

Rapid development of seismic imaging applications using Symbolic mathematics

Navjot Kukreja¹ M. Lange¹ M. Louboutin² F. Luporini¹ J.
Hueckelheim¹ P. Witte² C. Yount³ F. Herrmann² G. Gorman¹

October 2, 2017

¹Department of Earth Science and Engineering, Imperial College London, UK

²Seismic Lab. for Imaging and Modeling, The University of British Columbia, Canada

³Intel Corporation

Introduction

Devito

Verification

Performance

Vertical Integration

Conclusions

Use geophysics to understand the earth




Figure 1: Offshore seismic survey

Source: <http://www.open.edu/openlearn/science-maths-technology/science/environmental-science/earths-physical-resources-petroleum/content-section-3.2.1>

Acoustic wave equation with 2nd-order discretisation:

```
for ti in range(timesteps):  
    t0 = ti % 3  
    t1 = (ti + 1) % 3  
    t2 = (ti + 2) % 3  
    for i in range(1, nx-1):  
        for j in range(1, ny-1):  
            uxx = (u[t1, i+1, j] -2 * u[t1, i, j] + u[t1, i-1, j]) / dx2  
            uyy = (u[t1, i, j+1] -2 * u[t1, i, j] + u[t1, i, j-1]) / dy2  
            u[t2, i, j] = 2*u[t1, i, j] - u[t0, i, j] + dt * dt *  
                (uxx + uyy)/m[i, j]
```

12th-order acoustic wave equation:

```

for (int i4 = 0; i4<149; i4+=1) {
  for (int i1 = 6; i1<64; i1++) {
    for (int i2 = 6; i2<64; i2++) {
      for (int i3 = 6; i3<64; i3++) {
        u[i4][i1][i2][i3] = 6.01250601250601e-9F*(2.80896e+8F*damp[i1][i2][i3]*u[i4-2][i1][i2][i3]-3.3264e+8F*m[i1][i2][i3]*u[i4-2][i1][i2][i3]+6.6528e+8F*m[i1][i2][i3]*u[i4-1][i1][i2][i3]-2.12255421155556e+7F*u[i4-1][i1][i2][i3]-1.42617283950617e+2F*u[i4-1][i1][i2][i3-6]+2.464426666666667e+3F*u[i4-1][i1][i2][i3-5]-2.117866666666667e+4F*u[i4-1][i1][i2][i3-4]+1.25503209876543e+5F*u[i4-1][i1][i2][i3-3]-6.3536e+5F*u[i4-1][i1][i2][i3-2]+4.066304e+6F*u[i4-1][i1][i2][i3-1]+4.066304e+6F*u[i4-1][i1][i2][i3+1]-6.3536e+5F*u[i4-1][i1][i2][i3+2]+1.25503209876543e+5F*u[i4-1][i1][i2][i3+3]-2.117866666666667e+4F*u[i4-1][i1][i2][i3+4]+2.464426666666667e+3F*u[i4-1][i1][i2][i3+5]-1.42617283950617e+2F*u[i4-1][i1][i2][i3+6]-1.42617283950617e+2F*u[i4-1][i1][i2-6][i3]+2.464426666666667e+3F*u[i4-1][i1][i2-5][i3]-2.117866666666667e+4F*u[i4-1][i1][i2-4][i3]+1.25503209876543e+5F*u[i4-1][i1][i2-3][i3]-6.3536e+5F*u[i4-1][i1][i2-2][i3]+4.066304e+6F*u[i4-1][i1][i2-1][i3]+4.066304e+6F*u[i4-1][i1][i2+1][i3]-6.3536e+5F*u[i4-1][i1][i2+2][i3]+1.25503209876543e+5F*u[i4-1][i1][i2+3][i3]-2.117866666666667e+4F*u[i4-1][i1][i2+4][i3]+2.464426666666667e+3F*u[i4-1][i1][i2+5][i3]-1.42617283950617e+2F*u[i4-1][i1][i2+6][i3]-1.42617283950617e+2F*u[i4-1][i1-6][i2][i3]+2.464426666666667e+3F*u[i4-1][i1-5][i2][i3]-2.117866666666667e+4F*u[i4-1][i1-4][i2][i3]+1.25503209876543e+5F*u[i4-1][i1-3][i2][i3]-6.3536e+5F*u[i4-1][i1-2][i2][i3]+4.066304e+6F*u[i4-1][i1-1][i2][i3]+4.066304e+6F*u[i4-1][i1+1][i2][i3]-6.3536e+5F*u[i4-1][i1+2][i2][i3]+1.25503209876543e+5F*u[i4-1][i1+3][i2][i3]-2.117866666666667e+4F*u[i4-1][i1+4][i2][i3]+2.464426666666667e+3F*u[i4-1][i1+5][i2][i3]-1.42617283950617e+2F*u[i4-1][i1+6][i2][i3])/(1.688888888888889F*damp[i1][i2][i3]+2*m[i1][i2][i3]);
      }
    }
  }
}

