

VECTORIZATION AND LOCALITY OPTIMIZATIONS FOR SEISMIC IMAGING METHODS THROUGH AUTOMATED CODE GENERATION

F. Luporini¹, M. Lange¹, M. Louboutin², N. Kukreja¹, G. Gorman¹

¹Imperial College London,

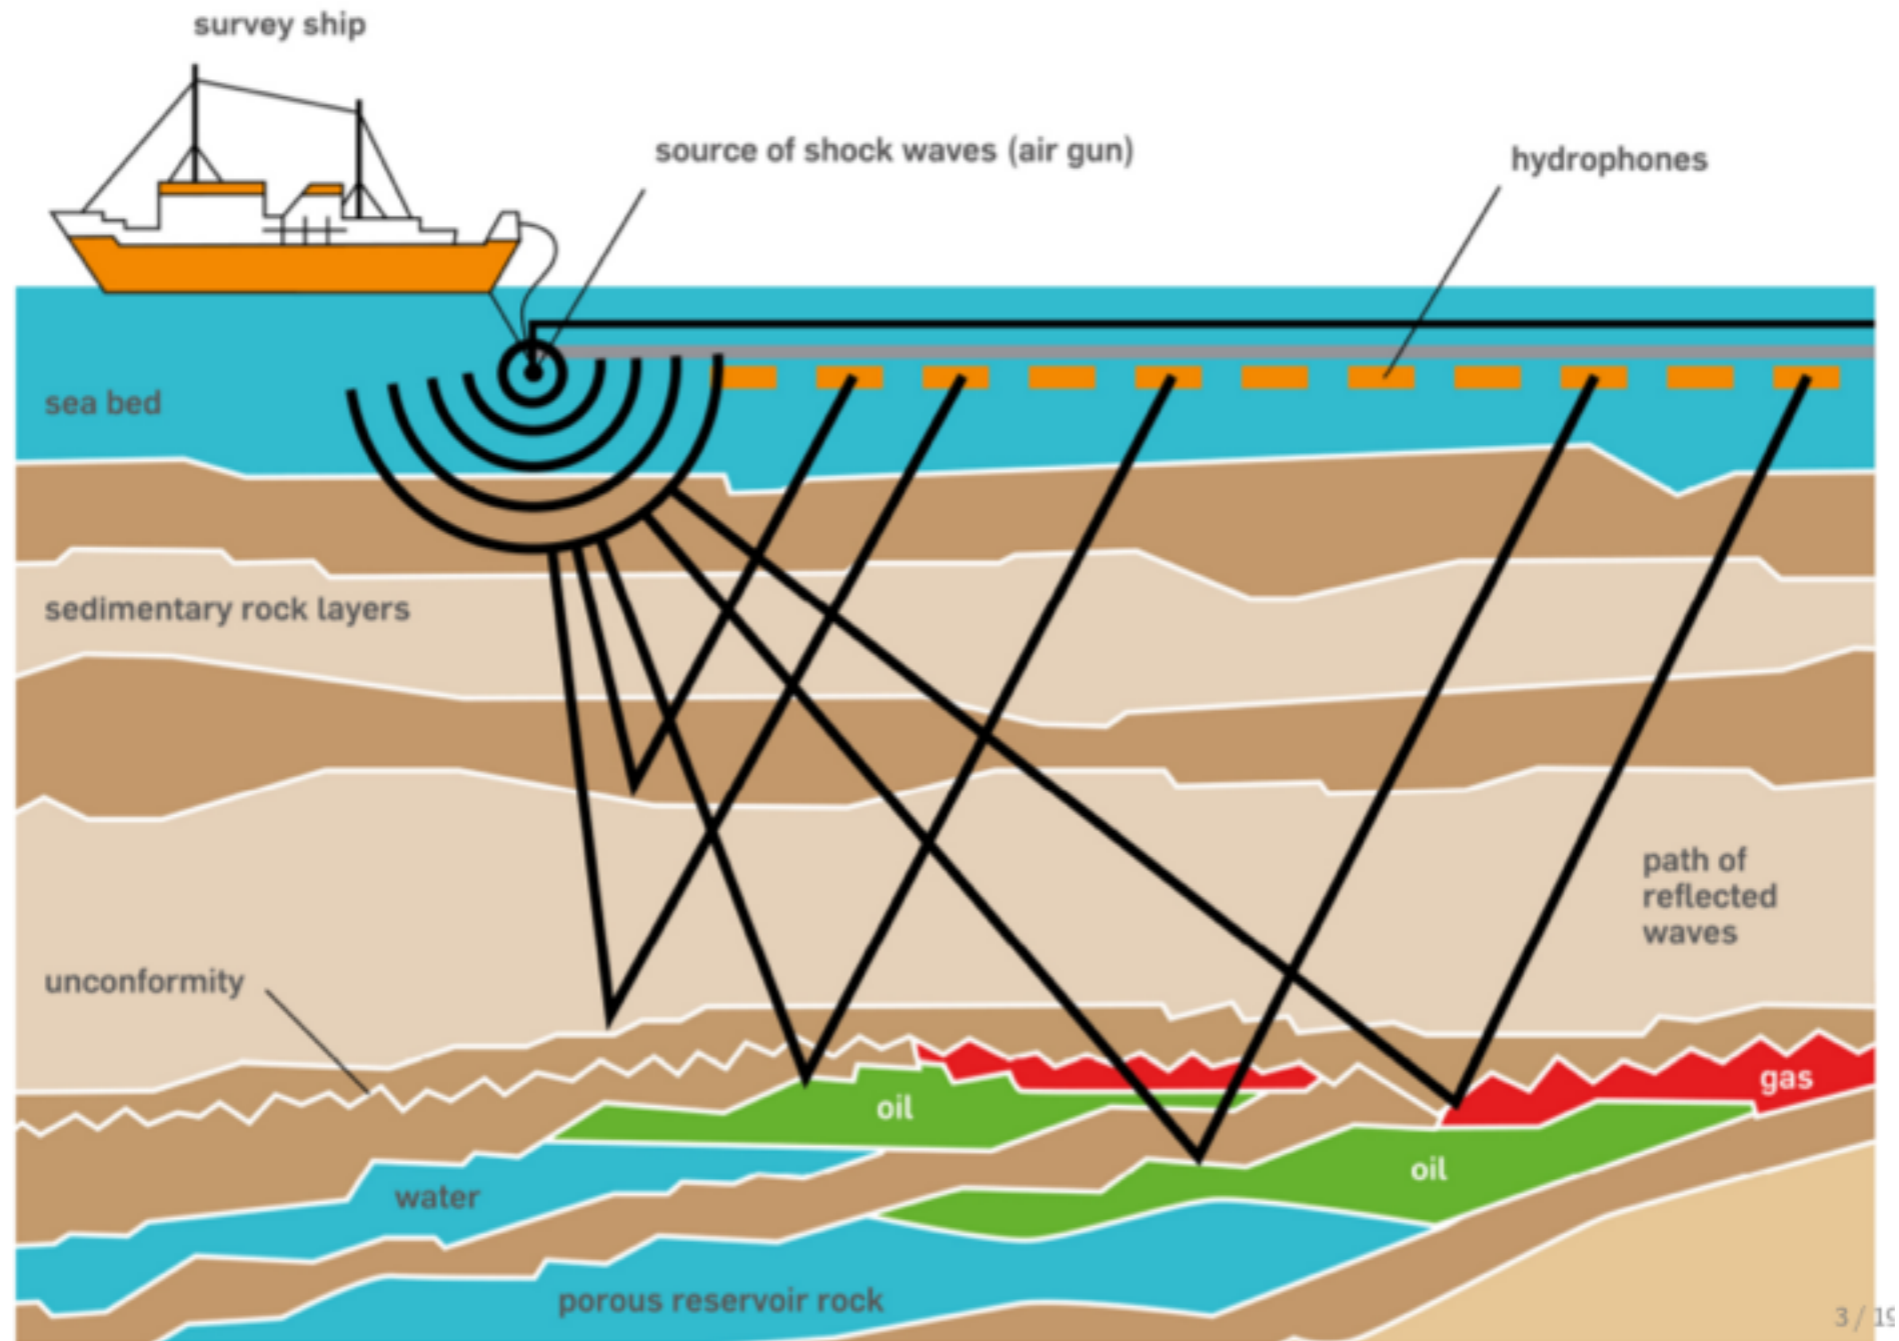
²The University of British Columbia (Seismic Lab for Imaging and Modeling)

2017 SIAM Conference on Computational Science and Engineering

MS: Efficiency of High-Order Methods on the 2nd Generation Intel Xeon Phi Processor

27/02/2017

Driving application: inversion problems for seismic imaging



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

- Challenging physics, many variants (e.g., wave equations),
- Big computational cost (wave simulation through subsurface)

So, it is not the “usual” Poisson equation that we aim to solve...

$$\frac{m}{\rho} \frac{d^2 p(x, t)}{dt^2} - (1 + 2\epsilon)(G_{\bar{x}\bar{x}} + G_{\bar{y}\bar{y}})p(x, t) - \sqrt{(1 + 2\delta)}G_{\bar{z}\bar{z}}r(x, t) = q,$$

$$\frac{m}{\rho} \frac{d^2 r(x, t)}{dt^2} - \sqrt{(1 + 2\delta)}(G_{\bar{x}\bar{x}} + G_{\bar{y}\bar{y}})p(x, t) - G_{\bar{z}\bar{z}}r(x, t) = q,$$

$$p(., 0) = 0,$$

$$\left. \frac{dp(x, t)}{dt} \right|_{t=0} = 0,$$

$$r(., 0) = 0,$$

$$\left. \frac{dr(x, t)}{dt} \right|_{t=0} = 0,$$

$$D_{x1} = \cos(\theta)\cos(\phi) \left. \frac{d}{dx} \right|_l + \cos(\theta)\sin(\phi) \left. \frac{d}{dy} \right|_l - \sin(\theta) \left. \frac{d}{dz} \right|_l$$

$$D_{x2} = \cos(\theta)\cos(\phi) \left. \frac{d}{dx} \right|_l + \cos(\theta)\sin(\phi) \left. \frac{d}{dy} \right|_l - \sin(\theta) \left. \frac{d}{dz} \right|_l$$

$$G_{\bar{x}\bar{x}} = \frac{1}{2} \left(D_{x1}^T \left(\frac{1}{\rho} \right) D_{x1} + D_{x2}^T \left(\frac{1}{\rho} \right) D_{x2} \right)$$

rotated second order differential operators

(incomplete)
specification of the
TTI (Tilted Transverse
Isotropy) forward
operator

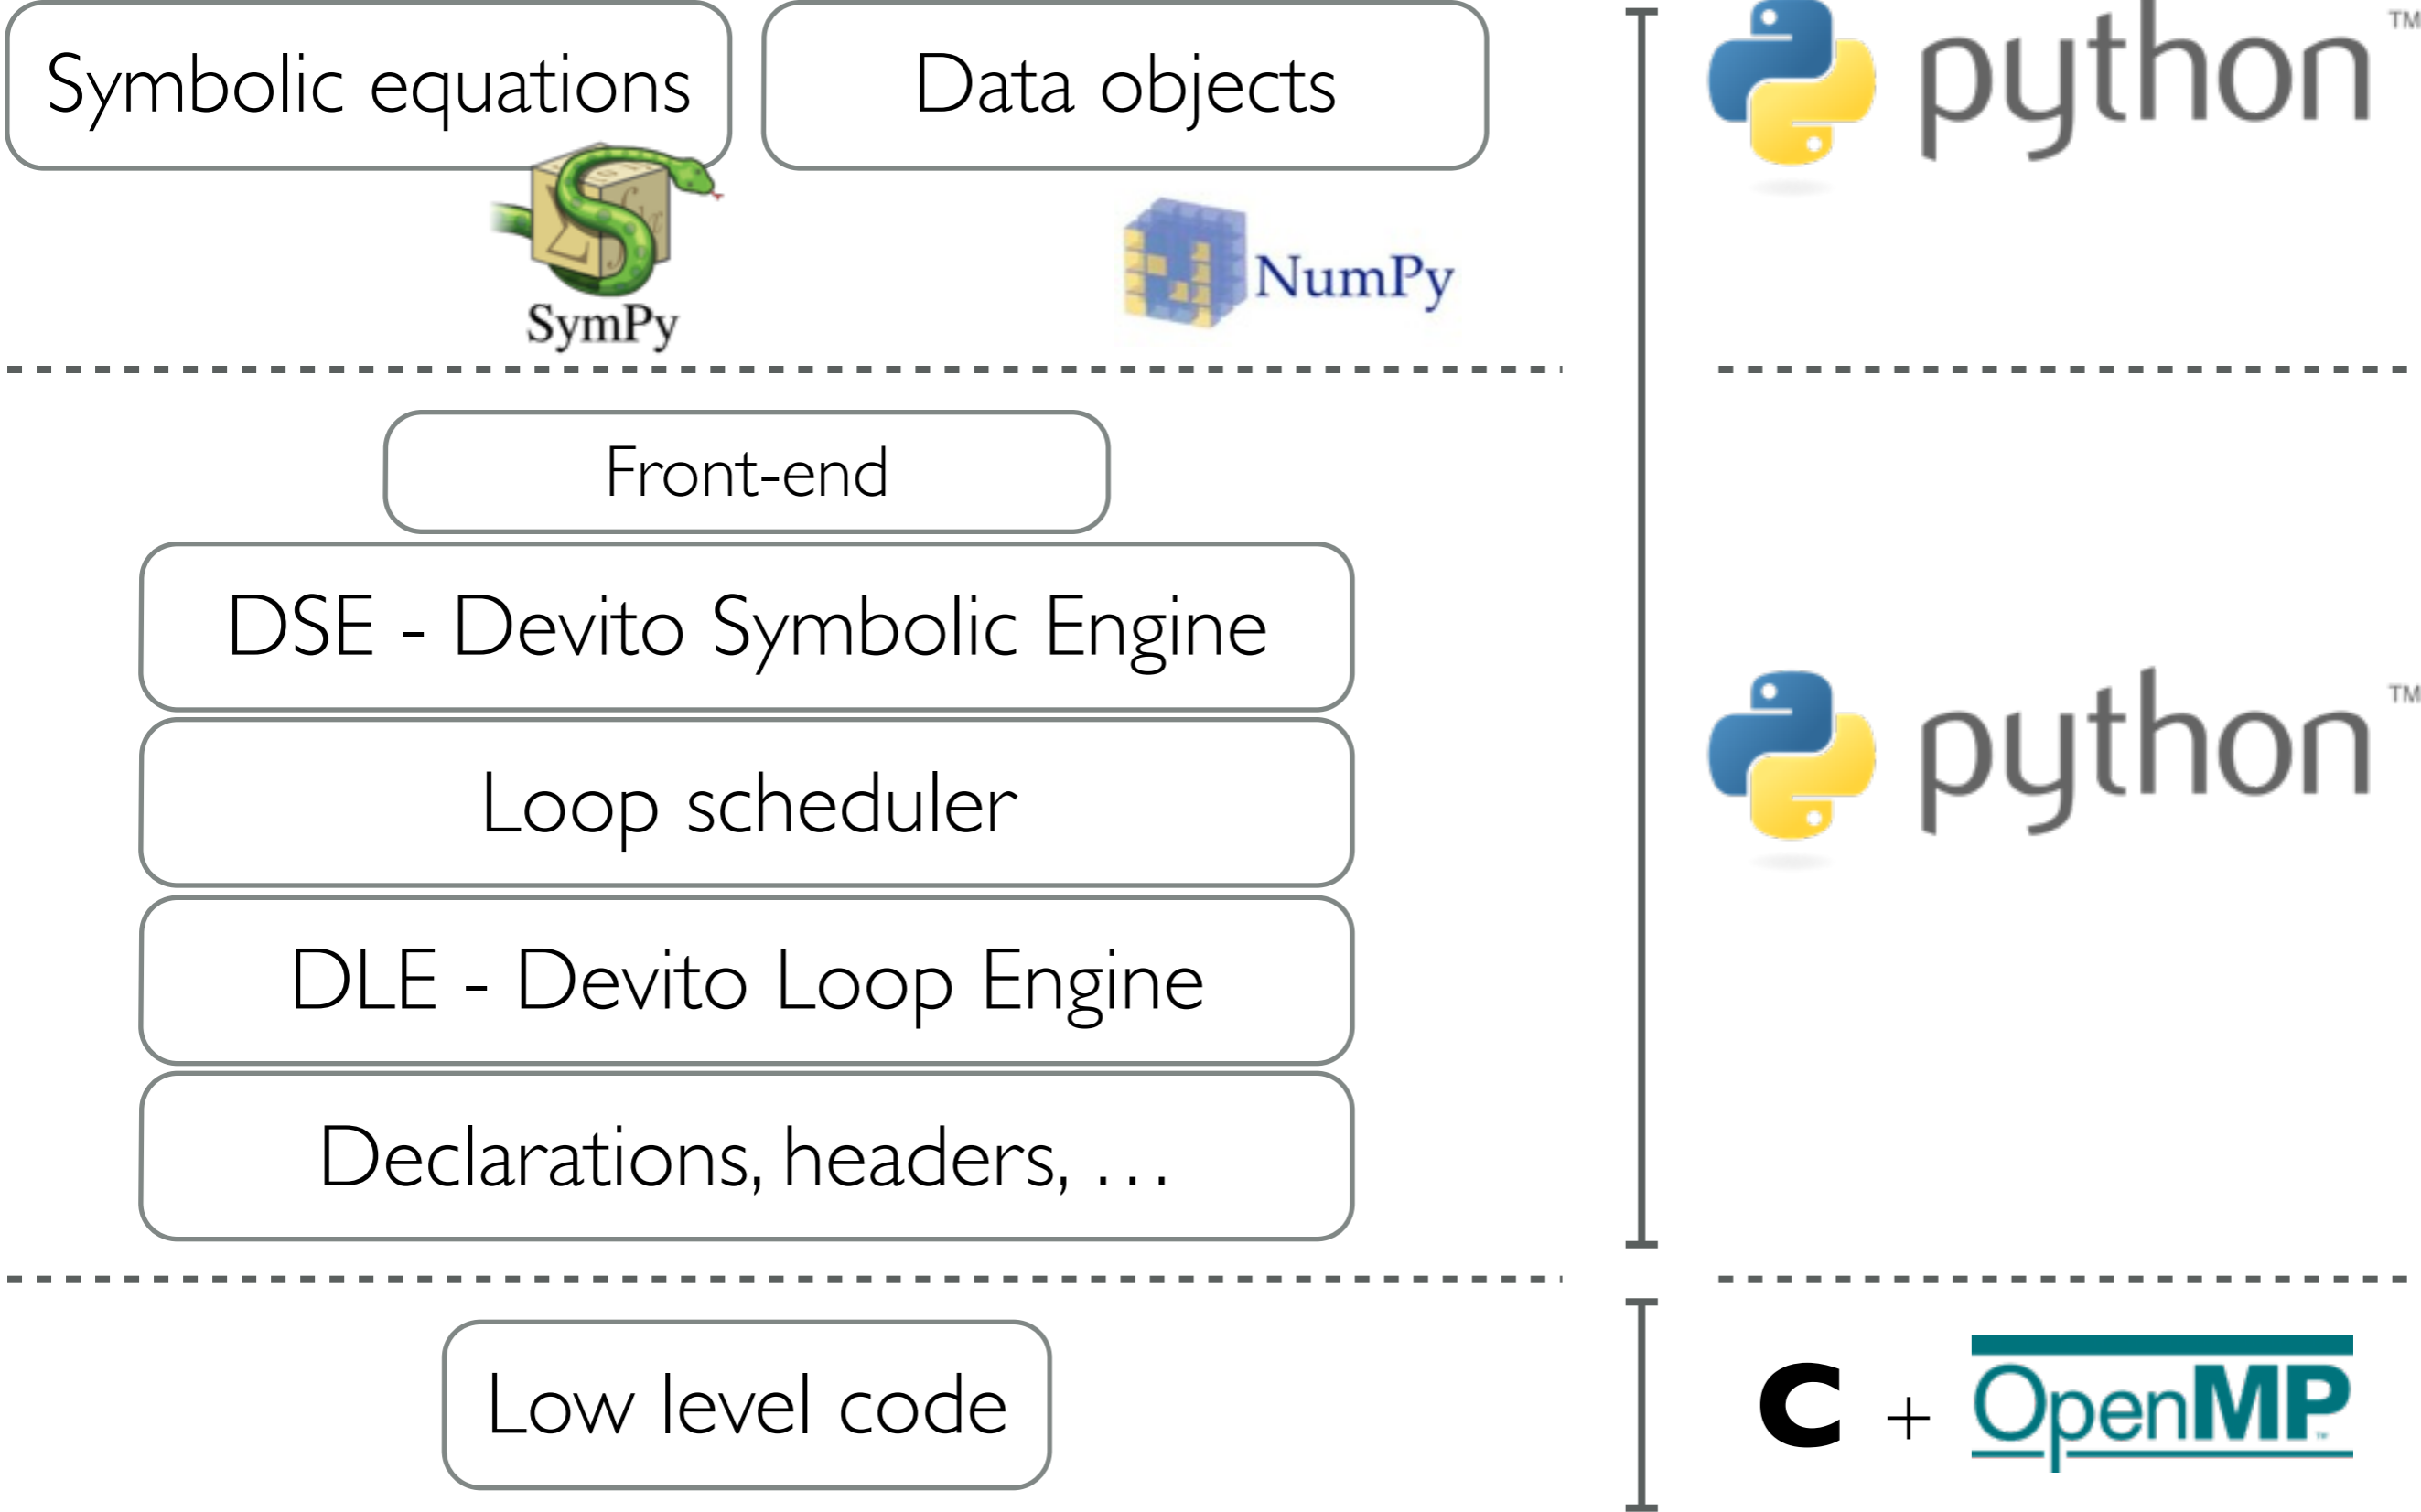
Why we need HPC implementations

- Huge number of floating-point operations: more than 6000 per loop iteration for a 16th order TTI operator
- Realistic 3D grids may have more than 10^9 grid points (e.g., 2.82 billions in SEAM benchmark)
- Often more than 3000 time steps
- Two operators: forward + adjoint
- Usually 30000 shots (“MPI level”)
- Around 15 Full-Waveform Inversion (FWI) iterations
- $\approx 6000 \times 2.82 \times 10^9 \times 3000 \times 2 \times 30000 \times 15 \approx 46$ billion TFLOPs
- ≈ 100 wall-clock days executing on the TACC Stampede (assuming Linpack-level performance)

