

# Devito: Towards a generic Finite Difference DSL using Symbolic Python

Michael Lange<sup>\*</sup>, Navjot Kukreja<sup>†</sup>, Mathias Louboutin<sup>‡</sup>, Fabio Luporini<sup>§</sup>, Felipe Vieira<sup>†</sup>,  
Vincenzo Pandolfo<sup>§</sup>, Paulius Veleško<sup>¶</sup>, Paulius Kazakas<sup>||</sup>, Gerard Gorman<sup>\*</sup>

<sup>\*</sup> Department of Earth Science and Engineering, Imperial College London, UK

<sup>†</sup> SENAI CIMATEC, Salvador, Brazil

<sup>‡</sup> Seismic Lab. for Imaging and Modeling, The University of British Columbia, Canada

<sup>§</sup> Department of Computer Science, Imperial College London, UK

<sup>¶</sup> College of Electrical and Computer Engineering, University of Oklahoma, USA

<sup>||</sup> Department of Computer Science, University of York, UK

**Abstract**—Domain specific languages (DSL) have been used in a variety of fields to express complex scientific problems in a concise manner and provide automated performance optimization for a range of computational architectures. As such DSLs provide a powerful mechanism to speed up scientific Python computation that goes beyond traditional vectorization and pre-compilation approaches, while allowing domain scientists to build applications within the comforts of the Python software ecosystem. In this paper we present Devito, a new finite difference DSL that provides optimized stencil computation from high-level problem specifications based on symbolic Python expressions. We demonstrate Devito’s symbolic API and performance advantages over traditional Python acceleration methods before highlighting its use in the scientific context of seismic inversion problems.

## I. INTRODUCTION

Python is one of the most popular high-level programming languages for academics and scientists due to its clean syntax, great expressiveness and a vast ecosystem of open source packages [1], [2], [3]. Its direct use in High Performance Computing (HPC), though, is complicated by the slow execution within the Python interpreter itself. Several mechanisms exist to overcome the interpreter overhead, ranging from static compilation via Cython [4] to generalised Just-In-Time (JIT) compilation via Numba [5].

However, for specific, often compute-intensive application areas such as solving partial differential equations (PDE) with the finite element method, domain specific languages (DSL) have also proven successful [6]. DSLs aim to decouple the problem specification from its low-level implementation to create a separation of concerns between domain scientists and HPC specialists, resulting in a direct payoff in productivity. Python is particularly useful for creating DSLs, since its runtime type inference and operator overloading via “magic functions” allows domain-specific constructs to be easily embedded in the language itself.

Moreover, due to their restriction to a single problem domain, code-generating DSLs can also augment the generated code to a particular computational architecture, allowing hardware-specific performance optimizations. In the age of

specialized HPC architectures, such as Intel<sup>®</sup> Xeon Phi<sup>™</sup> or GPU accelerators, this not only promises performance portability but also provides a valuable alternative to the continuing search for new all-encompassing parallel programming paradigms.

In this paper we present Devito<sup>1</sup>, a new finite difference DSL that utilises the SymPy [7] package to generate optimized parallel C code from high-level symbolic operator definitions. Devito is primarily targeted at generating fast wave propagation kernels for seismic inversion problems and therefore offers an abstraction hierarchy that enables not only the generation of fast stencil kernels, but also a range of domain-specific features, such as sparse point interpolation, required in a real scientific context. For this purpose Devito provides a two-level API that allows users to define operator kernels as a mixture of high-level symbolic equations and explicit C-like SymPy expressions to solve complex inversion problems.

## II. BACKGROUND

### A. Speeding up Python computation

Several methods exist to overcome the inherent interpreter overhead in Python, most of which focus on using pre-compiled functions to perform computationally intense loops. The most prominent example of this is the NumPy [8] array interface, which manages array data in contiguous blocks to avoid memory copies and provides vectorized calculations, where operations are applied element-wise over arrays using pre-compiled loops. The same technique is used by the SciPy package that adds a wealth of pre-compiled mathematical operators that can be applied in this vectorized way.

The effectiveness of NumPy vectorization, however, is limited by the need to create arrays with matching data layout, which often results in the “outer loop” being a pure Python loop. A more aggressive pre-compilation can be achieved with Cython [4], which allows users to annotate Python code with additional type information and invoke an explicit compilation

<sup>1</sup><https://github.com/opesci/devito>

step at install time at the expense of flexibility. Even more dynamic compilation can be achieved with Numba [5], a Just-In-Time (JIT) compiler that focuses on array-oriented computation and utilises LLVM to compile a subset of the Python language at runtime.

The Symbolic Python package SymPy used in Devito also provides JIT compilation via the `lambdify` utility function that compiles a symbolic expression into either a Python function that supports NumPy vectorization or pure C or Fortran kernels.

### B. DSLs embedded in Python

Interest in building generic DSLs for solving PDEs is not new with early attempts dating back as far as 1970 [9], [10], [11], [12]. More recently, two prominent finite element software packages, FEniCS [6] and Firedrake [13], have demonstrated the power of symbolic computation using the DSL paradigm. Both packages implement the Universal Form Language (UFL) [14] that allows scientists to express complex finite element problems symbolically in the weak form. Reducing the complexity of a problem to its symbolic definition then enables further symbolic abstraction, such as the automated generation of adjoint models, as demonstrated by Dolfin-Adjoint [15]. It is noteworthy that FEniCS, Firedrake and Dolfin-Adjoint all provide a native Python interface to the user, with the DSL embedded in the Python language itself.

The optimization of regular grid and stencil computations has also produced a vast range of libraries and DSLs that aim to ease the efficient automated creation of high-performance codes [16], [17], [18], [19], [20], [21], [22]. Most stencil DSLs, however, define their own custom input language, which limits their scope and applicability for practical applications. Since solving realistic problems often requires more than just a fast and efficient PDE solver, scientific applications can benefit from DSLs embedded in Python directly due to the native compatibility with the Python software ecosystem.

