

Optimised finite difference computation from symbolic equations

Michael Lange^{‡*}, Navjot Kukreja[‡], Fabio Luporini[‡], Mathias Louboutin[¶], Charles Yount[§], Jan Hückelheim[‡], Gerard J. Gorman[‡]

<https://youtu.be/KinmqFTEs94>



Abstract—Domain-specific high-productivity environments are playing an increasingly important role in scientific computing due to the levels of abstraction and automation they provide. In this paper we introduce Devito, an open-source domain-specific framework for solving partial differential equations from symbolic problem definitions by the finite difference method. We highlight the generation and automated execution of highly optimized stencil code from only a few lines of high-level symbolic Python for a set of scientific equations, before exploring the use of Devito operators in seismic inversion problems.

Index Terms—Finite difference, domain-specific languages, symbolic Python

Introduction

Domain-specific high-productivity environments are playing an increasingly important role in scientific computing. The level of abstraction and automation provided by such frameworks not only increases productivity and accelerates innovation, but also allows the combination of expertise from different specialised disciplines. This synergy is necessary when creating the complex software stack needed to solve leading edge scientific problems, since domain specialists as well as high performance computing experts are required to fully leverage modern computing architectures. Based on this philosophy we introduce Devito [Lange17], an open-source domain-specific framework for solving partial differential equations (PDE) from symbolic problem definitions by the finite difference method.

Symbolic computation, where optimized numerical code is automatically derived from a high-level problem definition, is a powerful technique that allows domain scientists to focus on algorithmic development rather than implementation details. For this reason Devito exposes an API based on Python (SymPy) [Meurer17] that allow users to express equations symbolically, from which it generates and executes optimized stencil code via just-in-time (JIT) compilation. Using latest advances in stencil compiler research, Devito thus provides domain scientists with the ability to quickly and efficiently generate high-performance

kernels from only a few lines of Python code, making Devito composable with existing open-source software.

While Devito was originally developed for seismic imaging workflows, the automated generation and optimization of stencil codes can be utilised for a much broader set of computational problems. Matrix-free stencil operators based on explicit finite difference schemes are widely used in industry and academic research, although they merely represent one of many approaches to solving PDEs [Baba16], [Liu09], [Rai91]. In this paper we therefore limit our discussion of numerical methods and instead focus on the ease with which these operators can be created symbolically. We give a brief overview of the design concepts and key features of Devito and demonstrate its API using a set of classic examples from computational fluid dynamics (CFD). Then we will discuss the use of Devito in an example of a complex seismic inversion algorithm to illustrate its use in practical scientific applications and to showcase the performance achieved by the auto-generated and optimised code.

Background

The attraction of using domain-specific languages (DSL) to solve PDEs via a high-level mathematical notation is by no means new and has led to various purpose-built software packages and compilers dating back to 1962 [Iverson62], [Cardenas70], [Umetani85], [Cook88], [VanEngelen96]. Following the emergence of Python as a widely used programming language in scientific research, embedded DSLs for more specialised domains came to the fore, most notably the FEniCS [Logg12] and Firedrake [Rathgeber16] frameworks, which both implement the unified Form Language (UFL) [Alnaes14] for the symbolic definition of finite element problems in the weak form. The increased level of abstraction that such high-level languages provide decouples the problem definition from its implementation, thus allowing domain scientists and mathematicians to focus on more advanced methods, such as the automation of adjoint models as demonstrated by Dolfin-Adjoint [Farrell13].

The performance optimization of stencil computation on regular cartesian grids for high-performance computing applications has also received much attention in computer science research [Datta08], [Brandvik10], [Zhang12], [Henretty13], [Yount15]. The primary focus of most stencil compilers or DSLs, however, is the optimization of synthetic problems which often limits their

* Corresponding author: michael.lange@imperial.ac.uk

‡ Imperial College London

¶ The University of British Columbia

§ Intel Corporation

applicability for practical scientific applications. The primary consideration here is that most realistic problems often require more than just a fast and efficient PDE solver, which entails that symbolic DSLs embedded in Python can benefit greatly from native interoperability with the scientific Python ecosystem.

Design and API

The primary objective of Devito is to enable the quick and effective creation of highly optimised finite difference operators for use in a realistic scientific application context. As such, its design is centred around composability with the existing Python software stack to provide users with the tools to dynamically generate optimised stencil computation kernels and to enable access to the full scientific software ecosystem. In addition, to accommodate the needs of "real life" scientific applications, a secondary API is provided that enables users to inject custom expressions, such as boundary conditions or sparse point interpolation routines, into the generated kernels.

The use of SymPy as the driver for the symbolic generation of stencil expressions and the subsequent code-generation are at the heart of the Devito philosophy. While SymPy is fully capable of auto-generating low-level C code for pre-compiled execution from high-level symbolic expressions, Devito is designed to combine these capabilities with automatic performance optimization based on the latest advances in stencil compiler technology. The result is a framework that is capable of automatically generating and optimising complex stencil code from high-level symbolic definitions.

