



Mat2Stencil: A Modular Matrix-Based DSL for Explicit and Implicit Matrix-Free PDE Solvers on Structured Grid

HUANQI CAO, Tsinghua University, China

SHIZHI TANG, Tsinghua University, China

QIANCHAO ZHU, Peking University, China

BOWEN YU, Tsinghua University, China

WENGUANG CHEN, Tsinghua University, China and Pengcheng Laboratory, China

Partial differential equation (PDE) solvers are extensively utilized across numerous scientific and engineering fields. However, achieving high performance and scalability often necessitates intricate and low-level programming, particularly when leveraging deterministic sparsity patterns in structured grids.

In this paper, we propose an innovative domain-specific language (DSL), Mat2Stencil, with its compiler, for PDE solvers on structured grids. Mat2Stencil introduces a structured sparse matrix abstraction, facilitating modular, flexible, and easy-to-use expression of solvers across a broad spectrum, encompassing components such as Jacobi or Gauss-Seidel preconditioners, incomplete LU or Cholesky decompositions, and multigrid methods built upon them. Our DSL compiler subsequently generates matrix-free code consisting of generalized stencils through multi-stage programming. The code allows spatial loop-carried dependence in the form of quasi-affine loops, in addition to the Jacobi-style stencil's embarrassingly parallel on spatial dimensions. We further propose a novel automatic parallelization technique for the spatially dependent loops, which offers a compile-time deterministic task partitioning for threading, calculates necessary inter-thread synchronization automatically, and generates an efficient multi-threaded implementation with fine-grained synchronization.

Implementing 4 benchmarking programs, 3 of them being the pseudo-applications in NAS Parallel Benchmarks with 6.3% lines of code and 1 being matrix-free High Performance Conjugate Gradients with 16.4% lines of code, we achieve up to 1.67 \times and on average 1.03 \times performance compared to manual implementations.

CCS Concepts: • **Computing methodologies** \rightarrow **Parallel programming languages**; *Shared memory algorithms*; • **Software and its engineering** \rightarrow **Domain specific languages**; *Source code generation*.

Additional Key Words and Phrases: domain-specific language, multi-stage programming, compiler, finite difference method, stencil, structured grid, polyhedral compilation, performance optimization

ACM Reference Format:

Huanqi Cao, Shizhi Tang, Qianchao Zhu, Bowen Yu, and Wenguang Chen. 2023. Mat2Stencil: A Modular Matrix-Based DSL for Explicit and Implicit Matrix-Free PDE Solvers on Structured Grid. *Proc. ACM Program. Lang.* 7, OOPSLA2, Article 246 (October 2023), 30 pages. <https://doi.org/10.1145/3622822>

Authors' addresses: [Huanqi Cao](mailto:caohq18@mails.tsinghua.edu.cn), caohq18@mails.tsinghua.edu.cn, Tsinghua University, Beijing, China; [Shizhi Tang](mailto:tsz19@mails.tsinghua.edu.cn), tsz19@mails.tsinghua.edu.cn, Tsinghua University, Beijing, China; [Qianchao Zhu](mailto:dysania@pku.edu.cn), dysania@pku.edu.cn, Peking University, Beijing, China; [Bowen Yu](mailto:yubowen@tsinghua.edu.cn), yubowen@tsinghua.edu.cn, Tsinghua University, Beijing, China; [Wenguang Chen](mailto:cwg@tsinghua.edu.cn), cwg@tsinghua.edu.cn, Tsinghua University, Beijing, China and Pengcheng Laboratory, Shenzhen, China.



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2023 Copyright held by the owner/author(s).

2475-1421/2023/10-ART246

<https://doi.org/10.1145/3622822>

1 INTRODUCTION

¹ Numerical solutions of partial differential equations (PDEs) are crucial in a wide variety of scientific computing applications, spanning from weather prediction [Skamarock et al. 2019] to seismic imaging [Mcmechan 2006]. Many of these applications employ structured grids, primarily Cartesian grids, for spatial discretization. These grids manifest as 2-dimensional or 3-dimensional arrays, with calculations executed through algorithm-driven traversal over the array elements. In the past, scientists employed general-purpose programming languages, such as FORTRAN and C/C++, to manually develop solvers for PDEs on structured grids. This process required extensive use of “for” loops and arithmetic operations, often leading to solvers encompassing thousands of lines of code. As multi-threading and single instruction multiple data (SIMD) technologies have advanced in modern computing, the manual development of such solvers is progressively becoming less practical and less effective.




Fig. 1. The different computation patterns of explicit and implicit solving algorithms, with a 1D heat equation (Equation (1)) as an example.

A natural direction to address this issue is to translate the operations on structured grids to manipulations of sparse matrices, which we term the *matrix-based* approach. The PDE solving procedure thus becomes multiplying sparse matrices to vectors (SpMV) and solving sparse linear systems according to the matrix. This approach thus typically employs general library routines [Intel 2023; Virtanen et al. 2020] to work on the user-provided matrices. However, there are two critical issues. First, constructing the sparse matrix for the PDE problem necessitates significant human effort due to the complexity of real-world equations and their solvers. Second, as these libraries serve general purposes, they must read the matrix and analyze it at runtime to cater to a broader range of applications, including unstructured grids. Consequently, they have limited ability to efficiently exploit structured sparsity patterns. Despite covering the domain well, the matrix-based approach suffers from unresolved programmability and performance issues.

However, *matrix-free* solvers that directly target the equations are typically handwritten in languages such as C/C++ or FORTRAN, making them labor-intensive to implement. Domain-specific languages (DSLs) and compilers have emerged as promising approaches to simplify programming and enhance performance. By eliminating matrices, the computations involved in solving partial differential equations (PDEs) are reduced to stencil operators, which iterate over grid points and compute new values based on their neighborhoods. A well-known class of stencil operators is the Jacobi-style stencils, which read and write different arrays at each time step and exhibit no spatial dependence, which is the dependence between different spatial locations at the same time step, as illustrated in Fig. 1a. In contrast, Seidel-style stencils read and write the same arrays, thereby introducing spatial dependence, as seen in Fig. 1b, which makes them challenging to automatically parallelize and optimize. While some DSLs, such as [Lengauer et al. 2020; Louboutin

¹We improved the writing of this section by prompting GPT-4 [OpenAI 2023], OpenAI’s large-scale language-generation model, with the draft written by humans. Upon generating, we reviewed, edited, and revised the language to ensure its fidelity and take ultimate responsibility for the content of this publication.

et al. 2019], have successfully simplified the programming of PDE equations and solvers by primarily targeting Jacobi-style stencils, they fail on many algorithms containing Seidel-style stencils due to limited abstractions. On the other hand, domain-specific compilers, such as [Bondhugula et al. 2017, 2008; Verdoolaege and Janssens 2017], support both types of stencil operators through advanced dependence analysis techniques but lack abstractions for PDEs, necessitating significant human effort to program a solver with nested loops.

In summary, we recognize several challenges in structured grid PDE solving: ❶ a labor-intensive process to express the equation and the solver; ❷ a conflict between expressiveness and the matrix-free property; ❸ the absence of effective automatic optimization for Seidel-style stencils. To address these challenges, we introduce a novel DSL called `MAT2STENCIL`, designed for both explicit and implicit solving algorithms for PDEs, while ensuring its expressiveness, modularity, and performance. Our contributions are as follows:

- We enable modular programming of equations and their solvers through a high-level abstraction of *structured sparse matrices*. It utilizes a set of simple matrices and their arithmetic operations to compose more complicated problem matrices. These pre-defined matrices reflect the various steps involved in discretizing partial differential equations (PDEs) and enable user-friendly problem construction in PDE-solving processes.
- We enhance expressiveness by introducing a low-level abstraction for structured sparse matrices based on *row functions*. It allows users to effortlessly define their custom sparse matrices as well as implement custom algorithms against any structured sparse matrices. The design is kept simple and compile-time resolvable, to enable further lowering to matrix-free implementations.
- We compile structured sparse matrices user code to matrix-free target code by expanding the sparse matrices during the first stage of our multi-stage programming infrastructure. In specific, we backtrack and discuss different cases of loop indices to fill the gap between compile-time (stage 1) and runtime (stage 2) control flow, generating matrix-free code consisting of regular quasi-affine loops.
- We develop a novel automatic parallelization technique, SewTh (abbrev. for **S**ewing the **T**hreads), based on polyhedral dependence analysis, to enable efficient multi-threaded execution of the more generalized stencil loops with spatial dependencies. This approach partitions the loop domain across threads and automatically generates fine-grained synchronization when dependencies cross thread boundaries.

Altogether, our new DSL is capable of expressing a much wider range of explicit and implicit PDE solvers. We evaluate our DSL by implementing four benchmarking programs from the NAS Parallel Benchmarks (NPB) [Bailey et al. 1991] and High Performance Conjugate Gradients (HPCG) [Dongarra et al. 2016], which includes Symmetric Gauss-Seidel, Symmetric Successive Over-Relaxation, and direct solving of multi-diagonal matrices (degenerated from Incomplete Lower-Upper decomposition) as the implicit solving kernels with spatial dependence. Our performance reaches up to $1.67\times$ performance in NPB compared with the original FORTRAN implementation, with only 6.3% lines of code; we also finish HPCG implementation in 16.4% lines of code compared to a manually matrix-free implementation. The geometric mean of our performance over all the programs is $1.03\times$ of the highly-optimized matrix-free human implementations.

2 BACKGROUND AND MOTIVATION

In this section, we discuss the current status and problems in differential equation solving, and then the background of techniques we will employ in this paper, including multi-stage programming and polyhedral analysis for automatic parallelization.

2.1 Solving Differential Equations on Structured Grids

In this subsection, we demonstrate existing approaches to solving differential equations on structured grids. We use the simple equation in Equation (1) as an example in which we want to solve $u(x, t)$.

$$\frac{\partial u}{\partial t} = \frac{\partial^2 u}{\partial x^2}, x \in [0, 1]; u(0, t) \text{ and } u(1, t) \text{ constant.} \quad (1)$$

To solve it numerically through Finite Difference [Strikwerda 2004], the variable t is discretized temporally to $T\Delta t$ and variable x spatially to $i\Delta x$ ($0 \leq i \leq n$, $\Delta x = 1/n$), in which T for timesteps and i for grid point number are both integers. We thus note $u(T\Delta t, i\Delta x)$ as $u_i^{(T)}$ in short. If we discretize $\partial u/\partial t$ as $(u_i^{(T+1)} - u_i^{(T)})/\Delta t$, we have the explicit numerical scheme in Equation (2), which is a iterative approach to compute u one by each timestep:

$$\begin{aligned} \frac{u_i^{(T+1)} - u_i^{(T)}}{\Delta t} &= \begin{cases} (u_{i-1}^{(T)} - 2u_i^{(T)} + u_{i+1}^{(T)})/(\Delta x)^2 & \text{if } 0 < i < n, \\ 0 & \text{otherwise} \end{cases} \\ u_i^{(T+1)} &= \begin{cases} \lambda u_{i-1}^{(T)} + (1 - 2\lambda)u_i^{(T)} + \lambda u_{i+1}^{(T)} & \text{if } 0 < i < n, \\ u_i^{(T)} & \text{otherwise} \end{cases} \quad \text{where } \lambda = \Delta t/(\Delta x)^2. \end{aligned} \quad (2)$$

However, such explicit schemes sometimes suffer from stability problems. Instead, if we discretize $\partial u/\partial t$ as $(u_i^{(T)} - u_i^{(T-1)})/\Delta t$, we then derive an implicit numerical scheme in Equation (3), which requires solving a linear system at each timestep:

$$\begin{aligned} \frac{u_i^{(T)} - u_i^{(T-1)}}{\Delta t} &= \begin{cases} (u_{i-1}^{(T)} - 2u_i^{(T)} + u_{i+1}^{(T)})/(\Delta x)^2 & \text{if } 0 < i < n \\ 0 & \text{otherwise} \end{cases} \\ \begin{cases} -\lambda u_{i-1}^{(T)} + (1 + 2\lambda)u_i^{(T)} - \lambda u_{i+1}^{(T)} = u_i^{(T-1)} & \text{for } 0 < i < n, \\ u_0^{(T)} = u_0^{(T-1)}, u_n^{(T)} = u_n^{(T-1)} \end{cases} & \quad \text{where } \lambda = \Delta t/(\Delta x)^2. \end{aligned} \quad (3)$$

We see that explicitly solving differential equations only requires straightforward computation, while an implicit method will need to solve systems of linear equations. To solve the sparse linear systems, direct methods (usually for multi-diagonal matrices) and iterative methods using Jacobi, Gauss-Seidel, Incomplete LU decomposition, etc., have been employed. Though in different patterns, the existing methods for both cases can be classified into two: one is matrix-based, generates sparse matrices, and uses general routines for sparse linear algebra; the other is matrix-free and targets the math formulas directly.

2.1.1 Matrix-Based Approach: General Sparse Linear Algebra. One general method to deal with the grids that PDEs are discretized onto involves sparse matrices. Regarding Equation (3), the linear system to be solved can be formulated in matrix form as follows:

$$\begin{bmatrix} 1 & & & & & \\ -\lambda & 1+2\lambda & -\lambda & & & \\ & \ddots & \ddots & \ddots & & \\ & & & -\lambda & 1+2\lambda & -\lambda \\ & & & & & 1 \end{bmatrix} \mathbf{u}^{(T)} = \mathbf{u}^{(T-1)} \quad (4)$$

Users may thus generate the sparse matrix by themselves and use general sparse linear algebra routines to solve their equation. Similarly, the explicit form in Equation (2) is formulated as a similar sparse matrix-vector multiplication, or *SpmV*. Thus, by implementing those algorithms against

arbitrary sparse matrices, library providers liberate domain experts from writing those complicated codes, only leaving the generating of sparse matrices.