### C. Abstraction layers in DSL design

While DSL-based approaches promise great flexibility and productivity, care needs to be taken when designing such domain-specific software stacks. In particular the choice of abstraction layers and interfaces is crucial, since future development and practical use can often be hampered by too narrow abstractions, while all-encompassing packages often become too generic, complicated and hard to maintain. The separation of the high-level symbolic problem definition layer from the low-level implementation details of the generated code is of crucial importance here, since it provides a separation of concerns between the domain scientists and HPC specialists. This not only increases productivity on both sides, but also ensures maintainability of the generated software stack.

Active libraries, such as OP2 [23] for unstructured mesh applications or OPS [24] for structured grids, aim to decouple the problem specification from the underlying execution by providing an abstraction of data movement following the access-execute paradigm [25]. A particular implementation of such an abstraction is provided by PyOP2 [26], which acts

as the underlying unstructured mesh traversal engine for the previously mentioned Firedrake finite element system. Similar mechanisms are provided by the Halide project [27], a domain specific language, compiler, and auto-tuning system that was originally designed for image processing pipelines

## III. DEVITO DESIGN AND API

In contrast to traditional Python-embedded DSLs, Devito does not define its own high-level language *per se*. Instead, it relies on SymPy to provide the building blocks for users to express a finite difference operator as a symbolic equation. Devito, however, annotates these objects with additional meta-data and provides new abstractions specific to defining finite difference problems. This allows Devito to automate the transformation of high-level symbolic expressions to optimized executable C code by maintaining and manipulating SymPy expression objects throughout the code generation process, while leveraging symbolic manipulation utilities provided by SymPy.

The automated code generation process is performed in multiple steps and follows the *Principle of Graceful Degradation*, which suggests that users should be able to circumvent limits to the top layer abstraction by utilising a lower-level API. A common example for this is sparse point interpolation, which is required by many scientific applications to extract solution values at specific points, but can not easily be expressed as a high-level PDE. For such cases Devito provides users with a lower level API to allow custom expressions to be added to the generated code in a C-like manner.

Fig. 1 highlights the overall design hierarchy of Devito and shows the individual layers of transformations. At the top, user can define data object that associate `sympy.Function` symbols with data buffers and use them to derive stencil equations in a concise symbolic format. The resulting stencil expressions are automatically expanded and transformed into explicit array accesses when the user builds a `devito.Operator` object from the SymPy expression. A secondary low-level kernel API is then utilised to generate optimized C code from the intermediate representation, and a set of compiler presets is used to perform Just-in-Time (JIT) compilation before executing the kernel operator over the given data.

### A. High-level API: Derivatives and equations

One of the core ideas behind Devito is to associate simulation data with symbolic objects to leverage the symbolic manipulation capabilities provided by SymPy during code generation. For this purpose Devito provides `DenseData` and `TimeData` objects that behave like `sympy.Function` objects, but also encapsulate data buffers and present them to the user as wrapped NumPy arrays. Importantly, Devito is thus able to manage data allocation internally to provide data caching on disk and first touch allocation to ensure data locality on NUMA systems.

To ensure associated meta-data is not lost during symbolic manipulation, Devito maintains its own symbol cache that “re-attaches” Devito-specific information whenever a new instance of a symbol is created. Since SymPy uses frequent

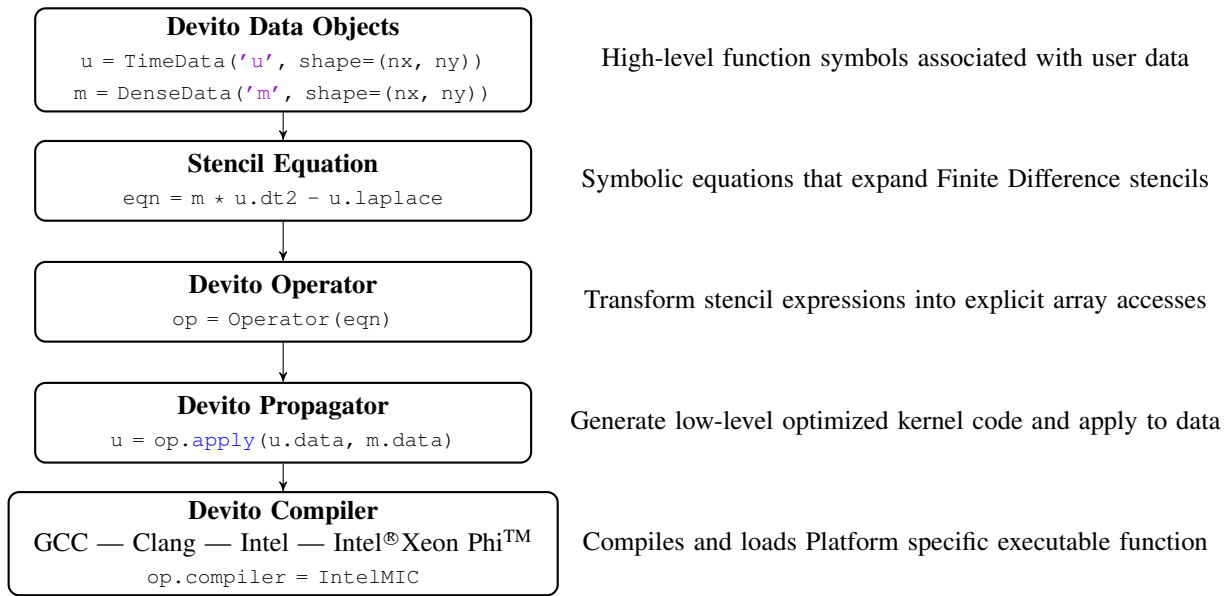


Fig. 1: The Devito code generation process utilises multiple abstraction layers that allow users to customise each individual transformation in the toolchain, ranging from high-level symbolic equation definitions to low-level compiler settings.

re-instantiation of `Function` objects to facilitate symbolic substitutions in expressions, this enables Devito symbols to behave opaquely during symbolic manipulation, for example when forming finite difference stencils. The following example demonstrates how a two-dimensional symbolic data object  $f(x, y)$  maintains the shape information of its associated data buffer (`f.data`), while expanding the first derivative in dimension  $x$  using SymPy's `as_finite_diff` utility.