The Devito API is based around two key concepts that allow users to express finite difference problems in a concise symbolic notation:

- **Symbolic data objects:** Devito's high-level symbolic objects behave like `sympy.Function` objects and provide a set of shorthand notations for generating derivative expressions, while also managing user data. The rationale for this duality is that many stencil optimization algorithms rely on data layout changes, mandating that Devito needs to be in control of data allocation and access.
- **Operator:** An `Operator` creates, compiles and executes a single executable kernel from a set of SymPy expressions. The code generation and optimization process involves various stages and accepts a mixture of high-level and low-level expressions to allow the injection of customised code.

Fluid Dynamics Examples

In the following section we demonstrate the use of the Devito API to implement two examples from classical fluid dynamics, before highlighting the role of Devito operators in a seismic inversion context. Both CFD examples are based in part on tutorials from the introductory blog "CFD Python: 12 steps to Navier-Stokes"¹ by the Lorena A. Barba group. We have chosen the examples in this section for their relative simplicity to concisely illustrate the capabilities and API features of Devito. For a more complete discussion on numerical methods for fluid flows please refer to [Peiro05].

1. <http://lorenabarba.com/blog/cfd-python-12-steps-to-navier-stokes/>

Linear Convection

We will demonstrate a basic Devito operator definition based on a linear two-dimensional convection flow (step 5 in the original tutorials)². The governing equation we are implementing here is:

$$\frac{\partial u}{\partial t} + c \frac{\partial u}{\partial x} + c \frac{\partial u}{\partial y} = 0 \quad (1)$$

A discretised version of this equation, using a forward difference scheme in time and a backward difference scheme in space might be written as

$$u_{i,j}^{n+1} = u_{i,j}^n - c \frac{\Delta t}{\Delta x} (u_{i,j}^n - u_{i-1,j}^n) - c \frac{\Delta t}{\Delta y} (u_{i,j}^n - u_{i,j-1}^n) \quad (2)$$

where the subscripts i and j denote indices in the space dimensions and the superscript n denotes the index in time, while Δt , Δx , Δy denote the spacing in time and space dimensions respectively.

The first thing we need is a function object with which we can build a timestepping scheme. For this purpose Devito provides so-called `TimeData` objects that encapsulate functions that are differentiable in space and time. With this we can derive symbolic expressions for the backward derivatives in space directly via the `u.dx1` and `u.dyl` shorthand expressions (the `1` indicates "left" or backward differences) and the shorthand notation `u.dt` provided by `TimeData` objects to derive the forward derivative in time.

```
from devito import *

c = 1.
u = TimeData(name='u', shape=(nx, ny))

eq = Eq(u.dt + c * u.dx1 + c * u.dyl)

[In] print eq
[Out] Eq(-u(t, x, y)/s + u(t + s, x, y)/s
      + 2.0*u(t, x, y)/h - 1.0*u(t, x, y - h)/h
      - 1.0*u(t, x - h, y)/h, 0)
```

The above expression results in a `sympy.Equation` object that contains the fully discretised form of Eq. 1, including placeholder symbols for grid spacing in space (h , assuming $\Delta x = \Delta y$) and time (s). These spacing symbols will be resolved during the code generation process, as described in the [code generation section](#). It is also important to note here that the explicit generation of the space derivatives `u_dx` and `u_dy` is due to the use of a backward derivative in space to align with the original example. A similar notation to the forward derivative in time (`u.dt`) will soon be provided.

In order to create a functional `Operator` object, the expression `eq` needs to be rearranged so that we may solve for the unknown $u_{i,j}^{n+1}$. This is easily achieved by using SymPy's `solve` utility and the Devito shorthand `u.forward` which denotes the furthest forward stencil point in a time derivative ($u_{i,j}^{n+1}$).

```
from sympy import solve

stencil = solve(eq, u.forward)[0]

[In] print(stencil)
[Out] (h*u(t, x, y) - 2.0*s*u(t, x, y)
      + s*u(t, x, y - h) + s*u(t, x - h, y))/h
```

The above variable `stencil` now represents the RHS of Eq. 2, allowing us to construct a SymPy expression that updates $u_{i,j}^{n+1}$ and build a `devito.Operator` from it. When creating this operator

2. http://nbviewer.jupyter.org/github/opesci/devito/blob/master/examples/cfd/test_01_convection_revisited.ipynb

we also supply concrete values for the spacing terms h and s via an additional substitution map argument `subs`.

```
op = Operator(Eq(u.forward, stencil),
             subs={h: dx, s:dt})

# Set initial condition as a smooth function
init_smooth(u.data, dx, dy)

op(u=u, time=100) # Apply for 100 timesteps
```

Using this operator we can now create a similar example to the one presented in the original tutorial by initialising the data associated with the symbolic function u , `u.data` with an initial flow field. However, to avoid numerical errors due to the discontinuities at the boundary of the original "hat function", we use the following smooth initial condition provided by [Krakos12], as depicted in Figure 1.

$$u_0(x,y) = 1 + u\left(\frac{2}{3}x\right) * u\left(\frac{2}{3}y\right)$$

The final result after executing the operator for 5s (100 timesteps) is depicted in Figure 2. The result shows the expected displacement of the initial shape, in accordance with the prescribed velocity ($c = 1.0$), closely mirroring the displacement of the "hat function" in the original tutorial. It should also be noted that, while the results show good agreement with expectations by visual inspection, they do not represent an accurate solution to the linear convection equation. In particular, the low order spatial discretisation introduces numerical diffusion that causes a decrease in the peak velocity. This is a well-known issue that could be addressed with more sophisticated solver schemes as discussed in [LeVeque92].