```

- Huge Problem sizes
 - Seismic surveys consist of thousands of individual experiments
 - Model wave propagation across large domains over thousands of timesteps
 - The code needs to be highly optimised for performance, including selecting appropriate discretisation for performance
- Performance optimisations are hardware-specific
 - Parallelisation/Vectorisation - Intrinsic
 - Cache-reuse (loop blocking)
 - Memory - alternate layouts, alignment, NUMA
 - Common sub-expression elimination
 - Others - Elemental functions/loop fission, Denormal numbers, streaming stores etc.
- Reverse mode requires the adjoint of the equation
 - Storing the entire forward wave-field in memory is prohibitively expensive
 - Solutions involving saving partial forward state and recomputing - e.g. checkpointing
- Various physical models
 - Inelastic - acoustic/VTI/TTI
 - Elastic

- Too much bespoke code being written
- Everyone who works on such a piece of code needs to be a polymath
- No suitable separation of concerns between the three kinds of expertise this code requires - Computer Science, Mathematics, Physics

Introduction

Devito

Verification

Performance

Vertical Integration

Conclusions

SymPy: Symbolic computer algebra system in pure Python¹

Enables automation of stencil generation

- Complex symbolic expressions as Python object trees
- Symbolic manipulation routines and interfaces
- Convert symbolic expressions to numeric functions
 - Python (NumPy) functions; C or Fortran kernels
- For a great overview see [A. Meurer's talk at SciPy 2016](#)


For specialised domains generating C code is not enough!

- Compiler-level optimization to leverage performance
- Stencil optimization is a research field of its own

Devito: Finite difference DSL based on SymPy

Devito generates highly optimized stencil code...

- OpenMP threading and vectorisation pragmas
- Cache blocking and auto-tuning
- Symbolic stencil optimisation



... from concise mathematical syntax

Example: acoustic wave equation with dampening

$$m \frac{\partial^2 u}{\partial t^2} + \eta \frac{\partial u}{\partial t} - \nabla^2 u = 0$$

can be written as

```
eqn = m * u.dt2 + eta * u.dt - u.laplace
```

```
from devito import TimeData, DenseData
from sympy import solve

shape = (10, 10)
space_order = 2
time_order = 2
m = DenseData(name='m', shape=shape, space_order=space_order)
u = TimeData(name='u', shape=shape, space_order=space_order, time_order=
            time_order)

eqn = m * u.dt2 - u.laplace
stencil=solve(eqn, u.forward)

[In] print(stencil)
[Out] [(h_x**2*h_y**2*(2*u(t, x, y) - u(t - s, x, y))*m(x, y) + h_x**2*s
**2*(-2*u(t, x, y) + u(t, x, y - h_y) + u(t, x, y + h_y)) + h_y**2*s
**2*(-2*u(t, x, y) + u(t, x - h_x, y) + u(t, x + h_x, y)))/(h_x**2*
h_y**2*m(x, y))]
```