Yet, generating the sparse matrix via hand-written codes can be annoying. With much more complex equations, the sparse matrix may require more than one page of math formulas to specify, not to mention the code. Besides the difficulty in constructing the matrices, the performance of the matrix-based approach is unsatisfying. Although general sparse routines are highly optimized, require extensive human efforts from expert programmers, and usually cannot be implemented efficiently by science domain researchers, these general sparse routines cannot compete with matrix-free approaches in structured problems, due to the lack of structural information at compile time.

2.1.2 Matrix-Free Approaches: Stencil by Hand or Automatically. Regarding the structured grids, it is possible to get rid of the matrices and perform computation directly, which we call “matrix-free”. Like in Equation (2), explicit approaches primarily involve embarrassingly parallel stencil computations in a single timestep. In each iteration, they read neighbor points in the input 2D or 3D tensors with a certain pattern, go through a certain function as its computation, and yield the corresponding value in output tensor(s). There is no data dependence between different grid points if we look at only one iteration. Due to its nature of independent computation, parallelizing such types of stencil computation is straightforward. We denote such stencils as “Jacobi-style” stencils considering that Jacobi iterations match such a computational pattern without purely spatial dependence. Since the computation is highly customized across applications due to different PDEs to solve, and optimization by hand requires massive human effort, many domain-specific languages (DSLs) have been designed to address the optimization of stencil computation, as well as programmability of explicit or Jacobi iteration-based implicit solvers.

While Jacobi-style stencil computations are widely researched, many implicit solvers require operators not covered by such DSLs. Such operators, while still operating neighbors on structured grids, have dependence carried purely by spatial loops on one or more dimensions, which we recognize as Seidel-style stencils, named after the most commonly used Gauss-Seidel iteration but not limited to that particular algorithm. For example, directly solving Equation (3) requires Gaussian eliminating forward and backward, introducing sequential dependence along the x axis. More complicated equations and numerical schemes may require iterative methods such as multi-grid solvers or conjugate gradients, which in turn requires routines like Gauss-Seidel iteration, incomplete lower-upper (ILU) factorization, fine-grained parallel ILU [Chow and Patel 2015], and so on.

Although the 1D case introduced here has to be executed sequentially, Seidel-style stencils with two or more dimensions are usually possible to get parallelized, but not as trivial as those Jacobi-style stencils, requiring more human efforts to optimize. For example, the block tri-diagonal pseudo-application in NAS Parallel Benchmarks (NPB-BT) is written in over 400 lines of Fortran code for solving the block tri-diagonal matrix in only one direction, and it has three directions to solve, involving over 1300 LOC. However, such routines are not yet supported by the DSLs targeting PDE solving. As the best among existing, ExaStencils is capable of implicitly solving differential equations, but only through Jacobian and forward Gauss-Seidel iterations in a multi-grid (MG) solver, which is just a particular case with limited computation patterns.

2.1.3 Discussion. Due to the difficulty of manually programming Seidel-style stencil operators, the matrix-based representation is usually preferred, especially in the case of implicit solvers. For example, the High Performance Conjugate Gradients (HPCG) Benchmark [Dongarra et al. 2016], a widely recognized effort for ranking top-level supercomputers, enforces sparse matrices and forbids matrix-free approaches in its valid implementations, indicating their concern about the generality of

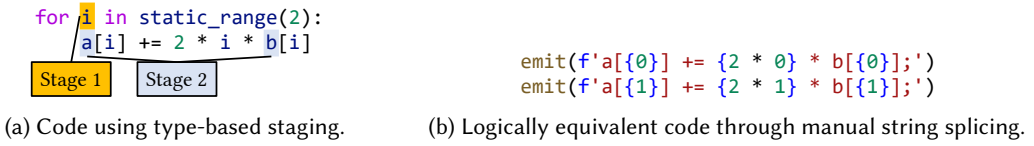


Fig. 2. A simple example for type-based multi-stage programming.

operator implementations. However, it turns out such prohibition limits performance drastically. It is reported that implementing the HPCG computation in a matrix-free form significantly improves the performance, by $4.67\times$ on the New Sunway supercomputer [Zhu et al. 2021]. When possible, HPC researchers still seek matrix-free approaches even for implicit approaches, e.g., the 2016 Gordon Bell Prize winner [Yang et al. 2016] designed and manually implemented a geometry-based ILU that maps the dependence to hardware-supported inter-core communication.

To this end, we seek a solution so that: ❶ the Seidel-style operators are naturally enabled, resolving the coverage issue of stencil DSLs; ❷ the construction of linear problems is easy for users, resolving the programmability issue of the matrix-based approach; ❸ the matrix-free property is retained at runtime, resolving the performance issue of the matrix-based approach. These challenges lead us to a new DSL working with sparse matrices but compiled into generalized stencils.

2.2 Multi-Stage Programming


Multi-stage programming, or staging, has been a promising method for building embedded domain-specific languages. In such practices of multi-stage programming, a program usually is split into two stages: stage 1 (at compile time) executes the user code partially to generate an optimized intermediate program, and stage 2 (at run time) executes the intermediate program to get the final result. In our design, programmers express the solvers as sparse matrices and their operations on stage 1, which generates matrix-free stage-2 code.

Lightweight Modular Staging (LMS) [Rompf and Odersky 2012] is one of such frameworks and has been adopted in multiple DSL works [Ofenbeck et al. 2013; Shaikhha et al. 2018; Sujeeth et al. 2014]. LMS separates the two stages through only types: `Rep[T]` refers to a T-typed value at the second stage. Since that LMS is a Scala library, we implement our own to build our embedded DSL in Python.

In such an approach, expressions to be delayed to the next stage are used together with normal values. The operators, including conventional ones (e.g. `__mul__(a, b)` for $a * b$) and indexed load/store (`__getitem__` and `__setitem__` for code like `a[...]`), are overloaded to make user code operating on the delayed expressions seamlessly. When the first stage runs, the codes generated by the overloaded operators for the second stage are spliced together for later compiling. The control flow structures are similarly overloaded to be able to generate control flow codes in the next stage. In particular, our backend, built on top of FreeTensor [Tang et al. 2022], transforms the Python AST to enable virtualized control flow, the details of which will be introduced in Section 5.1. A simple example is shown in Fig. 2.

2.3 Polyhedral Analysis for Automatic Parallelization and Optimization

Through multi-stage programming, we produce matrix-free code containing the generalized stencils, in the form of quasi-affine tensor programs. Such programs contain indexed for-loops and branches according to quasi-affine expressions, and access tensors, or multi-dimensional arrays, with quasi-affine indices as well. There have been many studies on optimizing such affine programs since the invention of FORTRAN, which can be summarized as *polyhedral analysis*. Program analysis and transformations for multithreading, vectorization, and better cache management are summarized



(a) A typical skewing parallel implementation. Colored bands are executed in parallel, constructing a skewed “time” axis T . Poor locality.

(b) A typical skewed tiling parallel implementation. Tiles with the same color are executed in parallel, while each tile is executed sequentially by one thread. High startup overhead.

Fig. 3. Automatically-produced polyhedral schedules for 2D 9-point Gauss-Seidel.

in [Padua and Wolfe 1986] and [Allen and Kennedy 2001]. As such methods can only be applied on nested loops, they can hardly be used directly on domain-specific user code, but they point out a path towards automatically parallelizing our generated code.

The most famous ones among those compilers are PPCG [Verdoolaage and Janssens 2017] and PLUTO [Bondhugula et al. 2008]. While capable of parallelizing computation in many different patterns, they need more specially-designed approaches to achieve efficient stencil computations, e.g. diamond tiling in PLUTO [Bondhugula et al. 2017]. However, previous stencil optimizations mainly target Jacobi-style stencils, leaving Seidel-style stencils with vanilla skewing and/or tiling. An example is shown in Fig. 3, which contains two typical methods to parallelize a 2D 9-point Gauss-Seidel loop, produced by the existing automatic approaches. The approach without tiling involves too many global synchronizations and has a poor memory locality. The approach with skewed tiling instead incurs huge starting and ending overhead. As such, they are not efficient at many of the implicit solvers our DSL wants to cover. The development of new techniques resolving the Seidel-style stencils is thus essential for us.

3 OVERVIEW OF MAT2STENCIL LANGUAGE AND COMPILER




Fig. 4. An overview of MAT2STENCIL workflow.

Based on the observation mentioned above, we propose MAT2STENCIL, our domain-specific language and its corresponding compiler for solving partial differential equations (PDEs) on structured grids. An overview is shown in Fig. 4. MAT2STENCIL offers a concise and flexible programming interface for a wide range of PDE solvers. Our structured sparse matrix abstraction enables the modular construction of sparse matrices through linear algebra operations on both predefined and user-defined simple matrices. Furthermore, our DSL separates the algorithms from the matrices, supporting operations ranging from simple stencil operators to complicated Gauss-Seidel iterations and Incomplete LU decomposition. Users write equation solver code that composes simple sparse matrices using arithmetic operations into complicated ones and invokes the algorithms against

them. The compile-time resolvable feature enables us to exploit sparsity patterns effectively. The language design will be covered in Section 4. As the first stage of the multi-stage programming frontend, we assemble matrix-free generalized stencil code in FreeTensor IR [Tang et al. 2022], by backtracking and discussing over compile-time undecidable branches. The staging techniques involved will be covered in Section 5.

After the first stage, the backend optimization of MAT2STENCIL, which we implemented as a FreeTensor extension, is responsible for optimizing and parallelizing the code before proceeding to stage 2 (or runtime). We propose a novel SewTh parallelizer (**S**ewing the **T**hreads), which empowers the MAT2STENCIL backend to automatically and efficiently parallelize Seidel-style stencils. Additionally, we use a PLUTO+-like [Bondhugula et al. 2016] algorithm to fuse the loops, and perform aggressive inlining of elementwise functions, to optimize memory traffic further. They will be discussed in Section 6.

4 MAT2STENCIL LANGUAGE DESIGN

In this section, we present the detailed design of the MAT2STENCIL language. We will start from the uppermost abstraction, which is on equation solvers, and then go down to the matrices abstraction and expression of linear algebra algorithms within MAT2STENCIL. Finally, we present the grammar of MAT2STENCIL.

4.1 Abstraction Over the Equation

In MAT2STENCIL, the solving procedure of a PDE is expressed with sparse linear algebra, in which the sparse matrices involved are structured and modular. As is mentioned in Section 2.1.1, constructing the sparse matrix can be problematic for complicated real-world solvers. Instead of requiring users to generate an entire matrix to solve with twisted problem generation code, we enable the assembly of such problem matrices through linear algebra expressions evaluated at compile-time. We present how MAT2STENCIL abstracts over the equation in Fig. 5 with an end-to-end example from the Scalar Penta-diagonal (SP) pseudo-application in the NAS Parallel Benchmark (NAS) before we dig into the details. Notice that instead of requiring the user to fill in the non-zeros of the matrices to solve one by one like in other matrix-based approaches, MAT2STENCIL only requires users to implement individual modules, including custom structured sparse matrices like `dssp4x` for $\epsilon h^4 (\partial^4 / \partial x^4)$ and `lambda_y`, `rho_n`, etc. as elementwise functions, then write linear expressions to combine them, which are close to the discretized equations.

$\left\{ \mathbf{I} - \Delta\tau \left[\frac{\partial}{\partial x} \Lambda_x(\mathbf{u}) + \frac{\partial^2}{\partial x^2} \rho N(\mathbf{u}) - \epsilon h^4 \frac{\partial^4}{\partial x^4} \right] \right\} \Delta \mathbf{u}_2 = \Delta \mathbf{u}_1$ $\left\{ \mathbf{I} - \Delta\tau \left[\frac{\partial}{\partial y} \Lambda_y(\mathbf{u}) + \frac{\partial^2}{\partial y^2} \rho Q(\mathbf{u}) - \epsilon h^4 \frac{\partial^4}{\partial y^4} \right] \right\} \Delta \mathbf{u}_4 = \Delta \mathbf{u}_3$ $\left\{ \mathbf{I} - \Delta\tau \left[\frac{\partial}{\partial z} \Lambda_z(\mathbf{u}) + \frac{\partial^2}{\partial z^2} \rho S(\mathbf{u}) - \epsilon h^4 \frac{\partial^4}{\partial z^4} \right] \right\} \Delta \mathbf{u}_6 = \Delta \mathbf{u}_5$	<pre style="font-family: monospace; font-size: 0.9em;"> ilu_solve(identity() - DELTA_T * Px * (Ax * Dx * diagonal(lambda_x(u)) + Dx * Dx * diagonal(rho_n(u)) - dssp4x), r) ... ilu_solve(identity() - DELTA_T * Py * (Ay * Dy * diagonal(lambda_y(u)) + Dy * Dy * diagonal(rho_q(u)) - dssp4y), r) ... ilu_solve(identity() - DELTA_T * Pz * (Az * Dz * diagonal(lambda_z(u)) + Dz * Dz * diagonal(rho_s(u)) - dssp4z), r) </pre>
(a) The equations to solve.	(b) Solver implementation.

Fig. 5. The NPB-SP specified problem and corresponded MAT2STENCIL implementation.