Fig. 1: Initial condition of `u.data` in the 2D convection example.




Fig. 2: State of `u.data` after 100 timesteps in convection example.

Laplace equation

The above example shows how Devito can be used to create finite difference stencil operators from only a few lines of high-level symbolic code. However, the previous example only required a single variable to be updated, while more complex operators might need to execute multiple expressions simultaneously, for example to solve coupled PDEs or apply boundary conditions as part of the time loop. For this reason `devito.Operator` objects can be constructed from multiple update expressions and allow multiple expression formats as input.

Nevertheless, boundary conditions are currently not provided as part of the symbolic high-level API. For exactly this reason, Devito provides a low-level, or "indexed" API, where custom SymPy expressions can be created with explicitly resolved grid accesses to manually inject custom code into the auto-generation toolchain. This entails that future extensions to capture different types of boundary conditions can easily be added at a later stage.

To illustrate the use of the low-level API, we will use the Laplace example from the original CFD tutorials (step 9), which implements the steady-state heat equation with Dirichlet and Neuman boundary conditions³. The governing equation for this problem is

$$\frac{\partial^2 p}{\partial x^2} + \frac{\partial^2 p}{\partial y^2} = 0 \quad (3)$$

The rearranged discretised form, assuming a central difference scheme for second derivatives, is

$$p_{i,j}^n = \frac{\Delta y^2 (p_{i+1,j}^n + p_{i-1,j}^n) + \Delta x^2 (p_{i,j+1}^n + p_{i,j-1}^n)}{2(\Delta x^2 + \Delta y^2)} \quad (4)$$

Using a similar approach to the previous example, we can construct the SymPy expression to update the state of a field p . For demonstration purposes we will use two separate function objects of type `DenseData` in this example, since the Laplace equation does not contain a time-dependence. The shorthand expressions `pn.dx2` and `pn.dy2` hereby denote the second derivatives in x and y .

```
# Create two separate symbols with space dimensions
p = DenseData(name='p', shape=(nx, ny),
              space_order=2)
pn = DenseData(name='pn', shape=(nx, ny),
              space_order=2)

# Define equation and solve for center point in `pn`
eq = Eq(a * pn.dx2 + pn.dy2)
stencil = solve(eq, pn)[0]
# The update expression to populate buffer `p`
eq_stencil = Eq(p, stencil)
```

Just as the original tutorial, our initial condition in this example is $p = 0$ and the flow will be driven by the boundary conditions

$$\begin{aligned} p &= 0 & \text{at } x &= 0 \\ p &= y & \text{at } x &= 2 \\ \frac{\partial p}{\partial y} &= 0 & \text{at } y &= 0, 1 \end{aligned}$$

To implement these BCs we can utilise the `.indexed` property that Devito symbols provide to get a symbol of type `sympy.IndexedBase`, which in turn allows us to use matrix indexing notation (square brackets) to create symbols of type `sympy.Indexed` instead of `sympy.Function`. This notation

allows users to hand-code stencil expressions using explicit relative grid indices, for example $p[x, y] - p[x-1, y] / h$ for the discretized backward derivative $\frac{\partial u}{\partial x}$. The symbols x and y hereby represent the respective problem dimensions and cause the expression to be executed over the entire data dimension, similar to Python's `:` operator.

The Dirichlet BCs in the Laplace example can thus be implemented by creating a `sympy.Eq` object that assigns either fixed values or a prescribed function, such as the utility symbol `bc_right` in our example, along the left and right boundary of the domain. To implement the Neumann BCs we again follow the original tutorial by assigning the second grid row from the top and bottom boundaries the value of the outermost row. The resulting SymPy expressions can then be used alongside the state update expression to create our `Operator` object.

```
# Create an additional symbol for our prescribed BC
bc_right = DenseData(name='bc_right', shape=(nx, ),
                    dimensions=(x, ))
bc_right.data[:] = np.linspace(0, 1, nx)

# Create explicit boundary condition expressions
bc = [Eq(p.indexed[x, 0], 0)]
bc += [Eq(p.indexed[x, ny-1], bc_right.indexed[x])]
bc += [Eq(p.indexed[0, y], p.indexed[1, y])]
bc += [Eq(p.indexed[nx-1, y], p.indexed[nx-2, y])]

# Build operator with update and BC expressions
op = Operator(expressions=[eq_stencil] + bc,
              subs={h: dx, a: 1.})
```

After building the operator, we can now use it in a time-independent convergence loop that minimizes the L^1 norm of p . However, in this example we need to make sure to explicitly exchange the role of the buffers `p` and `pn`. This can be achieved by supplying symbolic data objects via keyword arguments when invoking the operator, where the name of the argument is matched against the name of the original symbol used to create the operator.