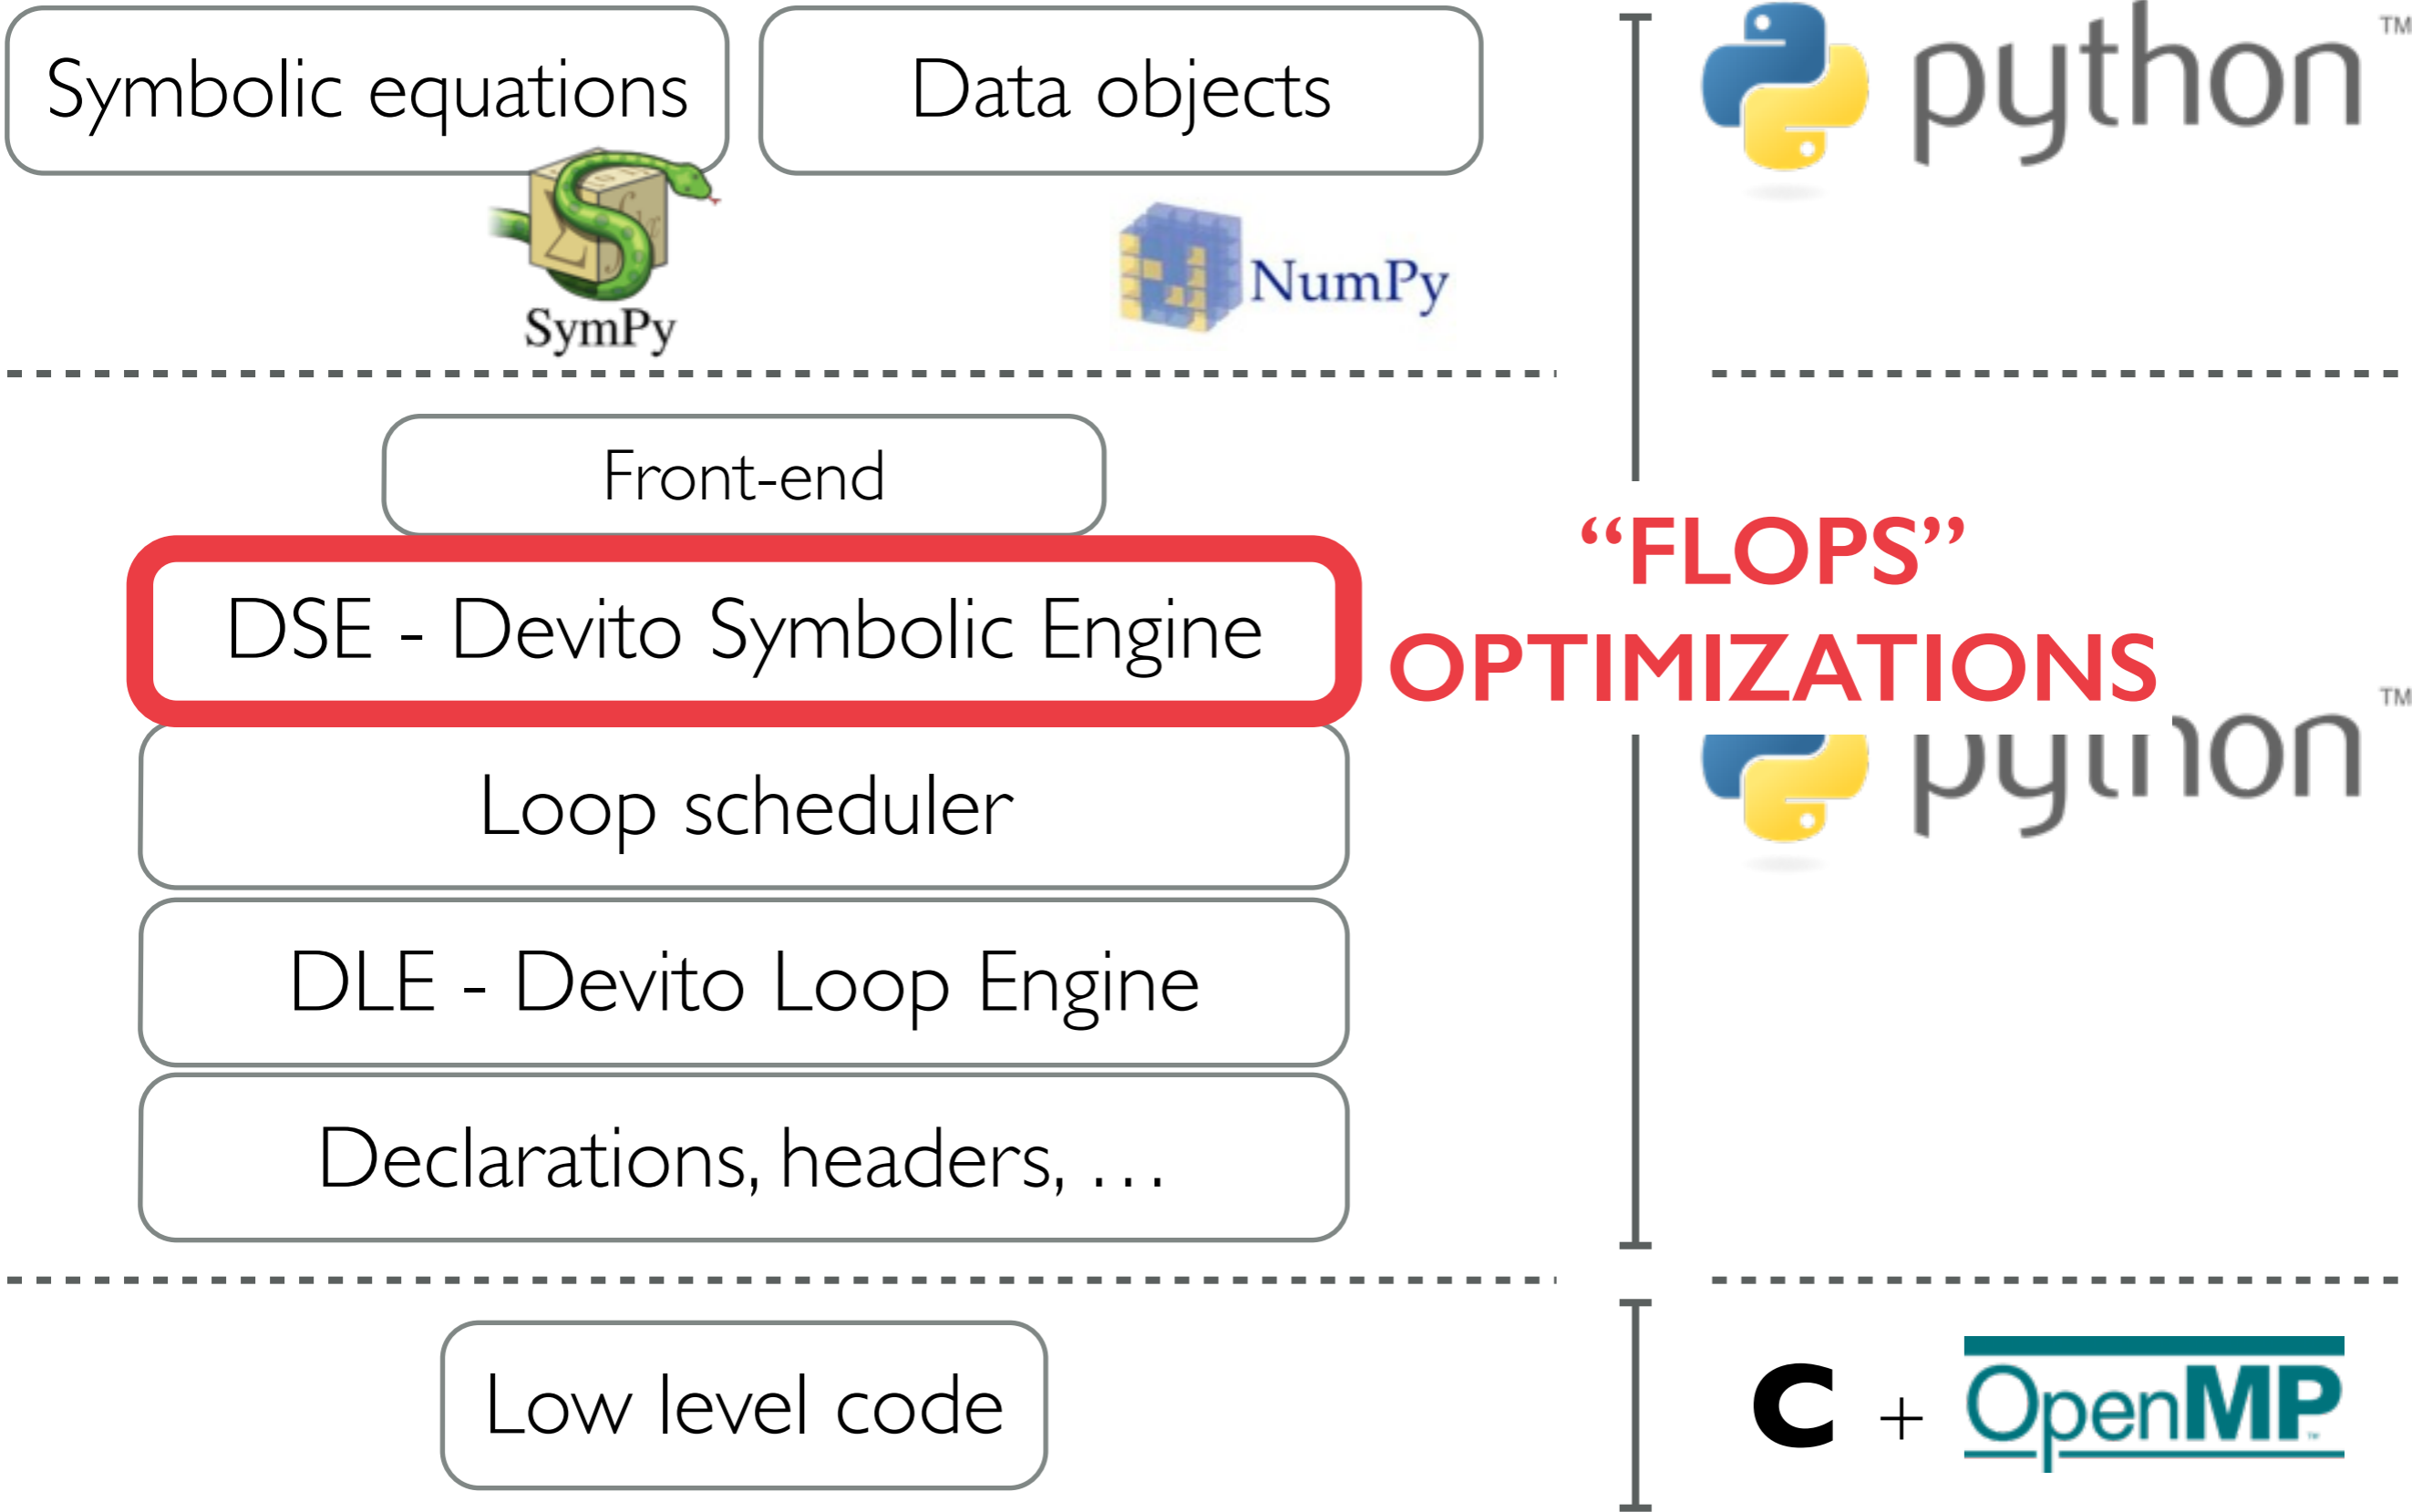
Devito: automated high performance finite difference

- **Real-world seismic imaging:**
 - Complex inversion methods (e.g., FWI)
 - Change of physics (e.g., acoustic, VTI, TTI — accuracy \Rightarrow complexity)
 - Change of discretization (FD schemes, up to very high order)
 - Boundary conditions, data acquisition, source/receivers modeling...
 - ...
- **Devito (\in OPESCI)**
 - Not “Yet another DSL for toy problems”: language + escape hatches
 - **Interdisciplinary research effort**
 - Used by geophysicists to write inversion operators
 - Based on actual compiler technology (you can write your own passes!)
- **This talk: the Devito compiler, its performance optimizations, application to real-world Acoustic and TTI operators**

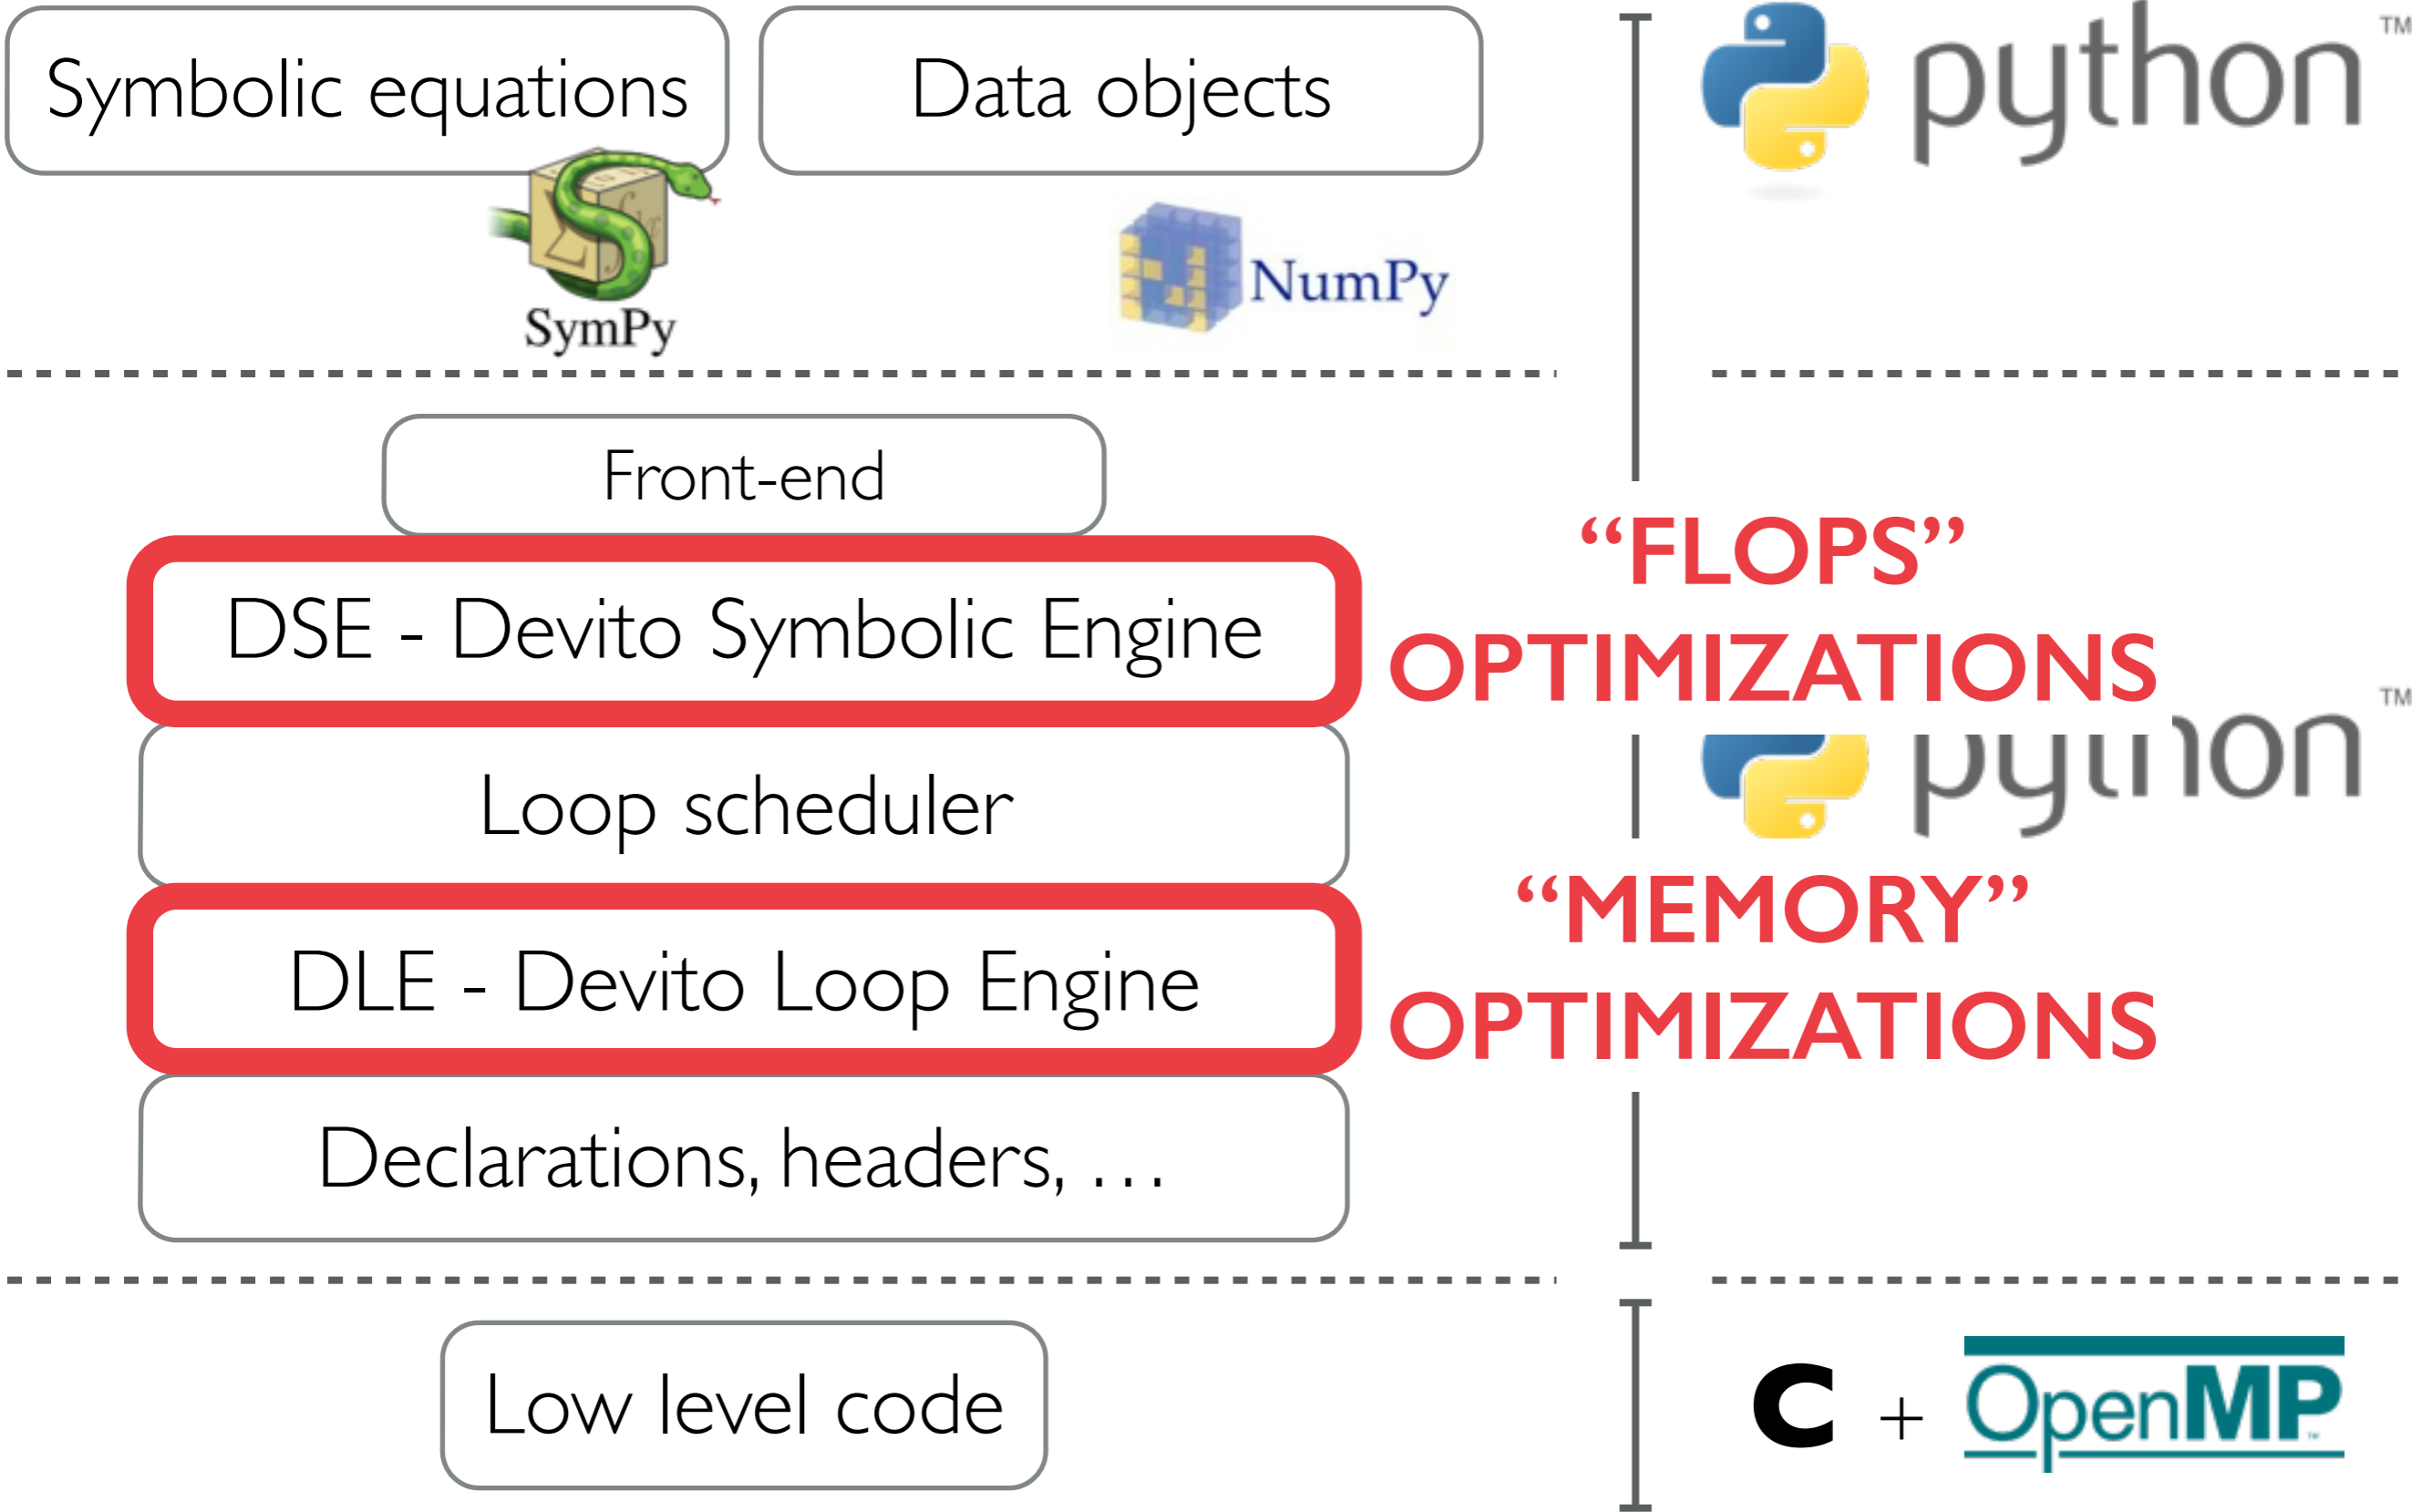
The compilation flow: from symbolics to HPC code