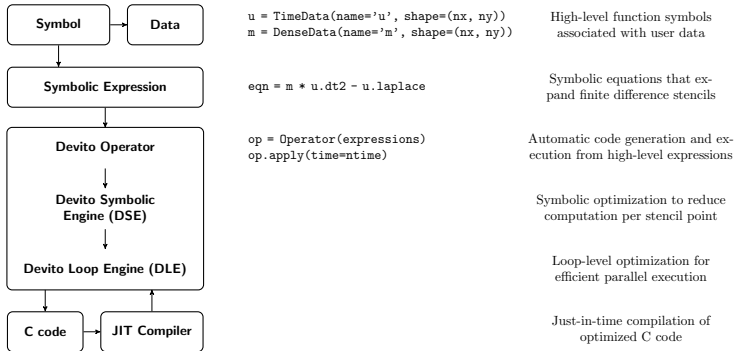


Figure 2: Overview of the Devito architecture and associated example workflow. Devito's top-level API allows users to generate symbolic stencil expressions from data-carrying function objects that can be used to for symbolic expressions vis SymPy. From this high-level definition an operator then generates, compiles and executes optimized high-performance C code.

Wave propagators in less than 20 lines

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

    # Inject wave as source term
    src_term = src.inject(field=u, expr=src * dt**2 / m)

    # Interpolate wavefield onto receivers
    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})
```

Wave propagators in less than 20 lines

```
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  
    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)
```

Reverse time migration in less than 100 lines

```
# Create the true and a smoothed model
m_true = Model(...)
m_smooth = Model(...)

# Create operators for forward and gradient
op_forward = forward(...)
op_gradient = forward(...)

# Create gradient field and loop over shots
grad = DenseData(name='grad', shape=model.shape)

for shot in shots:
    # Create receiver data from true model
    src = PointData(shot.source, ...)
    rec_true = PointData(shot.receiver.coordinates, ...)
    op.forward(src=src, rec=rec.true, m=m.true)

    # Run forward modelling operator with smooth model
    u = TimeData(name='u', shape=model.shape,
                 time_order=2, space_order=order)
    rec_smooth = PointData(shot.receiver.coordinates, ...)
    op.forward(u=u, src=src, rec=rec.smooth, m=m.smooth)