The convergence criterion for this example is defined as the relative error between two iterations and set to $\|p\|^1 < 10^{-4}$. The corresponding initial condition and the resulting steady-state solution, depicted in Figures 3 and 4 respectively, agree with the original tutorial implementation. It should again be noted that the chosen numerical scheme might not be optimal to solve steady-state problems of this type, since implicit methods are often preferred.

```
l1norm = 1
counter = 0
while l1norm > 1.e-4:
    # Determine buffer order
    if counter % 2 == 0:
        _p, _pn = p, pn
    else:
        _p, _pn = pn, p

    # Apply operator
    op(p=_p, pn=_pn)

    # Compute L1 norm
    l1norm = (np.sum(np.abs(_p.data[:])
                    - np.abs(_pn.data[:]))
              / np.sum(np.abs(_pn.data[:])))
    counter += 1
```

Seismic Inversion Example

The primary motivating application behind the design of Devito is the solution of seismic exploration problems that require highly optimised wave propagation operators for forward modelling and

3. http://nbviewer.jupyter.org/github/opesci/devito/blob/master/examples/cfd/test_05_laplace.ipynb




Fig. 3: Initial condition of $p.n.data$ in the 2D Laplace example.




Fig. 4: State of $p.data$ after convergence in Laplace example.

adjoint-based inversion. Obviously, the speed and accuracy of the generated kernels are of vital importance. Moreover, the ability to efficiently define rigorous forward modelling and adjoint operators from high-level symbolic definitions also implies that domain scientists are able to quickly adjust the numerical method and discretisation to the individual problem and hardware architecture [Louboutin17a].

In the following example we will show the generation of forward and adjoint operators for the acoustic wave equation and verify their correctness using the so-called *adjoint test* [Virieux09]⁴. This test, also known as *dot product test*, verifies that the implementation of an adjoint operator indeed computes the conjugate transpose of the forward operator.

The governing wave equation for the forward operator is defined as

$$m \frac{\partial^2 u}{\partial t^2} + \eta \frac{\partial u}{\partial t} - \nabla^2 u = q$$

where u denotes the pressure wave field, m is the square slowness, q is the source term and η denotes the spatially varying dampening factor used to implement an absorbing boundary condition.

On top of fast stencil operators, seismic inversion kernels also rely on sparse point interpolation to inject the modelled wave as a point source (q) and to record the pres-

sure at individual point locations. To accommodate this, Devito provides another symbolic data type `PointData`, which allows the generation of sparse-point interpolation expressions using the "indexed" low-level API. These symbolic objects provide utility routines `pt.interpolate(expression)` and `pt.inject(field, expression)` to create symbolic expressions that perform linear interpolation between the sparse points and the cartesian grid for insertion into `Operator` kernels. A separate set of explicit coordinate values is associated with the sparse point objects for this purpose in addition to the function values stored in the `data` property.

Adjoint Test

The first step for implementing the adjoint test is to build a forward operator that models the wave propagating through an isotropic medium, where the square slowness of the wave is denoted as m . Since m , as well as the boundary dampening function η , is re-used between forward and adjoint runs the only symbolic data object we need to create here is the wavefield u in order to implement and rearrange our discretised equation `eqn` to form the update expression for u . It is worth noting that the `u.laplace` shorthand notation used here expands to the set of second derivatives in all spatial dimensions, thus allowing us to use the same formulation for two-dimensional and three-dimensional problems.

In addition to the state update of u , we are also inserting two additional terms into the forward modelling operator:

- `src_term` injects a pressure source at a point location according to a prescribed time series stored in `src.data` that is accessible in symbolic form via the symbol `src`. The scaling factor in `src_term` is coded by hand but can be automatically inferred.
- `rec_term` adds the expression to interpolate the wavefield u for a set of "receiver" hydrophones that measure the propagated wave at varying distances from the source for every time step. The resulting interpolated point data will be stored in `rec.data` and is accessible to the user as a NumPy array.

```
def forward(model, m, eta, src, rec, order=2):
    # Create the wavefield function
    u = TimeData(name='u', shape=model.shape,
                 time_order=2, space_order=order)

    # Derive stencil from symbolic equation
    eqn = m * u.dt2 - u.laplace + eta * u.dt
    stencil = solve(eqn, u.forward)[0]
    update_u = [Eq(u.forward, stencil)]

    # Add source injection and receiver interpolation
    src_term = src.inject(field=u,
                          expr=src * dt**2 / m)
    rec_term = rec.interpolate(expr=u)

    # Create operator with source and receiver terms
    return Operator(update_u + src_term + rec_term,
                    subs={s: dt, h: model.spacing})
```

After building a forward operator, we can now implement the adjoint operator in a similar fashion. Using the provided symbols m and η , we can again define the adjoint wavefield v and implement its update expression from the discretised equation. However, since the adjoint operator needs to operate backwards in time there are two notable differences:

4. http://nbviewer.jupyter.org/github/opesci/devito/blob/master/examples/seismic/tutorials/test_01_modelling.ipynb

- The update expression now updates the backward stencil point in the time derivative $v_{i,j}^{n-1}$, denoted as `v.backward`. In addition to that, the `Operator` is forced to reverse its internal time loop by providing the argument `time_axis=Backward`
- Since the acoustic wave equation is self-adjoint without dampening, the only change required in the governing equation is to invert the sign of the dampening term `eta * u.dt`. The first derivative is an antisymmetric operator and its adjoint minus itself.