One may notice that derivative operators like $\partial/\partial x$ are replaced by some variables like D_x . These variables are predefined matrices that match the discretized differential and help the construction of matrices to solve. We observe that the sparse matrices to solve can usually be constructed from several simple matrices, and the construction procedure is tightly related to the differential formula. To explain this, we start with the first-order derivative. After discretization, the $\partial/\partial x$ differential operator is turned into a sparse matrix D , as is shown below. For simplicity, we write the expanded form of a column vector \mathbf{v} as a semicolon-separated list: $[v_0; \dots; v_n]$.

$$\begin{aligned} \mathbf{u} &\xrightarrow{\text{discretize}} [u_0; \dots; u_n] && \stackrel{\text{def}}{=} \mathbf{u} \\ \frac{\partial \mathbf{u}}{\partial x} &\xrightarrow{\text{discretize}} \left[\frac{u_1 - u_0}{\Delta x}; \dots; \frac{u_n - u_{n-1}}{\Delta x} \right] = \frac{1}{\Delta x} \begin{bmatrix} -1 & 1 & & & \\ & \ddots & \ddots & & \\ & & & -1 & 1 \end{bmatrix}_{(n-1) \times n} \mathbf{u} && \stackrel{\text{def}}{=} D_{(n)} \mathbf{u} \end{aligned}$$

Based on that, we similarly have $\frac{\partial^2 \mathbf{u}}{\partial x^2}$ discretized as $D_{(n-1)} D_{(n)} \mathbf{u}$, which can be simpler noted as $D^2 \mathbf{u}$ since the actual shape of each D can always be inferred from what vector it is multiplying to. Beside that, we further define $\mathbf{P} = [0; \mathbf{I}; 0]$, so that $\mathbf{P}\mathbf{u} = [0; \mathbf{u}; 0]$, which represents padding on the boundary. Using this notation, we can easily perform the discretization in Equation (3) without dealing with scalars:

$$\begin{aligned} \frac{\mathbf{u}^{(T)} - \mathbf{u}^{(T-1)}}{\Delta t} &= [0; D^2 \mathbf{u}^{(T)}; 0] = \mathbf{P} D^2 \mathbf{u}^{(T)} \\ \Rightarrow (\mathbf{I} - \Delta t \mathbf{P} D^2) \mathbf{u}^{(T)} &= \mathbf{u}^{(T-1)}. \end{aligned} \quad (5)$$

It can be easily verified that Equation (5) is equivalent to Equation (4). It naturally specifies how the matrix to solve is constructed, with a number of arithmetic operations on sparse matrices. Furthermore, since that higher dimensional \mathbf{u} will require flattening to a vector, D will be broadcasted into D_x , D_y , and D_z , reflecting the higher-dimensional differencing operations. Similarly, the neighborhood averaging A , which has occurred in Fig. 5b in broadcasted forms, is defined as

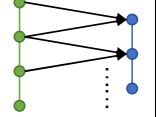
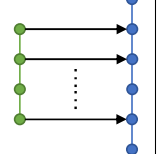
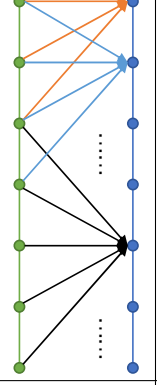
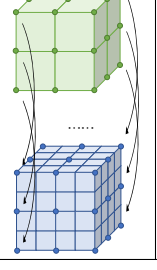
$$\mathbf{A} = \begin{bmatrix} 1/2 & 1/2 & & & \\ & 1/2 & 1/2 & & \\ & & \ddots & \ddots & \\ & & & 1/2 & 1/2 \end{bmatrix}, \text{ so that } \mathbf{A}\mathbf{u} = \left[\frac{u_0 + u_1}{2}; \dots; \frac{u_{n-1} + u_n}{2} \right].$$

However, computing the sparse matrices via these arithmetics at runtime can be costly, especially noticing that we need multiplications between sparse matrices. Given the compile-time known structure of the structured sparse matrices, we have the opportunity to fully resolve the matrix construction at compile-time. Besides, to generate efficient matrix-free compute kernels, the matrix-vector multiplication and sparse linear solvers need to cooperate with such structured matrices at compile time. Given the multi-stage programming infrastructure, the first stage (running in Python) is responsible to deal with this. The following subsections will introduce our design behind, and the staging procedure will be introduced in Section 5.

4.2 Sparse Matrices as Row Functions

In `MAT2STENCIL`, a sparse matrix is represented as a row function, which takes a row index and produces non-zeros in the specified row. Each non-zero is a pair of a column index and the corresponding matrix value, in the same way as in the adjacency list. The returned non-zero list is then consumed by sparse linear algebra procedures that define the computation, decoupled from the matrix. Speaking in the multi-stage programming framework, the row index, non-zero column

Table 1. Examples of sparse matrices in MAT2STENCIL.

Predefined Matrices Examples		
Diff		<pre>@spmat(lambda n: n - 1) def diff(col_domain: CartesianDomain, i: IdxExpr): return [(i, -1), (i + 1, 1)]</pre>
Padding		<pre>@spmat(lambda n: n + 2) def padding(col_domain: CartesianDomain, i: IdxExpr): if i > 0 and i < domain.size(0) + 1: return [(i - 1, 1)] return []</pre>
User-defined Matrices Examples		
NPB 4th order dissipation		<pre>@spmat() def dssp4(col_domain: CartesianDomain, i: IdxExpr): n = col_domain.size(0) if i == 0: return [(i, 5), (i + 1, -4), (i + 2, 1)] elif i == 1: return [(i - 1, -4), (i, 6), (i + 1, -4), (i + 2, 1)] elif i < n - 2: return [(i - 2, 1), (i - 1, -4), (i, 6), (i + 1, -4), (i + 2, 1)] elif i == n - 2: return [(i - 2, 1), (i - 1, -4), (i, 6), (i + 1, -4)] elif i == n - 1: return [(i - 2, 1), (i - 1, -4), (i, 5)]</pre>
HPCG Prolong		<pre>@spmat(lambda nx, ny, nz: [nx * 2, ny * 2, nz * 2]) def prolong(col_domain: CartesianDomain, z, y, x): if z % 2 == 0 and y % 2 == 0 and x % 2 == 0: return [(z // 2, y // 2, x // 2, 1.)] else: return []</pre>

indices, and values are all staged expressions that will be resolved at runtime, while the non-zero list itself is a compile-time Python list.

Some examples are shown in Table 1. A domain is assigned to both the row and column indices, representing their valid range starting from 0. Thus, the row function also takes a parameter `col_domain`, which contains the information on matrix size and will be described in detail in Section 4.2.1. Decorated by `@spmat`, which is a Python decorator function we provide, the row function is wrapped into a `MAT2STENCIL` matrix object of type `SpMat`. The decorator also optionally takes a lambda expression to compute the matrix size relatively, the detail of which will be discussed in Section 4.2.2. Besides, through simple if-else sequences, the special boundary in many structured sparse matrices can be well expressed, and so can periodical sparsity patterns that occur in, for example, prolongation in multi-grid solvers.

Our row function design imitates the adjacency list, corresponding the non-zero cells in the matrix to their rows. This abstraction suits our applications well: an unknown variable in a discretized linear equation system like Equation (3) is related to a limited and compile-time known number of other unknowns, which are usually neighbors on the grid. Thus in the matrix representing the linear systems to be solved, the non-zeros in a row are in the columns of the neighbors. It makes the row function a natural design: the existing approach through general sparse linear algebra routines also requires users to generate sparse matrices with similar code, but for the more complex end-to-end matrices instead of our simple modular ones.

Yet, during implementing applications, we observe that expressing a conventional matrix is not sufficient to maximally ease user programming. We extend the sparse matrix in `MAT2STENCIL` in the following aspects and expose optimization opportunities.

4.2.1 Multi-Dimensional Rows and Columns. Representing the structured problems as sparse matrices requires flattening the multi-dimensional unknowns. It adds burdens on user programming to write the flattened index of the rows and columns. Besides, such flattened access is not friendly to our backend optimization: to address the arbitrary dependence Seidel-style stencils may bring to us, the polyhedral analysis need complete information about the original multi-dimension structure to achieve good performance.

As such, we extend the sparse matrix to higher dimensions on the rows and columns. Such a “matrix” with m -D row domain and n -D column domain is thus a (m, n) -tensor in the terminology of General Relativity; yet for convenience of understanding and to match the code, we will continue to note it as a “matrix”, or `SpMat`. A matrix object then has two properties, `row_domain` and `col_domain`. Each domain object has a tuple of integers for its grid size; in our current implementation, only regular Cartesian grids are considered, but the domain type can be extended easily. An m -D point in the row domain corresponds to a row in the matrix consisting of non-zeros, each with a n -D point in the column domain and the value of the non-zero in the matrix. Correspondingly, a previously flattened vector becomes a multi-dimensional `Vector` either, with respect to the multi-dimensional space that the physical quantities span over. A `domain` is similarly attached as a property to each `Vector` object.

For example, in the prolongation matrix shown in Table 1 we are already using the multi-dimensional row index passed into the row function and returning multi-dimensional column indices of non-zeros in the row. If, for example, consumed by the `SpMV` procedure in Fig. 7, the returned column indices are later used to index the multi-dimensional `Vector` being multiplied against, thus performing the prolongation step mapping the coarser grid to a finer grid.

4.2.2 Automatically Deriving Row Domain from Column. A structured operation usually applies to many different grids, including grids of different sizes and dimensions. We notice that a given column domain for a structured sparse matrix will decide its row domain. In specific, a `SpMat` is logically right-multiplied to `Vectors`, thus the column domain is always decided first. As such, we make the `row_domain` of an `MAT2STENCIL` matrix always following its `col_domain`. Walking into details, `SpMat` does not require specific row or column domains in user codes. At each staging site that we decide its concrete shape, its `col_domain` will be set through a Python property setter. The setter calls a virtual function `set_col_domain` to set the column domain and compute the row domain. The `@spmat` generated `SpMat` optionally receives a lambda expression as its argument to map its sizes of column domain to row, which is used to implement its `set_col_domain`.

4.2.3 Inner Value as a Tensor. Multiple (but a small fixed number of) quantities are often attached to the same grid in real-world applications. In previous matrix-based approaches, these quantities are flattened along with the grid dimensions; after retaining the grid dimensions, we then need

to allow the SpMat and Vector to have tensors, instead of plain scalars, as its inner values. They thus have a fixed `inner_shape`, which describes the shape of their inner tensor values, as an extra parameter in constructing a MAT2STENCIL sparse matrix.

4.3 Linear Algebra on MAT2STENCIL

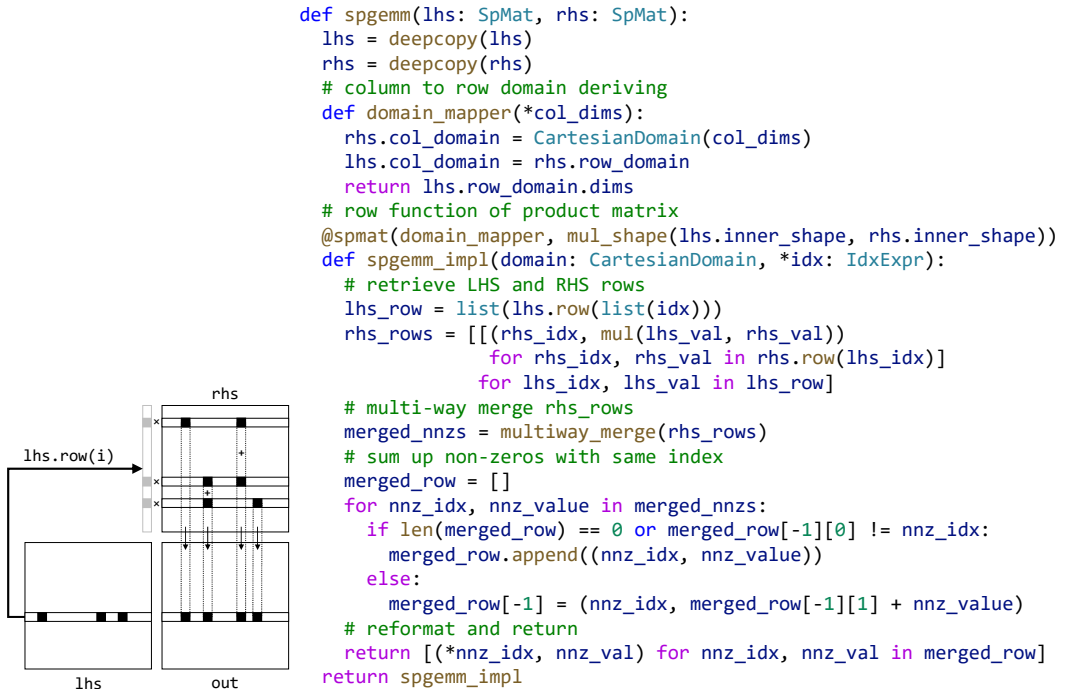
Through the programming interfaces of MAT2STENCIL, linear algebra routines can be easily implemented. Wrapped in a domain loop, such routines can retrieve lists of non-zeros in each row of MAT2STENCIL matrices, and perform computation according to their algorithm. Those sparse linear algebra routines are able and expected to be implemented in their simplest form, i.e. directly operating on the adjacency list, without the need for manual optimizations.

4.3.1 Sparse Matrix-Vector Multiplication (SpMV). The simplest example is SpMV in Fig. 7, multiplying a sparse matrix and a vector. It first adjusts the domains of the matrix and retrieves the expected domain of the product vector accordingly, the mechanism of which has been covered in Section 4.2.2.

Looping over the domain, it reads corresponding rows through the row method of the matrix object, which is a simple wrap of the above row function. It then multiplies the value from the matrix and vector for each non-zero and sums them up to get the final result, following the math of matrix-vector multiplication.

```
@ft.inline
def spmv(m: SpMat, v: Vector):
    m.col_domain = v.domain
    out = Vector(m.row_domain)
    for i in m.row_domain:
        out[i] = sum(nv * v[ni]
                    for ni, nv in m.row(i))
    return out
```

Fig. 7. Implementation of SpMV in MAT2STENCIL.



