

1 **A Performance Analysis of a Hybridized SBP-SAT**
2 **Finite-Difference Method for Large Scale Earth Science**
3 **Applications**

4 Alexandre Chen · Brittany A. Erickson · Jeremy
5 E. Kozdon

6
7 Received: date / Accepted: date

8 **Abstract** We present performance results from a new hybridized finite difference
9 method for the spatial discretization of partial differential equations. The method
10 is based on the standard Summation-By-Parts method with weak enforcement
11 of boundary and interface conditions through the Simultaneous-Approximation-
12 Term. We analyze the performance when applying the hybrid method to Poisson's
13 equation which arises in many steady-state physical problems, focusing on an
14 application in Earth science. When solving the resulting linear system we compare
15 direct and iterative solvers on both CPU and GPU, evaluating the performance on
16 meshes with different numbers of computational blocks. Our results demonstrate
17 the advantages of using the hybrid method in solving large-scale problems under
18 the restriction of system resources by utilizing techniques from parallel computing.

19 **1 Introduction and Background**

20 **1.1 Poisson's Equation and an Application in Earth Science**

21 Poisson's equation is a partial differential equation (PDE) of elliptic type that is
22 widely used in physics, fluid dynamics, mechanical engineering, and other fields to
23 study steady-state problems. The equation is given by

$$\nabla^2 \varphi = f, \tag{1}$$

Alexandre Chen and B. A. Erickson
Computer and Information Science
1202 University of Oregon
1477 E. 13th Ave.
Eugene, OR 97403-1202
E-mail: {chern,bae}@cs.uoregon.edu

J. E. Kozdon
Department of Applied Mathematics,
Naval Postgraduate School,
833 Dyer Road,
Monterey, CA 93943-5216
E-mail: jekozdon@nps.edu




Fig. 1: A 2D simplified model in an earthquake simulation. Figure from Erickson and Dunham (2014)

24 where f and ϕ are real or complex-valued functions on a Euclidean space, with f a
 25 given source function and φ is sought. When $f = 0$, we obtain Laplace's Equation.

26 In computational seismology, Poisson's equation arises when describing the
 27 2D antiplane problem as shown in figure 1, where the out-of-plane displacement
 28 $u = \phi$ is sought (Erickson and Dunham, 2014). As illustrated in the figure, a
 29 section in the $-yz$ plane is considered, containing symmetry with respect to the z
 30 axis. An 1D earthquake fault (an interface) is located along z -axis and is subject
 31 to a specified friction law. Tectonic motion is captured by setting the remote
 32 boundaries to be displaced at a slow plate rate of $\approx 32\text{mm/yr}$, which is enforced
 33 by applying Dirichlet boundary conditions. Although the boundary conditions are
 34 changing through the time at an extremely low rate, we have assumed a quasi-
 35 static response. A sequence of earthquakes nucleate at the fault in response to the
 36 remote tectonic loading. In this earthquake cycle simulation, the fault is a thin
 37 zone of crushed rock separating blocks of the Earth's crust. When an earthquake
 38 occurs on the fault, the rock on one side of the fault is displaced with respect to
 39 the other side, and this jump in displacement across the fault is known as slip. The
 40 fault length can be of several hundreds of kilometers, with frictional properties on
 41 the order of microns, which gives rise to large problems in simulation. Earth's free
 42 surface is at $z = 0$ and we also assume a free surface at depth, corresponding to
 43 Neumann boundary conditions. Due to the symmetry of the system, the problem
 44 can be further simplified by considering only one side of the fault. Once we have the
 45 numerical solution for one side of the fault, the other side can be easily obtained
 46 from symmetry properties.

47 The assumption of steady-state motion in the anti-plane scenario gives rise
 48 to the following anisotropic version of Poisson's equation in a two dimensional
 49 domain Ω :

$$-\nabla \cdot (\mathbf{b} \nabla \mathbf{u}) = \mathbf{f}, \text{ on } \Omega \quad (2a)$$

$$u = g_D, \text{ on } \partial\Omega_D \quad (2b)$$

$$\mathbf{n} \cdot \mathbf{b} \nabla \mathbf{u} = g_N, \text{ on } \partial\Omega_N \quad (2c)$$

$$\begin{cases} \{\{\mathbf{n} \cdot \mathbf{b} \nabla u\}\} = 0 \\ \llbracket u \rrbracket = \delta \end{cases} \text{ on } \partial\Gamma_I, \quad (2d)$$

50 where the field u is the material displacement. Here, $\mathbf{b}(x, y)$ is a matrix valued
 51 function that is symmetric positive definite and the scalar function $\mathbf{f}(x, y)$ is the

52 source function. The boundary conditions of the domain have been partitioned into
 53 a Dirichlet one and a Neumann one, namely, $\partial\Omega = \partial\Omega_D \cup \partial\Omega_N$ and $\partial\Omega_D \cap \partial\Omega_N = \emptyset$.
 54 At the Neumann boundaries, the vector \mathbf{n} is the outward pointing normal. We have
 55 also introduced an internal interface Γ_I , and along this interface, the \mathbf{b} -weighted
 56 normal derivative is taken to be continuous, with jumps allowed in the scalar field.
 57 Jumps allow us to apply this method to physical problems where displacements
 58 occur across an interface, e.g. earthquakes that occur along a fault which forms an
 59 interface in the solid Earth. $\{\{w\}\} = w^+ + w^-$ here denotes the sum of the scalar
 60 quantity on both sides of the interface and $[[w]] = w^+ - w^-$ is the difference across
 61 the interface. For system with jumps across the interface, we define a non-zero
 62 constant δ vector to depict this discontinuity. When there is no jump across the
 63 interface, δ is set to be zero.

64 1.2 Discretization with the SBP-SAT Scheme and Numerical Solution Methods

65 We mesh our domain with rectilinear grids in order to apply finite difference meth-
 66 ods. For real geographical domains where rectilinear meshing can not be applied,
 67 we can use coordinate transformations to transform the physical domain into a
 68 logical rectangular domain where we can apply rectilinear meshing. The solutions
 69 obtained from the logical domain can later be transformed back into solutions for
 70 the real domain using inverse transformation. For our study here, we consider the
 71 unit square for simplicity without losing generality. The detailed technique on co-
 72 ordinate transformation can be found in Kozdon, Erickson, et al. (2020) and won't
 73 be covered here.

74 Summation-By-Parts (SBP) finite difference methods have been proposed to
 75 solve problems with complex geometries such as the problem in this paper due
 76 to their desirable properties of high order accuracy and provable stability (Kreiss
 77 and Scherer, 1974; Kreiss and Scherer, 1977; Strand, 1994; Mattsson and Nord-
 78 ström, 2004). The inter-block coupling conditions can be enforced weakly using
 79 the Simultaneous-Approximation Term (SAT) method (Carpenter, Gottlieb, et al.,
 80 1994; Carpenter, Nordström, et al., 1999). The SAT term here is analogous to the
 81 penalty term in discrete Galerkin methods. More details on SBP-SAT method will
 82 be covered in 2.