Moreover, the role of the sparse point objects has now switched: Instead of injecting the source term, we are now injecting the previously recorded receiver values into the adjoint wavefield, while we are interpolating the resulting wave at the original source location. As the injection and interpolations are part of the kernel, we also insure that these two are adjoints of each other.

```
def adjoint(model, m, eta, srca, rec, order=2):
    # Create the adjoint wavefield function
    v = TimeData(name='v', shape=model.shape,
                 time_order=2, space_order=order)

    # Derive stencil from symbolic equation
    # Note the inversion of the dampening term
    eqn = m * v.dt2 - v.laplace - eta * v.dt
    stencil = solve(eqn, u.forward)[0]
    update_v = [Eq(v.backward, stencil)]

    # Inject the previous receiver readings
    rec_term = rec.inject(field=v,
                          expr=rec * dt**2 / m)

    # Interpolate the adjoint-source
    srca_term = srca.interpolate(expr=v)

    # Create operator with source and receiver terms
    return Operator(update_v + rec_term + srca_term,
                    subs={s: dt, h: model.spacing},
                    time_axis=Backward)
```

Having established how to build the required operators we can now define the workflow for our adjoint example. For illustration purposes we are using a utility object `Model` that provides the core information for seismic inversion runs, such as the values for `m` and the dampening term `eta`, as well as the coordinates of the point source and receiver hydrophones. It is worth noting that the spatial discretisation and thus the stencil size of the operators is still fully parameterisable.

```
# Create the seismic model of the domain
model = Model(...)

# Create source with Ricker wavelet
src = PointData(name='src', ntime=ntime,
                ndim=2, npoint=1)
src.data[0, :] = ricker_wavelet(ntime)
src.coordinates.data[:] = source_coords

# Create empty set of receivers
rec = PointData(name='rec', ntime=ntime,
                ndim=2, npoint=101)
rec.coordinates.data[:] = receiver_coords

# Create empty adjoint source symbol
srca = PointData(name='srca', ntime=ntime,
                 ndim=2, npoint=1)
srca.coordinates.data[:] = source_coords

# Create symbol for square slowness
m = DenseData(name='m', shape=model.shape,
              space_order=order)
```




Fig. 5: Shot record of the measured point values in `rec.data` after the forward run.

```
m.data[:] = model # Set m from model data

# Create dampening term from model
eta = DenseData(name='eta', shape=shape,
                space_order=order)
eta.data[:] = model.dampening

# Execute forward and adjoint runs
fwd = forward(model, m, eta, src, rec)
fwd(time=ntime)
adj = adjoint(model, m, eta, srca, rec)
adj(time=ntime)

# Test prescribed against adjoint source
adjoint_test(src.data, srca.data)
```

The adjoint test is the core definition of the adjoint of a linear operator. The mathematical correctness of the adjoint is required for mathematical adjoint-based optimizations methods that are only guaranteed to converged with the correct adjoint. The test can be written as:

$$\langle src, adjoint(rec) \rangle = \langle forward(src), rec \rangle$$

The adjoint test can be used to verify the accuracy of the forward propagation and adjoint operators and has been shown to agree for 2D and 3D implementations [Louboutin17b]. The shot record of the data measured at the receiver locations after the forward run is shown in Figure 5.

Automated code generation

The role of the `Operator` in the previous examples is to generate semantically equivalent C code to the provided SymPy expressions, complete with loop constructs and annotations for performance optimization, such as OpenMP pragmas. Unlike many other DSL-based frameworks, Devito employs actual compiler technology during the code generation and optimization process. The symbolic specification is progressively lowered to C code through a series of passes manipulating abstract syntax trees (AST), rather than working with rigid templates. This software engineering choice has an invaluable impact on maintainability, extensibility and composability.

Following the initial resolution of explicit grid indices into the low-level format, Devito is able to apply several types of automated performance optimization throughout the code generation pipeline, which are grouped into two distinct sub-modules:

- **DSE - Devito Symbolic Engine:** The first set of optimization passes consists of manipulating SymPy equations with the aim to decrease the number of floating-point operations performed when evaluating a single grid point. This initial optimization is performed following an initial analysis of the provided expressions and consists of sub-passes such as common sub-expressions elimination, detection and promotion of time-invariants, and factorization of common finite-difference weights. These transformations not only optimize the operation count, but they also improve the symbolic processing and low-level compilation times of later processing stages.
- **DLE - Devito Loop Engine:** After the initial symbolic processing Devito schedules the optimised expressions in a set of loops by creating an Abstract Syntax Tree (AST). The loop engine (DLE) is now able to perform typical loop-level optimizations in multiple passes by manipulating this AST, including data alignment through array annotations and padding, SIMD vectorization through OpenMP pragmas and thread parallelism through OpenMP pragmas. On top of that, loop blocking is used to fully exploit the memory bandwidth of a target architecture by increasing data locality and thus cache utilization. Since the effectiveness of the blocking technique is highly architecture-dependent, Devito can determine optimal block size through runtime auto-tuning.

Performance Benchmark

The effectiveness of the automated performance optimization performed by the Devito backend engines can be demonstrated using the forward operator constructed in the above example. The following performance benchmarks were run with for a three-dimensional grid of size $512 \times 512 \times 512$ with varying spatial discretisations resulting in different stencil sizes with increasing operational intensity (OI). The benchmark runs were performed on on a Intel(R) Xeon E5-2620 v4 2.1Ghz "Broadwell" CPU with a single memory socket and 8 cores per socket and the slope of the roofline models was derived using the Stream Triad benchmark [McCalpin95].