(a) Row-based algorithm.

(b) Implementation in MAT2STENCIL.

Fig. 6. Algorithm and implementation of SpGEMM.

4.3.2 Sparse Matrix Multiplication. With both operands expressed in `MAT2STENCIL`, it is possible to give their product still in `MAT2STENCIL` and being resolved all at compile stage through a row-wise `spgemm` as is shown in Fig. 6a.

Then, as shown in Fig. 6b, the row function of the resulting matrix first walks through the row of the specified index on the left-hand side (LHS), retrieving the non-zeros. It then uses the non-zeros' column indices to index the rows of the right-hand side (RHS) matrix. All the RHS rows are merged in order, multiplied with the LHS values accordingly, and summed up to one non-zero sequence, producing the result row. All this procedure happens in stage 1 and is thus expanded to plain matrix-free codes before stage 2.

Summing up and subtracting between `MAT2STENCIL` matrices is even simpler: they only require merging two rows from the two matrices, similar to combining the multiple rows from RHS in `spgemm`. We also implement them for `MAT2STENCIL` but do not include details here.

4.3.3 Enabling Spatial Dependence. These linear algebra routines can also contain spatial dependence carried by the domain loops, according to their different algorithms, enabling the implementation of more complicated routines required by PDE solvers. An example is in Fig. 8, in which the loops simultaneously read and write on the same `Vector` object and introduces matrix-defined dependence. The staging procedure will then produce their matrix-free form constructed of regular nested loops, without any indirect access typical in traditional sparse linear algebra, thanks to the staged domain loop indices (will be covered in Section 5). Our analysis and optimization mentioned in Section 6 will resolve the loop-carried dependence and exploit chances to parallelize.

To support our evaluated applications, we have implemented Incomplete Lower-Upper, Symmetric Successive Over-Relaxation, and Symmetric Gauss-Seidel, each with tens of lines of code. Users with science backgrounds will be able to implement their own routines in `MAT2STENCIL` in usually tens of lines, much easier than in traditional sparse algebra libraries.

4.4 Grammar of `MAT2STENCIL`

After describing our language design, we thereby demonstrate the grammar of the `MAT2STENCIL` language. As an embedded domain-specific language built upon type-based staging techniques, `MAT2STENCIL` involves various host language (Python) functions during the staging process, noted as `InType` \rightsquigarrow `OutType` and may have side effects. Many designs mentioned above work as Python components, and are thus invisible in the grammar, including the row function, structured sparse matrix objects, and linear algebra routines. Without hurting the core idea, it also omits details about “inner value as a tensor” from Section 4.2.3.

```
@ft.inline
def symgs(A: CartesianDomain,
          r: Vector, x: Vector):
    # forward sweep
    for i in x.domain:
        sx = r[i]
        diag = 0
        for nnz_idx, nnz_val in A.row(i):
            if nnz_idx == i:
                diag = nnz_val
            else:
                sx = sx + (-nnz_val) * x[nnz_idx]
        x[i] = sx / diag
    # backward sweep
    for i in reversed(x.domain):
        sx = r[i]
        diag = 0
        for nnz_idx, nnz_val in A.row(i):
            if nnz_idx == i:
                diag = nnz_val
            else:
                sx = sx + (-nnz_val) * x[nnz_idx]
        x[i] = sx / diag
```

Fig. 8. Example code for Symmetric Gauss-Seidel implemented against `MAT2STENCIL`.

$S_O ::=$	new_vector (domain : List[Int], vector : Name)	(creating vector on domain)
	new_scalar (scalar : Name)	(creating global scalar)
	store_scalar (scalar : Name, expr : E_O)	(store to global scalar)
	for (domain : List[Int], body : List[I] \rightsquigarrow S_I)	(runtime loop, dictionary order)
	ite_O (cond : B_O , then : S_O , else : Optional[S_O])	(runtime branch)
	$S_O; S_O$	(statements sequence)
$S_I ::=$	store (vector : Name, idx : List[I], expr : E_I)	(store an element to vector)
	reduce_add (scalar : Name, expr : E_I)	(reducing to global scalar)
	ite_I (cond : B_I , then : \rightsquigarrow S_I , else : Optional[\rightsquigarrow S_I])	(runtime branch)
	$S_I; S_I$	(statements sequence)
$E_O ::=$	Float	(lift from compile-time)
	$E_O + E_O$ $E_O - E_O$ $E_O \times E_O$ E_O / E_O ...	(floating point arithmetics)
	load_scalar (scalar : Name)	(load value of scalar)
$B_O ::=$	$E_O > E_O$ $E_O \geq E_O$ $E_O < E_O$ $E_O \leq E_O$	(floating point comparison)
	$B_O \wedge B_O$ $B_O \vee B_O$ $\neg B_O$	(boolean operations)
$E_I ::=$	E_O	(lift from global)
	$E_I + E_I$ $E_I - E_I$ $E_I \times E_I$ E_I / E_I ...	(floating point arithmetics)
	load (vector : Name, idx : List[I])	(load element from vector)
	I	lift from an integer index)
$I ::=$	Int $I + I$ $I - I$ $I * I$ $I \div I$ $I \bmod I$	(integer literal and arithmetics)
	Name	(index variable of domain loop)
$B_I ::=$	$I > I$ $I \geq I$ $I < I$ $I \leq I$ $I = I$ $I \neq I$	(integer comparison)
	$B_I \wedge B_I$ $B_I \vee B_I$ $\neg B_I$	(boolean operations)

Note that we support separate classes of statements and expressions for outside and inside domain loops. Domain loops cannot nest with each other; loading and storing the vectors only happens inside domain loops; and the branches inside domain loops, **ite_I** must be conditioned according to integer arithmetics of integer indices I ultimately introduced by domain loops. Global scalars are useful to aggregate values globally, e.g., dot products. The outer if-else-then **ite_O** is useful for iterative solver algorithms to control the workflow, e.g. check convergence. The aforementioned sparse linear algebra routines, invoked with our structured sparse matrices, are all taken place outside the domain loops and compose S_O nodes.

Then a sparse matrix will be a host language (Python) object of type:

```

spmat(domain_col2row : List[Int]  $\rightsquigarrow$  List[Int],
       row : (col_domain : List[Int], row_idx : List[ $I$ ])  $\rightsquigarrow$ 
       (eff :  $S_I$ , nnzs : List[(col_idx :  $I$ , value :  $E$ )]))

```

Which contains a function `domain_col2row`, mapping column domain to row domain, and a function `row`, for retrieving a row of non-zeros in the matrix symbolically.

5 MATRIX-FREE CODE GENERATION VIA STAGING

5.1 Customizable Control Flow Virtualization

Similar to LMS, our staging infrastructure works through control flow virtualization. With a decorator, it automatically transforms a function to its virtualized version: each for-loop and if-then-else branch is transformed into a corresponding virtual call, allowing the code generator to emit stage-2 code wrapping the loop or branch body. But unlike LMS, which directs all control flow to a central overload object defining the stage-2 language, we choose to make the virtual call point to the iterable or predicate expression itself to gain more customizability.

We demonstrate examples of the code transformation in Table 2. Walking into the details, if the iterable or predicate is of staged type (`StagedIterable` and `StagedPredicate` accordingly), the staging handler method of the staged object (`foreach` and `if_stmt` accordingly) is called to stage future codes. Otherwise, it automatically falls back to a stage-1 control flow. The logical operations, which are not overloadable per the design of Python, are virtualized similarly for `StagedPredicate`. This staging infrastructure has been merged into `FreeTensor` as the basis of its frontend.

It is natural to see that the domain loops iterating through all points in a domain have to be staged as loops in stage 2. Otherwise, the staged code size will bloat out as the problem grows, which is undesired since we will deal with large science problems. As such, we implement `StagedIterable` for the domain objects and emit the nested for-loops according to the domain dimensions in the staging handler. It also allows us to do more work in the staging handler behind the scenes to support the backtracking and discussing which we will introduce next.

5.2 Resolve Undecidable Branches through Backtracking and Discussing

As we have shown in Section 4.2, we want to enable taking branches with regard to row indices and manipulating stage-1 lists accordingly. Besides, in matrix multiplication, we need similar functionality to merge non-zero sequences according to the comparison results of column indices; more user programs also need such branches. They together raise a vital requirement on branching in stage 1 according to expressions related to the domain loop indices. Formally speaking, the

Table 2. Examples of code transformation for control flow virtualization. Standalone temporary function definitions are used in real implementation instead of lambda expressions to enable multi-line bodies.

User code	After control flow virtualization
<pre>for i in it: foo(i)</pre>	<pre>if isinstance(it, StagedIterable): it.foreach('i', lambda i: foo(i)) else: for i in it: foo(i)</pre>
pred1 and pred2	<pre>(pred1.logical_and(pred2) if isinstance(pred1, StagedPredicate) else (pred1 and pred2))</pre>
pred1 or pred2	<pre>(pred1.logical_or(pred2) if isinstance(pred1, StagedPredicate) else (pred1 or pred2))</pre>
not pred	<pre>(pred.logical_not() if isinstance(pred, StagedPredicate) else (not pred))</pre>
<pre>if pred: foo() else: bar()</pre>	<pre>if isinstance(pred, StagedPredicate): pred.if_stmt(lambda: foo(), lambda: bar()) else: if pred: foo() else: bar()</pre>

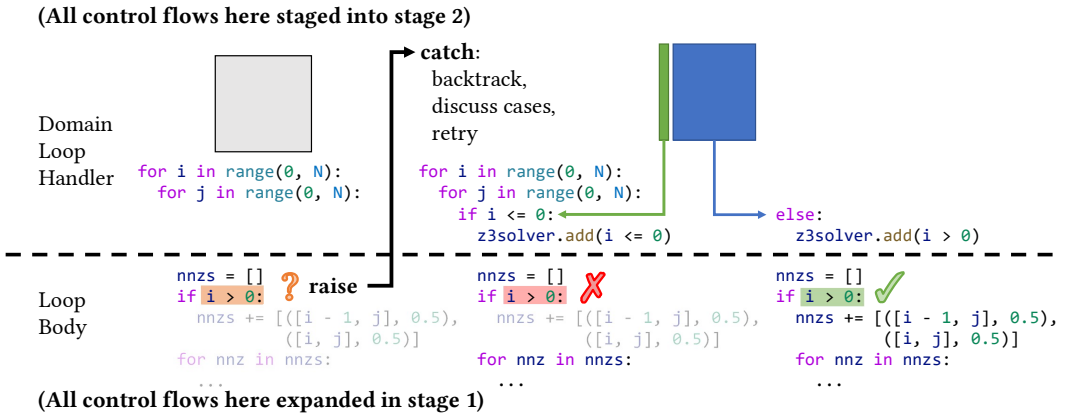


Fig. 9. The procedure of backtracking and discussing for resolving undecidable branches.

ite_I node shown in Section 4.4 accepts a runtime boolean expression B_I but coordinates with compile-time then/else bodies as Python functions, which is the time-reversing. However, notice that possible combinations of these branches are available at stage 1, we can split the original loop domain into pieces, making control flows deterministic in a single piece.

To achieve this, we backtrack on every undecidable branch and discuss them to be true or false. As is shown in Fig. 9, we backtrack through raising an exception when an undeterministic predicate is met in the ite_I if-else-then branch. Such branches must occur inside some domain loop, the stage 1 handler of which will catch the exception and retry staging the loop body twice, each with the predicate assumed to be true or false. The time-reversing branches thus become regular runtime branches. We implement the boolean expressions B_I as a sub-class of `StagedPredicate` and implement the checking and raising logic in its handler to achieve this.

It can be seen that as more branches are met, the number of cases easily grows exponentially with the number of branches. We observe that the predicates in `MAT2STENCIL` codes for structured solving turn out to be primarily duplicated; it is reasonable since the domain dimension is limited to no more than 3 in typical usages, and the branches are mostly related to boundaries or grid coarsening so the number of independent predicates on each dimension will be limited to only a handful. As such, we further utilize the Z3 [de Moura and Björner 2008] SMT solver to deduplicate the cases. Every time a predicate is taken to generate the loop body for its case, we append the assumption into the Z3 solver. Once a new predicate is met, we first check its truth value through Z3; if deterministic, we will not perform backtracking and instead directly take the branch. Applying such a scheme, we are able to generate clean case-discussing code for matrix-free algorithms.

Summarizing above, the backtrack and discuss algorithm is shown in Algorithm 1.

6 BACKEND OPTIMIZATION

In this section, we introduce our backend techniques to parallelize and optimize the generated matrix-free code. The most significant part is `SewTh`, a novel scheduling algorithm to efficiently parallelize arbitrary Seidel-style stencil operators. Unlike typical polyhedral schedulers, which linearly transform nested loops, optionally tile them, and selectively parallelize, `SewTh` additionally schedules fine-grained point-to-point synchronizations at certain iterations instead of performing global barrier synchronization. We also implement elementwise function inlining and loop fusion with linear transformation based on the `Pluto+` algorithm [Bondhugula et al. 2016].