83 In earthquake cycle simulations, earthquake nucleation and rupture propa-
 84 gation is simulated over thousands of years, where quasi-steady state problems
 85 are formed to depict slow and quiescent periods between earthquakes (Erickson
 86 and Dunham, 2014). In the steady-state regime, we need to solve elliptic partial
 87 differential equations, which will result in large linear systems of equations for
 88 realistically complex problems under the constraint of stability requirement for
 89 time stepping methods. In order to obtain stable solutions over long time scales,
 90 we need to apply a fine mesh in the spatial domain, and this is where large scale
 91 linear problems arise. This project is constructed around one key challenge: How
 92 can we obtain numerical results for a large linear system formed by SBP-SAT
 93 operators in order to study earthquake cycle simulations over long time scales?

94 In the terminology of algorithm, the time complexity refers to the number
 95 of steps required for an algorithm, and the space complexity refers to the total
 96 space taken by the algorithm to store the input data and intermediate results with
 97 respect to the input size. Solving a linear system of size $n \times n$ directly with an

98 LU decomposition or other factorization-based direct method is known to have
99 the time complexity of $\mathcal{O}(n^3)$ and space complexity of $\mathcal{O}(n^2)$. The computational
100 complexity mainly comes from the process of factorization which has $\mathcal{O}(n^3)$. After
101 obtaining the factorization, using forward/backward substitution to solve the sys-
102 tem has the time complexity of $\mathcal{O}(n^2)$. In time stepping methods, we can reuse the
103 factorization results during each linear solve, but the space complexity of factor-
104 ization has limited the size of the problem that system is capable of doing. Another
105 common issue is that the matrices that we form using the SBP-SAT methods can
106 be expressed in a sparse matrix fashion which reduces the cost of matrix storage
107 and operations. But this sparsity can be destroyed during the LU factorization,
108 and it relies on certain factorization algorithms that can preserve sparsity to obtain
109 optimal results.

110 One way to get around the size restriction imposed by the space complexity of
111 direct solvers is using iterative methods. Iterative methods convert the problem of
112 solving a linear system into a problem in optimization. A common example is the
113 Conjugate Gradient (CG) method which is a traditional iterative solver but more
114 known to computer scientists in recent years because it has been widely used in
115 machine learning to minimize the loss function defined with information entropy.
116 CG is particularly suitable for solving a linear system that has a positive defi-
117 nite (PD) left-hand side. For information on other iterative solvers, refer to Saad
118 (2003). Iterative solvers avoid the challenge of obtaining a factorization for sparse
119 matrices by using repeated matrix-vector products to obtain an approximate nu-
120 merical result, approaching the exact solution with proven asymptotic accuracy.
121 Given that in numerical methods, the accuracy is limited by the round-off error,
122 iterative solvers can provide a numerical result that matches the accuracy of the
123 result obtained from a direct solver. Iterative solvers have other desirable proper-
124 ties in that we can gain efficiency by loosening the accuracy constraint, obtaining
125 a less accurate result that is sufficient for the study in mind. The linear algebra
126 operations such as matrix-vector multiplications can be easily accelerated by pack-
127 ages such as Basic Linear Algebra Subprograms (BLAS) on both CPU and GPU
128 architecture. In 3, we will study different iterative implementations to compare
129 the accuracy, stability and performance for our problem in search for an optimal
130 iterative method for this specific problems.

131 Another way to bypass the limitation of using direct solvers on a large single
132 system is to use an new hybridization technique that was proposed for SBP-SAT
133 methods (Kozdon, Erickson, et al., 2020). This method reduces the system size
134 by writing the numerical method in a way that leverages the Schur complement
135 and eliminates degrees of freedom from within the element, leaving only degrees
136 of freedom on element boundaries. We set the values on all the interfaces to be
137 given input data known as the trace variables. These independent trace variables
138 along the faces of the blocks are introduced so the inter-block coupling penalty
139 terms can be expressed merely as a function of the trace variables. Hence the so-
140 lution in each block is uniquely determined by these trace variables. In this hybrid
141 method, the problem is broken into two pieces, a *local problem* and a *global problem*.
142 *local problem* refers to the solution within the block given trace data, and the
143 *global problem* refers to the value of the trace variable, given the block data. The
144 *local problem* and the *global problem* are connected via a Schur complement. This
145 is an extension of existing SBP-SAT scheme by introducing trace variables so we

146 can work on a domain with multiple blocks that share interfaces. Details will be
147 covered in section 4 when we discuss performance of the hybrid method.

148 1.3 Implementation

149 In this paper we implement algorithms solving Poisson’s equation with SBP-SAT
150 method and extended hybrid SBP-SAT method in the Julia programming language
151 with support of various open packages from linear algebra to GPU computing. Ju-
152 lia is a new programming language with an emphasis on scientific computing. It
153 is designed to solve the Two-Language-Problem that many researchers encounter
154 when developing a prototype in a high-level language such as Python or MATLAB
155 (for efficiency in development) and then implementing the code in a low-level lan-
156 guage such as C++ and FORTRAN for performance. Julia is a compiled language
157 leveraging the JIT compilation for performance. It supports dynamic notation
158 with multiple dispatch, which gives high code readability during development and
159 high performance in code execution. At the compilation level, Julia uses LLVM,
160 which generates an LLVM intermediate representation that can be used to work
161 with other languages/frameworks that are also using LLVM. The extensibility
162 from using LLVM as a compiler has been demonstrated with Julia’s capability of
163 leveraging GPU power for HPC (Besard et al., 2019).

164 Other nice features of Julia include metaprogramming from the legacy of the
165 LISP language. Metaprogramming allows us to write less code by reducing repeti-
166 tion. Unlike other languages or frameworks that are accessed in one programming
167 language but written in another programming language, many Julia packages are
168 written in Julia itself with core source code open-sourced. This makes cooperation
169 in Julia much more handy, and for this reason, the Julia language has become one
170 of the most fast-growing languages with professional active users in computational
171 science who help form a booming Julia ecosystem. However, being a relatively new
172 open language also means there is lack of official support when it comes to bugs.
173 Our implementations were limited by the compatibility issues of different pack-
174 ages that have to wait for the update from independent developers who wrote
175 these packages.

176 This rest of the paper is organized as follows: In Section 2, we provide a de-
177 tailed description of the method of block decomposition and forming of SBP-SAT
178 operators. In Section 3, we give performance evaluation of this problem on a single
179 domain with different implementations. Namely, we confirm convergence results of
180 SBP-SAT discretization with different orders of accuracy. We test time and space
181 resource consumption with different implementations. In Section 4, we cover the
182 key ideas of the hybrid method introduced in the previous section, with extensive
183 study on the performance of this novel method. We illustrate the promising aspects
184 of this method and describe several existing issues of the current implementation.
185 Solving these issues in the future implementation would be essential to further
186 leveraging the benefits of the hybrid SBP-SAT scheme.