The compilation flow: from symbolics to HPC code



The compilation flow: from symbolics to HPC code



Devito Symbolic Engine

A sequence of compiler passes to reduce FLOPS (no loops at this stage!)

Devito Symbolic Engine

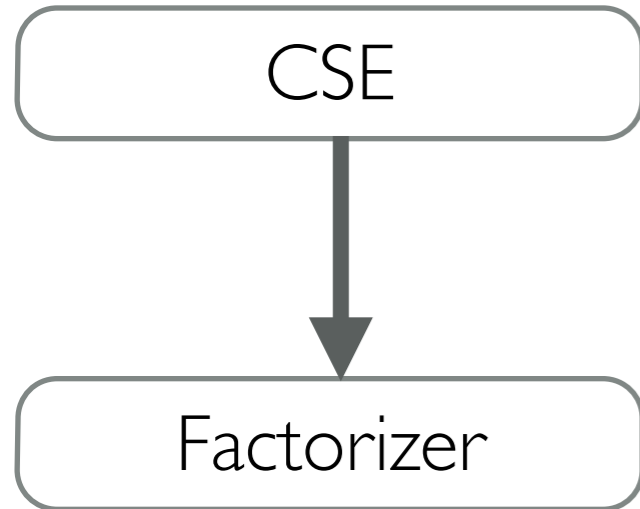
A sequence of compiler passes to reduce FLOPS (no loops at this stage!)

CSE

- Common sub-expressions elimination
 - C compilers do it already... but necessary for symbolic processing and compilation speed

Devito Symbolic Engine

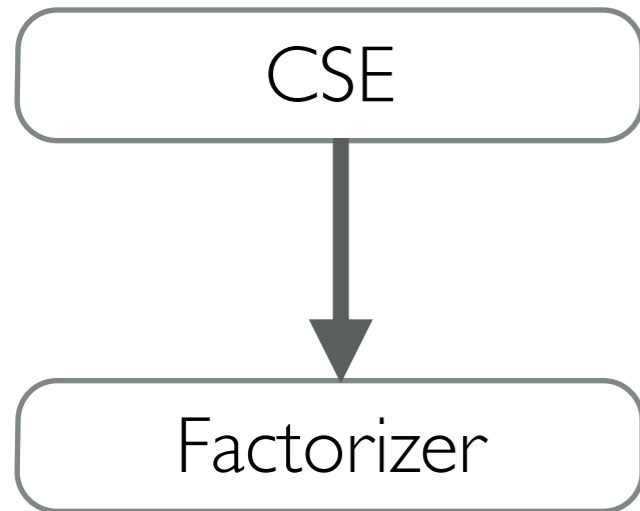
A sequence of compiler passes to reduce FLOPS (no loops at this stage!)



- Common sub-expressions elimination
 - C compilers do it already... but necessary for symbolic processing and compilation speed
- Heuristic re-factorization of recurrent terms
 - E.g., finite difference weights: $0.3*a + \dots + 0.3*b \Rightarrow 0.3*(a+b)$
 - Many possibilities (doesn't leverage domain properties yet!)

Devito Symbolic Engine

A sequence of compiler passes to reduce FLOPS (no loops at this stage!)



- Common sub-expressions elimination
 - C compilers do it already... but necessary for symbolic processing and compilation speed
- Heuristic re-factorization of recurrent terms
 - E.g., finite difference weights: $0.3*a + \dots + 0.3*b \Rightarrow 0.3*(a+b)$
 - Many possibilities (doesn't leverage domain properties yet!)

Factorizer impact:

TTI, space order 4: 1100 → 950

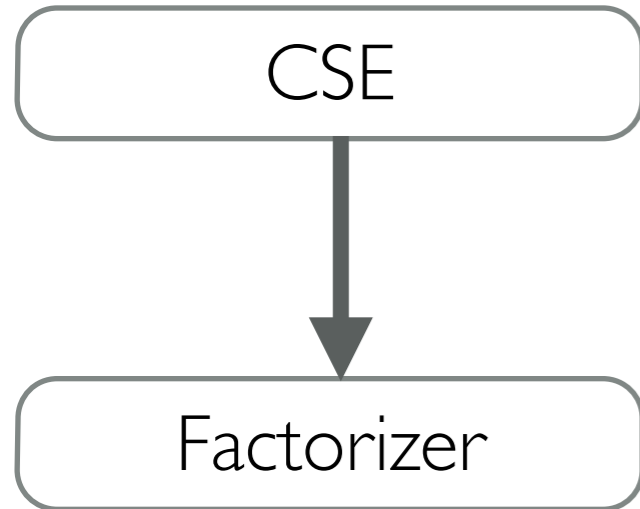
TTI, space order 8: 2380 → 2120

TTI, space order 12: 4240 → 3760

TTI, space order 16: 6680 → 5760

Devito Symbolic Engine

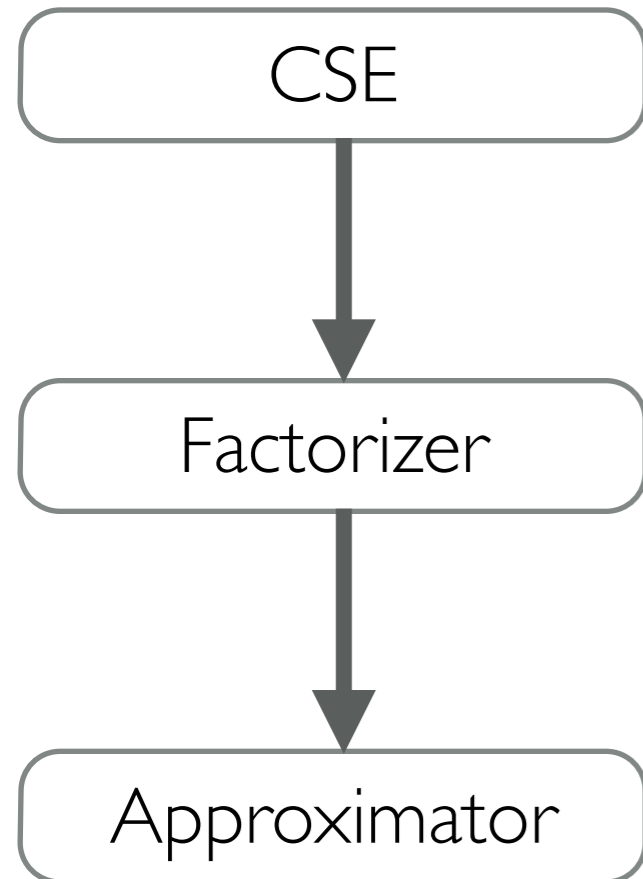
A sequence of compiler passes to reduce FLOPS (no loops at this stage!)



- Common sub-expressions elimination
 - C compilers do it already... but necessary for symbolic processing and compilation speed
- Heuristic re-factorization of recurrent terms
 - E.g., finite difference weights: $0.3*a + \dots + 0.3*b \Rightarrow 0.3*(a+b)$
 - Many possibilities (doesn't leverage domain properties yet!)

Devito Symbolic Engine

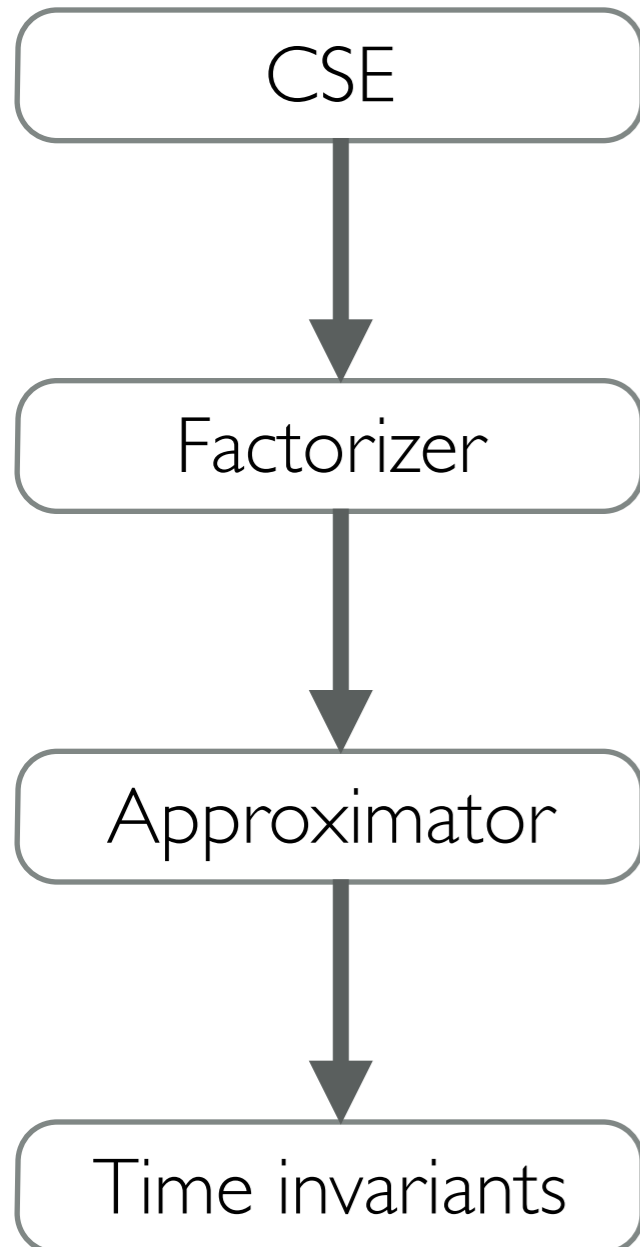
A sequence of compiler passes to reduce FLOPS (no loops at this stage!)



- Common sub-expressions elimination
 - C compilers do it already... but necessary for symbolic processing and compilation speed
- Heuristic re-factorization of recurrent terms
 - E.g., finite difference weights: $0.3*a + \dots + 0.3*b \Rightarrow 0.3*(a+b)$
 - Many possibilities (doesn't leverage domain properties yet!)
- Trigonometric functions are evil
 - Extremely costly
 - Therefore, approximation with e.g. Taylor polynomials
 - Vectorizable, quicker to compile

Devito Symbolic Engine

A sequence of compiler passes to reduce FLOPS (no loops at this stage!)



- Common sub-expressions elimination
 - C compilers do it already... but necessary for symbolic processing and compilation speed
- Heuristic re-factorization of recurrent terms
 - E.g., finite difference weights: $0.3*a + \dots + 0.3*b \Rightarrow 0.3*(a+b)$
 - Many possibilities (doesn't leverage domain properties yet!)
- Trigonometric functions are evil
 - Extremely costly
 - Therefore, approximation with e.g. Taylor polynomials
 - Vectorizable, quicker to compile
- Heuristic hoisting of time-invariant quantities
 - Currently, only Approximator's output (but pass is general)...
 - ... to minimize extra memory consumption
 - This is enhanced by the "aliases detection algorithm"

Devito Loop Engine

A sequence of compiler passes to introduce parallelism, SIMD vectorization and to improve data locality

Devito Loop Engine

A sequence of compiler passes to introduce parallelism, SIMD vectorization and to improve data locality

General Exploration General Exploration viewpoint

Collection Log Analysis Target Analysis Type Summary

Elapsed Time: 487.651s

Clockticks:	5,099,766,000,000
Instructions Retired:	4,963,350,000,000
CPI Rate:	1.027
MUX Reliability:	0.754
Front-End Bound:	1.7%
Back-End Bound:	77.5%
Memory Bound:	54.7%
L1 Bound:	31.0%
L2 Bound:	2.0%
L3 Bound:	3.3%
DRAM Bound:	24.7%
Memory Bandwidth:	24.1%
Memory Latency:	34.7%
Local DRAM:	95.0%
Remote DRAM:	0.0%
Remote Cache:	0.0%
Store Bound:	1.8%
Core Bound:	22.7%
Retiring:	20.5%
Total Thread Count:	205

DTLB Overhead: 3.9%

Loads Blocked by Store Forwarding: 0.0%

Lock Latency: 0.0%

Split Loads:	13.5%
4K Aliasing:	22.8%
FB Full:	22.1%

L1 Bound: 0.0%

Intel VTune, Broadwell E5-2620 v4, TTI space orders 4-8-12

MUX Reliability:	0.751
Front-End Bound:	2.0%
Back-End Bound:	79.0%
Memory Bound:	55.2%
L1 Bound:	31.4%
L2 Bound:	0.0%
L3 Bound:	4.0%
DRAM Bound:	24.5%
Memory Bandwidth:	15.3%
Memory Latency:	37.5%
Local DRAM:	19.3%
Remote DRAM:	0.0%
Remote Cache:	0.0%
Store Bound:	3.5%

DTLB Overhead: 3.9%

Loads Blocked by Store Forwarding: 0.0%

Lock Latency: 0.0%

Split Loads:	68.3%
4K Aliasing:	31.4%
FB Full:	89.2%

MUX Reliability:	0.752
Front-End Bound:	2.2%
Back-End Bound:	76.0%
Memory Bound:	53.1%
L1 Bound:	28.5%
L2 Bound:	0.0%
L3 Bound:	11.2%
DRAM Bound:	12.1%
Contested Accesses:	0.0%
Data Sharing:	0.0%
L3 Latency:	5.2%
SQ Full:	0.3%
Memory Bandwidth:	12.1%
Memory Latency:	40.1%

DTLB Overhead: 3.0%

Loads Blocked by Store Forwarding: 0.0%

Lock Latency: 0.0%

Split Loads:	38.7%
4K Aliasing:	36.6%
FB Full:	100.0%

Devito Loop Engine

A sequence of compiler passes to introduce parallelism, SIMD vectorization and to improve data locality

Cache opts

- Cache optimizations (mostly L1 cache)
 - Loop fission + elemental functions (register locality)
 - Padding + data alignment (split loads)
 - Work in progress: data layout transformations

Intel VTune, Broadwell E5-2620 v4, TTI space orders 4-8-12

Metric	Value
Elapsed Time	487.651s
Clockticks	5,099,766,000,000
Instructions Retired	4,963,350,000,000
CPI Rate	1.027
MUX Reliability	0.754
Front-End Bound	1.7%
Back-End Bound	77.5%
Memory Bound	54.7%
L1 Bound	31.0%
L2 Bound	2.0%
L3 Bound	3.3%
DRAM Bound	24.7%
Memory Bandwidth	24.1%
Memory Latency	34.7%
Local DRAM	95.0%
Remote DRAM	0.0%
Remote Cache	0.0%
Store Bound	1.8%
Core Bound	22.7%
Retiring	20.5%
Total Thread Count	205
Split Loads	13.5%
4K Aliasing	22.8%
FB Full	22.1%

MUX Reliability	0.751
Front-End Bound	2.0%
Back-End Bound	79.0%
Memory Bound	55.2%
L1 Bound	31.4%
L2 Bound	0.0%
L3 Bound	4.0%
DRAM Bound	24.5%
Memory Bandwidth	15.3%
Memory Latency	37.5%
Local DRAM	19.3%
Remote DRAM	0.0%
Remote Cache	0.0%
Store Bound	3.5%
Core Bound	22.7%
Retiring	20.5%
Total Thread Count	205
Split Loads	68.3%
4K Aliasing	31.4%
FB Full	89.2%

MUX Reliability	0.752
Front-End Bound	2.2%
Back-End Bound	76.0%
Memory Bound	53.1%
L1 Bound	28.5%
L2 Bound	0.0%
L3 Bound	11.2%
DRAM Bound	12.1%
Contested Accesses	0.0%
Data Sharing	0.0%
L3 Latency	5.2%
SQ Full	0.3%
Memory Bandwidth	12.1%
Memory Latency	40.1%
Store Bound	1.8%
Core Bound	22.7%
Retiring	20.5%
Total Thread Count	205
Split Loads	38.7%
4K Aliasing	36.6%
FB Full	100.0%

Devito Loop Engine

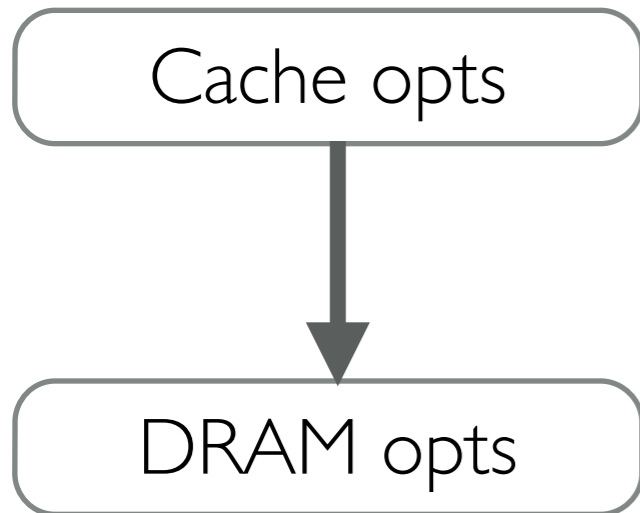
A sequence of compiler passes to introduce parallelism, SIMD vectorization and to improve data locality