    # Compute gradient update from the residual
    v = TimeData(name='v', shape=model.shape,
                 time_order=2, space_order=order)
    residual = rec_true.data[:] - rec_smooth.data[:]
    op.gradient(u=u, v=v, grad=grad, rec=residual, m=m.smooth)
```


- Test and verify in Python
- Acoustic Operators in < 20 lines
- TTI Operators in < 100 lines ¹
- RTM setup in < 100 lines
- Variable stencil order

¹Yu Zhang, Houzhu Zhang, and Guanquan Zhang. A stable tti reverse time migration and its implementation. *Geophysics*, 76(3):WA3–WA11, 2011

- Common sub-expression elimination - C compilers do it already but it is quicker to do it at the higher level
- Heuristic refactorisation e.g. $0.3 * a + \dots + 0.3 * b \Rightarrow 0.3 * (a + b)$
 - Impact: TTI, space order 16: 6680 5760
- Alias detection e.g. `sin(phi[i, j, k]), sin(phi[i-1, j-1, k-1])`
- Heuristic hoisting of time-invariants
- Loop fission + elemental functions (register locality)
- Padding + data alignment (split loads)
- Loop blocking in 1D/2D/3D (no time yet) - with autotuned block sizes
- SIMD vectorisation through pragmas (intrinsics not required)
- Thread parallelism through OpenMP pragmas
- YASK backend (WIP)

Introduction

Devito

Verification

Performance

Vertical Integration

Conclusions



Verification of the generated code:

- Extensive unit-testing already in place with continuous integration (Travis)
- Adjoint test ¹
 - For any $x \in \text{span}(P_s A^T P_r^T)$, $y \in \text{span}(P_r A^T P_s^T)$
 - $\langle P_r A^T P_s^T x, y \rangle - \langle x, P_s A^T P_r^T y \rangle = 0$
 - Passes with at-least 8 matching significant digits for 2D and 3D with 2,4,6,8,10,12th order discretization
- Gradient test
 - For a small model perturbation dm , $\phi_s(m + hdm) = \phi_s(m) + \mathcal{O}(h)$ and $\phi_s(m + hdm) = \phi_s(m) + h(J[m]^T \delta d)dm + \mathcal{O}(h^2)$
 - Passes at the level of the machine's accuracy
- Automatic formal self-verification ²

¹Louboutin, M., Lange, M., Luporini, F., Kukreja, N., Herrmann, F., Velesko, P. and Gorman, G. (2017). Code generation from symbolic finite-difference for geophysical exploration. Geoscientific Model Development. (under review)

²Huckelheim, J., Luo, Z., Luporini, F., Kukreja, N., Lange, M., Gorman, G., Siegel, S., Dwyer, M. and Hovland, P. (2017). Towards built-in verification in high-performance stencil code generation. Correctness 2017: First International Workshop on Software Correctness for HPC Applications. (accepted for presentation)