---

```

In [1]: from Devito import DenseData
In [2]: from sympy import as_finite_diff
In [3]: from sympy.abc import x, h

In [4]: f = DenseData(name='f', shape=(10, 12))
Out[4]: f(x, y)

In [5]: eqn = as_finite_diff(f.diff(x), [x, x+h])
Out[5]: -f(x, y)/h + f(h + x, y)/h

In [6]: f_h = eqn.args[0].args[1]
Out[6]: f(h + x, y)

In [7]: f_h.data.shape
Out[7]: (10, 12)

```

---

Moreover, Devito objects also provide shorthand notation for common finite difference formulations, such as first, second and cross derivatives in the space and time dimensions. The order of the stencil discretization is hereby defined on the symbolic object itself, which allows us to express the second derivative of  $f$  in  $x$  as:

---

```

In [1]: from devito import DenseData

In [2]: f = DenseData(name='f', shape=(10, 10),
                    space_order=2)
Out[2]: f(x, y)

In [3]: f.dx2
Out[3]: -2*f(x, y)/h**2 + f(-h + x, y)/h**2
        + f(h + x, y)/h**2

```

---

In the resulting stencil expression Devito inserts the symbol  $h$  to denote grid spacing in the  $x$  dimension. It is important to note here that this notation allows users to quickly increase the order of the stencil discretization by simply changing a single parameter during symbol construction.

Another utility property of Devito's symbolic data objects is the `f.laplace` operator that expands to  $(f.dx2 + f.dy2)$  for two-dimensional problems and  $(f.dx2 + f.dy2 + f.dz)$  for three-dimensional ones. This allows the dimensionless expression of complex equations, such as a simple wave equation, to be defined very concisely:

---

```

In [1]: from devito import DenseData
In [2]: from sympy import Eq

In [3]: u = TimeData(name='u', space_order=2,
                    time_order=2, shape=(10, 12))
In [4]: m = DenseData(name='m', space_order=2,
                    shape=(10, 12))

In [5]: eqn = Eq(m * u.dt2, u.laplace)
Out[5]: Eq((-2*u(t, x, y)/s**2
            + u(-s + t, x, y)/s**2
            + u(s + t, x, y)/s**2)*m(x, y),
            -4*u(t, x, y)/h**2
            + u(t, x, -h + y)/h**2
            + u(t, x, h + y)/h**2
            + u(t, -h + x, y)/h**2
            + u(t, h + x, y)/h**2)

```

---

The most important abstraction provided by Devito however is the `Operator` class that allows users to trigger the automated code generation and compilation cascade. Users simply need to provide an expanded stencil expression and optional compiler settings via `op = Operator(eqn, compiler=IntelCompiler)`. The resulting operator object can then be used to execute the generated kernel code via `op.apply(u, m)`.

## B. Low-level API: Expressions and loops

One of the key transformations performed during the code generation process is to resolve the grid spacing variables, for example `h`, into explicit grid indices. During this transformation Devito converts symbolic functions like  $f(x, y)$  into `sympy.Indexed` expressions of the form `f[x, y]`, before converting the dimension symbols to loop counters to yield `f[i1, i2]`. It is important to note here that the intermediate indexed representation with explicit dimensions is also accepted input when creating `Operator` instances, allowing users to input custom expressions with irregular grid indexing. The low-level API is also a useful tool for developers to rapidly test and develop additional higher-level abstractions.

Moreover, a new low-level API is currently under development that will allow the automatic inference of iteration spaces and loop constructs from the indexed stencil expression itself. The key to this are `Dimension` objects that act like index symbols (`x`, `y` or `t`) and define the iteration space over which to loop in the generated kernel code, while defining additional meta-information like boundary padding layers or buffer switching for time loops. This new API will allow the generic and automated construction of arbitrary loop nests via provided `Iteration` and `Expression` classes, enabling the rapid development of further mathematical abstractions in the high-level API. This will also provide the possibility of closer integration with dedicated access-execute frameworks (see section II-C) in the future to further leverage automated loop-level performance optimizations.

## C. Performance optimization

Devito provides a set of automated performance optimizations during code generation that allow user applications to fully utilise the target hardware without changing the model specification. The optimizations provided by Devito range from symbolic manipulations applied to the provided stencil expressions to hardware-specific implementation choices like loop blocking.

1) *Shared-memory Parallelism*: Devito provides thread-parallel execution through automated code annotation with OpenMP pragmas inserted at the appropriate locations in the code, with a single thread pool allocated outside of the primary timestepping loop. On multi-socket systems Devito performs so called first touch memory allocation on its symbolic data objects ensuring high data locality on NUMA architectures.

2) *Vectorization*: SIMD parallelism and vectorization are a crucial performance optimization for compute intensive stencil codes to fully exploit modern HPC architectures. Devito ensures efficient vectorization by inserting toolchain-specific compiler hints into the generated code and enforcing data alignment on page boundaries during allocation.