Cache opts

- Cache optimizations (mostly L1 cache)
 - Loop fission + elemental functions (register locality)
 - Padding + data alignment (split loads)
 - Work in progress: data layout transformations

Devito Loop Engine

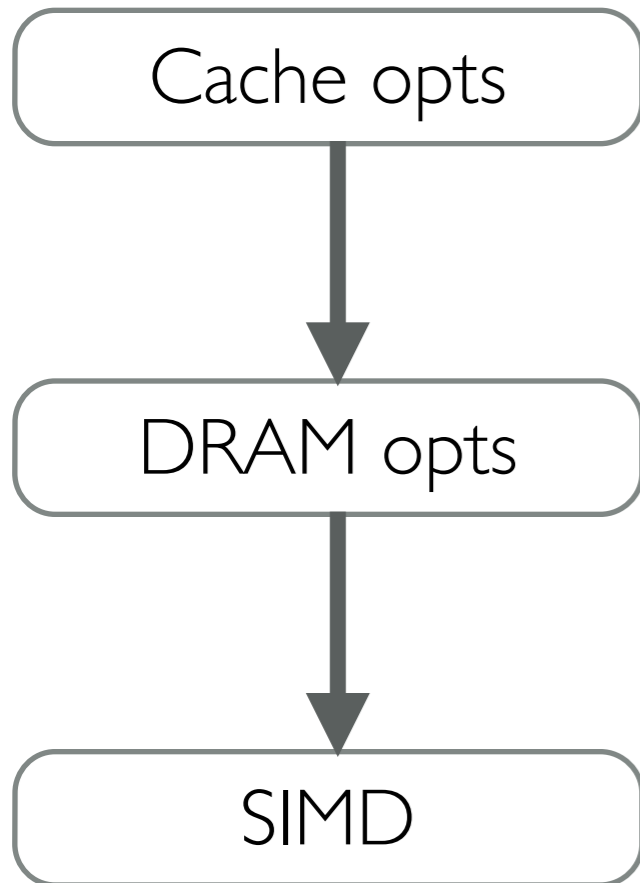
A sequence of compiler passes to introduce parallelism, SIMD vectorization and to improve data locality



- Cache optimizations (mostly L1 cache)
 - Loop fission + elemental functions (register locality)
 - Padding + data alignment (split loads)
 - Work in progress: data layout transformations
- DRAM optimizations: loop blocking
 - 1D, 2D, 3D supported (but no time loop)
 - Auto-tuning supported

Devito Loop Engine

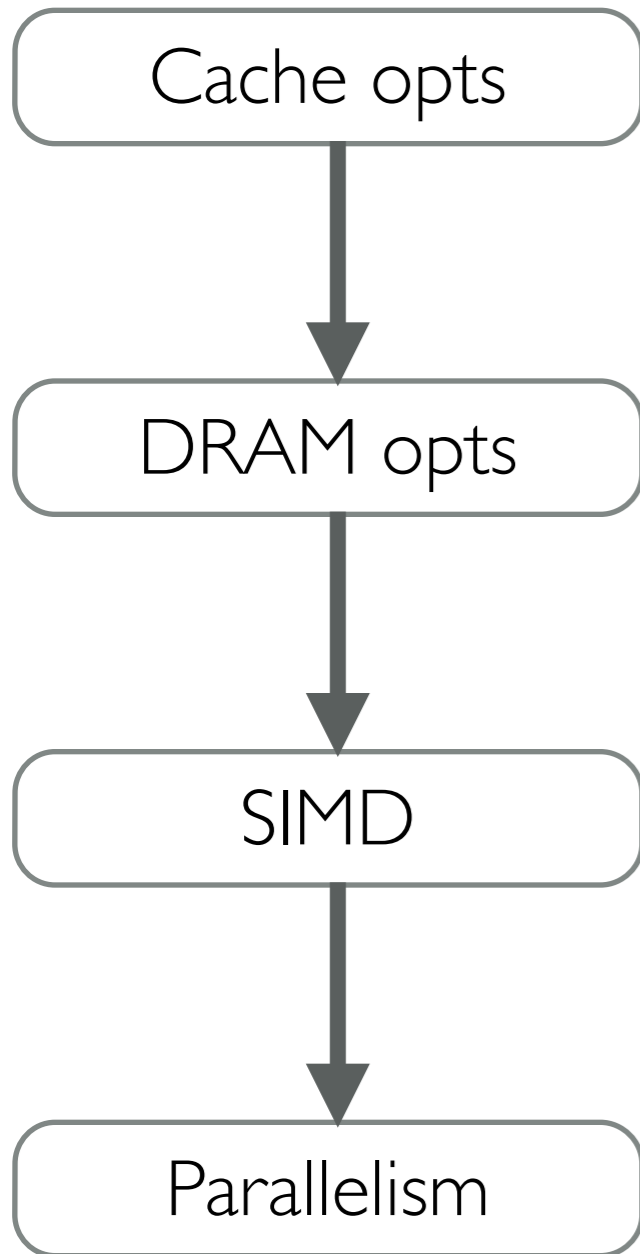
A sequence of compiler passes to introduce parallelism, SIMD vectorization and to improve data locality



- Cache optimizations (mostly L1 cache)
 - Loop fission + elemental functions (register locality)
 - Padding + data alignment (split loads)
 - Work in progress: data layout transformations
- DRAM optimizations: loop blocking
 - 1D, 2D, 3D supported (but no time loop)
 - Auto-tuning supported
- SIMD vectorization
 - Through compiler auto-vectorization
 - Why should I bother using intrinsics?
 - Various *#pragmas* introduced (e.g., *ivdep*, *alignment*, ...)

Devito Loop Engine

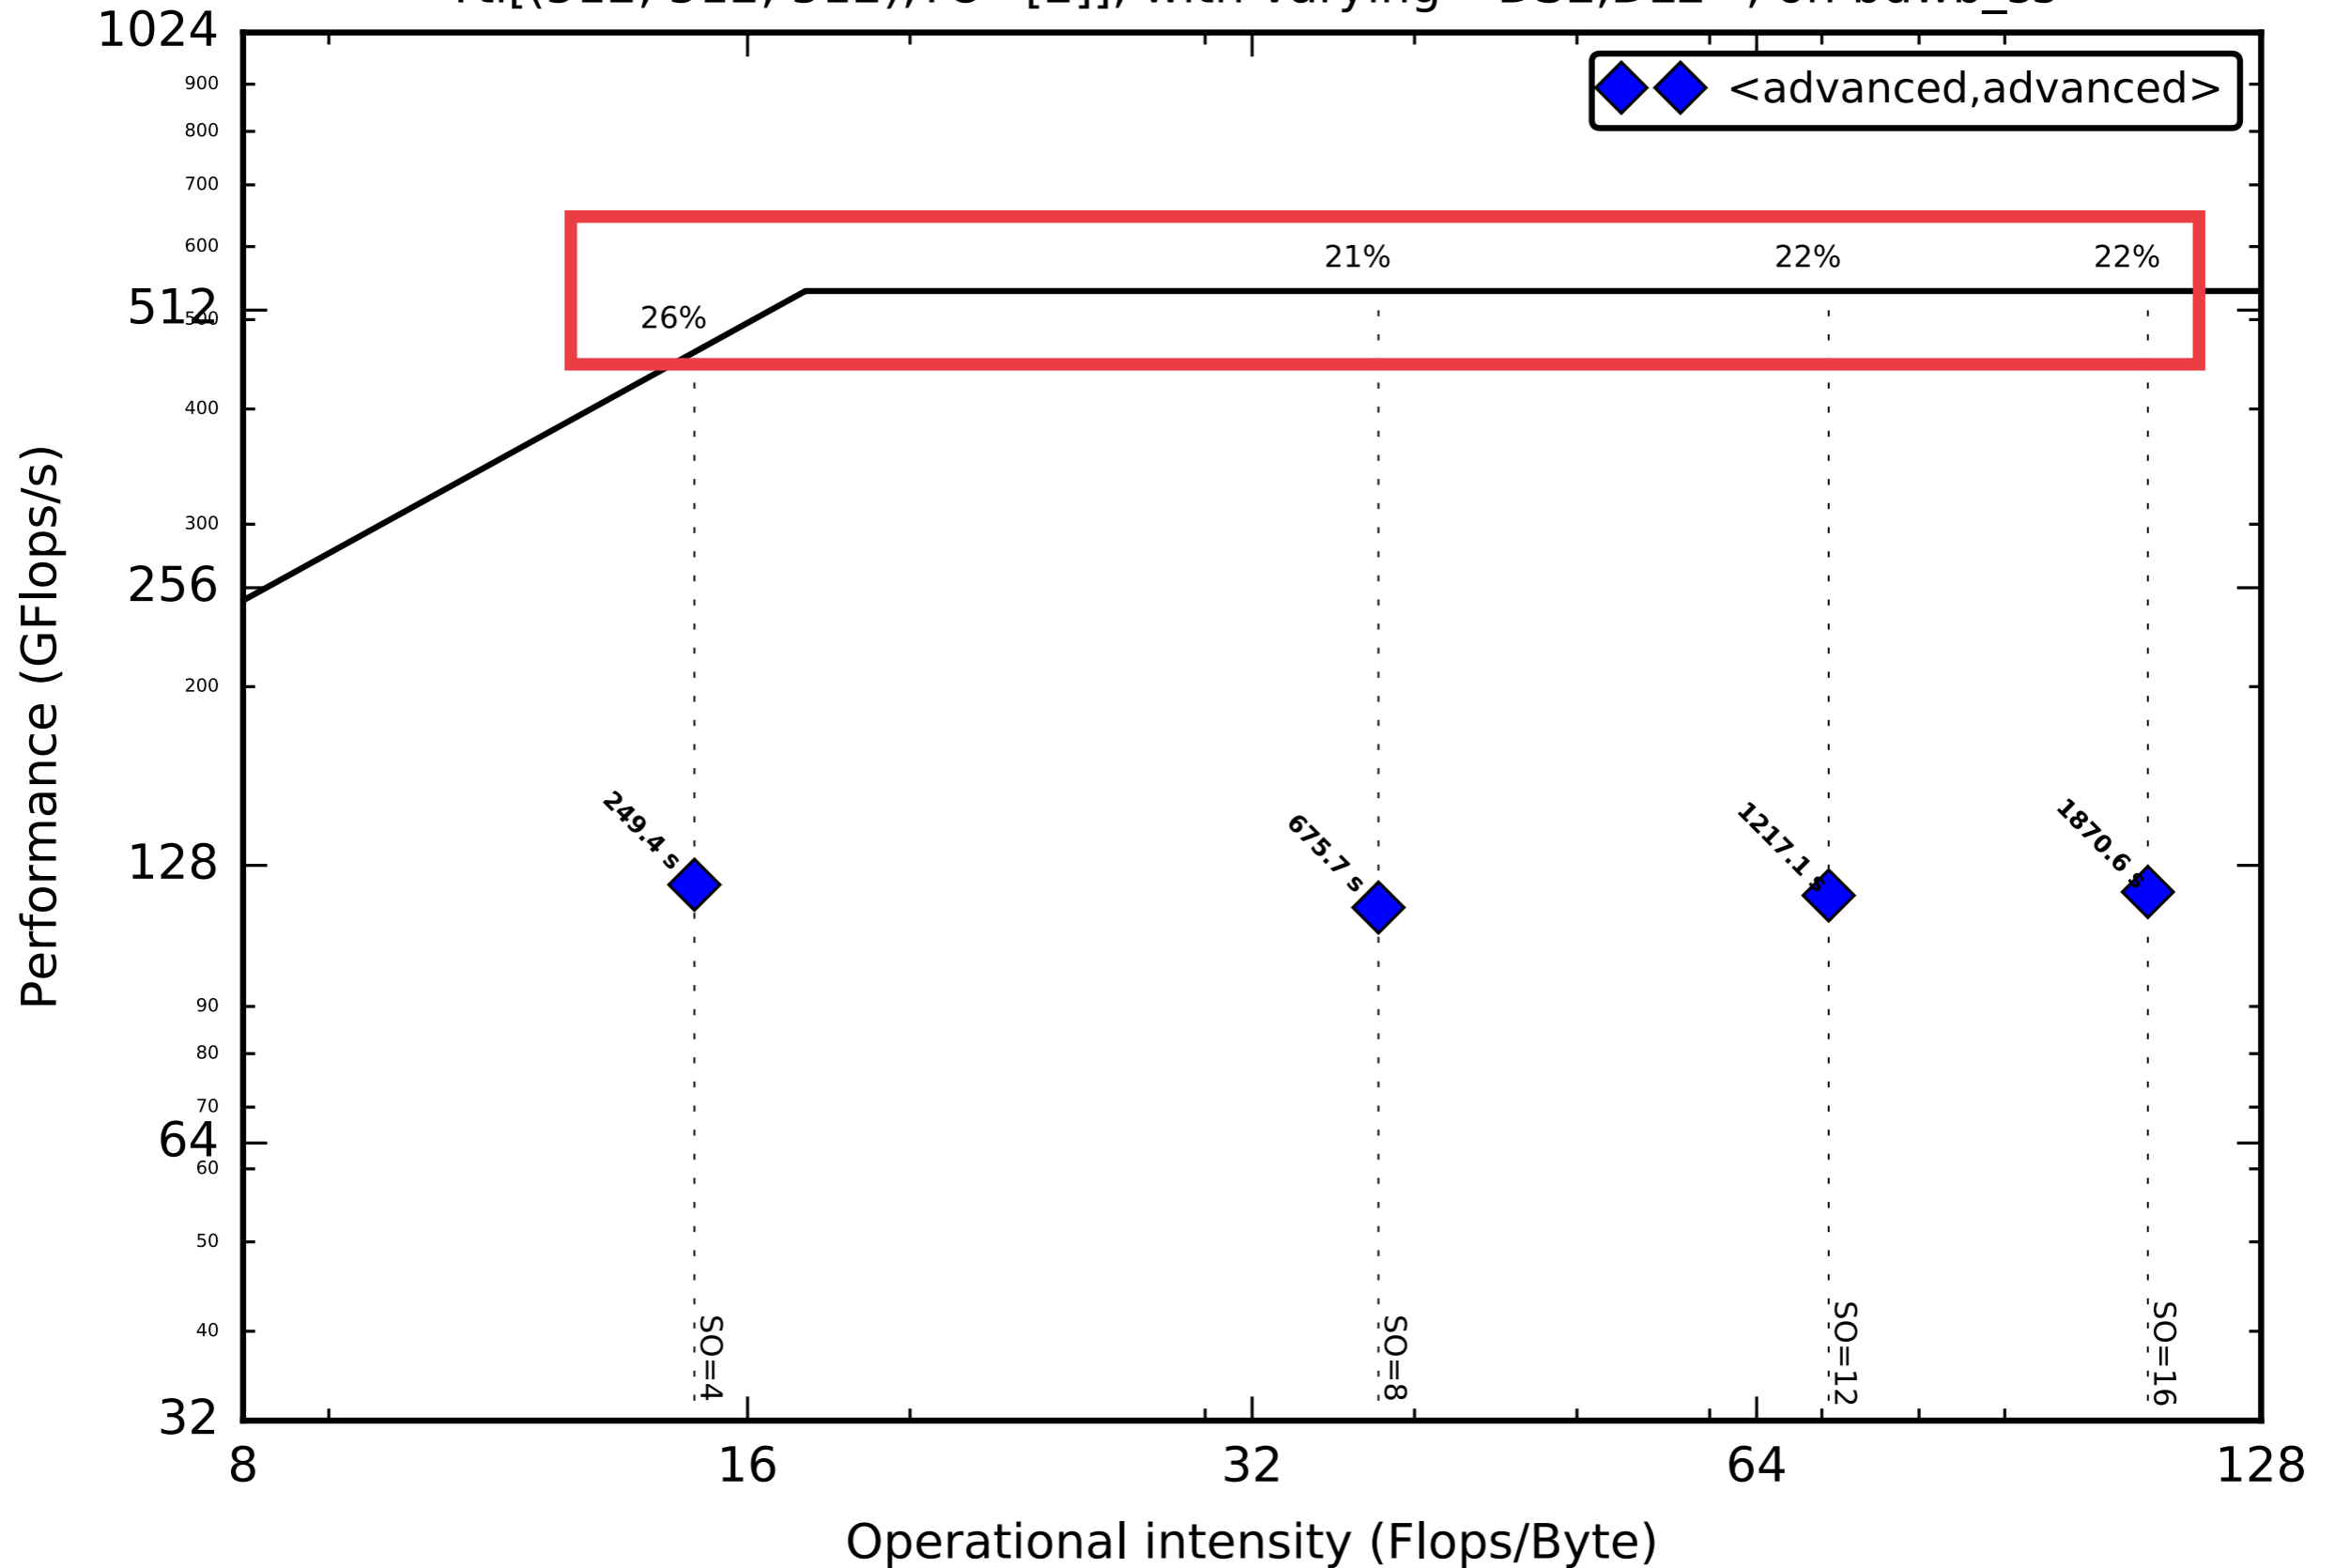
A sequence of compiler passes to introduce parallelism, SIMD vectorization and to improve data locality