Algorithm 1 Pseudo-code for the backtrack and discuss algorithm.

```

1: function  $S_O.FOR$ ( $domain : List[Int], body : List[I]$ )
2:   global  $solver \leftarrow Z3.Solver()$  ▷ create global solver
3:    $indices \leftarrow nil$  ▷ create list of loop indices
4:   for  $n \leftarrow domain$  do
5:      $i \leftarrow$  create new Name as  $I$ 
6:      $solver.add(0 \leq i < n)$  ▷ append constraint for dimension range
7:      $indices \leftarrow indices :: i$  ▷ append loop index of this dimension
8:   function RECURSINGBACKTRACK
9:     try
10:      return  $body(indices)$ 
11:     catch  $Undeterminant(cond)$ 
12:        $solver.push(cond)$  ▷ discuss with  $cond$  holds
13:        $true\_case \leftarrow RECURSINGBACKTRACK()$ 
14:        $solver.pop()$ 
15:        $solver.push(\neg cond)$  ▷ discuss with  $cond$  not holds
16:        $false\_case \leftarrow RECURSINGBACKTRACK()$ 
17:        $solver.pop()$ 
18:       return  $FT.ite(cond, true\_case, false\_case)$  ▷ assemble the wrapping if-else-then
19:
20:    $code \leftarrow RECURSINGBACKTRACK()$ 
21:   for  $i \leftarrow len(domain) - 1..0$  do ▷ assemble the domain for loops
22:      $code \leftarrow FT.for(indices[i], FT.slice(domain[i]), code)$ 
23:   return  $code$ 
24:
25: function  $S_I.ITE_I$ ( $cond : B_I, then : \rightsquigarrow S_I, else : Optional[\rightsquigarrow S_I]$ )
26:   global  $solver$  ▷ retrieve the global solver
27:   if  $solver.prove(cond)$  then ▷ use  $then$  when always true
28:     return  $then()$ 
29:   else if  $solver.prove(\neg cond)$  then ▷ use  $else$  if exists when always false
30:     return if  $else$  then  $else()$  else  $FT.nop$ 
31:   else ▷ undeterministic, raise exception to discuss the cases
32:     raise  $Undeterminant(cond)$ 

```

6.1 Sewing the Threads: A New Approach to Parallelize Seidel-Style Stencils

As is mentioned in Section 2.3, automatic parallelization of Seidel-style Stencils has not been explored extensively by existing research. To achieve both good memory locality and low overhead starting, our core idea is to assign part of the loop to each thread and perform fine-grained synchronization to satisfy the dependence across threads. Unlike the existing polyhedral code generation approaches synchronizing with global barriers, we generate point-to-point synchronizations based on single-producer, single-consumer (SPSC) lock-free queues. They are thus on-demand and introduce significantly lower overhead. With the large blocks assigned to each thread, it is rare to perform the synchronization. Such fine-grained synchronization sews the threads together, guaranteeing the correctness of produced code.

We cover the details of the work decomposition and what to synchronize, how we compute necessary synchronization from analyzed dependence and generate code for them, and how we

```

double x[14][14], r[12][12];
for (int iy = 0; iy < 12; iy++)
  for (int ix = 0; ix < 12; ix++) {
    x[iy+1][ix+1] = (r[iy][ix] +
      x[iy+0][ix+0] + x[iy+0][ix+1] + x[iy+0][ix+2] +
      x[iy+1][ix+0] +           + x[iy+1][ix+2] +
      x[iy+2][ix+0] + x[iy+2][ix+1] + x[iy+2][ix+2]) / 8.;
  }

```

(a) Input code of a 2D 9-point Gauss-Seidel.

```

double x[14][14], r[12][12];
#pragma omp parallel for
for (int thr_ix = 0; thr_ix < 4; thr_ix++)
  for (int iy = 0; iy < 12; iy++)
    for (int ix = 3 * thr_ix; ix < 3 * thr_ix + 3; ix++) {
      if ((ix == 2 + 3 * thr_ix && thr_ix <= 2 && iy > 0) ||
        (ix == 3 * thr_ix && thr_ix > 0))
        waitThread(thr_ix, -2 * thr_ix + ix - 1);
      x[iy+1][ix+1] = (r[iy][ix] +
        x[iy+0][ix+0] + x[iy+0][ix+1] + x[iy+0][ix+2] +
        x[iy+1][ix+0] +           + x[iy+1][ix+2] +
        x[iy+2][ix+0] + x[iy+2][ix+1] + x[iy+2][ix+2]) / 8.;
      if ((ix == 3 * thr_ix && thr_ix > 0 && iy <= 10) ||
        (ix == 2 + 3 * thr_ix && thr_ix <= 2))
        notifyThread(thr_ix, -2 * thr_ix + ix - 1);
    }

```

(b) Generated code by SewTh.


(c) The enforced partitioning schedule. The sequences of different colored arrows are in sequential order within threads. The arrows crossing thread boundaries are the **crossing dependence**. Dark arrows are the **compacted crossing dependence**, which are the only physically synchronized dependence.

Fig. 10. An example of SewTh that parallelizes a 2D 9-point Gauss-Seidel operator. Equivalent C++ code is presented instead of FreeTensor IR for better readability.

automatically prove the parallelism of our plan. We will use the 2D 9-point Gauss-Seidel presented in Fig. 10a as the example throughout this subsection. And we will continue to abbreviate $[u_1, u_2, \dots, u_n]$ as bold \mathbf{u} .

6.1.1 Enforced Partitioning to Threads. To exploit parallelism from the multidimensional loops with spatial dependence, we first assign certain partitioning to the loop space. We partition as few dimensions as possible in all inner loops without any skewing, while leaving the outermost loop and possibly some inner loops not partitioned. Assume the dimension sizes of the outer domain loops are described by N . Organizing the threads into a mesh, their partition sizes are N^t , of length 1 less than N , selected so that components of N^t are close to each other. Then we

enforce the function mapping an iteration \mathbf{x} to a thread to be $\mathbf{T}(\mathbf{x}) = [x_1 \div N_0^t, \dots, x_{i+1} \div N_i^t, \dots]$, in which $\mathbf{x} \in L = \{\mathbf{x} \mid \forall i, 0 \leq x_i < N_i\}$ (the looping range.) Within each thread, it still takes the lexicographical execution order, so that the intra-thread dependence naturally holds. An example is shown in Fig. 10c.

It is then only the dependence crossing thread boundaries, or *crossing dependence*, remained to be considered. The crossing dependence can be computed from the full dependence set, which is extracted from the program with the help of polyhedral dependence analysis and Presburger arithmetics. To start with, the full dependence set computed by FreeTensor dependence analysis is noted as $D = \{(\mathbf{a} \in L, \mathbf{b} \in L) \mid \mathbf{b} \text{ depends on } \mathbf{a}\}$. For all $(\mathbf{a}, \mathbf{b}) \in D$, \mathbf{b} must be computed later than \mathbf{a} in any valid schedule. The formal definition of crossing dependence is shown in Definition 6.1, which is also the Presburger arithmetics to automatically compute it.

Definition 6.1. The crossing dependence, in the form of a Presburger map from thread pair to dependencies, is computed as below:

$$D_{\times}(\mathbf{T}_a, \mathbf{T}_b) = \{(\mathbf{a}, \mathbf{b}) \in D \mid \mathbf{T}(\mathbf{a}) = \mathbf{T}_a \wedge \mathbf{T}(\mathbf{b}) = \mathbf{T}_b \wedge \mathbf{T}_a \neq \mathbf{T}_b\}.$$

Synchronizing the crossing dependence guarantees two depending iterations belonging to different threads to be executed in order. It is obvious that such synchronization will not cause deadlock: both the intra-thread execution order and inter-thread synchronization are following the lexicographical order of the multi-dimensional loop indices (also the execution order for sequential input program), making the combined synchronization graph acyclic.

6.1.2 Compact Dependence and Generate Code. We have addressed above what dependence should be covered by inter-thread synchronization. Yet, directly generating synchronizing code according to the crossing dependence has multiple downsides. Firstly, simply putting one point-to-point synchronization for each crossing dependence works fine, but introduces redundant synchronization. Besides, the possibly differently ordered sources and targets between a thread pair require the generated code to accurately synchronize against specific loop iterations instead of simply synchronizing against the thread. Taking the case in Fig. 10c, we attempt to avoid the explicit synchronization on the tinted crossing dependence.

We propose to compact the crossing dependence as follows. A crossing dependence is considered to be redundant if there is another crossing dependence between the same thread pair with a later or same source and an earlier or same target. Non-redundant crossing dependence is retained during compaction, i.e. compacted crossing dependence. A visual illustration is shown in Fig. 11: since the sequential order guarantees \mathbf{a}' happens later than \mathbf{a} , \mathbf{b} later than \mathbf{b}' , and another dependence $\mathbf{a}' \rightarrow \mathbf{b}'$ exists, $\mathbf{a} \rightarrow \mathbf{b}$ is marked redundant and not synchronized explicitly. The formal definition (which is also the algorithm) is listed in Definition 6.2.

Definition 6.2. The redundant crossing dependence is computed as:

$$D_{\times}^{-}(\mathbf{T}_a, \mathbf{T}_b) = \{(\mathbf{a}, \mathbf{b}) \in D_{\times}(\mathbf{T}_a, \mathbf{T}_b) \mid \exists(\mathbf{a}', \mathbf{b}') \in D_{\times}(\mathbf{T}_a, \mathbf{T}_b), (\mathbf{a} < \mathbf{a}' \wedge \mathbf{b}' \leq \mathbf{b}) \vee (\mathbf{a} = \mathbf{a}' \wedge \mathbf{b}' < \mathbf{b})\},$$

in which the comparisons are carried out in dictionary order. Then the compacted crossing dependencies are:

$$D_{\times}^{+}(\mathbf{T}_a, \mathbf{T}_b) = D_{\times}(\mathbf{T}_a, \mathbf{T}_b) - D_{\times}^{-}(\mathbf{T}_a, \mathbf{T}_b).$$

Then we conclude the following properties we want:




Fig. 11. Compact crossing dependence.

THEOREM 6.3 (COMPLETENESS). *With the sequential execution order of each thread and the compacted crossing dependence satisfied by point-to-point synchronization, all crossing dependence are guaranteed to hold. Formally speaking, we have*

$$\forall (\mathbf{a}, \mathbf{b}) \in D_{\times}(\mathbf{T}_a, \mathbf{T}_b), \exists (\mathbf{a}', \mathbf{b}') \in D_{\times}^+(\mathbf{T}_a, \mathbf{T}_b), \mathbf{a} \leq \mathbf{a}' \wedge \mathbf{b}' \leq \mathbf{b}.$$

PROOF. We will prove this theorem by contradiction.

Suppose for some $(\mathbf{a}_0, \mathbf{b}_0) \in D_{\times}(\mathbf{T}_a, \mathbf{T}_b), \forall (\mathbf{a}', \mathbf{b}') \in D_{\times}^+(\mathbf{T}_a, \mathbf{T}_b), \mathbf{a}_0 > \mathbf{a}' \vee \mathbf{b}' > \mathbf{b}_0$.

If $(\mathbf{a}_0, \mathbf{b}_0) \in D_{\times}^+(\mathbf{T}_a, \mathbf{T}_b)$, then taking $(\mathbf{a}', \mathbf{b}') = (\mathbf{a}_0, \mathbf{b}_0)$ contradicts the assumption.

Otherwise, according to Definition 6.2, we have $(\mathbf{a}_0, \mathbf{b}_0) \in D_{\times}^-(\mathbf{T}_a, \mathbf{T}_b)$. Thus,

$$\exists (\mathbf{a}_1, \mathbf{b}_1) \in D_{\times}(\mathbf{T}_a, \mathbf{T}_b), (\mathbf{a}_0, \mathbf{b}_0) \neq (\mathbf{a}_1, \mathbf{b}_1) \wedge \mathbf{a}_0 \leq \mathbf{a}_1 \wedge \mathbf{b}_1 \leq \mathbf{b}_0.$$

We can then repeat the process and get a sequence of different dependence $(\mathbf{a}_2, \mathbf{b}_2), (\mathbf{a}_3, \mathbf{b}_3), \dots$, until reaching some $(\mathbf{a}_n, \mathbf{b}_n) \in D_{\times}^+(\mathbf{T}_a, \mathbf{T}_b)$. The dependence set is finite, so we can always drain the $D_{\times}^-(\mathbf{T}_a, \mathbf{T}_b)$ and arrive at $D_{\times}^+(\mathbf{T}_a, \mathbf{T}_b)$, resulting in the contradiction above. \square

THEOREM 6.4 (CAN BLINDLY SYNCHRONIZE). *For every thread pair, the sources and targets of the compacted crossing dependence are in the same order, so blindly synchronizing against the thread will be equivalent to synchronizing against the accurate iteration. Formally speaking, we have*

$$\forall (\mathbf{a}_1, \mathbf{b}_1) \text{ and } (\mathbf{a}_2, \mathbf{b}_2) \in D_{\times}^+(\mathbf{T}_a, \mathbf{T}_b), \mathbf{a}_1 < \mathbf{a}_2 \Leftrightarrow \mathbf{b}_1 < \mathbf{b}_2.$$

PROOF. We will prove this theorem by contradiction. Suppose there exists $(\mathbf{a}_1, \mathbf{b}_1), (\mathbf{a}_2, \mathbf{b}_2)$ in $D_{\times}^+(\mathbf{T}_a, \mathbf{T}_b)$, that $\mathbf{a}_1 < \mathbf{a}_2 \Leftrightarrow \mathbf{b}_1 < \mathbf{b}_2$ does not hold. Then according to Definition 6.2 we have

$$\begin{aligned} & (\mathbf{a}_2 \leq \mathbf{a}_1 \wedge \mathbf{b}_1 < \mathbf{b}_2) \vee (\mathbf{a}_1 < \mathbf{a}_2 \wedge \mathbf{b}_2 \leq \mathbf{b}_1) \\ \Rightarrow & (\mathbf{a}_2, \mathbf{b}_2) \in D_{\times}^-(\mathbf{T}_a, \mathbf{T}_b) \vee (\mathbf{a}_1, \mathbf{b}_1) \in D_{\times}^-(\mathbf{T}_a, \mathbf{T}_b). \end{aligned}$$

Since $D_{\times}^+(\mathbf{T}_a, \mathbf{T}_b) \cap D_{\times}^-(\mathbf{T}_a, \mathbf{T}_b) = \emptyset$, this contradicts the assumption. \square

Then the compacted crossing dependence set between each thread pair can be computed as a quasi-affine map with the help of ILP libraries, ISL [Verdoolaege 2010] in our case, allowing us to generate code against it. A “wait” block is injected before the loop body so that it waits for each source thread that the current loop iteration depends on, and similarly, a “notify” block but inversed after the loop body. According to Theorem 6.3, the compaction ensures that only synchronizing the compacted dependence will be sufficient. Also, according to Theorem 6.4 the synchronization can ignore which loop iteration the other end is and blindly synchronize with the specific thread. We are thus possible to synchronize simply through efficient atomic counters, one for each thread pair. A code generation example is shown in Fig. 10b.