2 SBP-SAT introduction

2.1 One Dimensional SBP Operators

We discretize the domain $0 \leq x \leq 1$ with $N + 1$ evenly spaced grid points $x_i = ih, i = 0, \dots, N$ with spacing $h = 1/N$. We then project a function u onto the computational grid to be $u = [u_0, u_1, \dots, u_N]^T$. u is often taken to be the interpolant of u at grid points. We define grid basis vector e_j to be a vector with value 1 at grid point j and 0 for the rest. We only need e_0 and e_N to form projections at boundaries. Note that in general we have $u_j = e_j^T u$.

We apply the class of high-order accurate SBP finite difference methods for first order derivatives which were introduced in Kreiss and Scherer (1974) and Kreiss and Scherer (1977) and Strand (1994) as mentioned above. For second order derivatives, we apply Mattsson and Nordström (2004), with variable coefficients treated in Mattsson (2012). The exact form of definitions are given below.

Definition 1 (First Derivative) We define matrix D_x to be an SBP approximation to $\partial u / \partial x$ if it can be decomposed as $\mathbf{H}D_x = \mathbf{Q}$ with \mathbf{H} being symmetric positive definite and \mathbf{Q} satisfying $\mathbf{u}^T(\mathbf{Q} + \mathbf{Q}^T)\mathbf{v} = u_N v_N - u_0 v_0$.

Here, we only consider diagonal-norm SBP, i.e. finite difference operators where \mathbf{H} is a diagonal matrix and D_x is the standard central finite difference matrix in the interior which transitions to one-sided at boundaries. The condition of \mathbf{Q} defined above can be written as $\mathbf{Q} + \mathbf{Q}^T = e_N e_N^T - e_0 e_0^T$.

The reason why the operator D_x is called SBP because it mimics the integration-by-part property

$$\int_0^1 u \frac{\partial v}{\partial x} + \int_0^1 \frac{\partial u}{\partial x} v = uv \Big|_0^1, \quad (3)$$

in a discrete form

$$\mathbf{u}^T \mathbf{H} D_x \mathbf{v} + \mathbf{u}^T D_x^T \mathbf{H} \mathbf{v} = \mathbf{u}^T (\mathbf{Q} + \mathbf{Q}^T) \mathbf{v} = u_N v_N - u_0 v_0. \quad (4)$$

Following the same pattern of the first derivative, we can define the second derivative.

Definition 2 (Second Derivative) We define matrix $D_{xx}^{(c)}$ to be an SBP approximation to $\frac{\partial}{\partial x} \left(c \frac{\partial u}{\partial x} \right)$ if it can be decomposed as $\mathbf{H}D_{xx}^{(c)} = -\mathbf{A}^{(c)} + c_N e_N \mathbf{d}_N^T - c_0 e_0 \mathbf{d}_0^T$ where $\mathbf{A}^{(c)}$ is symmetric positive definite and $\mathbf{d}_0^T \mathbf{u}$ and $\mathbf{d}_N^T \mathbf{u}$ are approximations of the first derivative of u at the boundaries.

Similarly, the operator $D_{xx}^{(c)}$ mimics the integration-by-parts property

$$\int_0^1 u \frac{\partial}{\partial x} \left(c \frac{\partial v}{\partial x} \right) + \int_0^1 \frac{\partial u}{\partial x} c \frac{\partial v}{\partial x} = uc \frac{\partial v}{\partial x} \Big|_0^1, \quad (5)$$

in a discrete form

$$\mathbf{u}^T \mathbf{H} D_{xx}^{(c)} \mathbf{v} + \mathbf{u}^T \mathbf{A}^{(c)} \mathbf{v} = c_N u_N \mathbf{d}_N^T \mathbf{v} - c_0 u_0 \mathbf{d}_0^T \mathbf{v}. \quad (6)$$

As noted above, we only consider diagonal-norm SBP finite difference operators here. In the interior, the operators use the minimal bandwidth central difference

220 stencil and transition to one-sided at boundaries in a manner that preserves the
 221 SBP property.

222 It has been known that for SBP operators defined above, if the interior operator
 223 has accuracy $2p$, then the interior stencil bandwidth is $2p+1$ and the boundary
 224 operator has accuracy p . If we use operators with interior accuracy $2p = 2, 4,$ and
 225 6 , the expected global order of accuracy is the minimal of $2p$ and $p+2$ as evidenced
 226 by empirical study (Mattsson, Ham, et al. (2009) Virta and Mattsson (2014)) and
 227 proved for the Schrödinger equation (Nissen et al., 2013). We will verify the result
 228 in later sections.

229 2.2 Two Dimensional SBP Operators

230 Two-dimensional SBP operators can be developed by applying the one-dimensional
 231 SBP operators in a tensor product fashion. Here we describe the operators for a
 232 rectangular block $\hat{B} \in [0, 1] \times [0, 1]$. We discretize this domain similar to 1d case
 233 in each direction resulting in an $(N + 1) \times (N + 1)$ grid of points where grid point
 234 (i, j) is at $(r_i, s_j) = (ih, jh)$ for $0 \leq i, j \leq N$ with $h = 1/N$; For simplicity, we only
 235 consider the case where we have the same numbers of grid points in each direction.
 236 A more complex scenario where we have different numbers of grid points in each
 237 direction can be formed similarly, but we are not going to discuss in detail here.

238 The two dimensional SBP operators can be obtained from one dimensional SBP
 239 operators by taking Kronecker products with them in orders that are determined
 240 by directions in two dimensional space. The detailed technique can be found in
 241 Kozdon, Erickson, et al. (2020) and won't be repeated here. We should note that
 242 tensor products are used here mainly for the purpose of simplicity in theoretical
 243 analysis. In computer memory, data are stored in a one dimensional array. Hence,
 244 the Kronecker products here mainly affects the order where we read data from a
 245 one dimensional array.

246 2.3 SAT Terms

247 SAT terms weakly enforce boundary conditions penalizing the grid point at the
 248 boundary towards the boundary data. It has the following simplified form:

$$\mathbf{b} = \alpha * (\mu \mathbf{B} \mathbf{u} - \mathbf{g}). \quad (7)$$

249 Here, \mathbf{u} is the grid vector (the numerical approximation to the solution), \mathbf{g} is
 250 boundary condition for a particular interfaces. \mathbf{B} is an SBP operator that extracts
 251 boundary data from \mathbf{u} and it would contain information about boundary layouts
 252 and associated conditions. μ is the block-diagonal matrix associated with \mathbf{B} that
 253 needs to be compatible with boundary layouts. α is the penalty parameter in SAT
 254 term that is chosen under stability constraints from energy estimate. Finally, \mathbf{b}
 255 is the assembled vector that weakly enforces a certain boundary condition in fi-
 256 nite difference methods. More detailed examples of SAT terms in practice can be
 257 found in Erickson and Dunham (2014). Boundary conditions can be assembled by
 258 gradually adding \mathbf{b} terms to the RHS of the equation in a simple additive way.
 259 Compared to the traditional method of using injection or strong enforcement of
 260 boundary/interface conditions that would destroy the SBP property defined in

261 equations 4 and 6 in section 2, using SAT terms preserves strict stability mean-
 262 ing that the semi-discrete problem has the same asymptotic time-growth as the
 263 continuous problem (Mattsson, 2003).