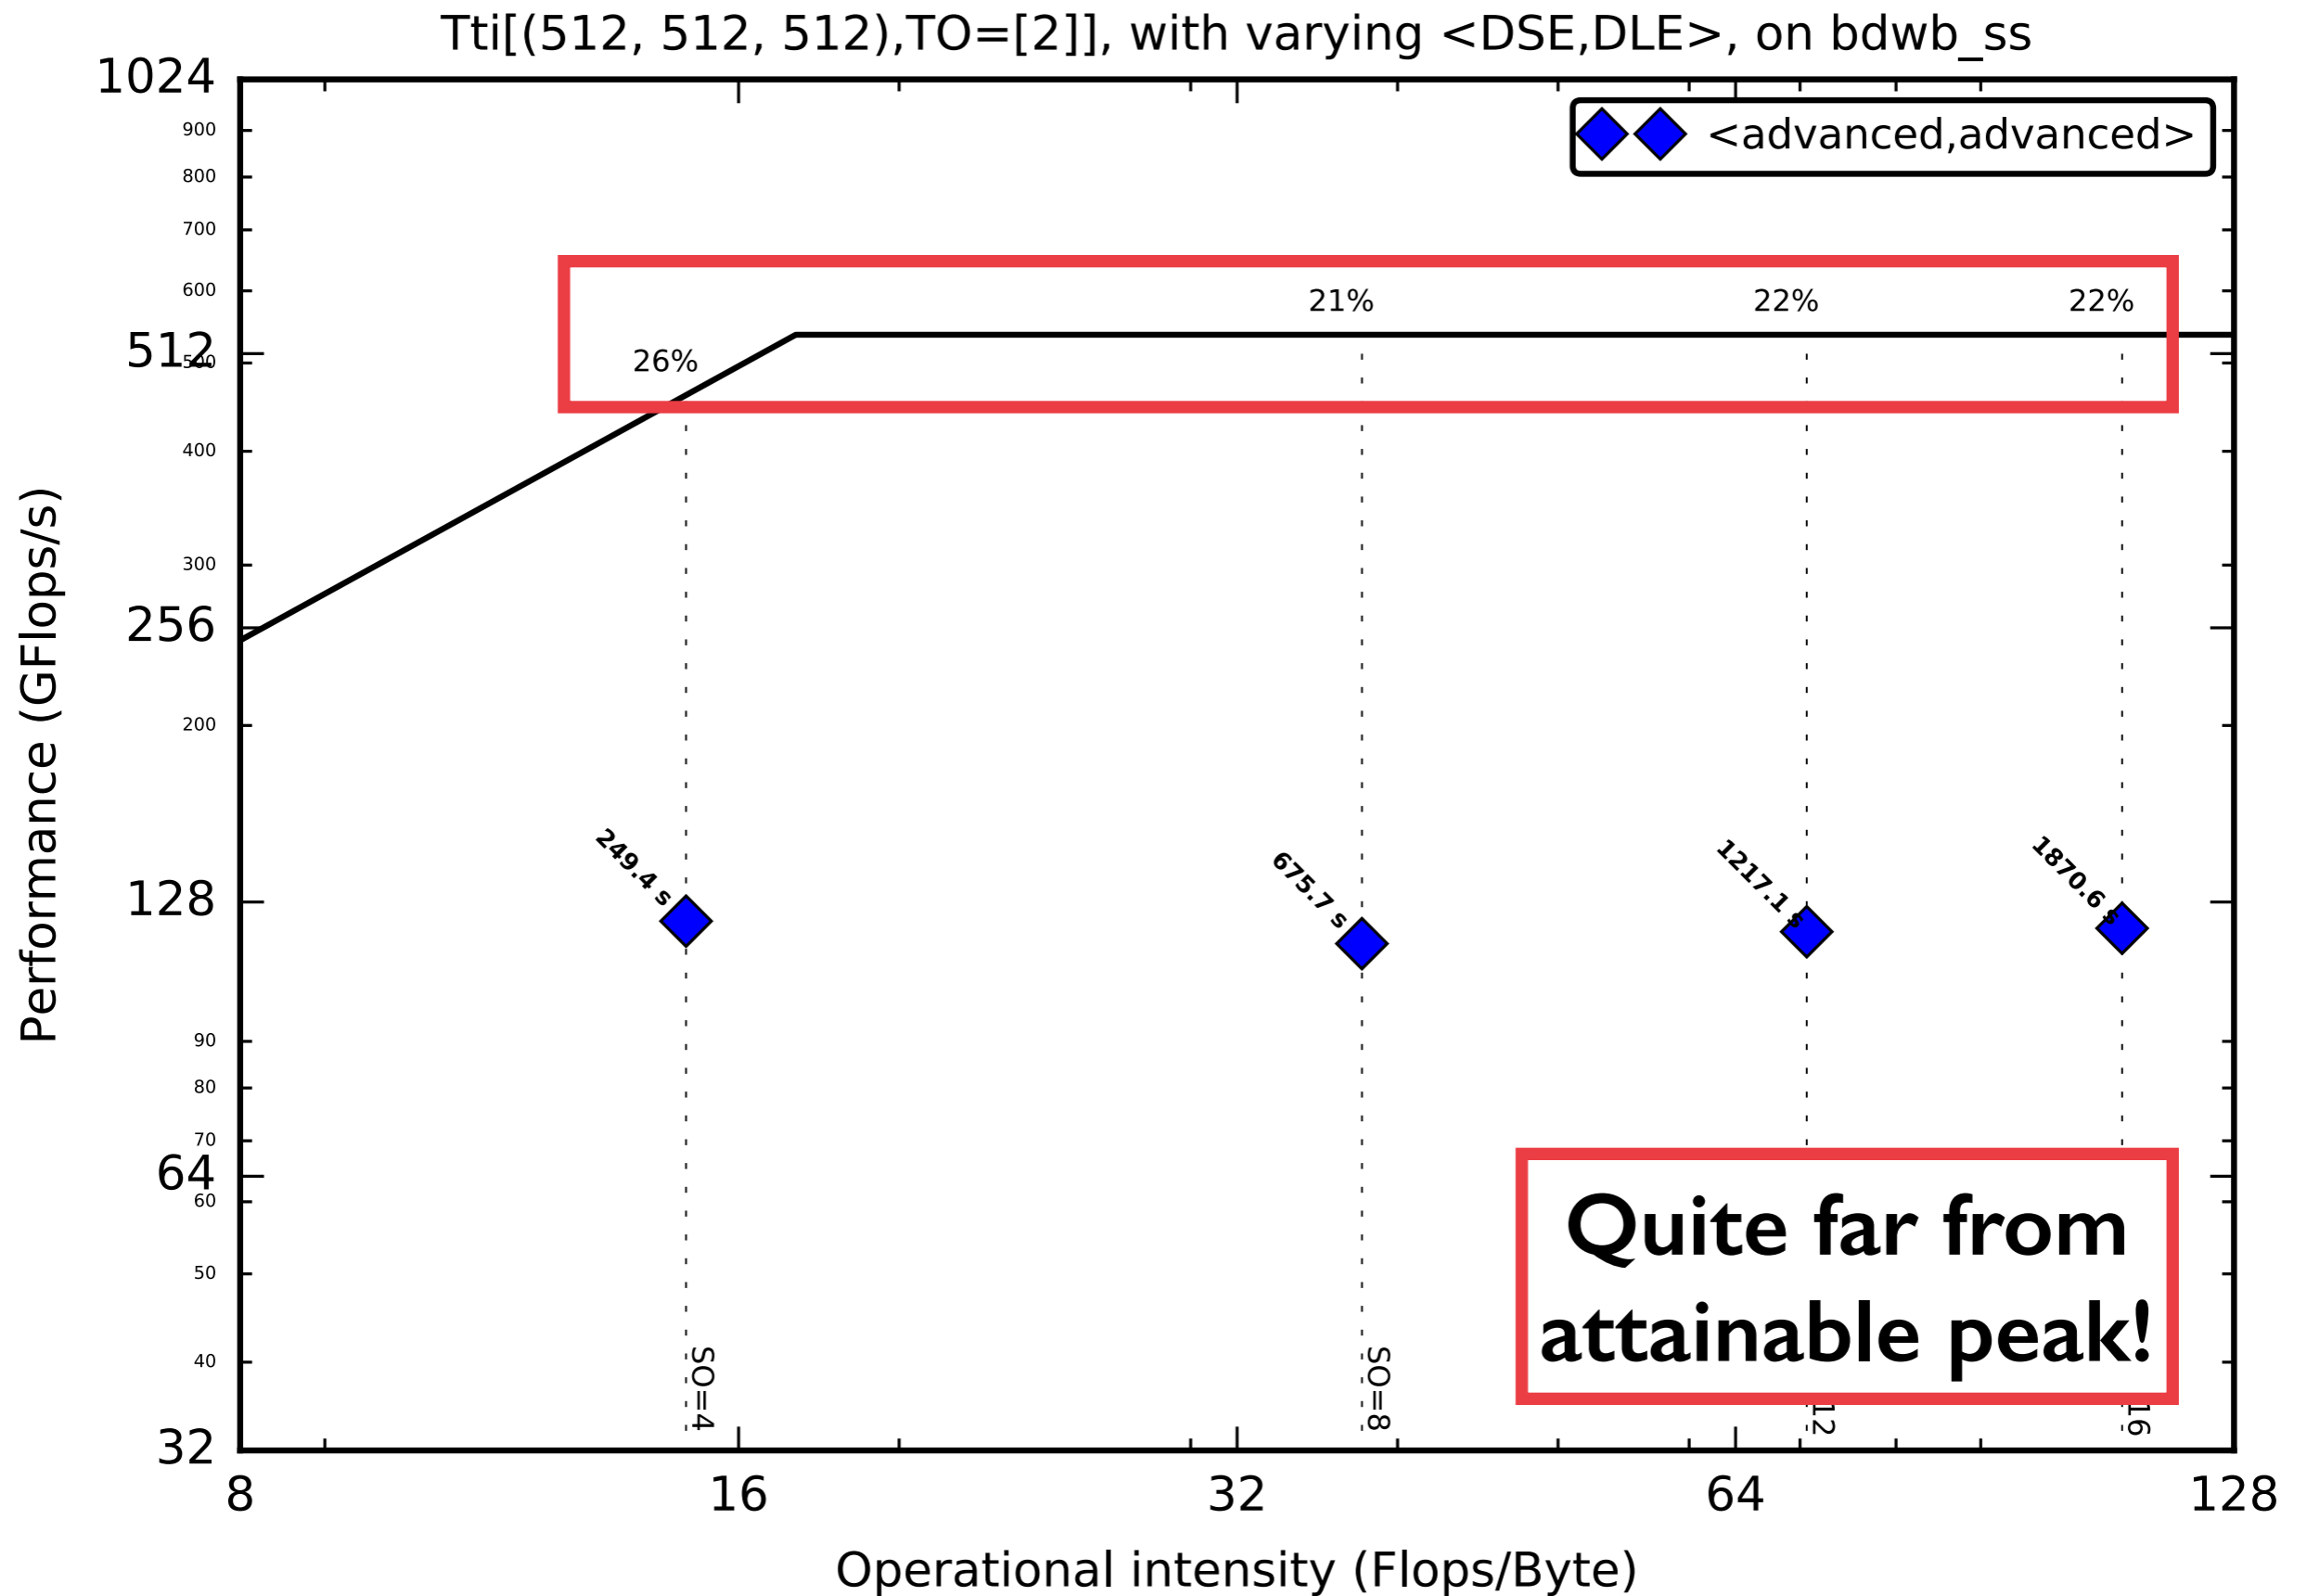
- Cache optimizations (mostly L1 cache)
 - Loop fission + elemental functions (register locality)
 - Padding + data alignment (split loads)
 - Work in progress: data layout transformations
- DRAM optimizations: loop blocking
 - 1D, 2D, 3D supported (but no time loop)
 - Auto-tuning supported
- SIMD vectorization
 - Through compiler auto-vectorization
 - Why should I bother using intrinsics?
 - Various `#pragmas` introduced (e.g., `ivdep`, `alignment`, ...)
- OpenMP
 - `#pragma collapse` clause on the Xeon Phi

TTI on Broadwell (8 threads, single socket)

Tti[(512, 512, 512),TO=[2]], with varying <DSE,DLE>, on bdwb_ss

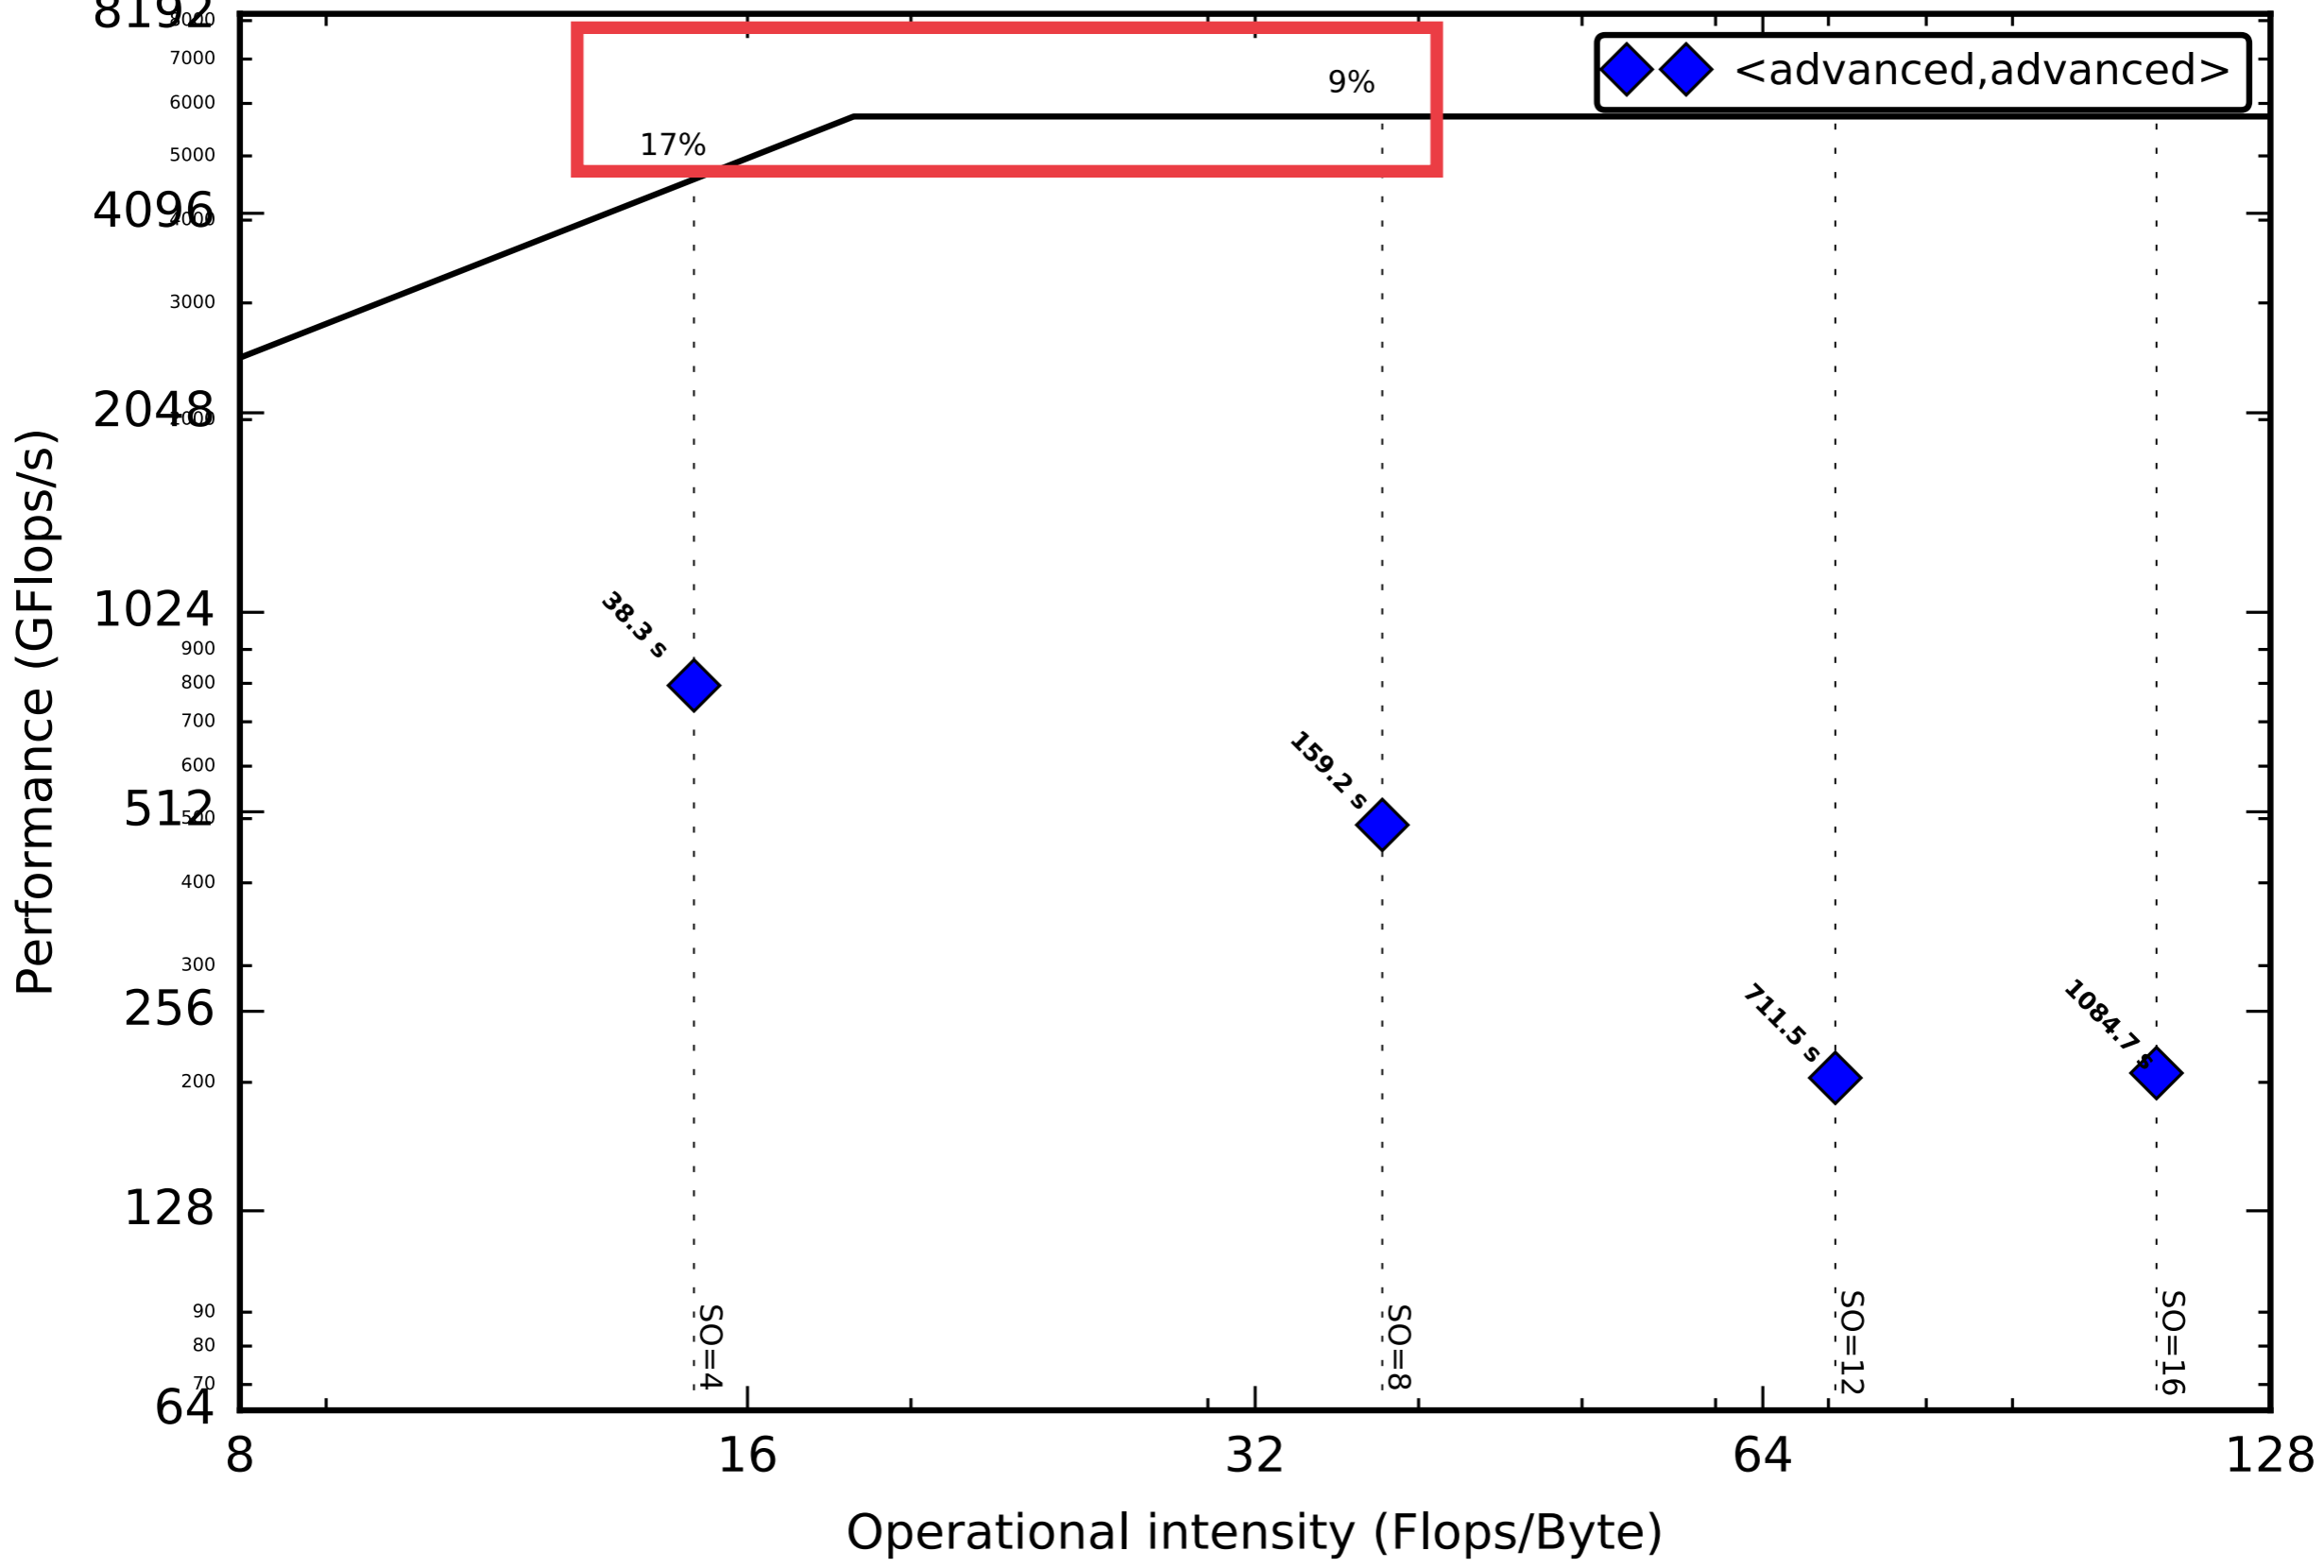


TTI on Broadwell (8 threads, single socket)