The first set of benchmark results, shown in Figure 6, highlights the performance gains achieved through loop-level optimizations. For these runs the symbolic optimizations were kept at a "basic" setting, where only common sub-expressions elimination is performed on the kernel expressions. Of particular interest are the performance gains achieved by increasing the loop engine mode from "basic" to "advanced", to insert loop blocking and explicit vectorization directives into the generated code. Due to the improved memory bandwidth utilization the performance increased to between 52% and 74% of the achievable peak. It is also worth noting that more aggressive optimization in the "speculative" DLE mode (directives for non-temporal stores and row-wise data alignment through additional padding) did not yield any consistent improvements due to the low OI inherent to the acoustic formulation of the wave equation and the subsequent memory bandwidth limitations of the kernel.

On top of loop-level performance optimizations, Figure 7 shows the achieved performance with additional symbolic optimizations and flop reductions enabled. While the peak performance shows only small effects from this set of optimizations due to the inherent memory bandwidth limitations of the kernel, it is interesting to note a reduction in operational intensity between

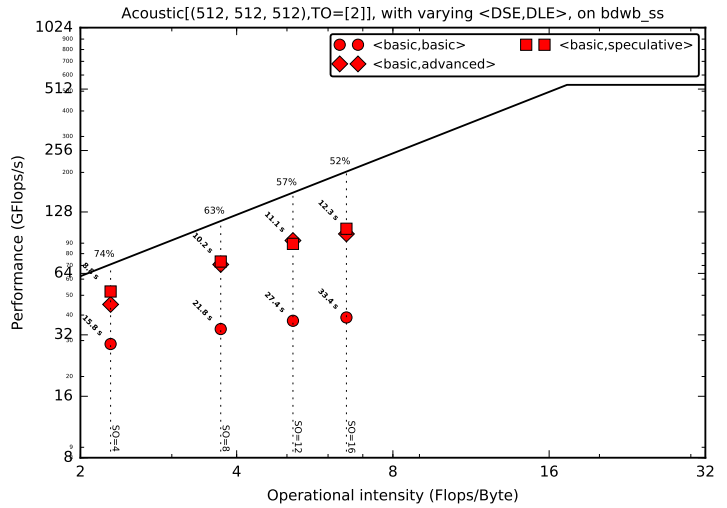


Fig. 6: Performance benchmarks for loop-level optimizations with different spatial orders (SO). The symbolic optimizations (DSE) have been kept at level 'basic', while loop optimisation levels (DLE) vary.

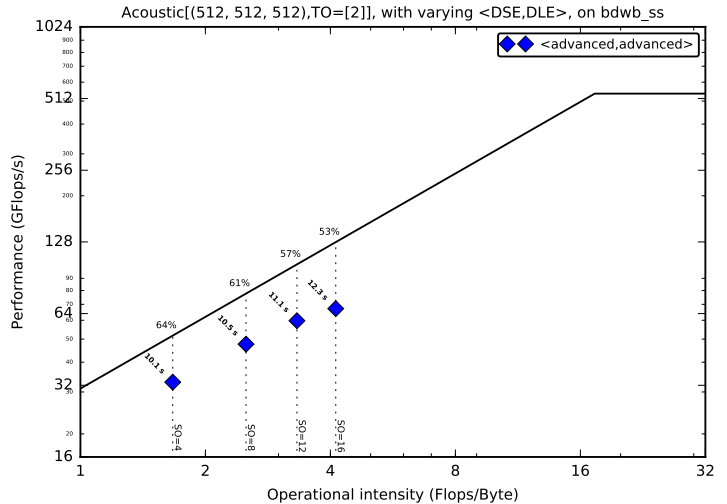


Fig. 7: Performance benchmarks with full symbolic and loop-level optimizations for different spatial orders (SO).

equivalent stencil sizes in Figures 6 and 7. This entails that, despite only marginal runtime changes, the generated code is performing less flops per stencil point, which is of vital importance for compute-dominated kernels with large OI [Louboutin17a].

Integration with YASK

As mentioned previously, Devito is based upon actual compiler technology with a highly modular structure. Each backend transformation pass is based on manipulating an input AST and returning a new, different AST. One of the reasons behind this software engineering strategy, which is clearly more challenging than a template-based solution, is to ease the integration of external tools, such as the YASK stencil optimizer [Yount16]. We are currently in the process of integrating YASK to complement the DLE, so that YASK may replace some (but not all) DLE passes.

The DLE passes are organized in a hierarchy of classes where each class represents a specific code transformation pipeline

based on AST manipulations. Integrating YASK becomes then a conceptually simple task, which boils down to three actions:

- 1) Adding a new transformation pipeline to the DLE.
- 2) Adding a new array type, to ease storage layout transformations and data views (YASK employs a data layout different than the conventional row-major format).
- 3) Creating the proper Python bindings in YASK so that Devito can drive the code generation process.

It has been shown that real-world stencil codes optimised through YASK may achieve an exceptionally high fraction of the attainable machine peak [Yount15], [Yount16]. Further, initial prototyping (manual optimization of Devito-generated code through YASK) revealed that YASK may also outperform the loop optimization engine currently available in Devito, besides ensuring seamless performance portability across a range of computer architectures. On the other hand, YASK is a C++ based framework that, unlike Devito, does not rely on symbolic mathematics and processing; in other words, it operates at a much lower level of abstraction. These observations, as well as the outcome of the initial prototyping phase, motivate the on-going Devito-YASK integration effort.