3) *Loop Blocking*: Loop or cache blocking is a well-known performance optimization for stencil codes that aims to increase data locality and thus cache utilisation by splitting iteration spaces (loops) into spatially aligned blocks. The effectiveness of the technique is highly dependent on the cache size of the target architecture. Devito therefore provides an interfaces for users to define block sizes explicitly, as well as

an “auto-tune” mode where a defined `Operator` object is used to identify the optimal block size through a brute-force search.

4) *Common sub-expression elimination*: On top of architecture-specific optimizations Devito also aims to leverage symbolic optimizations to speed up computation. An important example of this is common sub-expression elimination (CSE) that aims to reduce the amount of redundant computation within complex stencil expressions by factoring out common terms. While it may be argued that most modern optimizing compilers are able to do this, it is important to note here that Devito utilises SymPy’s internal CSE functionality to apply the technique at a much higher abstraction level.

As a result, processing times of symbolic expressions are reduced during the code generation phase, and the compilation time of the generated code itself is improved. This is important since complex stencils can increase compilation time to several hours if CSE is performed by the compiler, in particular for high-order seismic inversion kernels that often contain very large numbers of terms.

## IV. 2D DIFFUSION EXAMPLE

The use of symbolic notations to define finite difference operators strongly enhances the efficiency of model developers and allows well known problems to be expressed in a concise manner. In this section we will first demonstrate the use of symbolic Python notation in comparison to pure Python code on a simple two-dimensional diffusion example. We will also analyse the performance benefits of Devito operators over vectorized NumPy array notation, before demonstrating the construction of seismic inversion operators using Devito’s symbolic interfaces in the next section.

The general diffusion or heat equation in two dimensions is commonly notated as

$$\frac{\partial u}{\partial t} = \alpha \nabla^2 u \quad (1)$$

where  $u$  denotes the unknown function to solve for,  $\alpha$  is the diffusion coefficient and  $\nabla^2 u$  denotes the Laplace operator, as defined in section III-A. To keep the following examples short we will be using a forward Euler approach with constant grid spacing in  $x$  and  $y$ , while the the timestep size is computed as

$$\Delta t = \frac{\Delta x^2 \Delta y^2}{2\alpha(\Delta x^2 + \Delta y^2)} \quad (2)$$

where  $\Delta x$ ,  $\Delta y$  denotes the grid spacing and  $\Delta t$  is the timestep size.

### A. Pure Python

A first direct implementation of the described equation using an alternating buffer scheme for  $u$  in pure Python is shown in Listing 1. The implementation consists of three loops that apply a 5-point star stencil - a 3-point stencil in  $x$  ( $u_{xx}$ ) and a 3-point stencil in  $y$  ( $u_{yy}$ ) to compute the update. It is important to note here that the stencil is only applied to elements in the interior of the grid in order to avoid out-of-bound array accesses.

Listing 1: Diffusion code using pure Python loops.

---

```
_, nx, ny = u.shape
for t in range(timesteps):
    t0 = t % 2
    t1 = (t + 1) % 2
    for i in range(1, nx-1):
        for j in range(1, ny-1):
            uxx = (u[t0, i+1, j] - 2*u[t0, i, j]
                  + u[t0, i-1, j]) / dx2
            uyy = (u[t0, i, j+1] - 2*u[t0, i, j]
                  + u[t0, i, j-1]) / dy2
            u[t1, i, j] = u[t0, i, j]
                          + dt * a * (uxx + uyy)
```

---

Listing 2: Diffusion code with vectorized NumPy arrays.

---

```
for t in range(timesteps):
    t0 = t % 2
    t1 = (t + 1) % 2
    uxx = (u[t0, 2:, 1:-1] - 2*u[t0, 1:-1, 1:-1]
           + u[t0, :-2, 1:-1]) / dx2
    uyy = (u[t0, 1:-1, 2:] - 2*u[t0, 1:-1, 1:-1]
           + u[t0, 1:-1, :-2]) / dy2
    u[t1, 1:-1, 1:-1] = u[t0, 1:-1, 1:-1]
                       + a * dt * (uxx + uyy)
```

---

## B. Vectorized NumPy arrays

Since the loop performance overhead of the Python interpreter is well known, the most common optimization of simple arithmetic computations on array and grid data is to utilise NumPy vectorization. This applies arithmetic operations element-wise over all elements in an array using pre-compiled kernels. For stencil expressions it is therefore important to adjust the bounds of the array view that represents each stencil point, as demonstrated in Listing 2.

## C. Symbolic Python - *lambdified*

An alternative approach to specifying the loop kernel that implements a particular equation directly is to define it symbolically via SymPy, as demonstrated in Listing 3. The diffusion equation can thus be expressed in a manner that is mathematically concise, while SymPy utilities are used to expand the derivative stencils and create an executable function via `sympy.lambdify`. It is noteworthy that the `'numpy'` mode argument allows us to utilise vectorized operations on sub-arrays, similar to the pure NumPy variant.

## D. Symbolic Python - *Devito*

A similar symbolic approach is provided by Devito, where the expansion of the derivative stencils is automated by the provided symbolic data objects. As shown in Listing 4, this allows the stencil discretization to be adjusted via a single parameter change, while the equation definition and reorganisation of the stencil expression via `sympy.solve` is identical to the pure SymPy approach. The `devito.Operator` created from the expression is then used to generate low-level code and execute it for a given number of timesteps.

Listing 3: Diffusion code using symbolic SymPy objects.