264 The combined SBP-SAT approach has been extensively used in computational
 265 science for solving problems from natural sciences where physical interfaces are
 266 ubiquitous. In geophysics particularly, it can be used to solve earthquake problems
 267 where continental and oceanic crustal blocks are separated by faults or in multi-
 268 phase fluids with discontinuous properties (Kozdon, Dunham, et al., 2012; Erickson
 269 and Day, 2016; Karlstrom and Dunham, 2016; Lotto and Dunham, 2015).

270 2.4 Implementation in Julia

271 In the Julia implementation, we use SparseArrays.jl to form sparse matrices. We
 272 use LinearAlgebra.jl for Kronecker products and other linear algebra operations.
 273 We use CUDA.jl and CuArrays.jl for computations on CUDA supported
 274 GPUs. For iterative solvers, we use IterativeSolvers.jl as well as our own matrix-
 275 free version of the CG algorithm. This project has been completely done in Julia
 276 except the meshing part where we use the external meshing software Trelis ([https://](https://csimsoft.com/trelis)
 277 csimsoft.com/trelis).

278 3 Performance Study on A Single Domain

279 3.1 Comparison Between Direct Solver and Iterative Solver

280 We first study the performance of the hybrid SBP-SAT method on a single domain.
 281 We chose a unit square, where we have Dirichlet boundary conditions on the
 282 left and the right, and Neumann boundary conditions on the top and bottom.
 283 To study the accuracy and convergence of the method, we apply the method
 284 of manufactured solutions (MMS), see Roache (1998) for example. In the MMS
 285 technique, an analytic solution is assumed from which we can derive compatible
 286 boundary and source data. In our test, we manufactured a solution to have the
 287 following form:

$$u(x, y) = \sin(\pi x + \pi y), 0 \leq x \leq 1 \quad (8)$$

288 where x and y denote the x and y coordinates of a given point. From this fabricated
 289 solution, we can derive the conditions on the boundary and source function in
 290 interior, namely

$$\left\{ \begin{array}{ll} u_{xx} + u_{yy} + 2\pi^2 \sin(\pi x + \pi y) = 0, & 0 \leq x \leq 1 \quad (9a) \\ u = \sin(\pi y), & x = 0 \quad (9b) \\ u = -\sin(\pi y), & x = 1 \quad (9c) \\ -u_y = \cos(\pi x), & y = 0 \quad (9d) \\ u_y = -\cos(\pi x), & y = 1. \quad (9e) \end{array} \right.$$

291 Here u_y represents the first derivative of u with respect to y and u_{xx} represents
 292 the second derivative of u with respect to x . We don't have cross derivative terms
 293

in our simplified problem. The minus sign from the equation 9d above comes in because the normal vector at $y = 0$ points downward. We define the error $_h = \sqrt{(u_h - u)^T H (u_h - u)}$. Here u is the exact solution in equation 8 evaluated on the numerical grid, and u_h stands for numerical results from solving the system defined by and boundary conditions and source functions obtained from fabricated results u with governing equations in 9. u_h and u are stacked to be one-dimensional vectors. H is SBP operator that incorporates grid space information and is the Kronecker product of H_x and H_y defined in the 2D SBP operators. This definition of error $_h$ is the discrete-L2 error and is used for convergence tests.

We begin convergence tests using a direct solver on CPU. For different p values, we obtain the following accuracy results for convergence tests in table 1. We obtained expected convergence results for the 2nd order and 4th order SBP operators. For 6th order SBP operators, the convergence rates are close to 5.5. The reason why in higher order operators we don't observe convergence rate as the order of SBP operators is because the order of the accuracy is lower on boundaries as described in 2. Also we are reaching machine precision with $p = 6$, so the rate of convergence is also affected by this factor.

N	2nd Order		4th Order		6th Order	
	error $_N$	rate	error $_N$	rate	error $_N$	rate
2^4	1.735×10^{-3}		4.227×10^{-5}		1.139×10^{-5}	
2^5	4.319×10^{-4}	2.0013	2.117×10^{-6}	4.320	2.605×10^{-7}	5.451
2^6	1.079×10^{-4}	2.0003	1.095×10^{-7}	4.273	5.847×10^{-9}	5.477
2^7	2.696×10^{-5}	2.00007	5.956×10^{-9}	4.200	1.301×10^{-10}	5.489
2^8	6.740×10^{-6}	2.000017	3.401×10^{-10}	4.130	2.896×10^{-12}	5.489

Table 1: Error and convergence rates using the method of manufactured solutions.

Convergence results above have verified the correctness of our implementations. We can also verify this with numerical results from the iterative solvers with similar outcomes.

Although our ultimate goal is to solve this problem on a very large system, once convergence is verified, our next question is how can we solve this problem more efficiently while maintaining correct results. We fix $p = 2$ to reduce the number of variables in our study. Our linear system has a positive semi-definite (PD) left-hand-side (LHS). It is easy to verify that the CG method outperforms other iterative solvers that are designed to handle non-PSD cases such as MINRES (for indefinite matrices) or GMRES (for non-symmetric matrices when good preconditioning is available). We now compare the performance of a direct solver with the CG method on both CPU and GPU. The accuracy results are shown in table 2 to demonstrate all three methods succeed in producing correct results. We should note that by default GPU works with Float32 which normally has significantly higher peak FLOPS than Float64. For convergence and accuracy comparisons however, we chose Float64 on GPU to compare with Float64 on CPU32. Float32 on GPU still yields rather high accuracy (up to 2^{-9}) which can be sufficient enough (depending on our accuracy requirements) while obtaining optimal performance. We evaluate the performance here according to how long it takes to solve the lin-

Grid Amounts	Direct Solver	GPU Iterative	CPU Iterative
N	$\text{Log}_2(\text{error}_N)$	$\text{Log}_2(\text{error}_N)$	$\text{Log}_2(\text{error}_N)$
2^3	-7.140701	-7.140701	-7.140701
2^4	-9.170836	-9.170835	-9.170836
2^5	-11.17709	-11.177089	-11.177089
2^6	-13.178418	-13.178418	-13.178461
2^7	-15.178715	-15.178781	-15.178781

Table 2: Log Errors (base 2) Comparison of Direct Solver and Iterative Solver

ear system and how much memory is allocated. In order to achieve reproducible results, we use the zero vector as the initial guess for the CG method. In practice, using a random initialization normally gives 10x speed-up on our problem. We use the BenchmarkTools.jl package for performance evaluation.