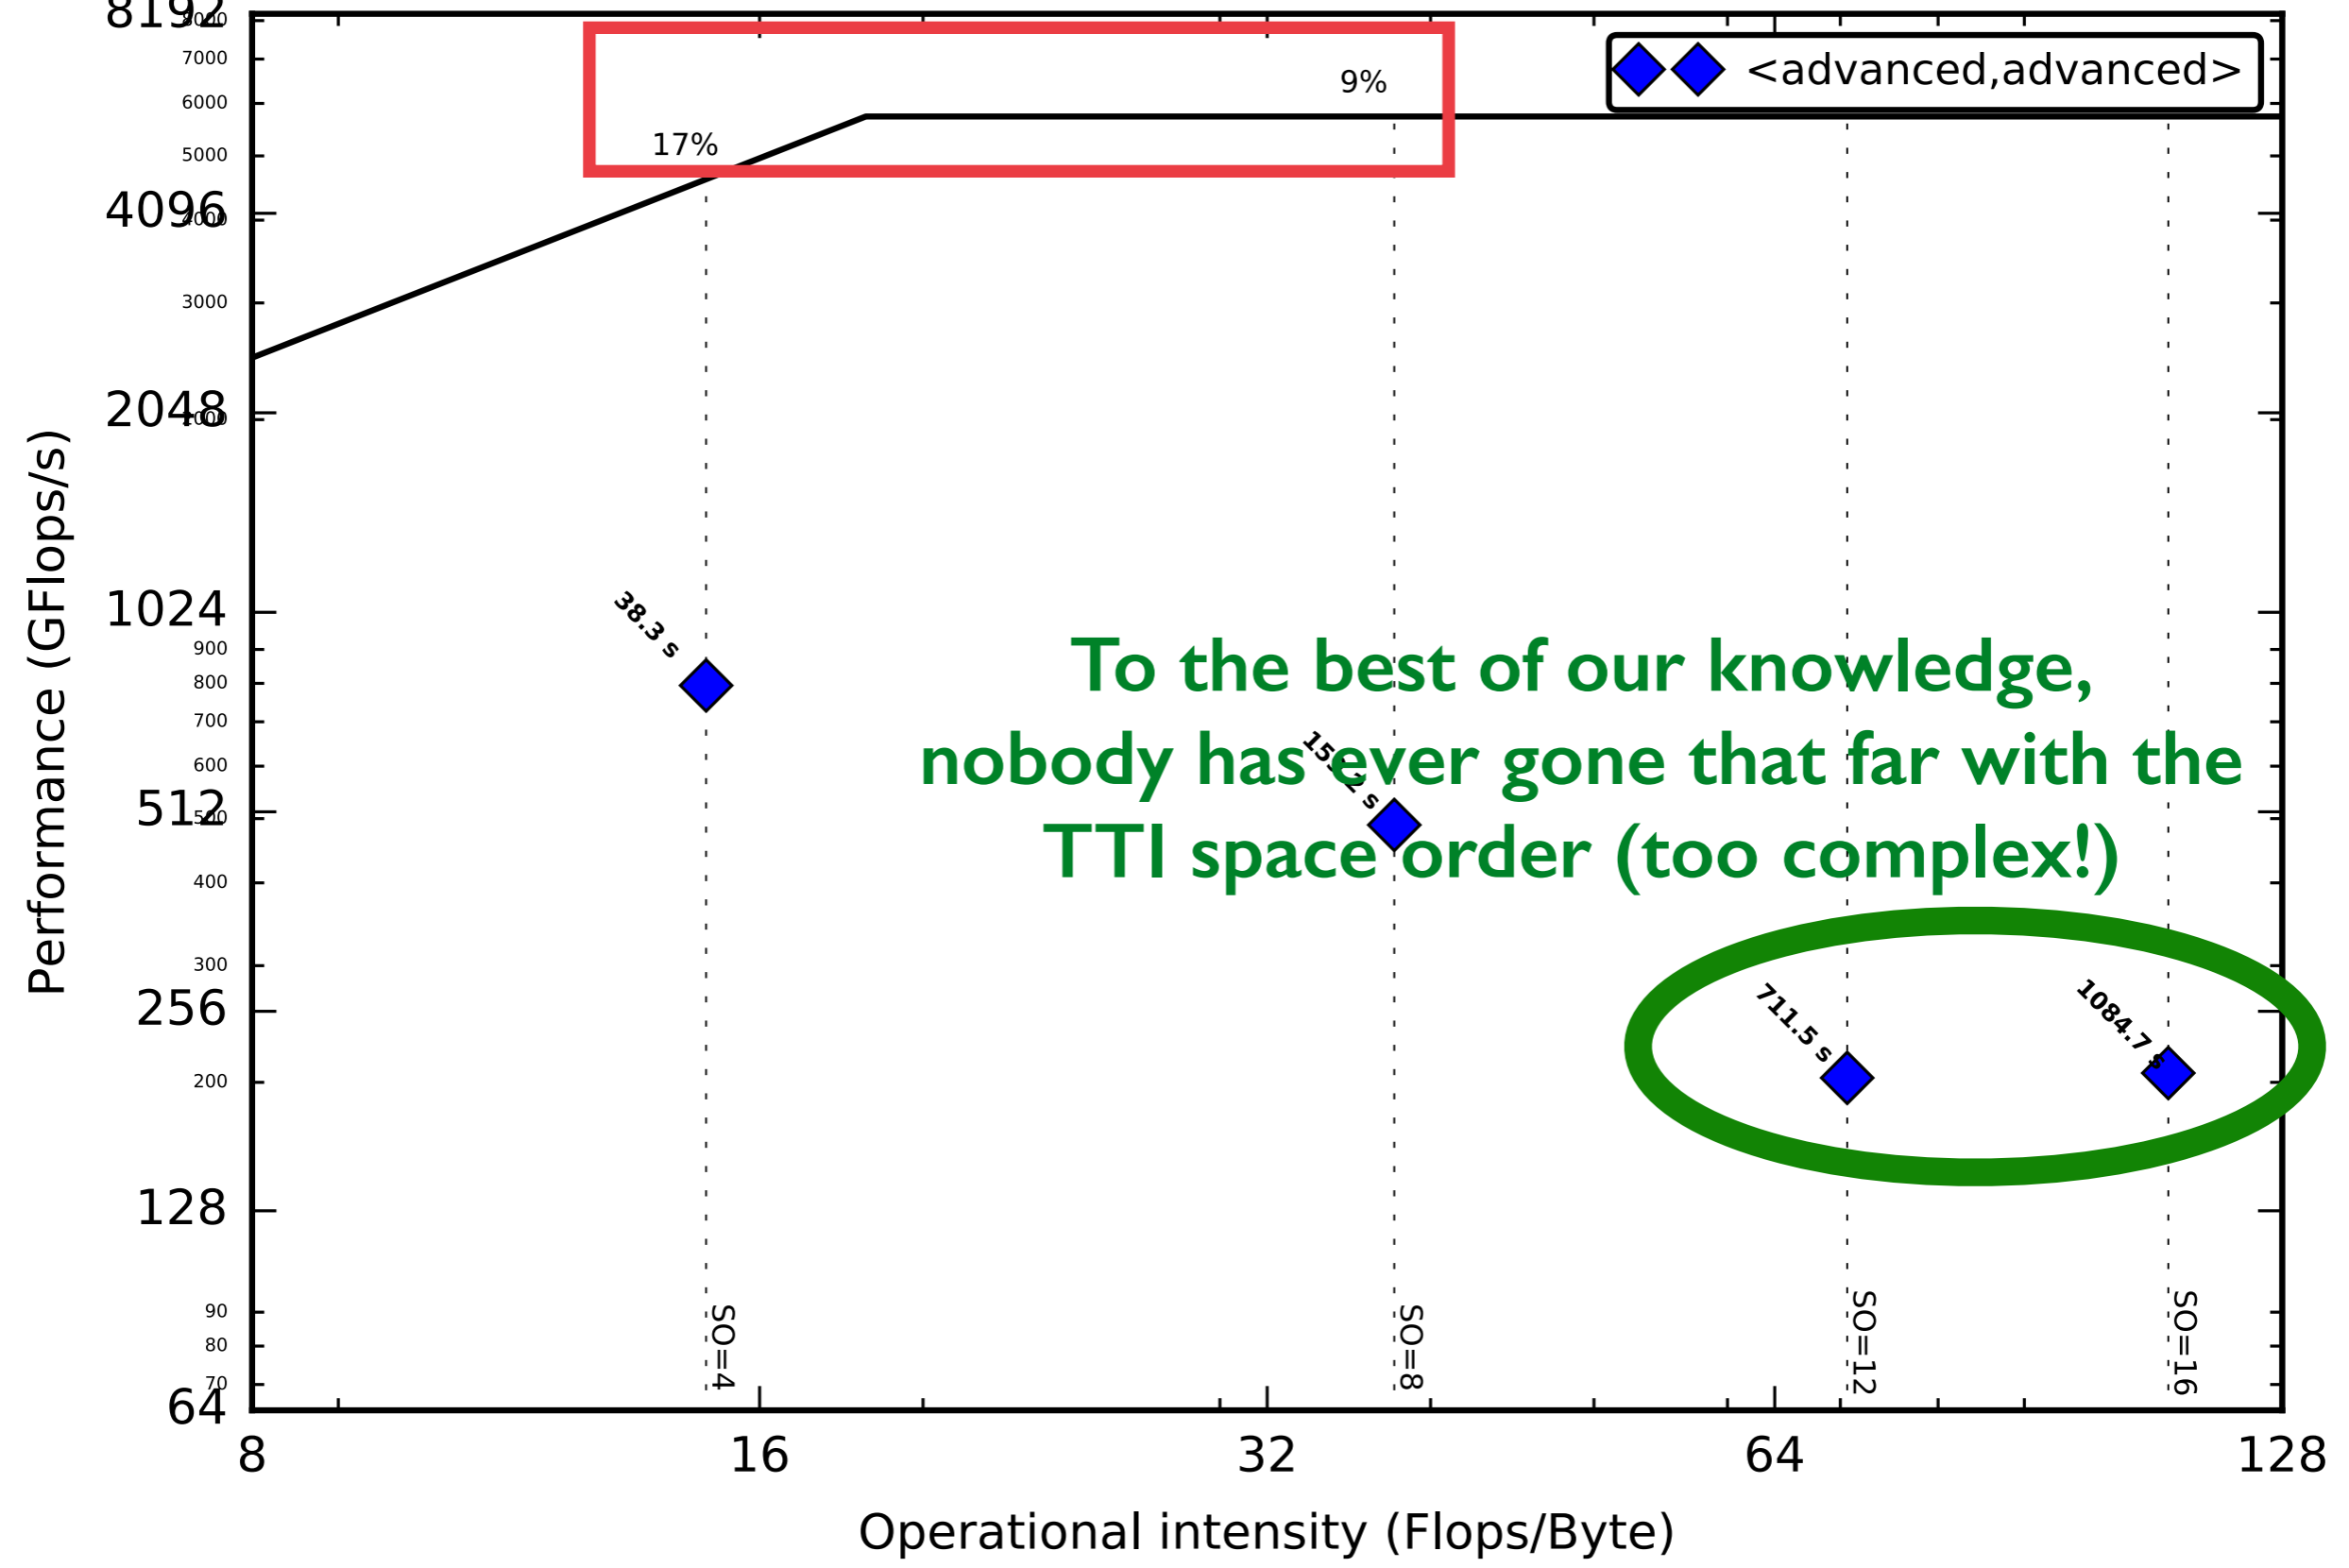
TTI on Xeon Phi (64 threads, cache mode, quadrant)

Tti[(512, 512, 512),TO=[2]], with varying <DSE,DLE>, on ekf_1



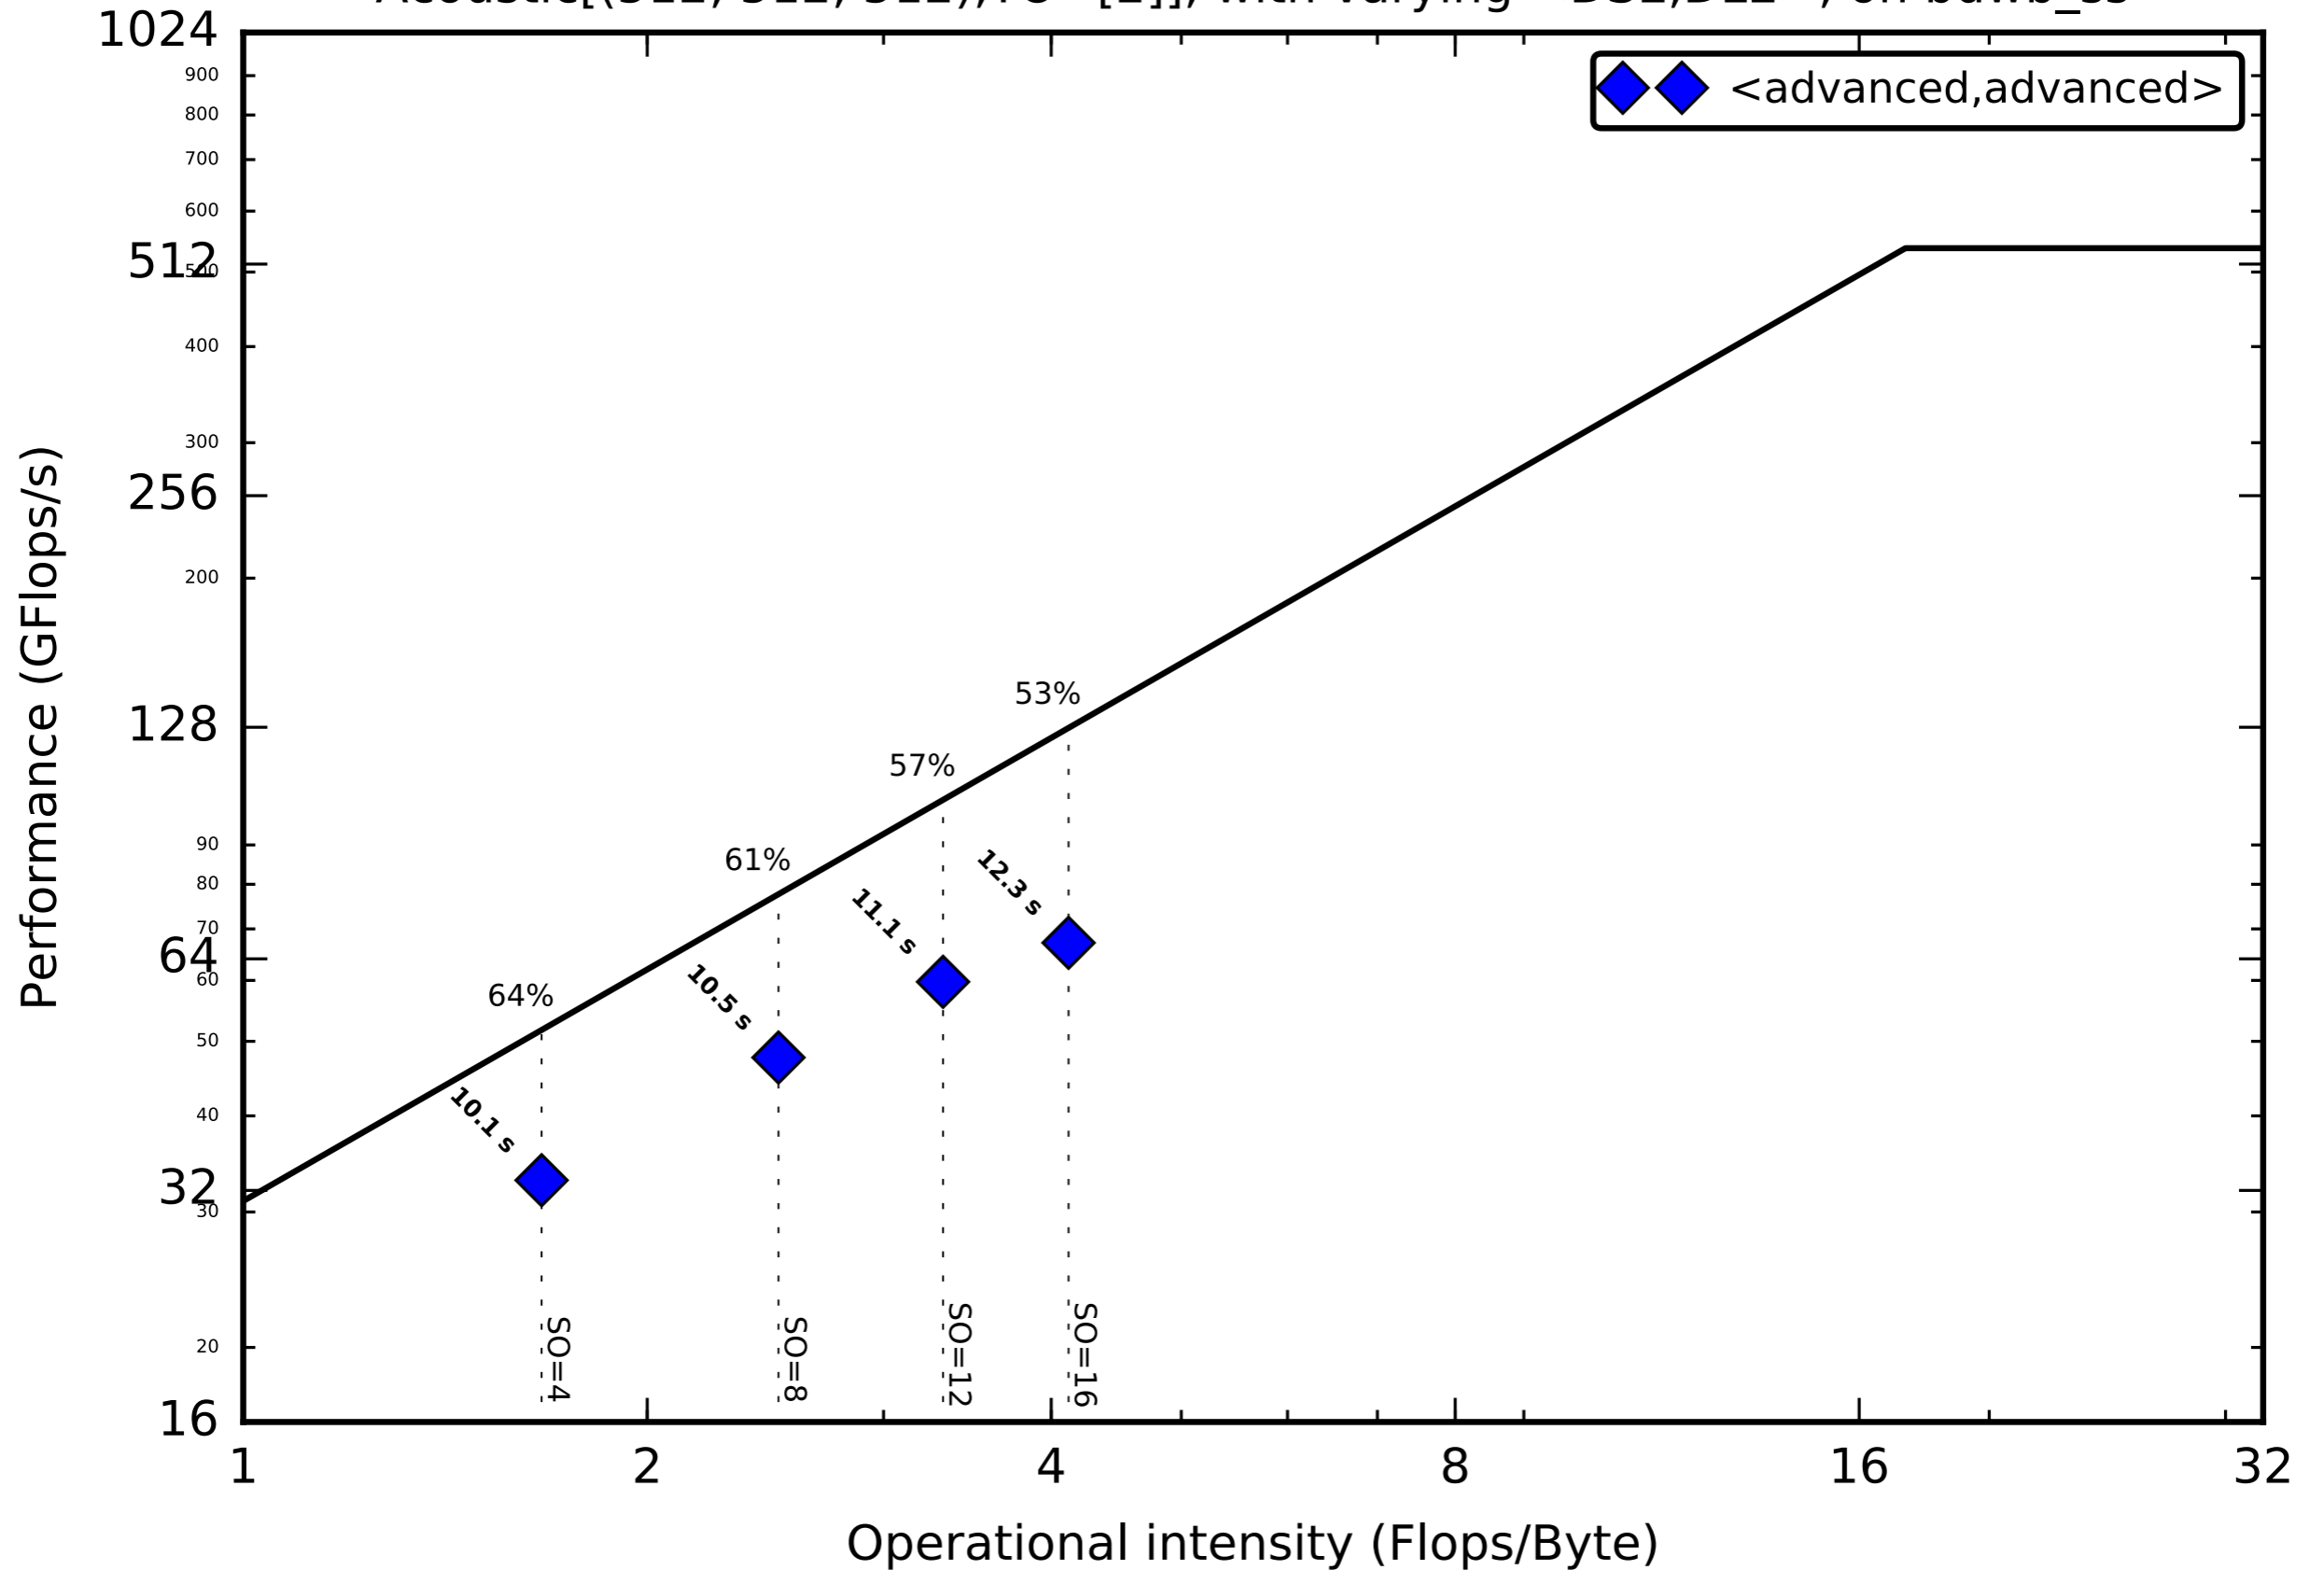
TTI on Xeon Phi (64 threads, cache mode, quadrant)