---

```
from sympy import Eq, Function, lambdify
from sympy.abc import x, y, t, s, h
p = Function('p')
dx2 = as_finite_diff(p(x, y, t).diff(x, x),
                    [x - h, x, x + h])
dy2 = as_finite_diff(p(x, y, t).diff(y, y),
                    [y - h, y, y + h])
dt = as_finite_diff(p(x, y, t).diff(t), [t, t+s])
eqn = Eq(dt, a * (dx2 + dy2))
stencil = solve(eqn, p(x, y, t + s))[0]
subs = (p(x, y, t), p(x+h, y, t), p(x-h, y, t),
        p(x, y+h, t), p(x, y-h, t), s, h, a)
kernel = lambdify(subs, stencil, 'numpy')
```

---

```
for ti in range(timesteps):
    t0 = ti % 2
    t1 = (ti + 1) % 2
    u[t1, 1:-1, 1:-1] = kernel(
        u[t0, 1:-1, 1:-1], u[t0, 2:, 1:-1],
        u[t0, :-2, 1:-1], u[t0, 1:-1, 2:]
        u[t0, 1:-1, :-2], dt, dx)
```

---

Listing 4: Diffusion code using the symbolic Devito API.

---

```
from devito import TimeData, Operator
from sympy.abc import s, h
u = TimeData(name='u', shape=(nx, ny),
             time_order=1, space_order=2)
u.data[0, :] = ui[:]
```

---

```
eqn = Eq(u.dt, a * (u.dx2 + u.dy2))
stencil = solve(eqn, u.forward)[0]
op = Operator(stencils=Eq(u.forward, stencil),
             subs={h: dx, s: dt}, nt=timesteps)
op.apply()
```

---

An example of the auto-generated C code, is shown in Fig. 5. The example demonstrates Devito's use of parallelisation and vectorization pragmas, as well as its internal resolution of alternating buffer access. Moreover it can be seen that Devito has replaced all numerical constants, such as the diffusion coefficient  $a$  and spacing parameters  $dx$  and  $dt$ , and inserted all known loop dimensions to aid compiler optimizations.

## E. Performance comparison

To demonstrate the performance advantages of custom code generation and JIT compilation as performed in Devito over traditional approaches to Python acceleration, we have compared the single core performance of the vectorized NumPy and SymPy implementations of the diffusion example. For this purpose the diffusion benchmark was run for 500 timesteps with a grid size of 0.001 on a single core of a Intel® Core i3-4030U CPU. The results presented in Fig. 2 show a clear performance advantage of the DSL approach, as well as highlighting a significant overhead of the vectorized SymPy implementation over the pure NumPy approach. It is important to note here that the Devito DSL is the only benchmarked approach that also encapsulates the timestepping loop in the compilation step and that further performance advantages can

Listing 5: Auto-generated C code to solve diffusion example.

```
extern ``C`` int Operator(float *u_vec)
{
    float (*u)[1000][1000] = (float (*)[1000][1000]) u_vec;
    {
        int t0;
        int t1;
        #pragma omp parallel
        for (int i3 = 0; i3<500; i3+=1)
        {
            #pragma omp single
            {
                t0 = (i3)%2;
                t1 = (t0 + 1)%2;
            }
            {
                #pragma omp for schedule(static)
                for (int i1 = 1; i1<999; i1++)
                {
                    #pragma omp simd aligned(u:64)
                    for (int i2 = 1; i2<999; i2++)
                    {
                        u[t1][i1][i2] = 2.5e-1F*u[t0][i1][i2-1]
                            + 2.5e-1F*u[t0][i1][i2+1]
                            + 2.5e-1F*u[t0][i1-1][i2]
                            + 2.5e-1F*u[t0][i1+1][i2];
                    }
                }
            }
        }
    }
}
```

be expected when shared-memory parallelism is utilised, as we will show in section V-D.




Fig. 2: Serial runtime comparison between NumPy, SymPy and Devito for the diffusion example.

## V. SEISMIC INVERSION EXAMPLE

The key motivating application behind the creation of the Devito finite difference DSL are seismic exploration problems including advanced modelling and adjoint-based inversion. These not only require a fast and accurate wave propagation model, but also efficient, easy-to-define and rigorous adjoint operators. In the following example we demonstrate the symbolic definition of modelling and adjoint modelling operations for the acoustic wave equation denoted as:

$$m \frac{\partial^2 u}{\partial t^2} + \eta \frac{\partial u}{\partial t} - \nabla^2 u = q \quad (3)$$

where  $u$  denotes the pressure wave field,  $m$  is the square slowness,  $q$  is the source term and  $\eta$  denotes the spatially

Listing 6: Example code to solve the wave equation

```
def forward(model, nt, dt, h, order=2):
    shape = model.shape
    m = DenseData(name="m", shape=shape,
                  space_order=order)
    m.data[:] = model # Set model data
    u = TimeData(name="u", shape=model.shape,
                 time_dim=nt, time_order=2,
                 space_order=order)
    eta = DenseData(name="eta", shape=shape,
                   space_order=order)

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

    op = Operator(stencils=Eq(u.forward, stencil),
                  nt=nt, subs={s: dt, h: h},
                  shape=shape, forward=True)
    # Source injection code omitted for brevity

    op.apply()
```

varying dampening factor used to implement an absorbing boundary condition.

### A. Acoustic wave equation

The implementation of the forward operator with an explicit time marching scheme is shown in Listing 6. First, the symbolic data objects are created, where  $u$  is defined as a time-varying function, whereas  $m$  and  $\eta$  are spatial grids only. Notably, the source term  $q$  is missing from the high-level PDE description, since the injection of the source requires sparse point-to-grid interpolation. Since interpolation does not fall within a finite difference abstraction, it is dealt with via separate expressions, as further explained in section V-B

It is also important to note here that the spatial order of the derivative stencils in the operator definition is provided as an argument and can thus be changed easily by the user, and that the entire formulation is independent of problem dimension. From the individual component symbols we create the wave equation in a purely symbolic form, leaving Devito's `dt`, `dt2` and `laplace` operators to perform the stencil expansion according to the spatial and temporal discretization order. The resulting expression is then rearranged using the SymPy `solve` routine to denote forward propagation, from which the corresponding `Operator` can be created and executed.