N	Direct Solver	GPU Iterative	CPU Iterative	Direct Solver After LU
2^3	91.387 μs	143.526 μs	1.112 μs	3.082 μs
2^4	313.194 μs	158.186 μs	4.249 μs	12.953 μs
2^5	1.180 ms	165.876 μs	10.934 μs	58.580 μs
2^6	5.799 ms	599.566 μs	37.629 μs	271.836 μs
2^7	32.527 ms	12.113 ms	144.675 μs	1.308 ms

Table 3: Time Comparison of Direct Solver and Iterative Solver

N	Direct Solver	GPU Iterative	CPU Iterative	Direct Solver After LU
2^3	68.99 KiB	5.64 KiB	2.97 KiB	1.47 KiB
2^4	250.24 KiB	5.67 KiB	8.13 KiB	4.91 KiB
2^5	988.21 KiB	5.75 KiB	26.91 KiB	17.41 KiB
2^6	4.03 MiB	5.75 KiB	100.27 KiB	66.31 KiB
2^7	22.27 MiB	2.93 MiB	651.58 KiB	260.31 KiB

Table 4: Memory Comparison of Direct Solver and Iterative Solver

After we verified the correctness, we compare the performance of the direct and iterative methods in terms of time and memory. For iterative methods, we use built-in CG method in IterativeSolvers package. For hyperparameters in CG, we use default ones. The maximum number of iterations is chosen to be the size of the linear system and the tolerance is set to be the square root of machine epsilon for a given floating point type. The results are given in table 3 and 4 respectively. Within the context of a time stepping method, the cost of the LU factorization is overhead cost; once it's obtained, it can be reused in direct solver. Therefore, to compare the actual performance of the iterative solver with direct solver, excluding the cost of LU factorization is an important consideration when working with a time-dependent problem. However we should note that the size of the system that we can solve is limited by the memory cost from LU decomposition. As we can see from the table 4 and 3, for our system, even with factorization cost removed,

347 the iterative solver on CPU still out-performs the direct solver both in time and
348 memory for the same problem. This nice speedup from using the CG iterative
349 solver is most likely coming from the fact that our matrix is PD. Another nice
350 property of our system is the narrow-band sparsity, which requires significantly
351 less computational cost with iterative solvers compared to having factorization
352 that cannot preserve sparsity. Both iterative solvers on CPU and GPU would use
353 parallel framework such as BLAS and CUDA to accelerate, which is less practical in
354 factorization. All these three contribute to the fact that even though factorization
355 results can be reused, solving this system with direct method is still much more
356 expensive than using iterative solvers.

357 3.2 Further Speedup of Iterative Solvers

358 One nice feature of using an iterative solver is the access to matrix-free methods
359 where instead of forming matrix A , we compute the product Ax . We do this by
360 writing a function $f(x)$ that takes vector x as input and directly modifies the entries
361 of x to achieve the same results of Ax . A naive example would be when A is identity
362 matrix, we can use $f(x) = x$ that export x itself as output which has the same
363 effect as multiplying x by identity matrix A in Ax . A matrix-free implementation is
364 extremely handy when the matrix A is sparse and is close to diagonal or tridiagonal
365 with very low bandwidth such as we have in the hybridized SBP-SAT method.
366 Using matrix free methods reduces the memory allocation from forming matrix A .
367 And in-situ operations on x are faster in some cases by reducing the cost of forming
368 intermediate results and storing them in matrix multiplications. One thing we need
369 to be careful about is that Julia's hidden pointer mechanism makes it prone to
370 data contamination when we tried to use some of Julia's notations for fast I/O.
371 The solution is similar to writing Julia in C++ where you need to determine
372 the data containers needed in advance and allocate them in memory, then reuse
373 these containers in matrix-free functions. Our implementation showed that matrix-
374 free functions achieve speedup of several times to ten times compared to sparse-
375 matrix counterparts with zero garbage collector (GC) times. The zero overhead
376 is extremely important because for large matrix operations, GC would determine
377 the maximum time to finish a task, which would cause the volatility in run-time
378 behavior depending on system load. For ideal parallelization, we would expect less
379 volatility in order to have nice static load balancing in designing parallel scheme
380 with optimal performance. Our matrix-free method can be easily parallelized using
381 built-in multi-threading macros in Julia.

382 To demonstrate the speed-up of multi-threading in Julia, we compared two
383 versions of the same matrix-free functions, one with multi-threading and one in
384 serial. The only difference is that in the multi-threaded matrix-free function, we
385 add the `@threads` macro before the for loops. We tested speedup from multi-
386 threading with respect to different system sizes. In our local environment, we set
387 the number of threads in Julia to be 4. The results are shown in the table 5. Multi-
388 threading has fixed overhead which makes it more expensive when our system is
389 small, but as the size of the problem increases, this overhead is negligible and we
390 can see significant speed-up from even naive implementation using the built-in
391 macro. To further explore the speedup we can achieve by throwing more threads,
392 we tested the case where $N = 10000$ on the Talapas server (<https://datascience>).

393 uoregon.edu/talapas-supercomputer) . Using 32 threads, we achieve 10x speedup.
 394 As the system size increases, We are expected to see more speedup from the multi-
 395 threading technique.

N	Serial		Multithreading	
	Time	Memory	Time	Memory
100	56.014 μs	79.83 KiB	119.677 μs	87.33 KiB
1000	2.965 ms	7.64 MiB	2.173 ms	7.65 MiB
10000	809.077 ms	763.09 MiB	423.203 ms	763.10 MiB

Table 5: Performance of Multi-threading in Matrix-free Operators

396 The speedup from matrix-free operators themselves associated with feasible
 397 multi-threading in matrix-free methods make it more appealing when using iterative
 398 solvers compared to sparse-matrix formulations. However, one big challenge is
 399 to optimize the memory allocation and multi-threading in a compound functions
 400 that calls external functions. Our matrix-free CG method out-performs existing
 401 sparse-matrix CG implementations, but the speed-up is not ideal to what we
 402 expected from the speedup from consisting functions within our CG. Further op-
 403 timization would be needed to further leveraging the power of matrix-free method
 404 and parallel processing including multi-threading.

405 4 Hybridized SBP Scheme

406 One of the main goals of this work is to explore performance gains using the newly
 407 proposed hybridized SBP-SAT scheme from Kozdon, Erickson, et al. (2020). In
 408 the finite element literature, a hybrid method is the method where one unknown
 409 is a function on the interior of the elements and the unknown is function on the
 410 trace of the elements (Ciarlet, 2002, page 421). For SBP methods particularly, we
 411 write the method in terms of local problems and associated global problem. In the
 412 local problems, for each block $B \in \mathcal{B}$, a grid of blocks \mathcal{B} , the trace of the solution
 413 (i.e., the boundary and interface data) is assumed and each set of equations (2)
 414 is solved locally over B . In the global problem the solution traces for each $B \in \mathcal{B}$
 415 are coupled. This technique will result in a linear system of the form

$$\begin{bmatrix} \bar{M} & \bar{F} \\ \bar{F}^T & \bar{D} \end{bmatrix} \begin{bmatrix} \bar{u} \\ \bar{\lambda} \end{bmatrix} = \begin{bmatrix} \bar{g} \\ \bar{g}_\delta \end{bmatrix}. \quad (10)$$

416 Here \bar{u} is the approximate solution to (2) at all the grid points and $\bar{\lambda}$ are the
 417 trace variables along internal interfaces; trace variables that are associated with
 418 boundary conditions can be eliminated. The matrix \bar{M} is block diagonal consist-
 419 ing symmetric positive definite blocks for each $B \in \mathcal{B}$, \bar{D} is diagonal, and the
 420 matrix \bar{F} is sparse and incorporates the coupling conditions. The right-hand side
 421 vector \bar{g} incorporates both boundary data (g_D, g_N) and source terms whereas \bar{g}_δ
 422 incorporates the interface data δ .