Comparison with a reference implementation - IWAVE ¹



¹ Symes, W.W., 2015. IWAVE structure and basic use cases. THE RICE INVERSION PROJECT, p.85.

Introduction

Devito

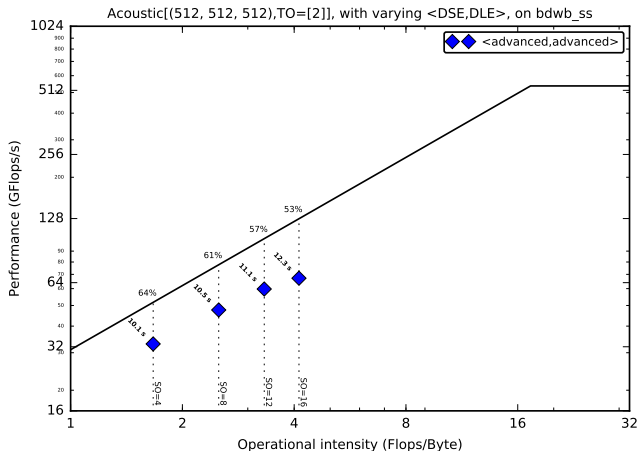
Verification

Performance


Vertical Integration

Conclusions

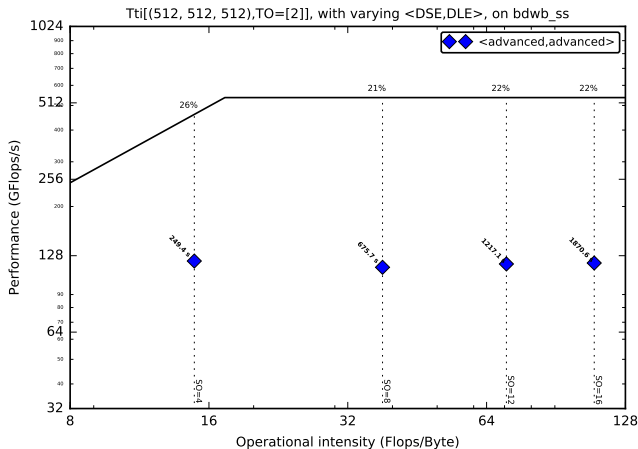
- Performance of acoustic forward operator
- Intel® Xeon™ E5-2620 v4 2.1Ghz Broadwell (8 cores)
- Model size $512 \times 512 \times 512$, $t_n = 250$



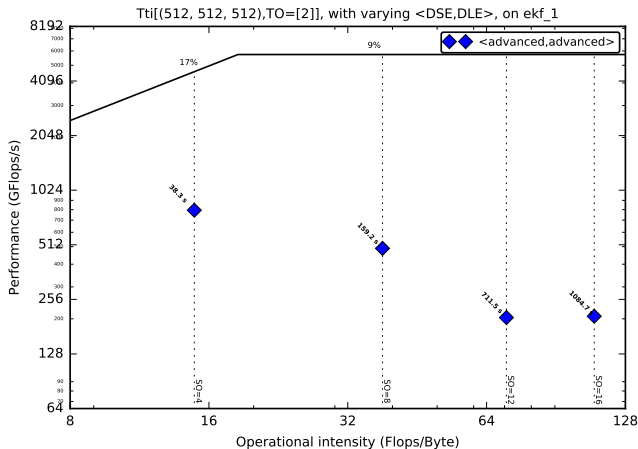
- Performance of acoustic forward operator
- Intel® Xeon Phi™ 7650 Knightslanding (68 cores) Quadrant Mode
- Model size $512 \times 512 \times 512$, $t_n = 3000$



- Performance of TTI forward operator
- Intel® Xeon™ E5-2620 v4 2.1Ghz Broadwell (8 cores)
- Model size $512 \times 512 \times 512$, $t_n = 250$



- Performance of TTI forward operator
- Intel® Xeon Phi™ 7650 Knightslanding (68 cores) Quadrant mode
- Model size $512 \times 512 \times 512$, $t_n = 3000$



Introduction

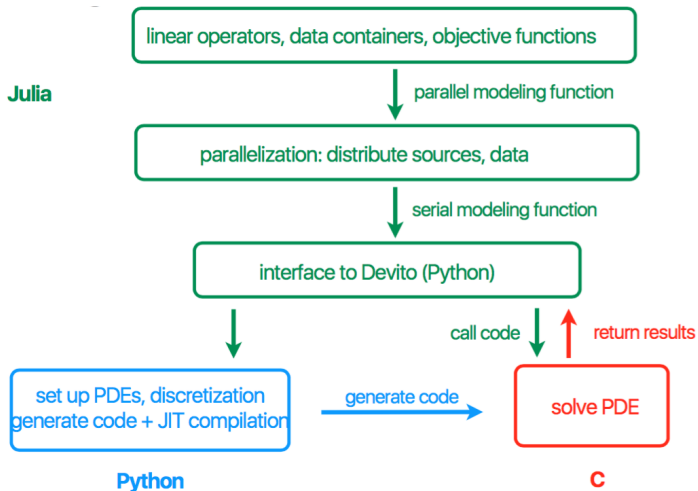
Devito

Verification

Performance

Vertical Integration

Conclusions



1

¹ Philipp Witte, Mathias Louboutin, and Felix J. Herrmann. Large-scale workflows for wave-equation based inversion in Julia. In *Domain-Specific Abstractions for Full-Waveform Inversion at SIAM CSE*, 2017

Introduction

Devito

Verification

Performance

Vertical Integration

Conclusions

- Devito: A finite difference DSL for seismic imaging
 - Symbolic problem description (PDEs) via SymPy
 - Low-level API for kernel customization
 - Automated performance optimization
- Devito is driven by real-world scientific problems
 - Bring the latest in performance optimization closer to real science
- Future work:
 - Yask Backend
 - MPI parallelism for larger models
 - Checkpointing
 - Better boundary conditions

Publications

- N. Kukreja, M. Louboutin, F. Vieira, F. Luporini, M. Lange, and G. Gorman. Devito: automated fast finite difference computation. WOLFHPC 2016
- M. Lange, N. Kukreja, M. Louboutin, F. Luporini, F. Vieira, V. Pandolfo, P. Velesko, P. Kazakas, and G. Gorman. Devito: Towards a generic Finite Difference DSL using Symbolic Python. PyHPC 2016
- M. Louboutin, M. Lange, N. Kukreja, F. Herrmann, and G. Gorman. Performance prediction of finite-difference solvers for different computer architectures. Submitted to Computers and Geosciences, 2016

Web links

- www.opesci.org
- github.com/opesci



EPSRC

Pioneering research
and skills