A very similar symbolic definition can be used to define the adjoint operator for the wave equation, as demonstrated in Listing 7. The key difference to the forward implementation is that the dampening term  $\eta * u.dt$  is subtracted rather than added, and that we rearrange the stencil expression for execution backward in time. The shorthand notation `u.forward` and `u.backward` hereby denote the highest and lowest stencil point in the second-order time discretization stencil,  $t + s$  and  $t - s$  respectively. These can be used to ensure the correct time marching direction once the symbolic function offsets  $s$  and  $h$  have been resolved to explicit grid indices in the final stencil expression.

Listing 7: Example code to solve the adjoint wave equation

```

def adjoint(model, nt, dt, h, spc_order=2):
    shape = model.shape
    m = DenseData(name="m", shape=shape,
                  space_order=order)
    m.data[:] = model
    v = TimeData(name='v', shape=shape,
                 time_dim=nt, time_order=2,
                 space_order=spc_order)
    η = DenseData(name="η", shape=shape,
                  space_order=order)

    # Derive stencil from symbolic equation
    eqn = m * v.dt2 - v.laplace - η * v.dt
    stencil = solve(eqn, v.backward)[0]

    op = Operator(stencils=Eq(u.backward, stencil),
                  nt=nt, subs={s: dt, h: h},
                  shape=shape, forward=False)
    # Receiver interpolation omitted for brevity

    op.apply()

```

Listing 8: Example code for point-to-grid interpolation.

```

eqns = [Eq(u[t, x, y], u[t, x, y] + expr),
        Eq(u[t, x+1, y], u[t, x+1, y] + expr),
        Eq(u[t, x, y+1], u[t, x, y+1] + expr),
        Eq(u[t, x+1, y+1], u[t, x+1, y+1] + expr)]
it = Iteration(eqns, index=p, limits=source.shape)

```

### B. Sparse point interpolation

One key restriction of the high-level PDE notation used in the previous examples is the inability to express interpolation between irregularly spaced sparse points and the grid. This feature is frequently required to inject source terms by point-to-grid interpolation and sample the resulting wave field at non-aligned sparse points via grid-to-point interpolation.

Listing 8 demonstrates how to construct a set of expressions to perform point-to-grid interpolation, where each grid value enclosing the sparse point is updated according to a source term. Most notably, the update to the grid value in the wave field  $u$  is created using the low-level indexed variable format, allowing us to define the explicit offsets needed to update every enclosing grid point with the relevant source expression `expr`. A similar set of expressions can be generated to perform the inverse interpolation from grid values to sparse points, but an example of this is omitted here for brevity. Grid-to-point interpolation, however, is of particular importance for seismic inversion problems, since the behaviour of sparse points changes between the forward and the adjoint runs.

Since there are potentially many sources in a wave field model, we can also construct an additional loop around the expression set to iterate over all sparse points contained in `source`. For this, Devito provides an additional low-level `Iteration` class that adds a user-defined loop construct into the auto-generated code. This enables custom iterations over source points and their associated coordinate data, allowing users to create coordinate-dependent symbolic expressions that define the interpolation for each reference cell point.

### C. Validation

We pointed out the requirement to have rigorous adjoints in order to use our framework in seismic inversion. To validate our implementation we use the simple definition of an adjoint as a test. For a given matrix  $A$ , its adjoint is defined as the matrix  $A^T$  such that:

$$\forall(x, y), \quad \langle Ax, y \rangle - \langle x, A^T y \rangle = 0, \quad (4)$$

$$\frac{\langle Ax, y \rangle}{\langle x, A^T y \rangle} = 1, \quad (5)$$

where  $\langle \cdot, \cdot \rangle$  denotes the inner product. In the case of the wave equation,  $A$  represent the modelling operator while  $A^T$  represents the adjoint modelling operator, and from the definition in Equation 4 we can compare the output of our kernels to verify the accuracy of the adjoint. We validate our implementation with this test in a two-dimensional and three-dimensional setting for multiple spacial discretization order in Table I .

### D. Performance

The performance of the finite difference code generated by Devito is demonstrated on two target architectures, a Intel®Xeon™E5-2690v2 architecture with 10 physical cores running at 3GHz and a Intel®Xeon Phi™accelerator card, in Fig. 3. The model used for the benchmark is a gradient test with forward and adjoint operators on a three-dimensional grid with  $201 \times 201 \times 70$  grid points and 40 PML grid points on each side, resulting in a computational grid of size  $281 \times 281 \times 150$ . The grid size is  $15m$  and the source term is a Ricker wavelet at  $10Hz$ . The wave field is modelled for 1 second with spatial discretizations of varying order from 2 to 16. The results indicate that Devito is able to utilise both architectures to a high degree of efficiency, while maintaining the ability to increase accuracy by switching to higher order stencil discretizations dynamically.

## VI. DISCUSSION

In this paper we present the Devito finite difference DSL and showcase its multi-level Python API to efficiently generate complex mathematical operators from SymPy expressions. We highlight Devito’s practical applicability in creating forward and adjoint operators for seismic inversion operators from concise symbolic definitions and verify their correctness by showing that the generated code does in fact behave mathematically like the adjoint of the equation. Moreover, we demonstrate that the optimized code generated by Devito is capable of achieving a significant percentage of peak performance on a traditional CPU, as well as a Intel®Xeon Phi™architecture.

When showcasing Devito’s symbolic Python API we also compare it to traditional approaches to accelerating Python computation and show significant speed-ups on even a single core CPU. On top of the evident performance benefits, a more concise formulation of simple mathematical operators can be achieved using Devito, which increases the productivity of domain scientists. Conversely, our total reliance on the

TABLE I: Adjoint test for different discretization orders in 2D and 3D. The test is computed on a two layer model.