6.1.3 Modeling the Parallelism. The above work decomposition and synchronization together do not necessarily imply sufficient parallelism. It is possible that most threads are busy waiting and only some of the threads are working all the time under inappropriate parameters. We here propose a model for the parallelism lower bound of the generated code.

Existing methods have been utilizing *hyperplanes* to determine the time axis that enables parallelism. Recall the parallelizable bands in Fig. 3, these are an example of a valid setting of parallelizable hyperplanes: since the nested loop has 2 dimensions, the hyperplanes become 1D lines. Inspired by those methods, we design a set of thread partitioning hyperplanes corresponding to the enforced partitioning proposed above. Say we have the loop dimensions




Fig. 12. A plane from eq. (6).

$\mathbf{x} = [x_0, x_1, \dots, x_n]$, $0 \leq x_i < N_i$, and their inner partition sizes $[N_1^t, N_2^t, \dots, N_n^t]$. Then we define the thread partitioning hyperplanes as

$$\phi_t(\mathbf{x}) = \sum_{i=0}^n \prod_{j=i+1}^n N_j^t x_i = c. \quad (6)$$

Such a partitioning hyperplane ensures contiguous execution on each thread, following the thread-local lexicographic execution order as specified in Section 6.1.1. A visual example is shown in Fig. 12. Each cube cell represents a loop iteration and chessboard-colored ranges are thread partitions. The yellow cells are the loop iterations executed in parallel if the plane satisfies all loop dependence.

If such a hyperplane setting satisfies all dependence, then all loop iterations in the same hyperplane and assigned to different threads will be parallelizable. Formally speaking, we check whether for any dependence pair $(\mathbf{a}, \mathbf{b}) \in D$, $\phi_t(\mathbf{a}) < \phi_t(\mathbf{b})$, as a sufficient condition for the parallelism. If the check passed, the hyperplanes will reveal a possible parallel execution plan, and serve as a lower bound for the parallelism of our generated parallel implementation, which is already fully parallelized except at the short pipeline start and exit.

Take the two dimensional examples in Fig. 13, in which $N_1^t = 3$, so that $\phi_t([y, x]) = 3y + x$. If the thread partitioning hyperplanes successfully satisfy all dependence as in Fig. 13a, our plan is guaranteed to provide near-optimal parallelism (fully parallel except at the beginning and the end). Instead in Fig. 13b, the thread partitioning hyperplanes do not satisfy all dependence: for the brown dependence, $\phi_t([y + 1, x - 3]) = \phi_t([y, x])$; in such circumstances, our plan will yield degraded parallelism.

There is a great chance that the loop dependence in a seidel-style stencil operator satisfies this condition. Given a large grid and usually tens of threads, the thread partitions will be much larger than the stencil pattern, the size of which is determined by the order of partial differential equation and usually a small constant. The hyperplanes can then cover all iterations in the near neighborhood and earlier than the current iteration, satisfying all possible dependence. For example, in Fig. 13b, the hyperplanes fail only after the long dependence is added. The dependence is too long such that it spreads over an entire thread partition. With larger loop spaces and stencils still in such sizes, it will be impossible. Actually, in all evaluation cases requiring Seidel-style stencils in Section 7 including micro-kernels and end-to-end solvers, the thread partitioning hyperplanes satisfy all dependence in every kernel.

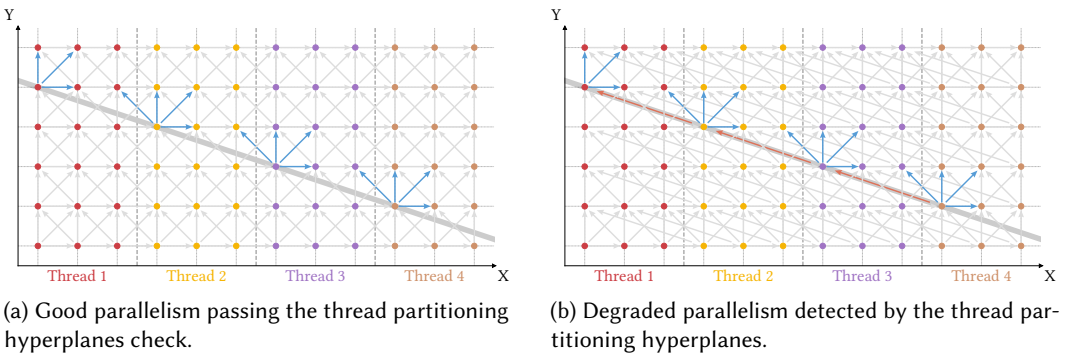


Fig. 13. Two examples for checking parallelism with the thread partitioning hyperplanes.

6.2 Memory Optimizations

In addition to basically parallelizing everything, including Jacobi and Seidel-style stencils, we introduce memory optimizations as below. While simple to work with, these optimizations are critical to many real-world applications, since the stencils involved in PDE solving are usually memory-bound, making reducing memory traffic important.

6.2.1 Aggressive Inlining of Pointwise Operations. In Section 4.1 we have mentioned elementwise functions, which allow the user to easily build pointwise operations. Using such components will create intermediate variables that can be easily inlined and eliminated. We perform such inlining whenever possible so that the memory traffic is minimized. The inlining is automatically accomplished by recognizing all elementwise domain loops and calling the FreeTensor schedule inline.

6.2.2 Loop Fusion with Affine Transformation. Loop fusion plays a critical role in memory intensive compute kernels, which is usually the case for the stencils in PDE solving. While previous DSLs can perform loop fusion only based on operator-level dependence information, our added support for customizable spatial loop dependence prohibits such simple approaches. Besides, the modern multigrid methods involve finer and coarser grids, and the loops on them can only be fused with strides. As such, we implement automatic loop fusion also based on polyhedral dependence analysis, which iteratively selects one axis from each of the two adjacent loop nests, performs affine transformation computed similar to the PLUTO+ [Bondhugula et al. 2016] algorithm, and fuses the two loops; the procedure is repeated until no level of loops can be fused in the loop nests. At the program level, the loop nests are iteratively fused, scanned in the forward order, and only fuse loops that do not sacrifice parallelism after fusion. The implementation is accomplished by ourselves instead of invoking PLUTO+ directly.

We do not perform tiling additionally since SewTh already exposes a tiled execution order. Other than that, we also automatically vectorize sequential loops with best efforts through the corresponding FreeTensor schedule `vectorize`; FreeTensor determines if a loop is vectorizable also through polyhedral analysis.

7 EVALUATION

In this section, we evaluate our domain-specific language design and performance through end-to-end benchmarking and backend optimization assessment. We set up our experiments on a dual-socket server, each socket being one Intel Xeon Gold 6126 CPU, with 12 physical cores or 24 logical cores per CPU, counting 24 physical or 48 logical together. The server has 565 GB of DDR4-2666 main memory. The server runs Ubuntu 18.04.4 LTS with Linux kernel 4.15.0. We use GCC 12.2.0 to compile the programs, using `g++` for C++ and `gfortran` for FORTRAN, including our generated codes and baseline manual implementations. All parallelism is through OpenMP, including reference codes and our implementation.

7.1 End-to-End Benchmarking

We evaluate MAT2STENCIL against official manual implementations of all the three pseudo-applications in the NAS Parallel Benchmark (NPB) [Bailey et al. 1991], BT, SP, and LU, which are three solvers of a Computational Fluid Dynamics (CFD) problem, and a matrix-free variant of High Performance Conjugate Gradient (HPCG) [Dongarra et al. 2016], which is a widely adopted performance benchmark solving Poisson equation with a multigrid solver.

We use NPB 3.4.2 as the baseline for the NPB evaluation and a hand-written matrix-free HPCG from [Zhu et al. 2021] for the HPCG evaluation. Existing stencil DSLs are not capable of expressing

these solvers due to the presence of Seidel-style stencils in them, thus we do not include such baselines. We neither compare with the reference HPCG implementation, given it is matrix-based according to the specification and such comparison will be unfair². For NPB, we run the Class B, C, D, and E problems (grid size 102^3 , 162^3 , 408^3 , and 1020^3) on all three pseudo-applications, the sizes of which are specified in the NPB Specification. For HPCG, we run on grids of size 768^3 , 1024^3 , 1280^3 , and 1536^3 . We show the computed grid points per second, which is the grid size divided by time consumption.




Fig. 14. Performance results on NAS Parallel Benchmarks (NPB-BT, NPB-SP, NPB-LU), and HPCG. **Higher is better.**

The evaluation result is shown in Fig. 14. In the three pseudo-applications of NPB, we achieve an average (geometric mean) of 1.25 \times , 0.94 \times , and 0.93 \times performance, compared to the manual matrix-free implementation, respectively. In HPCG, we achieve an average (geometric mean) of 1.04 \times performance. The performance is competitive at all scales. We also observe that larger grids yield relatively better performance in our implementation, possibly due to that the Python-native interface overhead is less significant at a larger scale.

7.1.1 Comparing the Code Length. We evaluate the human effort required to program matrix-free solvers manually or using our MAT2STENCIL by comparing the lines of code for different implementations in Table 3. MAT2STENCIL implementation takes 1044 lines of Python code for the majority of work and 329 lines of C++ code as a FreeTensor plugin implementing some auxiliary passes. Based on that, we implement all three NPB pseudo-applications in a total of 750 lines of code, with 485 lines among them to be shared across applications. It takes only 6.3% of the total code length compared to the total of three NAS-provided manually matrix-free implementations. Notably, even considering the lines of code for our compiler, we still involve significantly less code. In addition, the matrix-free HPCG implementation through MAT2STENCIL takes 16.4% of the code length in hand-written matrix-free implementation. MAT2STENCIL significantly releases the burden of programming efficient matrix-free solvers.

7.1.2 Evaluating Compile time. To demonstrate the compilation cost of MAT2STENCIL, we compare the end-to-end cost for all four evaluated workloads on their largest problems. The time breakdown is shown in Table 4. We observe that the stage-1 time, including the backtracking and discussing, takes a minor part of the compilation time. The majority of compile time is in the backend optimizations, which is reasonable since polyhedral schedulers are known to be costly. However, given larger and more long-running problems than the benchmarks, the compilation will take a smaller

²We test a 48-process reference HPCG run on a total of 768^3 sized grid, which is $192 \times 192 \times 256$ local grid on each process, because the reference HPCG code does not support multithreaded SymGS. It is worth noting that although it achieves full parallelization through domain decomposition, it is at the cost of worse convergence. Even so, it takes 355.8 seconds per run (50 iterations), which is significantly slower than 64.2 seconds in MAT2STENCIL, being 5.5 \times longer.

Table 3. Lines of code in evaluated cases under different implementations.

App. Impl.	NPB-BT	NPB-SP	NPB-LU	HPCG
Manual	4288	3079	3068	914
MAT2STENCIL	$408^\dagger + 24$	$408^\dagger + 188^\ddagger$	$408^\dagger + 41$	150

[†] Common codes for all NPB pseudo-applications; mostly elementwise functions with long math formulas.

[‡] Longer than others due to its unique extra math formulas, mentioned in Section 4.1.

fraction of the total time even if the application runs only once, making MAT2STENCIL practical for real applications.

Table 4. Comparing the compile and running time (in seconds) of MAT2STENCIL-based programs. The running times listed are at the maximum tested scale.

	Stages	NPB-BT	NPB-SP	NPB-LU	HPCG
Compile	Stage 1 Execution	1.27	5.52	1.51	0.67
	Stage 2 Optimize	280.89	1072.42	579.42	83.49
	Stage 2 GCC	22.63	23.80	24.98	33.77
	Run (Stage 2)	10,414.35	8479.75	6753.30	539.53

7.2 Assessing Backend Optimizations

7.2.1 Comparing with Other Polyhedral Compilers. While our targeted implicit solvers with Seidel-style stencils are not covered by existing stencil DSLs, the polyhedral compilers naturally support such computation and can reasonably parallelize them. We thus compare MAT2STENCIL with PPCG [Verdoolaege and Janssens 2017] and PLUTO [Bondhugula et al. 2008] on microbenchmarks to demonstrate the extra effectiveness of our optimization pipeline in such circumstances. The sequential C++ codes fed into PPCG and PLUTO are generated by the MAT2STENCIL frontend, without going through our optimizations. PPCG is invoked with outer-coincidence enforced³, which is vital for it to parallelize the test cases with loop-carried dependence. We enable diamond tiling in PLUTO to get the best performance on it. We use PLUTO but not PLUTO+ because PLUTO is better maintained and PLUTO+, as an improving fork of PLUTO, focuses on enabling negative coefficients in the linear transformations, which is not useful in our microbenchmarks and will not bring performance improvement. The microbenchmarks first include symmetric Gauss-Seidel (SymGS) and Incomplete LU (ILU) preconditioners with a 3D 27-point stencil on a 1536^3 grid. We also evaluate multiple Jacobi iterations (JacIter) with a 2D 9-point stencil on a 4096^2 grid, which only involves temporal dependence and is well-optimized by previous stencil optimizations.