423 Using the Schur complement we can transform (10) to (11a) and (11b). The
 424 problem size in (10) is significantly reduced since the number of trace variables is




Fig. 2: Sample 2D meshes with different block numbers

425 significantly smaller than the number of solution variables. Because $\bar{\mathbf{M}}$ is block
 426 diagonal here, the inverse of a block diagonal matrix can be obtained in a decoupled
 427 manner for each $B \in \mathcal{B}$. Thus there is a trade-off between the number of blocks
 428 and the size of system (11a), since for a fixed resolution increasing the number of
 429 blocks means that $\bar{\mathbf{M}}$ will be more efficiently factored but the size of (11a) will
 430 increase through the introduction of additional trace variables.

$$\left(\bar{\mathbf{D}} - \bar{\mathbf{F}}^T \bar{\mathbf{M}}^{-1} \bar{\mathbf{F}}\right) \bar{\boldsymbol{\lambda}} = \bar{\mathbf{g}}_\delta - \bar{\mathbf{F}}^T \bar{\mathbf{M}}^{-1} \bar{\mathbf{g}}, \quad (11a)$$

$$\left(\bar{\mathbf{M}} - \bar{\mathbf{F}} \bar{\mathbf{D}}^{-1} \bar{\mathbf{F}}^T\right) \bar{\mathbf{u}} = \bar{\mathbf{g}} - \bar{\mathbf{F}} \bar{\mathbf{D}}^{-1} \bar{\mathbf{g}}_\delta. \quad (11b)$$

431 The big picture of the hybridized method is described above, more detailed
 432 formulation of how local problems and global problems are formed can be found
 433 in Kozdon, Erickson, et al. (2020). For the problem that we are studying, we
 434 want to see the trade-off between the number of the blocks and the number of
 435 trace variables, and how this trade-off would affect the performance of solving
 436 Poisson's equation 9 on a given domain. For evaluation, we used and modified the
 437 code in <https://github.com/bfam/HybridSBP> for our test. We generate different
 438 square meshes that have unit length with decreasing grid spacing using built-in
 439 mesh refinement in the Trellis meshing software. For a unit square domain, we
 440 can choose different numbers of blocks with different local meshing levels within
 441 each block to achieve the same number of total grid points in each direction. To
 442 demonstrate this, we use two different meshes with the same number of total grid
 443 points in each direction in figure 2. The first mesh has only one block with no
 444 trace variables, and each block has $2^2 + 1$ grid points in each direction. The second
 445 mesh has four blocks with internal boundaries marked in red, and each block has
 446 $2^1 + 1$ grid points. Trace variables along these internal boundaries can be obtained
 447 from solving a global problem using Schur complement as described previously.

448 We use mesh refinement to generate unit mesh with the number of blocks in
 449 each direction ranging from 1 to 2^{10} in Trellis. Within each block, we start with
 450 2^4 and gradually start mesh refinement by increasing the number of grid points
 451 in each direction and creating the respective local operators. Note that this is

Number of Grid Points	Block Size Of Given Mesh				
	1 x 1 (LU Factorization)	2 x 2	4 x 4	8 x 8	16 x 16
N					
2^4	✓				
2^5	✓	✓			
2^6	✓	✓	✓		
2^7	✓	✓	✓	✓	
2^8	✓	✓	✓	✓	
2^9	✓	✓	✓	✓	
2^{10}	✓	✓	✓	✓	
2^{11}	✓	LoadError	LoadError	✓	✓
2^{12}	MallocError	×	×	?	?

Table 6: The largest system that can be solved with different number of blocks

452 different than in mesh refinement software where increasing the number of blocks
 453 also increases the number of interfaces. Mesh refining in Julia only applies to local
 454 operators, with number of interfaces fixed. The number of trace variables would
 455 increase because each interface now have more trace variables as a result of local
 456 block mesh refinement. The first question we are trying to answer is whether the
 457 hybrid-method with multiple blocks can help us solve a problem with more grid
 458 points in each direction without using iterative solvers under the constraint of
 459 available memory. We choose 6th order operators for comparison. We requested
 460 128 GB memory on the Talapas server and explored the largest problems we can
 461 solve with respect to different block sizes using the hybrid method with a direct
 462 solvers. We also added results from a single block as benchmark. The results are
 463 given in table 6

464 In this table, $N+1$ refers to the total number of grid points in each direc-
 465 tion. The MallocError in the table refers to the memory allocation errors that
 466 we encountered in the LU decomposition. The LoadError in the table refers to
 467 the error of loading the Cholesky factorization which is used in the hybridized
 468 method. They are denoted differently but the reasons are all associated with the
 469 limitation of memory for factroizations. We are limited by the size of the matrices
 470 that we can factorize. Note that our non-hybrid implementation on single block
 471 previously used LU factorization instead of Cholesky factorization. The results
 472 between a single block with multi-blocks can be inconsistent. But we can see as
 473 we increase the number of the blocks, we are able to solve this problem on a larger
 474 system. The reason that we couldn't finish working on the server is limited by the
 475 job length that we submitted. The calculation for 2^{12} block was terminated for
 476 both cases. No factorization error was reported after successfully obtaining results
 477 for 2^{11} . This means that given enough time, we are expected to see the system
 478 being solved with 2^{12} grid points along each side on an 8×8 block. This is further
 479 confirmed that even though both the 8×8 grid of blocks and 16×16 grid of blocks
 480 can be used to solve a system with $2^{11} + 1$ total grid points in each direction, the
 481 memory allocations on the latter one is much less than the previous one because
 482 the most computationally expensive part is factorization, and in hybrid scheme,
 483 this is associated with local block sizes. The number 8 in 8×8 here refers the
 484 number of the grid points is $2^8 + 1$. If there is no mesh refinement within each
 485 grid, the associated linear system has the size of 2^{16} by 2^{16} .

N	Grid Size Of Given Mesh								
	4 x 4			8 x 8			16 x 16		
	N_p^{Vol}	N_p^{Tr}	N_p^{Vol}/N_p^{Tr}	N_p^{Vol}	N_p^{Tr}	N_p^{Vol}/N_p^{Tr}	N_p^{Vol}	N_p^{Tr}	N_p^{Vol}/N_p^{Tr}
2^6	4624	408	11.33	5184	1008	5.14	6400	2400	2.67
2^7	17424	792	22	18496	1904	9.71	20736	4320	4.8
2^8	67600	1560	43.33	69696	3696	18.86	73984	8160	9.07
2^9	266256	3096	86	270400	7280	37.14	278784	15840	17.6
2^{10}	1056784	6168	171.33	1065024	14448	73.71	1081600	31200	34.67
2^{11}	4210704	12312	342	4227136	28784	146.86	4260096	61920	68.8

Table 7: Number of volume points and trace variables with different number of blocks