Discussion

In this paper we present the finite difference DSL Devito and demonstrate its high-level API to generate two fluid dynamics operators and a full seismic inversion example. We highlight the relative ease with which to create complex operators from only a few lines of high-level Python code while utilising highly optimised auto-generated C kernels via JIT compilation. On top of purely symbolic top-level API based on SymPy, we show how to utilise Devito's secondary API to inject custom expressions into the code generation toolchain to implement Dirichlet and Neumann boundary conditions, as well as the sparse-point interpolation routines required by seismic inversion operators.

Moreover, we demonstrate that Devito-generated kernels are capable of exploiting modern high performance computing architectures by achieving a significant percentage of machine peak. Devito's code-generation engines achieve this by automating well-known performance optimizations, as well as domain-specific optimizations, such as flop reduction techniques - all while maintaining full compatibility with the scientific software stack available through the open-source Python ecosystem.

Limitations and Future Work

The examples used in this paper have been chosen for their relative simplicity in order to concisely demonstrate the current features of the Devito API. Different numerical methods may be used to solve the presented examples with greater accuracy or achieve more realistic results. Nevertheless, finite difference methods play an important role and are widely used in academic and industrial research due to the relative ease of implementation, verification/validation and high computational efficiency, which is of particular importance for inversion methods that require fast and robust high-order PDE solvers.

The interfaces provided by Devito are intended to create high-performance operators with relative ease and thus increase user productivity. Several future extensions are planned to enhance the high-level API to further ease the construction of more complex operators, including explicit abstractions for symbolic boundary

conditions, perfectly matched layer (PML) methods and staggered grids. Devito's secondary low-level API and use of several intermediate representations are intended to ease the gradual addition of new high-level features.

Moreover, the addition of YASK as an alternative backend will not only provide more advanced performance optimisation, but also an MPI infrastructure to allow Devito to utilise distributed computing environments. Further plans also exist for integration with linear and non-linear solver libraries, such as PETSc, to enable Devito to handle implicit formulations.

Acknowledgements

This work was financially supported in part by EPSRC grant EP/L000407/1 and the Imperial College London Intel Parallel Computing Centre. This research was carried out as part of the SINBAD project with the support of the member organizations of the SINBAD Consortium. Part of this work was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics and Computer Science programs under contract number DE-AC02-06CH11357.

REFERENCES

- [Alnaes14] M. S. Alnaes, A. Logg, K. B. Ølgaard, M. E. Rognes, and G. N. Wells, "Unified Form Language: a domain-specific language for weak formulations of partial differential equations", *ACM Transactions on Mathematical Software (TOMS)*, vol. 40, no. 2, p. 9, 2014. <https://dx.doi.org/10.1145/2566630>
- [Baba16] Y. Baba and V. Rakov, "The Finite-Difference Time Domain Method for Solving Maxwell's Equations", in "Electromagnetic Computation Methods for Lightning Surge Protection Studies", 2016, pp. 43–72, Wiley, ISBN 9781118275658. <http://dx.doi.org/10.1002/9781118275658.ch3>
- [Brandvik10] T. Brandvik and G. Pullan, "Sblock: A framework for efficient stencil-based pde solvers on multi-core platforms", in "Proceedings of the 2010 10th IEEE International Conference on Computer and Information Technology", IEEE Computer Society, 2010, pp. 1181–1188. <http://dx.doi.org/10.1109/CIT.2010.214>
- [Cardenas70] Cárdenas, A. F. and Karplus, W. J.: PDEL — a language for partial differential equations, *Communications of the ACM*, 13, 184–191, 1970.
- [Cook88] Cook Jr, G. O.: ALPAL: A tool for the development of large-scale simulation codes, Tech. rep., Lawrence Livermore National Lab., CA (USA), 1988.
- [Datta08] K. Datta, M. Murphy, V. Volkov, S. Williams, J. Carter, L. Oliker, D. Patterson, J. Shalf, and K. Yelick, "Stencil computation optimization and auto-tuning on state-of-the-art multicore architectures", in *Proceedings of the 2008 ACM/IEEE Conference on Supercomputing*, IEEE Press, 2008, pp. 4:1–4:12. <http://dl.acm.org/citation.cfm?id=1413370.1413375>
- [Farrell13] Farrell, P. E., Ham, D. A., Funke, S. W., and Rognes, M. E.: Automated Derivation of the Adjoint of High-Level Transient Finite Element Programs, *SIAM Journal on Scientific Computing*, 35, C369–C393, 2013. <http://dx.doi.org/10.1137/120873558>
- [Henretty13] T. Henretty, R. Veras, F. Franchetti, L.-N. Pouchet, J. Ramanujam, and P. Sadayappan, "A stencil compiler for short-vector simd architectures," in *Proceedings of the 27th International ACM Conference on International Conference on Supercomputing*, ACM, 2013, pp. 13–24. <http://doi.acm.org/10.1145/2464996.2467268>
- [Iverson62] Iverson, K.: *A Programming Language*, Wiley, 1962.
- [Krkos12] J.A. Krkos, "Unsteady Adjoint Analysis for Output Sensitivity and Mesh Adaptation", PhD thesis, 2012. <https://dspace.mit.edu/handle/1721.1/77133>
- [Lange17] Lange, M., Luporini, F., Louboutin, M., Kukreja, N., Pandolfo, V., Kazakas, P., Velesko, P., Zhang, S., Peng, P., and Gorman, G. Dylan McCormick. 2017, June 7. *opesci/devito: Devito-3.0.1*. Zenodo. <https://doi.org/10.5281/zenodo.823172>