Tti[(512, 512, 512),TO=[2]], with varying <DSE,DLE>, on ekf_1



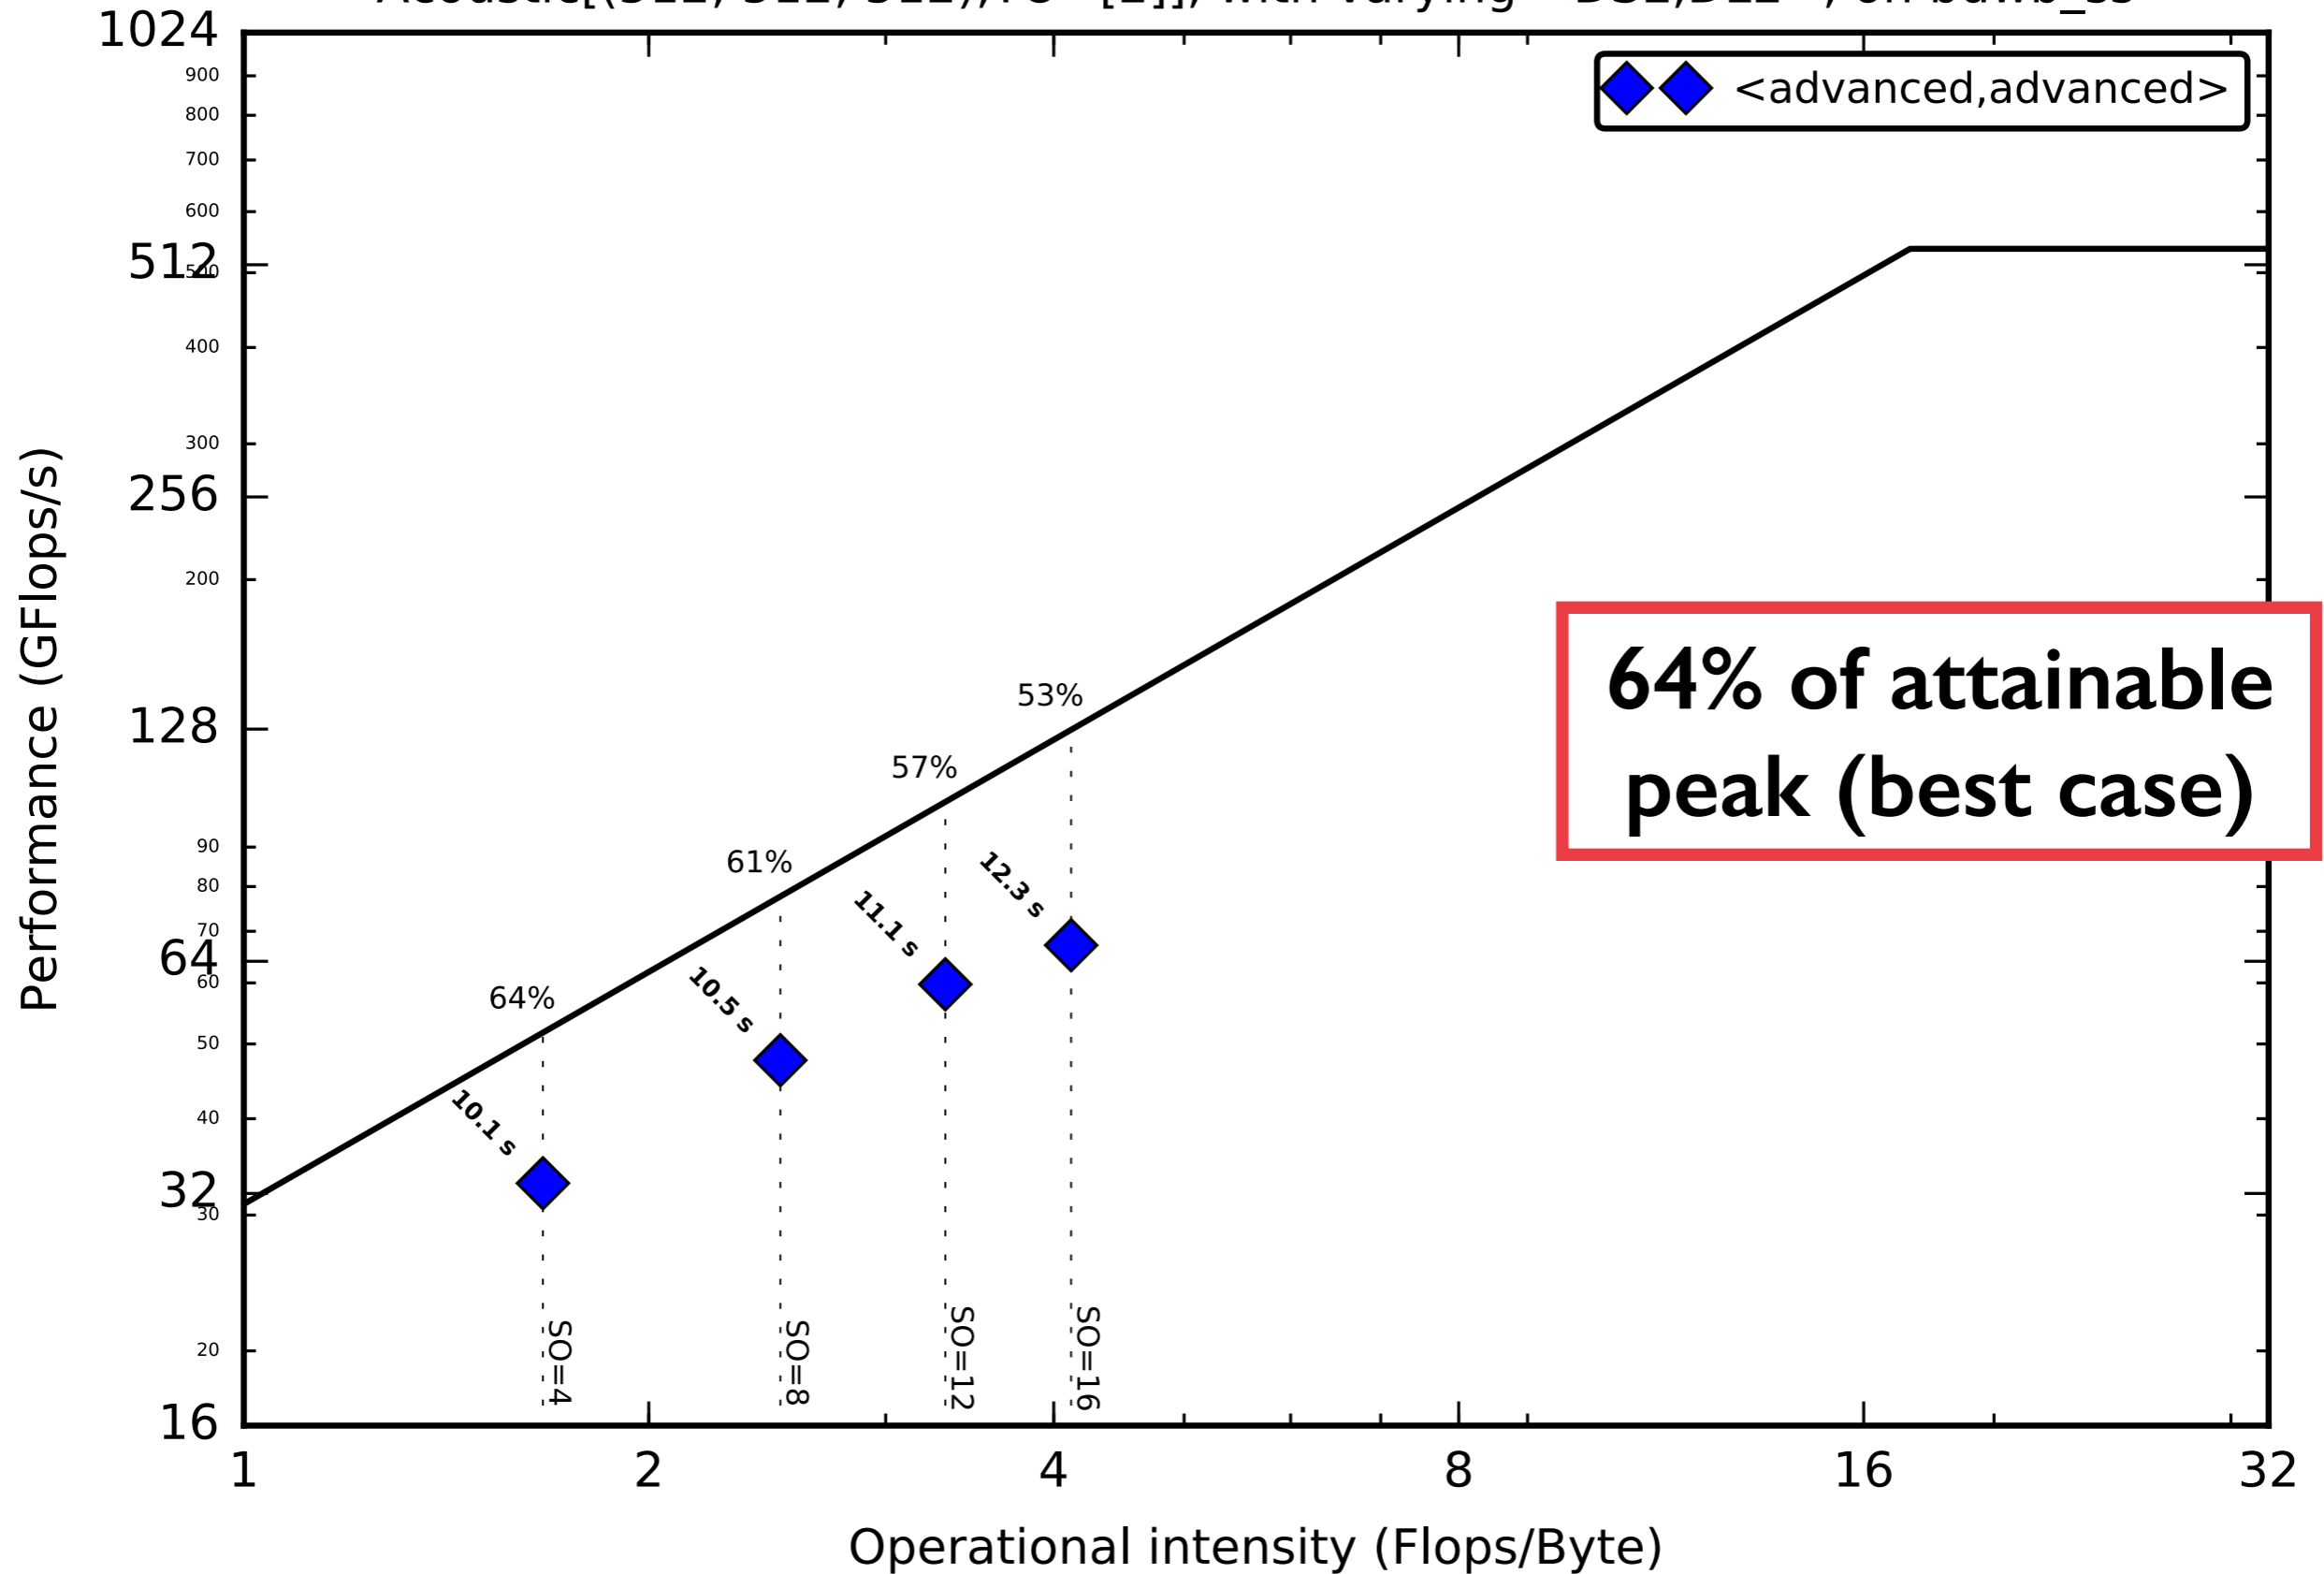
Acoustic on Broadwell

Acoustic[(512, 512, 512),TO=[2]], with varying <DSE,DLE>, on bdwb_ss



Acoustic on Broadwell

Acoustic[(512, 512, 512),TO=[2]], with varying <DSE,DLE>, on bdwb_ss



Conclusions and resources

- Devito (part of OPESCI): towards an efficient and sustainable finite difference DSL
- Driven/inspired by real-world seismic imaging
- Based on actual compiler technology
- Performance: promising, but still quite a lot to do
- Future: plug in backends such as YASK

Useful links

- <http://www.opesci.org>
- <https://github.com/opesci/devito>



Appendix


Experimentation details

- Compiler
 - ICC 17 -xHost -O3 (-O2 no difference)
 - -xMIC-AVX512 on Xeon Phi
- OpenMP
 - Single socket (still no support for NUMA issue through first touch)
 - Thread pinning
 - Numactl
- Intel(R) Xeon(R) E5-2620 v4 2.1 Ghz “Broadwell” (8 cores per socket)
- Intel(R) XeonPhi(R) 7650
 - 68 cores (used only 64)
 - Quadrant mode (still no support for NUMA)
 - Tried 1, 2, 4 threads per core. Shown 1 thread (no critical differences)
 - Cache mode performs equivalently to Flat mode when datasets fit in MCDRAM
- Roofline calculations available at: <https://gist.github.com/FabioLuporini/12485f08576674d8452fec8673d6f26e>
 - Memory bandwidth: STREAM
 - CPU peak: pen & paper
 - Operational intensity: source-level analysis (automated through Devito)



Summary of experimentation

	Intel Broadwell	Intel Xeon Phi
Acoustic forward		
TTI forward		




Summary of experimentation

	Intel Broadwell	Intel Xeon Phi
Acoustic forward		
TTI forward		





Summary of experimentation

	Intel Broadwell	Intel Xeon Phi
Acoustic forward		
TTI forward		





Summary of experimentation

	Intel Broadwell	Intel Xeon Phi
Acoustic forward		
TTI forward		

Summary of experimentation

	Intel Broadwell	Intel Xeon Phi
Acoustic forward		
TTI forward		

Summary of experimentation

	Intel Broadwell	Intel Xeon Phi
Acoustic forward		
TTI forward		

Devito Loop Engine (example output)

```
int Kernel(float *restrict damp_vec, float *restrict delta_vec, float *restrict epsilon_vec, float *restrict m_vec, float
*restrict phi_vec, ..., const int x_size, const int y_size, const int z_size, const int x_block_size, const int y_block_size,
struct profiler *timings)
{
    ...

    // PADDED BUFFERS
    float (*pu)[532][532][536];
    ...
    posix_memalign((void**)&pti3, 64, sizeof(float[532][532][536]));

    // TIME INVARIANTS
    for (int x = 0; x < x_size; x += 1)
    {
        for (int y = 0; y < y_size; y += 1)
        {
            #pragma noline
            f_2_0(phi_vec,x_size,x,y_size,y,z_size,(float*) pti0,(float*) pti1,(float*) pti2,(float*) pti3,theta_vec);
        }
    }

    for (int time = 0; time < time_size; time += 1)
    {
        // NEXT SLIDE
    }
    ...
}
```

Devito Loop Engine (example output)

```
#pragma omp parallel
```

```
{  
  /* Flush denormal numbers to zero in hardware */  
  __MM_SET_DENORMALS_ZERO_MODE(__MM_DENORMALS_ZERO_ON);  
  __MM_SET_FLUSH_ZERO_MODE(__MM_FLUSH_ZERO_ON);  
  #pragma omp for schedule(static,1)  
  for (int x_block = 4; x_block < x_size - (x_size - 8)%(x_block_size) - 4; x_block += x_block_size)  
  {  
    for (int y_block = 4; y_block < y_size - (y_size - 8)%(y_block_size) - 4; y_block += y_block_size)  
    {  
      double ptemp276[536] __attribute__((aligned(64)));  
      double ptemp278[536] __attribute__((aligned(64)));  
      // MORE PADDED TEMPORARIES
```

Loop blocking loops

```
      for (int x = x_block; x < x_block + x_block_size; x += 1) Intra-block loops
```

```
      {  
        for (int y = y_block; y < y_block + y_block_size; y += 1)  
        {  
          #pragma noline  
          f_2_1((float*) ptemp276,z_size,(float*) pu,t_size,x_size,x,y_size,y,t1);  
          ...  
          #pragma noline  
          f_2_119((float*) pts48,z_size,(float*) pts49,(float*) pts50,(float*) pv,t_size,x_size,x,y_size,y,t2);  
        }  
      }  
    }  
  }  
}
```

SIMD dimension
within these
functions

```
// BLOCKING REMAINDER LOOPS
```

```
// MODEL SOURCES AND RECEIVERS
```

```
#pragma noline  
f_2_477(m_vec,x_size,y_size,z_size,(float*) pu,t_size,(float*) pv,src_vec,time_size,src_coords_vec,d_size,time,t2);  
#pragma noline  
f_2_478((float*) pu,t_size,x_size,y_size,z_size,(float*) pv,rec_vec,time_size,rec_coords_vec,d_size,time,t2);  
...
```


DLE significantly benefited from Intel VTune

General Exploration General Exploration viewpoint (change)