486 The reason why we are seeing significantly long run-time from the code is
487 because in the original implementation of the hybrid method, there is no parallel
488 mechanism involved. Ideally, since each block is decoupled after solving the global
489 problem, the factorization for each block can be done independently. Also in the
490 hybrid method, many interior blocks have the same boundary conditions, which
491 results in identical SBP operators on the LHS for these blocks. Therefore there is
492 no need to do factorization for each matrix on the LHS. We only need to compute
493 a factorization for each distinctive local matrix and reuse the result when we
494 encounter an identical one. This is also not implemented in the code. The work of
495 performance improvement is crucial to not just being able to solve a larger problem,
496 but also solving it faster. We will continue exploring the optimization of the hybrid
497 method. Nevertheless, with larger block-sizes, we can still solve the problem faster
498 with more blocks. For direct solve after obtaining factorization results, we can
499 also use multi-threading for acceleration. We set number of threads to be a fixed
500 value 6, and we evaluate performance on meshes with different block numbers. To
501 avoid the impact of jobs being assigned to different node with different capacity,
502 we tested performance on a local computer with i-9500f CPU and fixed 16 Gb
503 memory. The result is given in the table 8. For comparison we also list the number
504 of volume points and trace variables in table 7. They are defined as:

$$N_p^{vol} = (N_l + 1)^2 N_b$$

$$N_p^{tr} = (N_l + 1) N_l$$

505 Here N_b represents the number of blocks, N_l refers to the number of grid points
506 in each direction in the local blocks. N_l represents the number of interfaces. The
507 ratio N_p^{tr}/N_p^{vol} would increase for the same N as we increase the number of blocks,
508 which can be confirmed with data in table 7.

509 From table 8, we can see as we increase the number of blocks, for the same N ,
510 more numbers of blocks mean more local systems to solve, but the total time to
511 solve all these local system is less. The computational complexity of direct solvers
512 after factorization is $O(N^2)$. (Here N means the size of the system, not number
513 of grid points in each direction). For each N , doubling the number of blocks in
514 each direction would reduce the size of the local system to 1/4. The cost of solving
515 each local problem is 1/16 of the previous one. But we also have 4 times more
516 local systems. The total cost should be reduced by 1/4. We can see the time to
517 do the direct solve after LU decomposition is faster, but the speedup is not 4x.

Number of Grid Points	Grid Size Of Given Mesh		
N	4 x 4	8 x 8	16 x 16
2^6	0.00410 s		
2^7	0.00270 s	0.00510 s	
2^8	0.01370 s	0.00920 s	0.00840 s
2^9	0.05480 s	0.02920 s	0.021300 s
2^{10}	0.29040 s	0.19320 s	0.11650 s
2^{11}	LoadError	1.12210 s	0.89510 s

Table 8: Time for direct solvers after factorization with different number of blocks

Number of Grid Points	Grid Size Of Given Mesh		
N	4 x 4	8 x 8	16 x 16
2^6	1251.687673025973		
2^7	0.2722459250042488	8076.298905171896	
2^8	0.006891702238438522	0.006643733883857528	11821.804240653208
2^9	0.00016887738908282545	0.0001429235518197867	0.0017873144253085962
2^{10}	3.932605990678488e-6	2.7931022138805148e-6	3.6928926337184484e-6
2^{11}	LoadError	7.434960489253316e-8	8.87865211871948e-8

Table 9: Errors for direct solvers with different number of blocks

518 The reason could be from implementation side, or it is coming from the fact that
 519 we are solving a sparse system instead of a dense system. Further experiments are
 520 needed to answer this question.

521 We need to further verify that using different number of blocks in hybrid
 522 method can give results with errors at the same level if not identical. We listed
 523 results in table 9. Errors for the coarsest mesh on different blocks are significantly
 524 larger than the rest, this is because of the test function we used in Kozdon, Erick-
 525 son, et al., 2020. As we continue further mesh refinement in each local block, we
 526 would obtain significantly reduced errors that are also converging.

527 The results are shown in table 10. The cost for forming the global problem
 528 and solving for trace variables, although not optimized with parallel scheme, also
 529 favors using more blocks for a given N that is sufficiently large. But for smaller N ,
 530 further increasing the number blocks increases the number of trace variables in the
 531 global problem, hence solving the global problem becomes the bottle-neck for the
 532 whole problem. This can be seen when $N=2^8$ and 2^9 , using 16×16 blocks results
 533 in worse performance compared to using 8×8 block. With optimized parallel
 534 implementation leveraging the advantage of having more decoupled local systems,
 535 we would expect for a very large system size that the hybrid SBP method would
 536 outperform the non-hybrid SBP method on a single domain. The fact that each
 537 local system is decoupled makes this problem embarrassingly parallel, which is
 538 extremely suitable for architecture that is better for computationally intensive
 539 tasks rather than I/O intensive tasks.

540 5 Conclusions

541 We tested the SBP-SAT method in solving Poisson's equation on a single domain.
 542 Our results show that the Conjugate Gradient method on CPU outperforms Con-

Number of Grid Points	Grid Size Of Given Mesh		
N	4 x 4	8 x 8	16 x 16
2^6	0.202 <i>s</i>		
2^7	1.460 <i>s</i>	0.005 <i>s</i>	
2^8	3.961 <i>s</i>	2.374 <i>s</i>	4.650 <i>s</i>
2^9	34.330 <i>s</i>	19.740 <i>s</i>	29.408 <i>s</i>
2^{10}	307.085 <i>s</i>	183.804 <i>s</i>	167.213 <i>s</i>
2^{11}	LoadError	1522.338 <i>s</i>	1170.490 <i>s</i>

Table 10: Time for solving global problems with different number of blocks

543 jugate Gradient methods on GPU and direct solvers on CPU even excluding fac-
544 torization costs. We weren't able to test direct methods on GPU in our previous
545 test due to the lack of properly implemented Julia interface to CUDA libraries.
546 GPU iterative solvers have high start-up costs which make them not suitable for
547 solving Poisson equations on a small system. We were limited by the incompat-
548 ibility of Julia packages in GPU computing and linear algebra and we failed to
549 test further on a large system. But computations on GPU are more scalable for
550 linear algebra operations and also more ideal for computationally intensive jobs
551 that are bottle-necked by I/O through-put between GPU and CPU. We are able
552 to test more thoroughly with recent updates in Julia and related packages that fix
553 the existing issues that limit our previous implementation.

554 We also tested the newly proposed hybrid SBP scheme. Our result shows that
555 hybrid SBP method can work on a larger system compared to traditional non-
556 hybridized SBP method under the constraint of system memory. Our empirical
557 results also show that the hybrid method achieves better performance on a system
558 with more blocks. For certain values of N , using more grid points at the start
559 improves performance, but further increasing the number of blocks reduces the
560 performance. To explore the trade-off relationship between the number of trace
561 variables and the local problem size quantitatively, we need to test our problem
562 on more grids with different numbers of blocks.

563 We bench-marked existing implementations to identify several issues that limit
564 the size of the system that we are able to solve. We tested different methods to
565 improve the current implementation. Results from our study form the foundation
566 of utilizing the new finite difference method whose hybrid structure is ideally
567 suited for GPU parallelization to achieve performance gains. This is much needed
568 for simulations on a large problem which would give rise to a extremely large linear
569 system that can't be solved with existing method and implementation, in order to
570 understand the behaviors of earthquakes in real scenario.