Order	Dimension	$\langle Ax, y \rangle$	$\langle x, A^T y \rangle$	Difference	ratio
2nd order	2D:	373323.7976	373323.7975434	6.07169350e-05	1.0
4th order	2D:	340158.1486	340158.1485252	0.00012756	1.0
6th order	2D:	341557.3948	341557.3947399	0.00014287	1.0
8th order	2D:	358240.8513	358240.8511931	0.00016741	1.0
10th order	2D:	393488.5561	393488.5559269	0.00023841	1.0
12th order	2D:	439561.4005	439561.4002034	0.00035794	1.0
2nd order	3D:	2.17496552	2.174965534979	-1.23030883e-08	0.99999999
4th order	3D:	3.64447937	3.644479393901	-2.13132316e-08	0.99999999
6th order	3D:	3.78730372	3.787303745086	-2.22477072e-08	0.99999999
8th order	3D:	3.80286229	3.802862312852	-2.23545817e-08	0.99999999
10th order	3D:	3.80557957	3.805579596334	-2.23736993e-08	0.99999999
12th order	3D:	3.80318675	3.803186773144	-2.23587757e-08	0.99999999


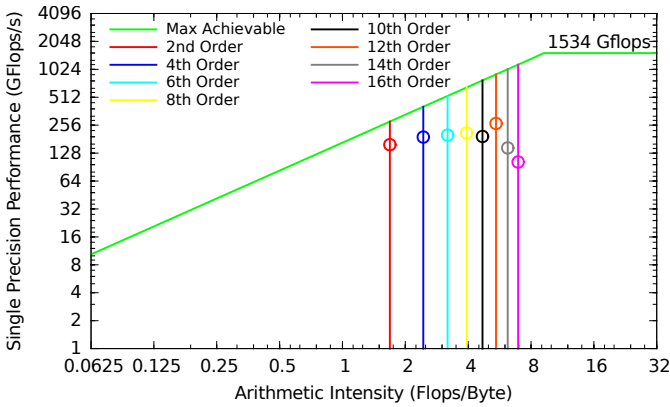
(a) Intel<sup>®</sup>Xeon<sup>™</sup>E5-2690v2 10C @ 3GHz(b) Intel<sup>®</sup>Xeon Phi<sup>™</sup>

Fig. 3: Roofline plots for two target systems

SymPy package demonstrates the importance of the Python programming language and its vast software ecosystem in creating performance-oriented DSLs and utilising them to build actual scientific applications.

The use of Devito’s second-layer API to create custom point-grid interpolation routines further illustrates the complexities that often prevent stencil DSLs from being used in practice. In the context of seismic inversion, sparse point interpolation is the most important feature that other finite difference DSLs lack. In addition, Devito’s full compatibility with other scientific Python packages through the NumPy array API entails that it is easy to implement a full-scale inversion workflow in Python - including additional features

such as MPI parallelism between multiple inversion operators.

## VII. FUTURE WORK

Even though the primary target application for Devito is seismic inversion, we aim to develop Devito towards becoming a generic finite difference DSL. The diffusion example presented in section IV-D shows that the current abstraction generalises well to problems beyond the wave equation and its additional low-level API provides an important degree of flexibility when adding new features. Additional finite difference mechanisms, such as staggered grids and new boundary condition types are of great importance here to generalise Devito’s applicability to other application areas.

Additional improvements to the high and low-level symbolic APIs are also planned that include customizable dimension symbols and the automated derivation of iteration spaces. This additional automation will in turn enable further integration with optimising compilers and other external stencil frameworks and potentially even allow the automated use of polyhedral compilation tools.

The seismic inversion operators demonstrated in this paper, and in particular the sparse point interpolation features discussed in section V-B, are still in an early development stage. The seismic inversion example presented here uses the acoustic isotropic wave equation, although, for real world applications, more complex physical models involving anisotropy and elastic wave propagation are commonly seen. In the future we plan to implement these complex physical models using Devito, expanding Devito’s repertoire of features wherever necessary.

## VIII. ACKNOWLEDGEMENTS

This work was financially supported in part by the Natural Sciences and Engineering Research Council of Canada Collaborative Research and Development Grant DNOISE II (CDRP J 375142-08), EPSRC grant EP/L000407/1, the Imperial College London Intel Parallel Computing Centre, British Gas, SENAI CIMATEC and the MCTI (Federal Govt of Brazil) scholarship MCTI/FINEP/CNPq 384214/2015-0. This research was carried out as part of the SINBAD II project with the support of the member organizations of the SINBAD Consortium.