The result is shown in Fig. 15. PPCG consistently performs worse than MAT2STENCIL, even on Jacobi iterations; it is reasonable since PPCG only performs linear transformations. PLUTO times out during transpiling the ILU code, taking more than one hour. Other than that, MAT2STENCIL outperforms PLUTO by $3.14\times$ on SymGS, demonstrating the effectiveness of our SewTh parallelizer. Yet, MAT2STENCIL achieves only 41.4% performance of PLUTO on JacIter, showing that the diamond tiling gives PLUTO outstanding performance on JacIter, but is not available for the Seidel-style stencils. However, such temporal tiling methods will not help in complicated implicit solvers, since too many loops with their own spatial dependence are attending a single timestep.

³Through the command-line flag `-isl-schedule-outer-coincidence`.




Fig. 15. Comparing with PPCG and PLUTO. The performance percentages achieved by PPCG and PLUTO compared to our MAT2STENCIL are presented alongside. PLUTO times out during compiling ILU. **Lower is better.**

7.2.2 Ablation Study on Memory Optimizations. To analyze the performance gains of the memory optimizations, we conduct ablation tests on previous benchmarking cases: CLASS D for NPB pseudo-applications, and grid size 1024^3 for HPCG. We start from SewTh-only and incrementally introduce inlining and loop fusion. Together with the manually implemented baseline, the performance numbers are shown in Fig. 16.




Fig. 16. Results of ablation study on memory optimizations. **ST** is SewTh-only, **+inl.** adds inlining for elementwise ops, **+fuse** then adds loop fusion. **Man.** is manually optimized implementation. **Lower is better.**

It can be seen that both inlining and loop fusion helps a lot in NPB pseudo-applications. The nonlinearity in their solved Navier-Stokes equation results in a number of elementwise functions and Jacobi-style stencil operators, which will introduce noticeable memory traffic and footprint if not inlined. Besides, for BT and SP, we observe that the direct solving happens in the three distinct axes, each sweeping forward and backward. Fusing the forward and backward sweeping loops greatly reduces the size of intermediate tensors by removing two unrelated dimensions, thus significantly improving the performance; the same optimization is manually implemented in the baseline. Such loop fusion, requiring permuting and reverting the nested loops, is only possible to be automated with a smart enough loop fuser that can select axes and perform affine transforms automatically.

Instead, SewTh-parallelized HPCG already yields good performance matching the manual implementation, and further optimizations help little. The reason is that HPCG solves a simple Poisson equation through the complicated multigrid conjugate gradient with the SymGS preconditioner and thus ❶ requires no elementwise functions, and ❷ incurs most of the time cost in SymGS.

8 DATA-AVAILABILITY STATEMENT

The implementation of MAT2STENCIL and baselines corresponded to the evaluation results are available at [Cao 2023] together with the running scripts. No external dataset is used in the evaluation.

MAT2STENCIL will also be open-sourced at <https://github.com/thu-pacman/Mat2Stencil>.

9 RELATED WORK

Solving Differential Equations through Matrix-based Approaches. Numerically solving differential equations has been widely treated as sparse linear algebra computation. Sparse matrix-vector multiplication (SpMV) [Merrill and Garland 2016; Pinar and Heath 1999; Williams et al. 2010] are well-studied, covering explicit solving and Jacobi iterations. Other routines are also available in sparse linear algebra, e.g. SuperLU contains a general-purpose efficient ILU implementation [Li and Shao 2011]. SciPy [Virtanen et al. 2020] is a commonly used Python library for scientific computing and covers many such routines.

While those approaches decouple the equation and solvers, they are too general to be highly optimized against structured grids. Borrowing their idea of representing differential equation solving in sparse linear algebra, we propose our MAT2STENCIL based on matrix abstraction, enabling easy algorithm programming while persisting the performance benefits of matrix-free codes. It is also worth mentioning that [Augustine et al. 2019] takes a totally different approach to optimize a solver already using sparse linear routines, by automatically analyzing the pattern of the sparse matrix and reconstructing matrix-free codes with the best effort.

Solving Differential Equations with Matrix-free Approaches. The best-studied area of matrix-free approaches is Jacobi-style stencil computations, including with or without time iterations [Habich et al. 2009; Kamil et al. 2006; Nguyen et al. 2010; Williams et al. 2007]. Early automatic optimization techniques include [Datta et al. 2008; Krishnamoorthy et al. 2007], which work on manually written nested loops. In order to provide easier interfaces for programming stencils, researchers have come up with many stencil libraries and DSLs, including [Huang et al. 2019; Lengauer et al. 2020; Louboutin et al. 2019; Pieper et al. 2021; Tang et al. 2011; Zhang et al. 2017]. Among those, Devito [Louboutin et al. 2019] and ExaStencils [Lengauer et al. 2020] present a set of program representations that is pretty much close to the math formula of the differential equation, being the easiest ones in programming.

However, none of these works fully enables programming implicit solvers, including direct solving of multi-diagonal matrices, (Symmetric) Gauss-Seidel iteration, (Symmetric) Successive Under/Over-Relaxation, Incomplete Lower-Upper [Saad 2003], Fine-Grained Parallel Incomplete Lower-Upper [Chow and Patel 2015], etc. The great diversity of such algorithms makes it hard for DSLs to cover all of them. Matrix-free approaches on implicit solvers have been requiring manual optimizations, for example, in [Zhu et al. 2021] and [Yang et al. 2016].

[Essadki et al. 2023] introduces automatic code generation for “in-place stencils”, which is noted as Seidel-style stencils in this paper and helps the construction of implicit solvers, but not includes advanced high-level abstractions required for easier programming of implicit solvers. Their methodology will also partially degenerate to existing polyhedral approaches for some stencil shapes, e.g. 9 point on 2-D or 27 point on 3-D, in which case our SewTh will perform significantly better. We cannot compare with their work in evaluation due to absence of their implementation.

Multi-stage Programming. Our embedded DSL implementation is inspired by Lightweight Modular Staging (LMS) [Rompf and Odersky 2012], which represents multi-staged programming through types in Scala. We adopt their type-based design to achieve user-transparent staging. There has also been a Python implementation for it named Snek [Decker 2019] that is similar to our staging infrastructure in design, but it does not perfectly feed our needs, thus we implement our own with small tweaks to make the control flow virtualization customizable for different staged types. Other approaches, e.g. the quote-splice in Scala 3.0 [Stucki et al. 2018] and lift-annotated MetaML [Taha 1999], will not fit us well since we do not expect DSL users to get in touch with the staging details.

Tensor Programs. Our backend generates generalized stencils containing quasi-affine loops that manipulate multi-dimensional arrays, or tensors. Our optimizations originate from various previous works on these types of programs. Many studies have been taken on optimizing tensor programs since the invention of FORTRAN. Compiling optimization techniques, especially *polyhedral compilation*, targeting general tensor programs, can be applied automatically [Allen and Kennedy 2001; Bondhugula et al. 2016, 2017, 2008; Fautrier 1992; Padua and Wolfe 1986; Verdoolaege and Janssens 2017], but the resulting performance is far from optimal for the important Seidel-style stencils. Nevertheless, compilers implementing such optimizations can be used as a backend to optimize further our solver programs generated by MAT2STENCIL and we have evaluated against some of them. Among recent works, FreeTensor [Tang et al. 2022] is a language and compiler to optimize such a tensor program. FreeTensor provides a flexible Python interface to construct the tensor program, which is ideal as the target language of our multi-stage programming. It further provides a set of code-transforming schedules, as performance-tuning primitives. We use FreeTensor as our backend and our proposed optimizations are implemented as a FreeTensor plugin.

10 DISCUSSION

While MAT2STENCIL has proposed a set of optimizations for the Seidel-style stencils involved in implicit solvers, the optimizations for conventional Jacobi-style stencils are left behind. It thus leaves huge potential on such optimizations, such as temporal blocking (when possible, it is sometimes forbidden by backward iterations), cache-aware tiling, etc. Applying such techniques to Seidel-style stencils is a challenging problem left to be solved.

Other than multithreading on CPUs, the latest advancement in accelerators, especially GPGPUs (general-purpose graphics processing units), has gained more and more attention. The SewTh does not fit into their programming model, which does not allow arbitrary inter-block synchronization. Algorithmic adjustments are essential as a result, involving techniques like multi-color reordering; incorporating such changes, new compilation techniques need to be developed to automatically parallelize Seidel-style stencils to make them efficiently run on GPGPUs. Automatic distributed code generation is also possible through polyhedral methods analyzing loop-carried dependences, which has been explored previously [Baghdadi et al. 2019].

Though our current abstraction focuses on Cartesian grids, other types, e.g., star-shaped or hexagonal ones, are possible to be supported with reconsidered predefined matrices. Speaking even further, semi-structured grids such as block-structured grids might also benefit from our method by combining the unstructured matrix-based representation at the coarse level and our sparse matrix representation for the finer grids, requiring more programming optimization designs. Besides, our cases only cover the Finite Difference method solvers and new matrices would be required to address Finite Volume and Finite Elements methods.

11 CONCLUSION

We propose an innovative domain-specific language (DSL), MAT2STENCIL, for solving partial differential equations (PDEs) on structured grids. MAT2STENCIL introduces a structured sparse matrix abstraction, facilitating modular, flexible, and easy-to-use expression of solvers across a broad spectrum of PDEs, including explicit and implicit solving algorithms. We evaluate our DSL by implementing four benchmarking programs from the NAS Parallel Benchmarks and High Performance Conjugate Gradients, achieving up to 1.67× and on average 1.03× performance compared to manually matrix-free codes, with significantly less code.

ACKNOWLEDGMENTS

This work is supported in part by NSFC U20B2044 and The Major Key Project of PCL.

REFERENCES

- Randy Allen and Ken Kennedy. 2001. *Optimizing Compilers for Modern Architectures: A Dependence-based Approach*. Morgan Kaufmann, San Francisco, CA, USA.
- Travis Augustine, Janarthanan Sarma, Louis-Noël Pouchet, and Gabriel Rodríguez. 2019. Generating piecewise-regular code from irregular structures. In *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2019, Phoenix, AZ, USA, June 22-26, 2019*, Kathryn S. McKinley and Kathleen Fisher (Eds.). ACM, New York, NY, USA, 625–639. <https://doi.org/10.1145/3314221.3314615>
- Riyadh Baghdadi, Jessica Ray, Malek Ben Romdhane, Emanuele Del Sozzo, Abdurrahman Akkas, Yunming Zhang, Patricia Suriana, Shoaib Kamil, and Saman P. Amarasinghe. 2019. Tiramisu: A Polyhedral Compiler for Expressing Fast and Portable Code. In *IEEE/ACM International Symposium on Code Generation and Optimization, CGO 2019, Washington, DC, USA, February 16-20, 2019*, Mahmut Taylan Kandemir, Alexandra Jimborean, and Tipp Moseley (Eds.). IEEE, 193–205. <https://doi.org/10.1109/CGO.2019.8661197>
- David H. Bailey, Eric Barszcz, John T. Barton, D. S. Browning, Robert L. Carter, Leonardo Dagum, Rod A. Fatoohi, Paul O. Frederickson, T. A. Lasinski, Robert Schreiber, Horst D. Simon, V. Venkatakrishnan, and Sisira Weeratunga. 1991. The Nas Parallel Benchmarks. *Int. J. High Perform. Comput. Appl.* 5, 3 (1991), 63–73. <https://doi.org/10.1177/109434209100500306>
- Uday Bondhugula, Aravind Acharya, and Albert Cohen. 2016. The Pluto+ Algorithm: A Practical Approach for Parallelization and Locality Optimization of Affine Loop Nests. *ACM Trans. Program. Lang. Syst.* 38, 3 (2016), 12:1–12:32. <https://doi.org/10.1145/2896389>
- Uday Bondhugula, Vinayaka Bandishti, and Irshad Pananilath. 2017. Diamond Tiling: Tiling Techniques to Maximize Parallelism for Stencil Computations. *IEEE Trans. Parallel Distributed Syst.* 28, 5 (2017), 1285–1298. <https://doi.org/10.1109/TPDS.2016.2615094>
- Uday Bondhugula, Albert Hartono, J. Ramanujam, and P. Sadayappan. 2008. A practical automatic polyhedral parallelizer and locality optimizer. In *Proceedings of the ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation, Tucson, AZ, USA, June 7-13, 2008*, Rajiv Gupta and Saman P. Amarasinghe (Eds.). ACM, 101–113. <https://doi.org/10.1145/1375581.1375595>
- Huanqi Cao. 2023. *Artifact of Mat2Stencil: A Modular Matrix-Based DSL for Explicit and Implicit Matrix-Free PDE Solvers on Structured Grid*. <https://doi.org/10.5281/zenodo.8149701>
- Edmond Chow and Aftab Patel. 2015. Fine-Grained Parallel Incomplete LU Factorization. *SIAM J. Sci. Comput.* 37, 2 (2015). <https://doi.org/10.1137/140968896>
- Kaushik Datta, Mark Murphy, Vasily Volkov, Samuel Williams, Jonathan Carter, Leonid Oliker, David A. Patterson, John Shalf, and Katherine A. Yelick. 2008. Stencil computation optimization and auto-tuning on state-of-the-art multicore architectures. In *Proceedings of the ACM/IEEE Conference on High Performance Computing, SC 2008, November 15-21, 2008, Austin, Texas, USA*. IEEE/ACM, 4. <https://doi.org/10.1109/SC.2008.5222004>
- Leonardo Mendonça de Moura and Nikolaj S. Bjørner. 2008. Z3: An Efficient SMT Solver. In *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2008, Budapest, Hungary, March 29-April 6, 2008. Proceedings (Lecture Notes in Computer Science, Vol. 4963)*, C. R. Ramakrishnan and Jakob Rehof (Eds.). Springer, 337–340. https://doi.org/10.1007/978-3-540-78800-3_24
- James Decker. 2019. *Implementation of Lightweight Modular Staging (LMS) in Python*. Retrieved November 10, 2022 from <https://github.com/jmd1011/snek-LMS>
- Jack J. Dongarra, Michael A. Heroux, and Piotr Luszczyk. 2016. High-performance conjugate-gradient benchmark: A new metric for ranking high-performance computing systems. *Int. J. High Perform. Comput. Appl.* 30, 1 (2016), 3–10. <https://doi.org/10.1177/1094342015593158>
- Mohamed Essadki, Bertrand Michel, Bruno Maugars, Oleksandr Zinenko, Nicolas Vasilache, and Albert Cohen. 2023. Code Generation for In-Place Stencils. In *Proceedings of the 21st ACM/IEEE International Symposium on Code Generation and Optimization, CGO 2023, Montréal, QC, Canada, 25 February 2023- 1 March 2023*, Christophe Dubach, Derek Bruening, and Ben Hardekopf (Eds.). ACM, 2–13. <https://doi.org/10.1145/3579990.3580006>
- Paul Feautrier. 1992. Some efficient solutions to the affine scheduling problem. I. One-dimensional time. *Int. J. Parallel Program.* 21, 5 (1992), 313–347. <https://doi.org/10.1007/BF01407835>
- Johannes Habich, T. Zeiser, Georg Hager, and Gerhard Wellein. 2009. Enabling temporal blocking for a lattice Boltzmann flow solver through multicore-aware wavefront parallelization.
- X. Huang, X. Huang, D. Wang, Q. Wu, Y. Li, S. Zhang, Y. Chen, M. Wang, Y. Gao, Q. Tang, Y. Chen, Z. Fang, Z. Song, and G. Yang. 2019. OpenArray v1.0: a simple operator library for the decoupling of ocean modeling and parallel computing. *Geoscientific Model Development* 12, 11 (2019), 4729–4749. <https://doi.org/10.5194/gmd-12-4729-2019>
- Intel. 2023. Intel oneAPI Math Kernel Library. <https://www.intel.com/content/www/us/en/developer/tools/oneapi/onemkl.html>