571 Acknowledgement

572 This work benefited from access to the University of Oregon high performance
573 computer, Talapas.

574 **References**

- 575 Besard, T., C. Foket, and B. De Sutter (2019). “Effective extensible programming:
576 unleashing julia on gpus”. In: *IEEE Transactions on Parallel and Distributed*
577 *Systems* 30.4, pp. 827–841.
- 578 Carpenter, M. H., D. Gottlieb, and S. Abarbanel (1994). “Time-stable boundary
579 conditions for finite-difference schemes solving hyperbolic systems: methodol-
580 ogy and application to high-order compact schemes”. In: *Journal of Computa-*
581 *tional Physics* 111.2, pp. 220–236. DOI: [10.1006/jcph.1994.1057](https://doi.org/10.1006/jcph.1994.1057).
- 582 Carpenter, M. H., J. Nordström, and D. Gottlieb (1999). “A stable and conserva-
583 tive interface treatment of arbitrary spatial accuracy”. In: *Journal of Compu-*
584 *tational Physics* 148.2, pp. 341–365. DOI: [10.1006/jcph.1998.6114](https://doi.org/10.1006/jcph.1998.6114).
- 585 Ciarlet, P. G. (2002). *The Finite Element Method for Elliptic Problems*. Society for
586 Industrial and Applied Mathematics. DOI: [10.1137/1.9780898719208](https://doi.org/10.1137/1.9780898719208).
- 587 Erickson, B. A. and S. M. Day (2016). “Bimaterial effects in an earthquake cycle
588 model using rate-and-state friction”. In: *Journal of Geophysical Research: Solid*
589 *Earth* 121, pp. 2480–2506. DOI: [10.1002/2015JB012470](https://doi.org/10.1002/2015JB012470).
- 590 Erickson, B. A. and E. M. Dunham (2014). “An efficient numerical method for
591 earthquake cycles in heterogeneous media: alternating subbasin and surface-
592 rupturing events on faults crossing a sedimentary basin”. In: *Journal of Geo-*
593 *physical Research: Solid Earth* 119.4, pp. 3290–3316. DOI: [10.1002/2013JB010614](https://doi.org/10.1002/2013JB010614).
- 594 Karlstrom, L. and E. M. Dunham (2016). “Excitation and resonance of acoustic-
595 gravity waves in a column of stratified, bubbly magma”. In: *Journal of Fluid*
596 *Mechanics* 797, pp. 431–470. DOI: [10.1017/jfm.2016.257](https://doi.org/10.1017/jfm.2016.257).
- 597 Kozdon, J. E., E. M. Dunham, and J. Nordström (2012). “Interaction of waves
598 with frictional interfaces using summation-by-parts difference operators: weak
599 enforcement of nonlinear boundary conditions”. In: *Journal of Scientific Com-*
600 *puting* 50, pp. 341–367. DOI: [10.1007/s10915-011-9485-3](https://doi.org/10.1007/s10915-011-9485-3).
- 601 Kozdon, J. E., B. A. Erickson, and L. C. Wilcox (2020). “Hybridized summation-
602 by-parts finite difference methods”. In: *ArXiv* abs/2002.00116.
- 603 Kreiss, H. and G. Scherer (1974). “Finite element and finite difference methods
604 for hyperbolic partial differential equations”. In: *Mathematical aspects of finite*
605 *elements in partial differential equations; Proceedings of the Symposium*. Madison,
606 WI, pp. 195–212. DOI: [10.1016/b978-0-12-208350-1.50012-1](https://doi.org/10.1016/b978-0-12-208350-1.50012-1).
- 607 — (1977). *On the existence of energy estimates for difference approximations for hy-*
608 *perbolic systems*. Tech. rep. Department of Scientific Computing, Uppsala Uni-
609 versity.
- 610 Lotto, G. C. and E. M. Dunham (2015). “High-order finite difference modeling
611 of tsunami generation in a compressible ocean from offshore earthquakes”. In:
612 *Computational Geosciences* 19.2, pp. 327–340. DOI: [10.1007/s10596-015-9472-0](https://doi.org/10.1007/s10596-015-9472-0).
- 613 Mattsson, K. (2012). “Summation by parts operators for finite difference approxi-
614 mations of second-derivatives with variable coefficients”. In: *Journal of Scientific*
615 *Computing* 51.3, pp. 650–682. DOI: [10.1007/s10915-011-9525-z](https://doi.org/10.1007/s10915-011-9525-z).
- 616 Mattsson, K., F. Ham, and G. Iaccarino (2009). “Stable boundary treatment for
617 the wave equation on second-order form”. In: *Journal of Scientific Computing*
618 41.3, pp. 366–383. DOI: [10.1007/s10915-009-9305-1](https://doi.org/10.1007/s10915-009-9305-1).
- 619 Mattsson, K. and J. Nordström (2004). “Summation by parts operators for finite
620 difference approximations of second derivatives”. In: *Journal of Computational*
621 *Physics* 199.2, pp. 503–540. DOI: [10.1016/j.jcp.2004.03.001](https://doi.org/10.1016/j.jcp.2004.03.001).

- 622 Mattsson, K. (Feb. 2003). “Boundary procedures for summation-by-parts opera-
623 tors”. In: *Journal of Scientific Computing* 18.1, pp. 133–153.
- 624 Nissen, A., G. Kreiss, and M. Gerritsen (2013). “High order stable finite difference
625 methods for the Schrödinger equation”. In: *Journal of Scientific Computing* 55.1,
626 pp. 173–199. DOI: [10.1007/s10915-012-9628-1](https://doi.org/10.1007/s10915-012-9628-1).
- 627 Roache, P. (1998). *Verification and Validation in Computational Science and Engi-
628 neering*. 1st ed. Albuquerque, NM: Hermosa Publishers.
- 629 Saad, Y. (2003). *Iterative Methods for Sparse Linear Systems*. Second. Society for
630 Industrial and Applied Mathematics. DOI: [10.1137/1.9780898718003](https://doi.org/10.1137/1.9780898718003). eprint:
631 <https://epubs.siam.org/doi/pdf/10.1137/1.9780898718003>. URL: [https:
632 //epubs.siam.org/doi/abs/10.1137/1.9780898718003](https://epubs.siam.org/doi/abs/10.1137/1.9780898718003).
- 633 Strand, B. (1994). “Summation by parts for finite difference approximations for
634 d/dx ”. In: *Journal of Computational Physics* 110.1, pp. 47–67. DOI: [10.1006/
635 jcph.1994.1005](https://doi.org/10.1006/jcph.1994.1005).
- 636 Virta, K. and K. Mattsson (2014). “Acoustic wave propagation in complicated
637 geometries and heterogeneous media”. In: *Journal of Scientific Computing* 61.1,
638 pp. 90–118. DOI: [10.1007/s10915-014-9817-1](https://doi.org/10.1007/s10915-014-9817-1).