerford, A. Bolis, G. Rocco, G. Mengaldo, D. D. Grazia, S. Yakovlev, J.-E. Lombard, D. Ekelschot, B. Jordi, H. Xu, Y. Mohamied, C. Eskilsson, B. Nelson, P. Vos, C. Biotto, R. Kirby, S. Sherwin, Nektar++: An open-source spectral/hp element framework, *Computer Physics Communications* 192 (2015) 205–219.

- [46] E. J. Ching, B. Bornhoft, A. Lasemi, M. Ihme, Quail: A lightweight open-source discontinuous galerkin code in python for teaching and prototyping, *SoftwareX* 17 (2022) 100982.
- [47] M. Giacomini, R. Sevilla, A. Huerta, HDGLab: An Open-Source Implementation of the Hybridisable Discontinuous Galerkin Method in MATLAB, *Archives of Computational Methods in Engineering* 28 (3) (2020) 1941–1986.
- [48] A. Dedner, R. Klöfkorn, M. Nolte, M. Ohlberger, A generic interface for parallel and adaptive discretization schemes: abstraction principles and the dune-fem module, *Computing* 90 (3-4) (2010) 165–196.
- [49] J. Schöberl, C++11 implementation of finite elements in ngsolve, Tech. rep., Institute for Analysis and Scientific Computing, Vienna University of Technology - TU Wien (2014).
- [50] C. Prud'homme, V. Chabannes, Stephane Veys, A. Ancel, T. Metivet, R. Hild, Jbwahl, C. Daversin-Catty, G. Dollé, Tarabay, Lsala, LANTZT, Doyeux, Trophime, Abdoulaye SAMAKE, B. Vanthong, M. ISMAIL, V. Huber, Kyoshe Winstone, Schenone, T. Saigre, , Philippe, D. Prada, Lberti, Prj-, D. Barbier, Clayrc, , Yacine, J. Veysset, feelpp/feelpp: Feel++ v109 (2021).
- [51] B. Reuter, FESTUNG: A MATLAB / GNU octave toolbox for the discontinuous galerkin method. part IV: Generic problem framework and model-coupling interface, *Communications in Computational Physics* 28 (2) (2020) 827–876.
- [52] A. Klöckner, T. Warburton, J. S. Hesthaven, Solving wave equations on unstructured geometries, in: *GPU Computing Gems Jade Edition*, Elsevier, 2012, pp. 225–242.
- [53] F. Witherden, A. Farrington, P. Vincent, PyFR: An open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach, *Computer Physics Communications* 185 (11) (2014) 3028–3040.
- [54] R. Alexander, Diagonally implicit Runge-Kutta methods for stiff ODEs, *SIAM J. Numer. Anal.* 14 (1977) 1006–1021.
- [55] C. Geuzaine, J.-F. Remacle, Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities, *International Journal for Numerical Methods in Engineering* 79 (11) (2009) 1309–1331.

- [56] M. Skroch, S. J. Owen, M. L. Staten, R. W. Quadros, B. Hanks, B. Clark, T. Hensley, C. Ernst, R. Morris, C. McBride, C. Stimpson, J. Perry, M. Richardson, K. Merkley, CUBIT: Geometry and Mesh Generation Toolkit. 16.02 User Documentation, Tech. rep., Sandia National Laboratories (2022).
- [57] The CGAL Project, CGAL User and Reference Manual, 5.4 Edition, CGAL Editorial Board, 2022.
- [58] P.-O. Persson, G. Strang, A simple mesh generator in MATLAB, *SIAM Review* 46 (2) (2004) 329–345.
- [59] H. Si, TetGen, a Delaunay-Based Quality Tetrahedral Mesh Generator, *ACM Transactions on Mathematical Software* 41 (2) (2015) 1–36.
- [60] C. Dapogny, C. Dobrzynski, P. Frey, Three-dimensional adaptive domain remeshing, implicit domain meshing, and applications to free and moving boundary problems, *Journal of Computational Physics* 262 (2014) 358–378.
- [61] P. Cignoni, M. Callieri, M. Corsini, M. Dellepiane, F. Ganovelli, G. Ranzuglia, Meshlab: an open-source mesh processing tool (2008).
- [62] G. Nicolas, T. Fouquet, Adaptive mesh refinement for conformal hexahedral meshes, *Finite Elements in Analysis and Design* 67 (2013) 1–12.
- [63] P.-O. Persson, J. Peraire, Curved mesh generation and mesh refinement using Lagrangian solid mechanics, in: 47th AIAA Aerospace Sciences Meeting, no. 949, 2009.
- [64] A. Gargallo-Peiró, X. Roca, J. Peraire, J. Sarrate, Optimization of a regularized distortion measure to generate curved high-order unstructured tetrahedral meshes, *International Journal for Numerical Methods in Engineering* 103 (5) (2015) 342–363.
- [65] W. Moses, V. Churavy, Instead of rewriting foreign code for machine learning, automatically synthesize fast gradients, in: H. Laroche, M. Ranzato, R. Hadsell, M. F. Balcan, H. Lin (Eds.), *Advances in Neural Information Processing Systems*, Vol. 33, Curran Associates, Inc., 2020, pp. 12472–12485.
- [66] W. S. Moses, V. Churavy, L. Paehler, J. Hüchelheim, S. H. K. Narayanan, M. Schanen, J. Doerfert, Reverse-mode automatic differentiation and optimization of GPU kernels via Enzyme, in: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ACM, 2021.

- [67] C. Hader, H. F. Fasel, Direct numerical simulations of hypersonic boundary-layer transition for a flared cone: Fundamental breakdown, *Journal of Fluid Mechanics* 869 (2019) 341–384.
- [68] B. C. Chynoweth, S. P. Schneider, C. Hader, H. Fasel, A. Batista, J. Kuehl, T. J. Juliano, B. M. Wheaton, History and progress of boundary-layer transition on a Mach-6 flared cone, *Journal of Spacecraft and Rockets* 56 (2) (2019) 333–346.
- [69] K. J. Baumeister, K. L. Kreider, Scattering cross section of sound waves by the modal element method, in: *Winter Annual Meeting by the American Society of Mechanical Engineers*, 1994.
- [70] F. J. Poulin, G. R. Flierl, The nonlinear evolution of barotropically unstable jets, *Journal of Physical Oceanography* 33 (10) (2003) 2173–2192.
- [71] W. M. van Rees, A. Leonard, D. Pullin, P. Koumoutsakos, A comparison of vortex and pseudo-spectral methods for the simulation of periodic vortical flows at high Reynolds numbers, *Journal of Computational Physics* 230 (8) (2011) 2794–2805.