## REFERENCES

- [1] T. E. Oliphant, "Python for scientific computing," *Computing in Science Engineering*, vol. 9, no. 3, pp. 10–20, May 2007.
- [2] K. J. Millman and M. Aivazis, "Python for scientists and engineers," *Computing in Science Engineering*, vol. 13, no. 2, pp. 9–12, March 2011.
- [3] F. Perez, B. E. Granger, and J. D. Hunter, "Python: An ecosystem for scientific computing," *Computing in Science Engineering*, vol. 13, no. 2, pp. 13–21, March 2011.
- [4] S. Behnel, R. Bradshaw, C. Citro, L. Dalcin, D. S. Seljebotn, and K. Smith, "Cython: The best of both worlds," *Computing in Science Engineering*, vol. 13, no. 2, pp. 31–39, March 2011.
- [5] S. K. Lam, A. Pitrou, and S. Seibert, "Numba: A llvm-based python jit compiler," in *Proceedings of the Second Workshop on the LLVM Compiler Infrastructure in HPC*, ser. LLVM '15. New York, NY, USA: ACM, 2015, pp. 7:1–7:6. [Online]. Available: <http://doi.acm.org/10.1145/2833157.2833162>
- [6] A. Logg, K.-A. Mardal, and G. Wells, *Automated Solution of Differential Equations by the Finite Element Method: The FEniCS Book*. Springer Publishing Company, Incorporated, 2012.
- [7] A. Meurer, C. P. Smith, M. Paprocki, O. Čertík, M. Rocklin, A. Kumar, S. Ivanov, J. K. Moore, S. Singh, T. Rathnayake *et al.*, "SymPy: Symbolic computing in python," PeerJ Preprints, Tech. Rep., 2016.
- [8] S. van der Walt, S. C. Colbert, and G. Varoquaux, "The numpy array: A structure for efficient numerical computation," *Computing in Science Engineering*, vol. 13, no. 2, pp. 22–30, March 2011.
- [9] A. F. Cárdenas and W. J. Karplus, "Pdela language for partial differential equations," *Communications of the ACM*, vol. 13, no. 3, pp. 184–191, 1970.
- [10] G. O. Cook Jr, "Alpal: A tool for the development of large-scale simulation codes," Lawrence Livermore National Lab., CA (USA), Tech. Rep., 1988.
- [11] Y. Umetani, "Deqsol a numerical simulation language for vector/parallel processors," *Proc. IFIP TC2/WG22, 1985*, vol. 5, pp. 147–164, 1985.
- [12] 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.
- [13] F. Rathgeber, D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. McRae, G.-T. Bercea, G. E. Markall, and P. H. Kelly, "Firedrake: automating the finite element method by composing abstractions," *Submitted to ACM TOMS*, 2015.
- [14] M. S. Alnæs, 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.
- [15] P. E. Farrell, D. A. Ham, S. W. Funke, and M. E. Rognes, "Automated derivation of the adjoint of high-level transient finite element programs," *SIAM Journal on Scientific Computing*, vol. 35, no. 4, pp. C369–C393, 2013. [Online]. Available: <http://dx.doi.org/10.1137/120873558>
- [16] 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*, ser. CIT '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 1181–1188. [Online]. Available: <http://dx.doi.org/10.1109/CIT.2010.214>
- [17] 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*, ser. ICS '13. New York, NY, USA: ACM, 2013, pp. 13–24. [Online]. Available: <http://doi.acm.org/10.1145/2464996.2467268>
- [18] 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*, ser. CGO '12. New York, NY, USA: ACM, 2012, pp. 155–164. [Online]. Available: <http://doi.acm.org/10.1145/2259016.2259037>
- [19] 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*, ser. SC '08. Piscataway, NJ, USA: IEEE Press, 2008, pp. 4:1–4:12. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1413370.1413375>
- [20] M. Christen, O. Schenk, and H. Burkhart, "Patus: A code generation and autotuning framework for parallel iterative stencil computations on modern microarchitectures," in *Proceedings of the 2011 IEEE International Parallel & Distributed Processing Symposium*, ser. IPDPS '11. Washington, DC, USA: IEEE Computer Society, 2011, pp. 676–687. [Online]. Available: <http://dx.doi.org/10.1109/IPDPS.2011.70>
- [21] K. A. Hawick and D. P. Playne, "Simulation software generation using a domain-specific language for partial differential field equations," in *11th International Conference on Software Engineering Research and Practice (SERP13)*, no. CSTN-187. Las Vegas, USA: WorldComp, 22-25 July 2013, p. SER3829.
- [22] R. Membarth, F. Hannig, J. Teich, and H. Köstler, "Towards domain-specific computing for stencil codes in hpc," in *High Performance Computing, Networking, Storage and Analysis (SCC), 2012 SC Companion*. IEEE, 2012, pp. 1133–1138.
- [23] G. R. Mudalige, M. B. Giles, I. Reguly, C. Bertolli, and P. H. J. Kelly, "Op2: An active library framework for solving unstructured mesh-based applications on multi-core and many-core architectures," in *Innovative Parallel Computing (InPar), 2012*, May 2012, pp. 1–12.
- [24] I. Z. Reguly, G. R. Mudalige, M. B. Giles, D. Curran, and S. McIntosh-Smith, "The ops domain specific abstraction for multi-block structured grid computations," in *Proceedings of the Fourth International Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing*, ser. WOLFHPC '14. Piscataway, NJ, USA: IEEE Press, 2014, pp. 58–67. [Online]. Available: <http://dx.doi.org/10.1109/WOLFHPC.2014.7>
- [25] L. W. Howes, A. Likhomotov, A. F. Donaldson, and P. H. Kelly, "Deriving efficient data movement from decoupled access/execute specifications," in *Proceedings of the 4th International Conference on High Performance Embedded Architectures and Compilers*, ser. HiPEAC '09. Berlin, Heidelberg: Springer-Verlag, 2009, pp. 168–182. [Online]. Available: [http://dx.doi.org/10.1007/978-3-540-92990-1\\_14](http://dx.doi.org/10.1007/978-3-540-92990-1_14)
- [26] F. Rathgeber, G. R. Markall, L. Mitchell, N. Lorient, D. A. Ham, C. Bertolli, and P. H. Kelly, "Pyop2: A high-level framework for performance-portable simulations on unstructured meshes," in *High Performance Computing, Networking Storage and Analysis, SC Companion*. Los Alamitos, CA, USA: IEEE Computer Society, 2012, pp. 1116–1123.
- [27] J. Ragan-Kelley, C. Barnes, A. Adams, S. Paris, F. Durand, and S. Amarasinghe, "Halide: A language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines," in *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI '13. New York, NY, USA: ACM, 2013, pp. 519–530. [Online]. Available: <http://doi.acm.org/10.1145/2491956.2462176>