- Shoab Kamil, Kaushik Datta, Samuel Williams, Leonid Oliker, John Shalf, and Katherine A. Yelick. 2006. Implicit and explicit optimizations for stencil computations. In *Proceedings of the 2006 workshop on Memory System Performance and Correctness, San Jose, California, USA, October 11, 2006*, Antony L. Hosking and Ali-Reza Adl-Tabatabai (Eds.). ACM, 51–60. <https://doi.org/10.1145/1178597.1178605>
- Sriram Krishnamoorthy, Muthu Manikandan Baskaran, Uday Bondhugula, J. Ramanujam, Atanas Rountev, and P. Sadayappan. 2007. Effective automatic parallelization of stencil computations. In *Proceedings of the ACM SIGPLAN 2007 Conference on Programming Language Design and Implementation, San Diego, California, USA, June 10-13, 2007*, Jeanne Ferrante and Kathryn S. McKinley (Eds.). ACM, 235–244. <https://doi.org/10.1145/1250734.1250761>
- Christian Lengauer, Sven Apel, Matthias Boltén, Shigeru Chiba, Ulrich Rüdé, Jürgen Teich, Armin Größlinger, Frank Hannig, Harald Köstler, Lisa Claus, Alexander Grebhahn, Stefan Groth, Stefan Kronawitter, Sebastian Kuckuk, Hannah Rittich, Christian Schmitt, and Jonas Schmitt. 2020. ExaStencils: Advanced Multigrid Solver Generation. In *Software for Exascale Computing - SPPEXA 2016-2019*, Hans-Joachim Bungartz, Severin Reiz, Benjamin Uekermann, Philipp Neumann, and Wolfgang E. Nagel (Eds.). Lecture Notes in Computational Science and Engineering, Vol. 136. Springer, 405–452. https://doi.org/10.1007/978-3-030-47956-5_14
- XiaoYe S. Li and Meiyue Shao. 2011. A Supermodal Approach to Incomplete LU Factorization with Partial Pivoting. *ACM Trans. Math. Softw.* 37, 4 (2011), 43:1–43:20. <https://doi.org/10.1145/1916461.1916467>
- M. Louboutin, M. Lange, F. Luporini, N. Kukreja, P. A. Witte, F. J. Herrmann, P. Velesko, and G. J. Gorman. 2019. Devito (v3.1.0): an embedded domain-specific language for finite differences and geophysical exploration. *Geoscientific Model Development* 12, 3 (2019), 1165–1187. <https://doi.org/10.5194/gmd-12-1165-2019>
- George Mcmechan. 2006. Migration by extrapolation of time-dependent boundary values. *Geophysical Prospecting* 31 (04 2006), 413 – 420. <https://doi.org/10.1111/j.1365-2478.1983.tb01060.x>
- Duane Merrill and Michael Garland. 2016. Merge-based parallel sparse matrix-vector multiplication. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2016, Salt Lake City, UT, USA, November 13-18, 2016*, John West and Cherri M. Pancake (Eds.). IEEE Computer Society, 678–689. <https://doi.org/10.1109/SC.2016.57>
- Anthony D. Nguyen, Nadathur Satish, Jatin Chhugani, Changkyu Kim, and Pradeep Dubey. 2010. 3.5-D Blocking Optimization for Stencil Computations on Modern CPUs and GPUs. In *Conference on High Performance Computing Networking, Storage and Analysis, SC 2010, New Orleans, LA, USA, November 13-19, 2010*. IEEE, 1–13. <https://doi.org/10.1109/SC.2010.2>
- Georg Ofenbeck, Tiark Rompf, Alen Stojanov, Martin Odersky, and Markus Püschel. 2013. Spiral in scala: towards the systematic construction of generators for performance libraries. (2013), 125–134. <https://doi.org/10.1145/2517208.2517228>
- OpenAI. 2023. GPT-4 Technical Report. arXiv:2303.08774 [cs.CL]
- David A. Padua and Michael Wolfe. 1986. Advanced Compiler Optimizations for Supercomputers. *Commun. ACM* 29, 12 (1986), 1184–1201. <https://doi.org/10.1145/7902.7904>
- Andreas Pieper, Georg Hager, and Holger Fehske. 2021. A domain-specific language and matrix-free stencil code for investigating electronic properties of Dirac and topological materials. *Int. J. High Perform. Comput. Appl.* 35, 1 (2021). <https://doi.org/10.1177/1094342020959423>
- Ali Pinar and Michael T. Heath. 1999. Improving Performance of Sparse Matrix-Vector Multiplication. In *Proceedings of the ACM/IEEE Conference on Supercomputing, SC 1999, November 13-19, 1999, Portland, Oregon, USA*. ACM, 30. <https://doi.org/10.1145/331532.331562>
- Tiark Rompf and Martin Odersky. 2012. Lightweight modular staging: a pragmatic approach to runtime code generation and compiled DSLs. *Commun. ACM* 55, 6 (2012), 121–130. <https://doi.org/10.1145/2184319.2184345>
- Yousef Saad. 2003. *Iterative methods for sparse linear systems*. SIAM. <https://doi.org/10.1137/1.9780898718003>
- Amir Shaikhha, Yannis Klonatos, and Christoph Koch. 2018. Building Efficient Query Engines in a High-Level Language. *ACM Trans. Database Syst.* 43, 1 (2018), 4:1–4:45. <https://doi.org/10.1145/3183653>
- C. Skamarock, BogumiŁa Klemp, Jimy Dudhia, O. Gill, Zhiqian Liu, Judith Berner, Wei Wang, G. Powers, Greg Duda, Dale M. Barker, and Xiangyu Huang. 2019. A Description of the Advanced Research WRF Model Version 4.
- John C. Strikwerda. 2004. *Finite Difference Schemes and Partial Differential Equations, Second Edition*. Society for Industrial and Applied Mathematics. <https://doi.org/10.1137/1.9780898717938> arXiv:<https://epubs.siam.org/doi/pdf/10.1137/1.9780898717938>
- Nicolas Stucki, Aggelos Biboudis, and Martin Odersky. 2018. A practical unification of multi-stage programming and macros. In *Proceedings of the 17th ACM SIGPLAN International Conference on Generative Programming: Concepts and Experiences, GPCE 2018, Boston, MA, USA, November 5-6, 2018*, Eric Van Wyk and Tiark Rompf (Eds.). ACM, 14–27. <https://doi.org/10.1145/3278122.3278139>
- Arvind K. Sujeeth, Kevin J. Brown, HyoukJoong Lee, Tiark Rompf, Hassan Chafi, Martin Odersky, and Kunle Olukotun. 2014. Delite: A Compiler Architecture for Performance-Oriented Embedded Domain-Specific Languages. *ACM Trans. Embed. Comput. Syst.* 13, 4s (2014), 134:1–134:25. <https://doi.org/10.1145/2584665>

- Walid Taha. 1999. *Multi-Stage Programming: Its Theory and Applications*. Ph.D. Dissertation. Halmstad University, Sweden. <https://urn.kb.se/resolve?urn=urn:nbn:se:hh:diva-15052>
- Shizhi Tang, Jidong Zhai, Haojie Wang, Lin Jiang, Liyan Zheng, Zhenhao Yuan, and Chen Zhang. 2022. FreeTensor: a free-form DSL with holistic optimizations for irregular tensor programs. In *PLDI '22: 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation, San Diego, CA, USA, June 13 - 17, 2022*, Ranjit Jhala and Isil Dillig (Eds.). ACM, 872–887. <https://doi.org/10.1145/3519939.3523448>
- Yuan Tang, Rezaul Alam Chowdhury, Bradley C. Kuszmaul, Chi-Keung Luk, and Charles E. Leiserson. 2011. The pochoir stencil compiler. In *SPAA 2011: Proceedings of the 23rd Annual ACM Symposium on Parallelism in Algorithms and Architectures, San Jose, CA, USA, June 4-6, 2011 (Co-located with FCRC 2011)*, Rajmohan Rajaraman and Friedhelm Meyer auf der Heide (Eds.). ACM, 117–128. <https://doi.org/10.1145/1989493.1989508>
- Sven Verdoolaege. 2010. *isl: An Integer Set Library for the Polyhedral Model*. In *Mathematical Software - ICMS 2010, Third International Congress on Mathematical Software, Kobe, Japan, September 13-17, 2010. Proceedings (Lecture Notes in Computer Science, Vol. 6327)*, Komei Fukuda, Joris van der Hoeven, Michael Joswig, and Nobuki Takayama (Eds.). Springer, 299–302. https://doi.org/10.1007/978-3-642-15582-6_49
- Sven Verdoolaege and Gerda Janssens. 2017. Scheduling for PPCG. <https://doi.org/10.13140/RG.2.2.28998.68169>
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. 2020. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* 17 (2020), 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- Samuel Williams, Nathan Bell, Jee Whan Choi, Michael Garland, Leonid Oliker, and Richard Vu. 2010. Sparse Matrix-Vector Multiplication on Multicore and Accelerators. In *Scientific Computing with Multicore and Accelerators*, Jakub Kurzak, David A. Bader, and Jack J. Dongarra (Eds.). CRC Press / Taylor & Francis, 83–109. <https://doi.org/10.1201/b10376-8>
- Samuel Williams, John Shalf, Leonid Oliker, Shoab Kamil, Parry Husbands, and Katherine A. Yelick. 2007. Scientific Computing Kernels on the Cell Processor. *Int. J. Parallel Program.* 35, 3 (2007), 263–298. <https://doi.org/10.1007/s10766-007-0034-5>
- Chao Yang, Wei Xue, Haohuan Fu, Hongtao You, Xinliang Wang, Yulong Ao, Fangfang Liu, Lin Gan, Ping Xu, Lanning Wang, Guangwen Yang, and Weimin Zheng. 2016. 10M-core scalable fully-implicit solver for nonhydrostatic atmospheric dynamics. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2016, Salt Lake City, UT, USA, November 13-18, 2016*, John West and Cherri M. Pancake (Eds.). IEEE Computer Society, 57–68. <https://doi.org/10.1109/SC.2016.5>
- Nathan Zhang, Michael B. Driscoll, Charles Markley, Samuel Williams, Protonu Basu, and Armando Fox. 2017. Snowflake: A Lightweight Portable Stencil DSL. In *2017 IEEE International Parallel and Distributed Processing Symposium Workshops, IPDPS Workshops 2017, Orlando / Buena Vista, FL, USA, May 29 - June 2, 2017*. IEEE Computer Society, 795–804. <https://doi.org/10.1109/IPDPSW.2017.89>
- Qianchao Zhu, Hao Luo, Chao Yang, Mingshuo Ding, Wanwang Yin, and Xinhui Yuan. 2021. Enabling and scaling the HPCG benchmark on the newest generation Sunway supercomputer with 42 million heterogeneous cores. In *International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2021, St. Louis, Missouri, USA, November 14-19, 2021*, Bronis R. de Supinski, Mary W. Hall, and Todd Gamblin (Eds.). ACM, 57. <https://doi.org/10.1145/3458817.3476158>

Received 2023-04-14; accepted 2023-08-27