- [LeVeque92] LeVeque, R. J., "Numerical Methods for Conservation Laws", Birkhauser-Verlag (1992).
- [Liu09] Y. Liu and M. K. Sen, "Advanced Finite-Difference Method for Seismic Modeling," *Geohorizons*, Vol. 14, No. 2, 2009, pp. 5-16.
- [Logg12] Logg, A., Mardal, K.-A., Wells, G. N., et al.: Automated Solution of Differential Equations by the Finite Element Method, Springer, doi:10.1007/978-3-642-23099-8, 2012.
- [Louboutin17a] Louboutin, M., Lange, M., Herrmann, F. J., Kukreja, N., and Gorman, G.: Performance prediction of finite-difference solvers for different computer architectures, *Computers Geosciences*, 105, 148—157, <https://doi.org/10.1016/j.cageo.2017.04.014>, 2017.
- [Louboutin17b] M. Louboutin, M. Lange, F. Luporini, N. Kukreja, F. Herrmann, P. Velesko, and G. Gorman: Code generation from symbolic finite-difference for geophysical exploration. In preparation for Geoscientific Model Development (GMD), 2017.
- [McCalpin95] McCalpin, J. D., "Memory Bandwidth and Machine Balance in Current High Performance Computers", IEEE Computer Society Technical Committee on Computer Architecture (TCCA) Newsletter, December 1995.
- [Meurer17] Meurer A, Smith CP, Paprocki M, Čertík O, Kirpichev SB, Rocklin M, Kumar A, Ivanov S, Moore JK, Singh S, Rathnayake T, Vig S, Granger BE, Muller RP, Bonazzi F, Gupta H, Vats S, Johansson F, Pedregosa F, Curry MJ, Terrel AR, Roučka Š, Saboo A, Fernando I, Kulal S, Cimrman R, Scopatz A. (2017) SymPy: symbolic computing in Python. *PeerJ Computer Science* 3:e103 <https://doi.org/10.7717/peerj-cs.103>
- [Peiro05] J. Peiró, S. Sherwin, "Finite Difference, Finite Element and Finite Volume Methods for Partial Differential Equations", in "Handbook of Materials Modeling, pp. 2415—2446, ISBN 978-1-4020-3286-8, 2005. http://dx.doi.org/10.1007/978-1-4020-3286-8_127.
- [Rai91] M. M. Rai and P. Moin. 1991. "Direct simulations of turbulent flow using finite-difference schemes", *J. Comput. Phys.* 96, 1 (October 1991), 15-53. [http://dx.doi.org/10.1016/0021-9991\(91\)90264-L](http://dx.doi.org/10.1016/0021-9991(91)90264-L)
- [Rathgeber16] Rathgeber, F., Ham, D. A., Mitchell, L., Lange, M., Luporini, F., McRae, A. T. T., Bercea, G., Markall, G. R., and Kelly, P. H. J.: "Firedrake: automating the finite element method by composing abstractions", *ACM Trans. Math. Softw.*, 43(3):24:1–24:27, 2016. <http://dx.doi.org/10.1145/2998441>.
- [Umetani85] Umetani, Y.: DEQSOL A numerical Simulation Language for Vector/Parallel Processors, Proc. IFIP TC2/WG22, 1985, 5, 147–164, 1985.
- [VanEngelen96] R. Van Engelen, L. Wolters, and G. Cats, "Ctadel: A generator of multi-platform high performance codes for pde-based scientific applications," in Proceedings of the 10th international conference on Supercomputing. ACM, 1996, pp. 86–93.
- [Virieux09] Virieux, J. and Operto, S., "An overview of full-waveform inversion in exploration geophysics", *GEOPHYSICS*, 74, WCC1–WCC26, 2009. <http://dx.doi.org/10.1190/1.3238367>
- [Yount15] C. Yount, "Vector Folding: Improving Stencil Performance via Multi-dimensional SIMD-vector Representation," 2015 IEEE 17th International Conference on High Performance Computing and Communications, 2015 IEEE 7th International Symposium on Cyberspace Safety and Security, and 2015 IEEE 12th International Conference on Embedded Software and Systems, New York, NY, 2015, pp. 865-870. <https://doi.org/10.1109/HPCC-CSS-ICSS.2015.27>
- [Yount16] C. Yount, J. Tobin, A. Breuer and A. Duran, "YASK — Yet Another Stencil Kernel: A Framework for HPC Stencil Code-Generation and Tuning," 2016 Sixth International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing (WOLFHPC), Salt Lake City, UT, 2016, pp. 30-39. <https://doi.org/10.1109/WOLFHPC.2016.08>
- [Zhang12] Y. Zhang and F. Mueller, "Auto-generation and auto-tuning of 3d stencil codes on gpu clusters," in Proceedings of the Tenth International Symposium on Code Generation and Optimization, ACM, 2012, pp. 155–164. <http://doi.acm.org/10.1145/2259016.2259037>