Collection Log Analysis Target Analysis Type Summary Bottom-up

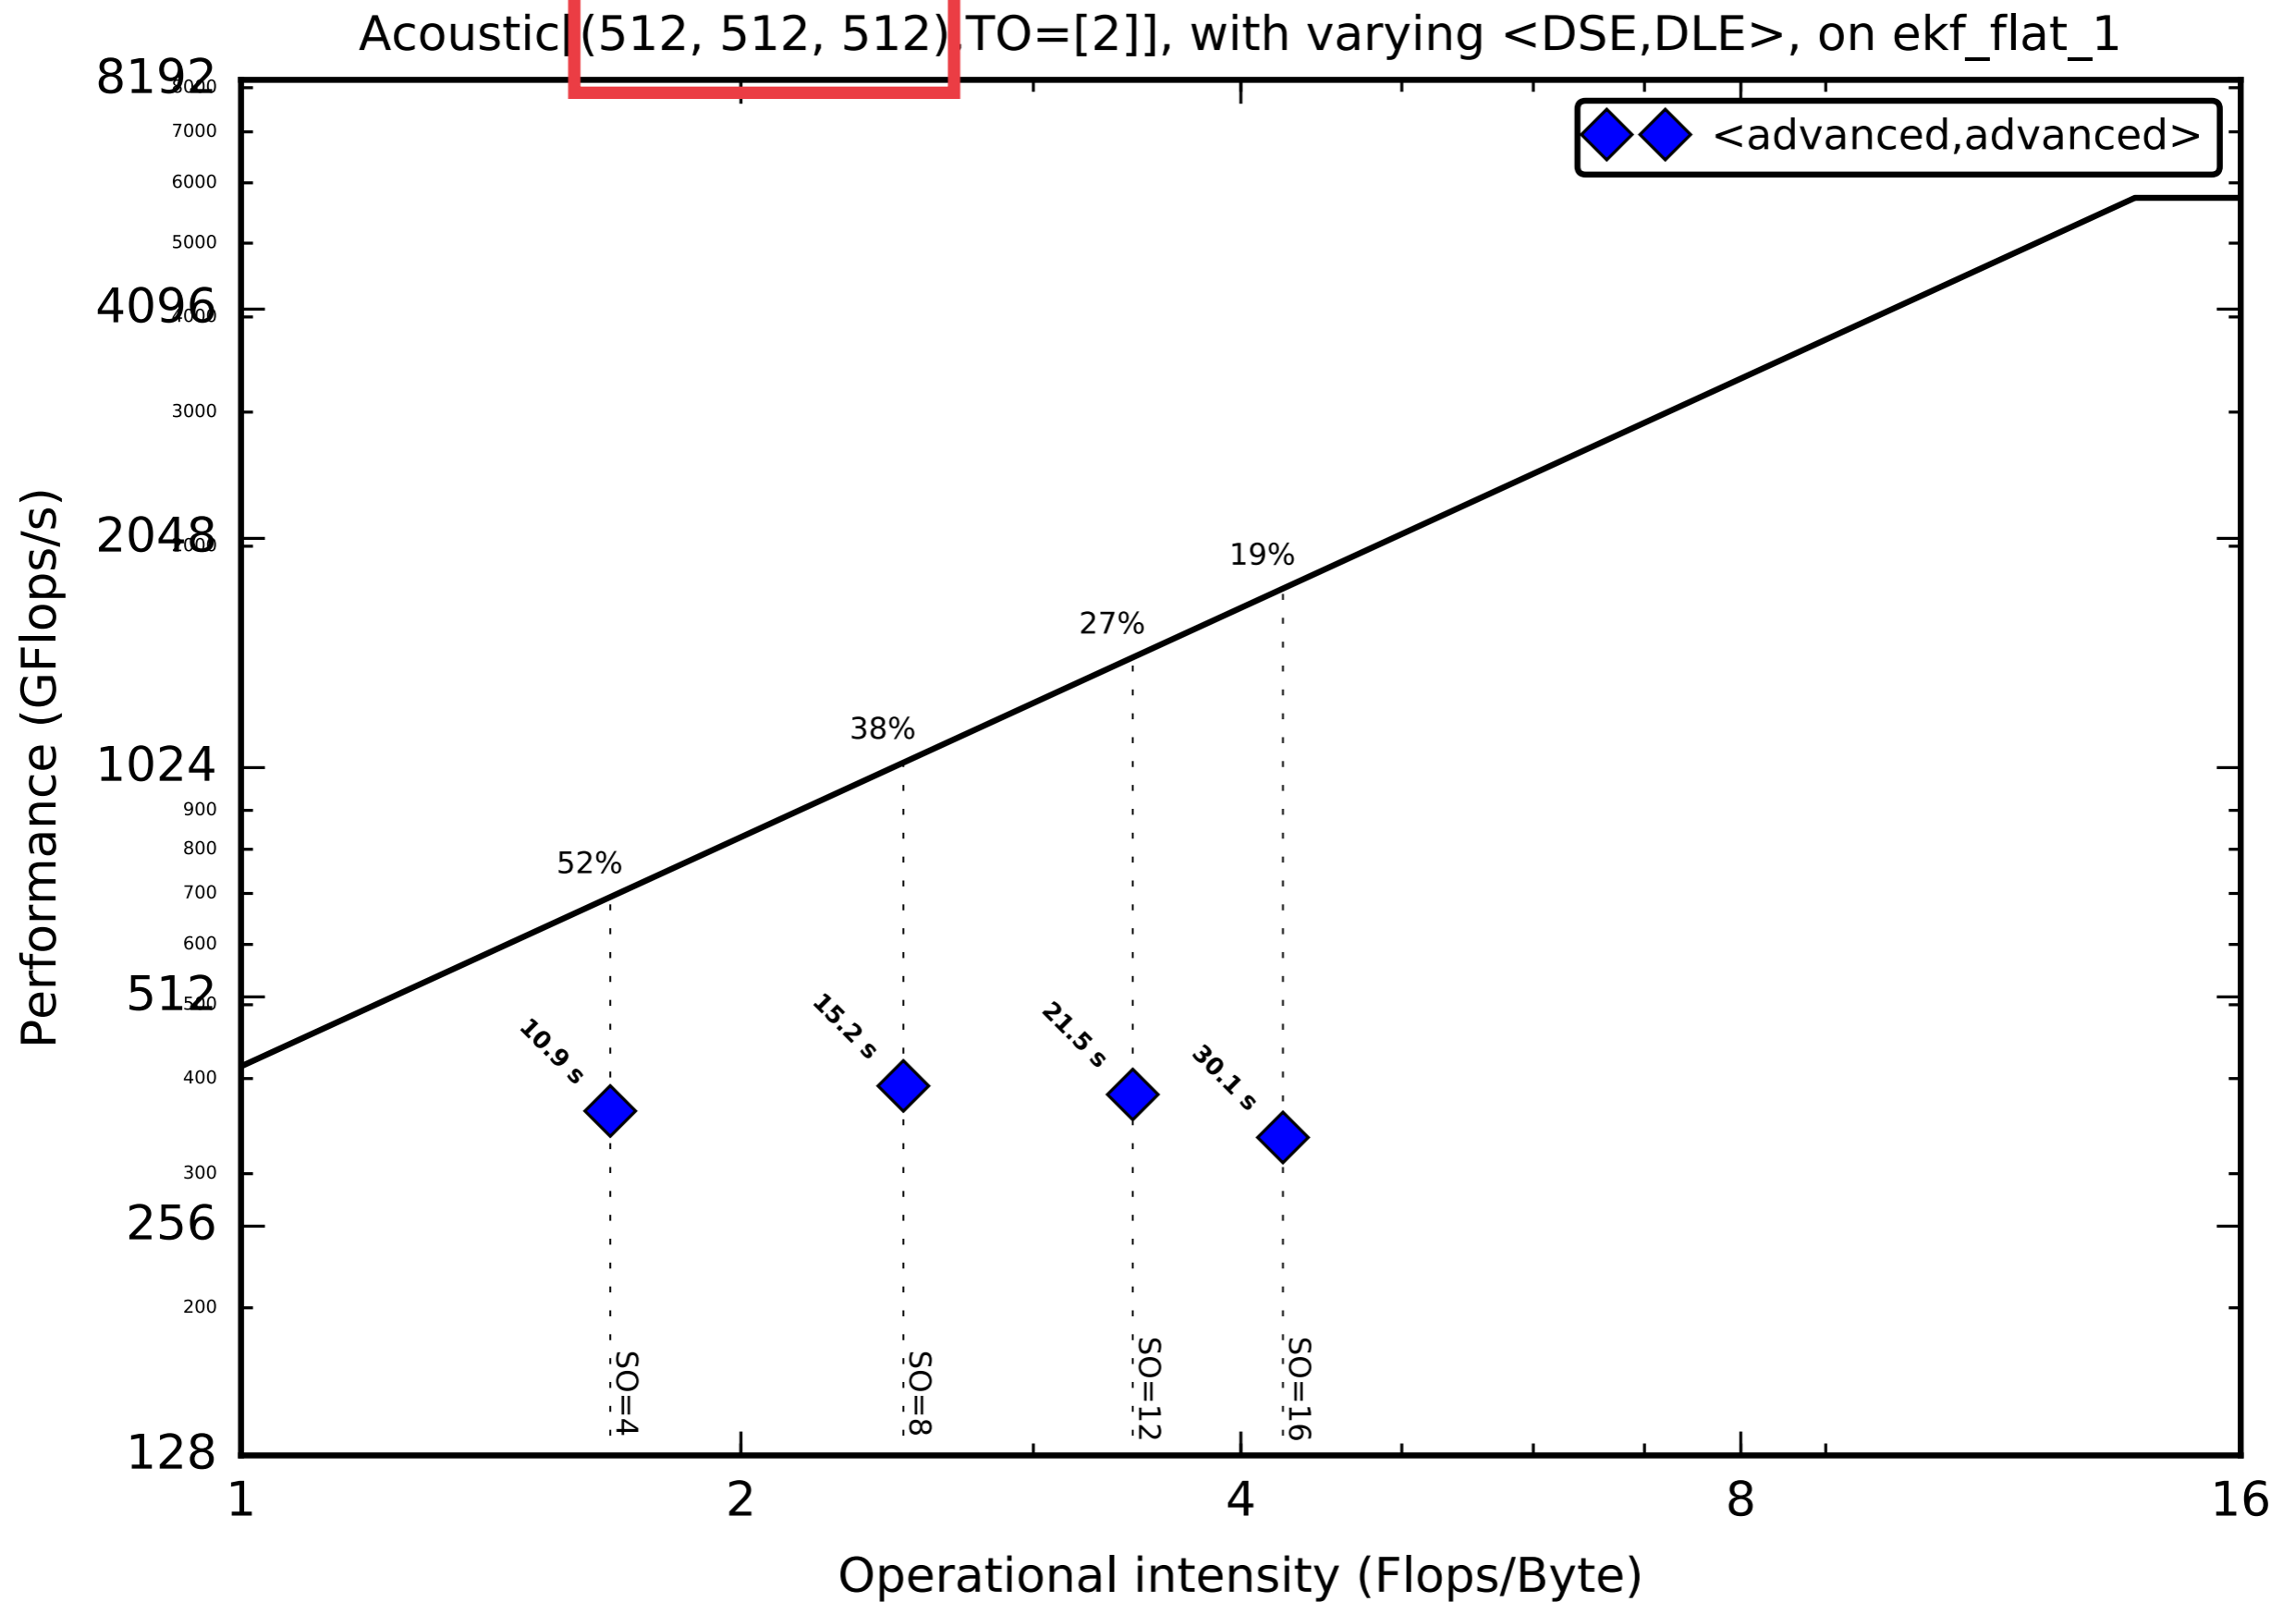
Elapsed Time [?]: 270.961s

Clockticks:	2,667,380,800,000	
Instructions Retired:	485,917,600,000	
CPI Rate [?] :	5.489	
MUX Reliability [?] :	1.000	
Front-End Bound [?] :	7.5%	of Pipeline Slots
Bad Speculation [?] :	3.1%	of Pipeline Slots
Back-End Bound [?] :	79.3%	of Pipeline Slots
Memory Latency:		
L1 Hit Rate [?] :	61.6%	
L2 Hit Rate [?] :	95.1%	
L2 Hit Bound [?] :	60.3%	of Clockticks
L2 Miss Bound [?] :	41.6%	of Clockticks
UTLB Overhead [?] :	0.3%	of Clockticks
SIMD Compute-to-L1 Access Ratio [?] :	0.964	
SIMD Compute-to-L2 Access Ratio [?] :	2.612	
Contested Accesses (Intra-Tile) [?] :	0.0%	
Page Walk [?] :	0.2%	of Clockticks
Memory Reissues:		
Split Loads [?] :	67.4%	
Split Stores [?] :	45.4%	
Loads Blocked by Store Forwarding [?] :	0.0%	
Retiring [?] :	10.1%	of Pipeline Slots

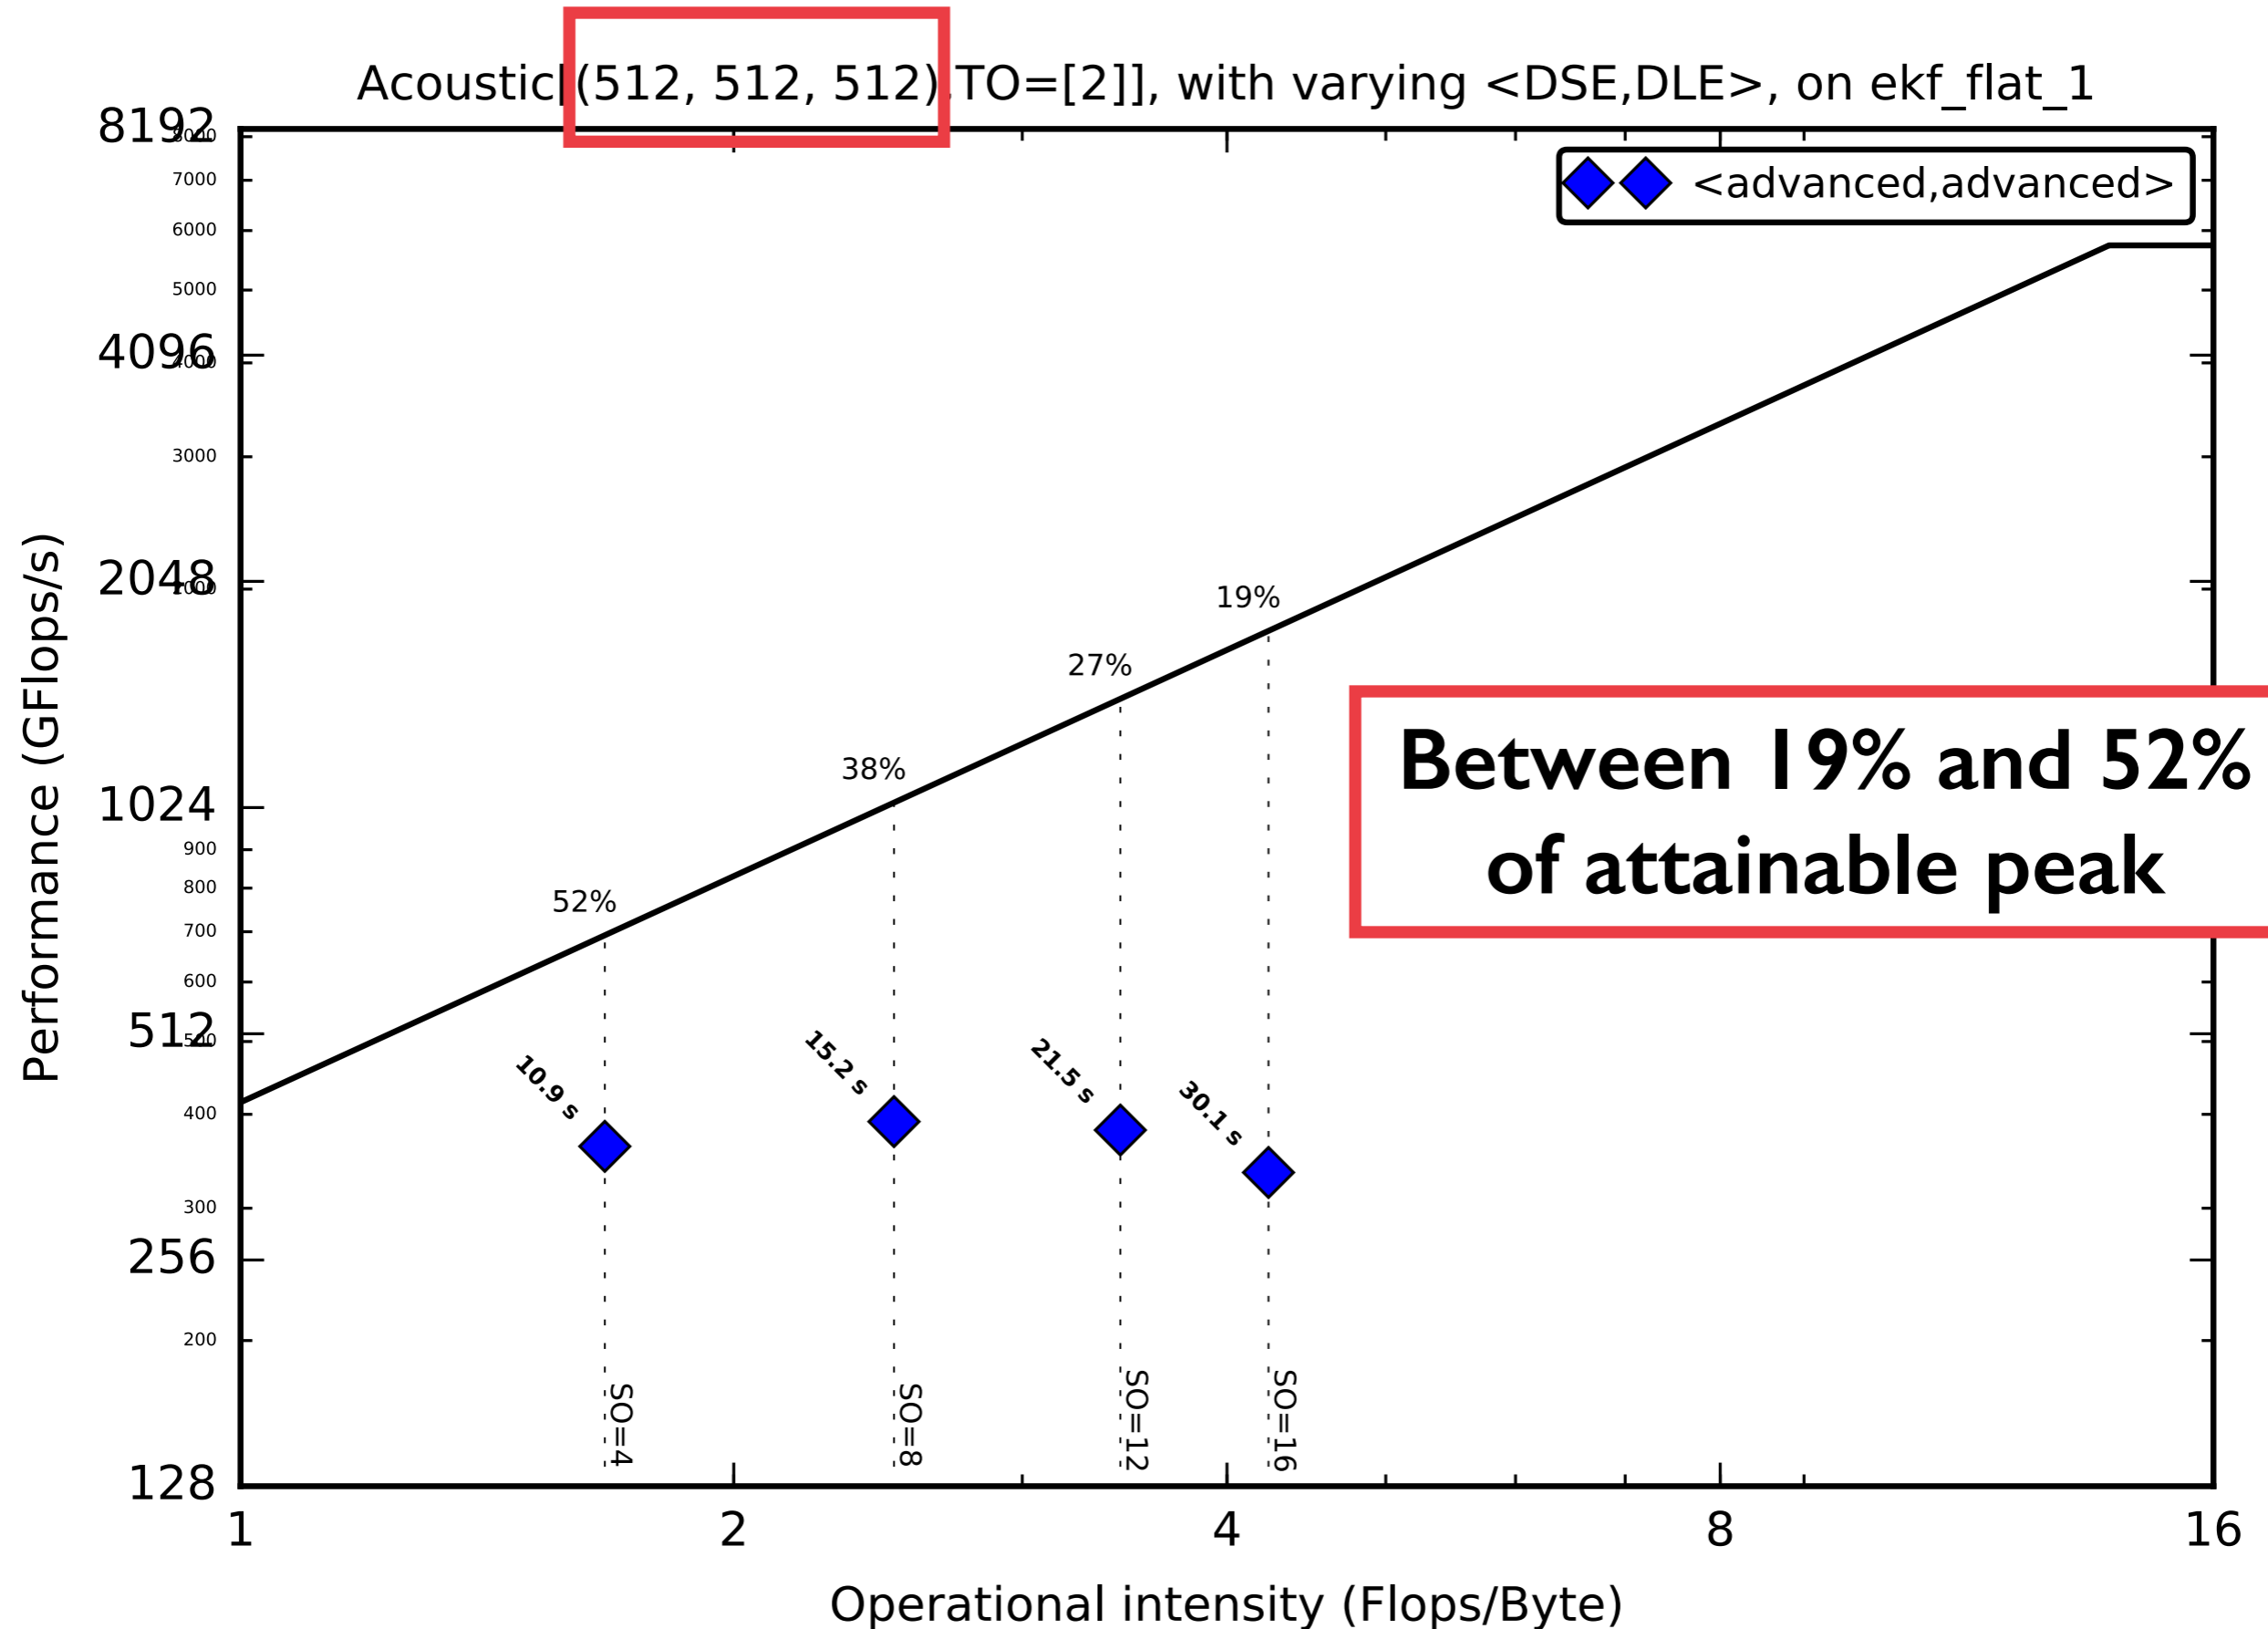
KNL 7650

TTI space order 4

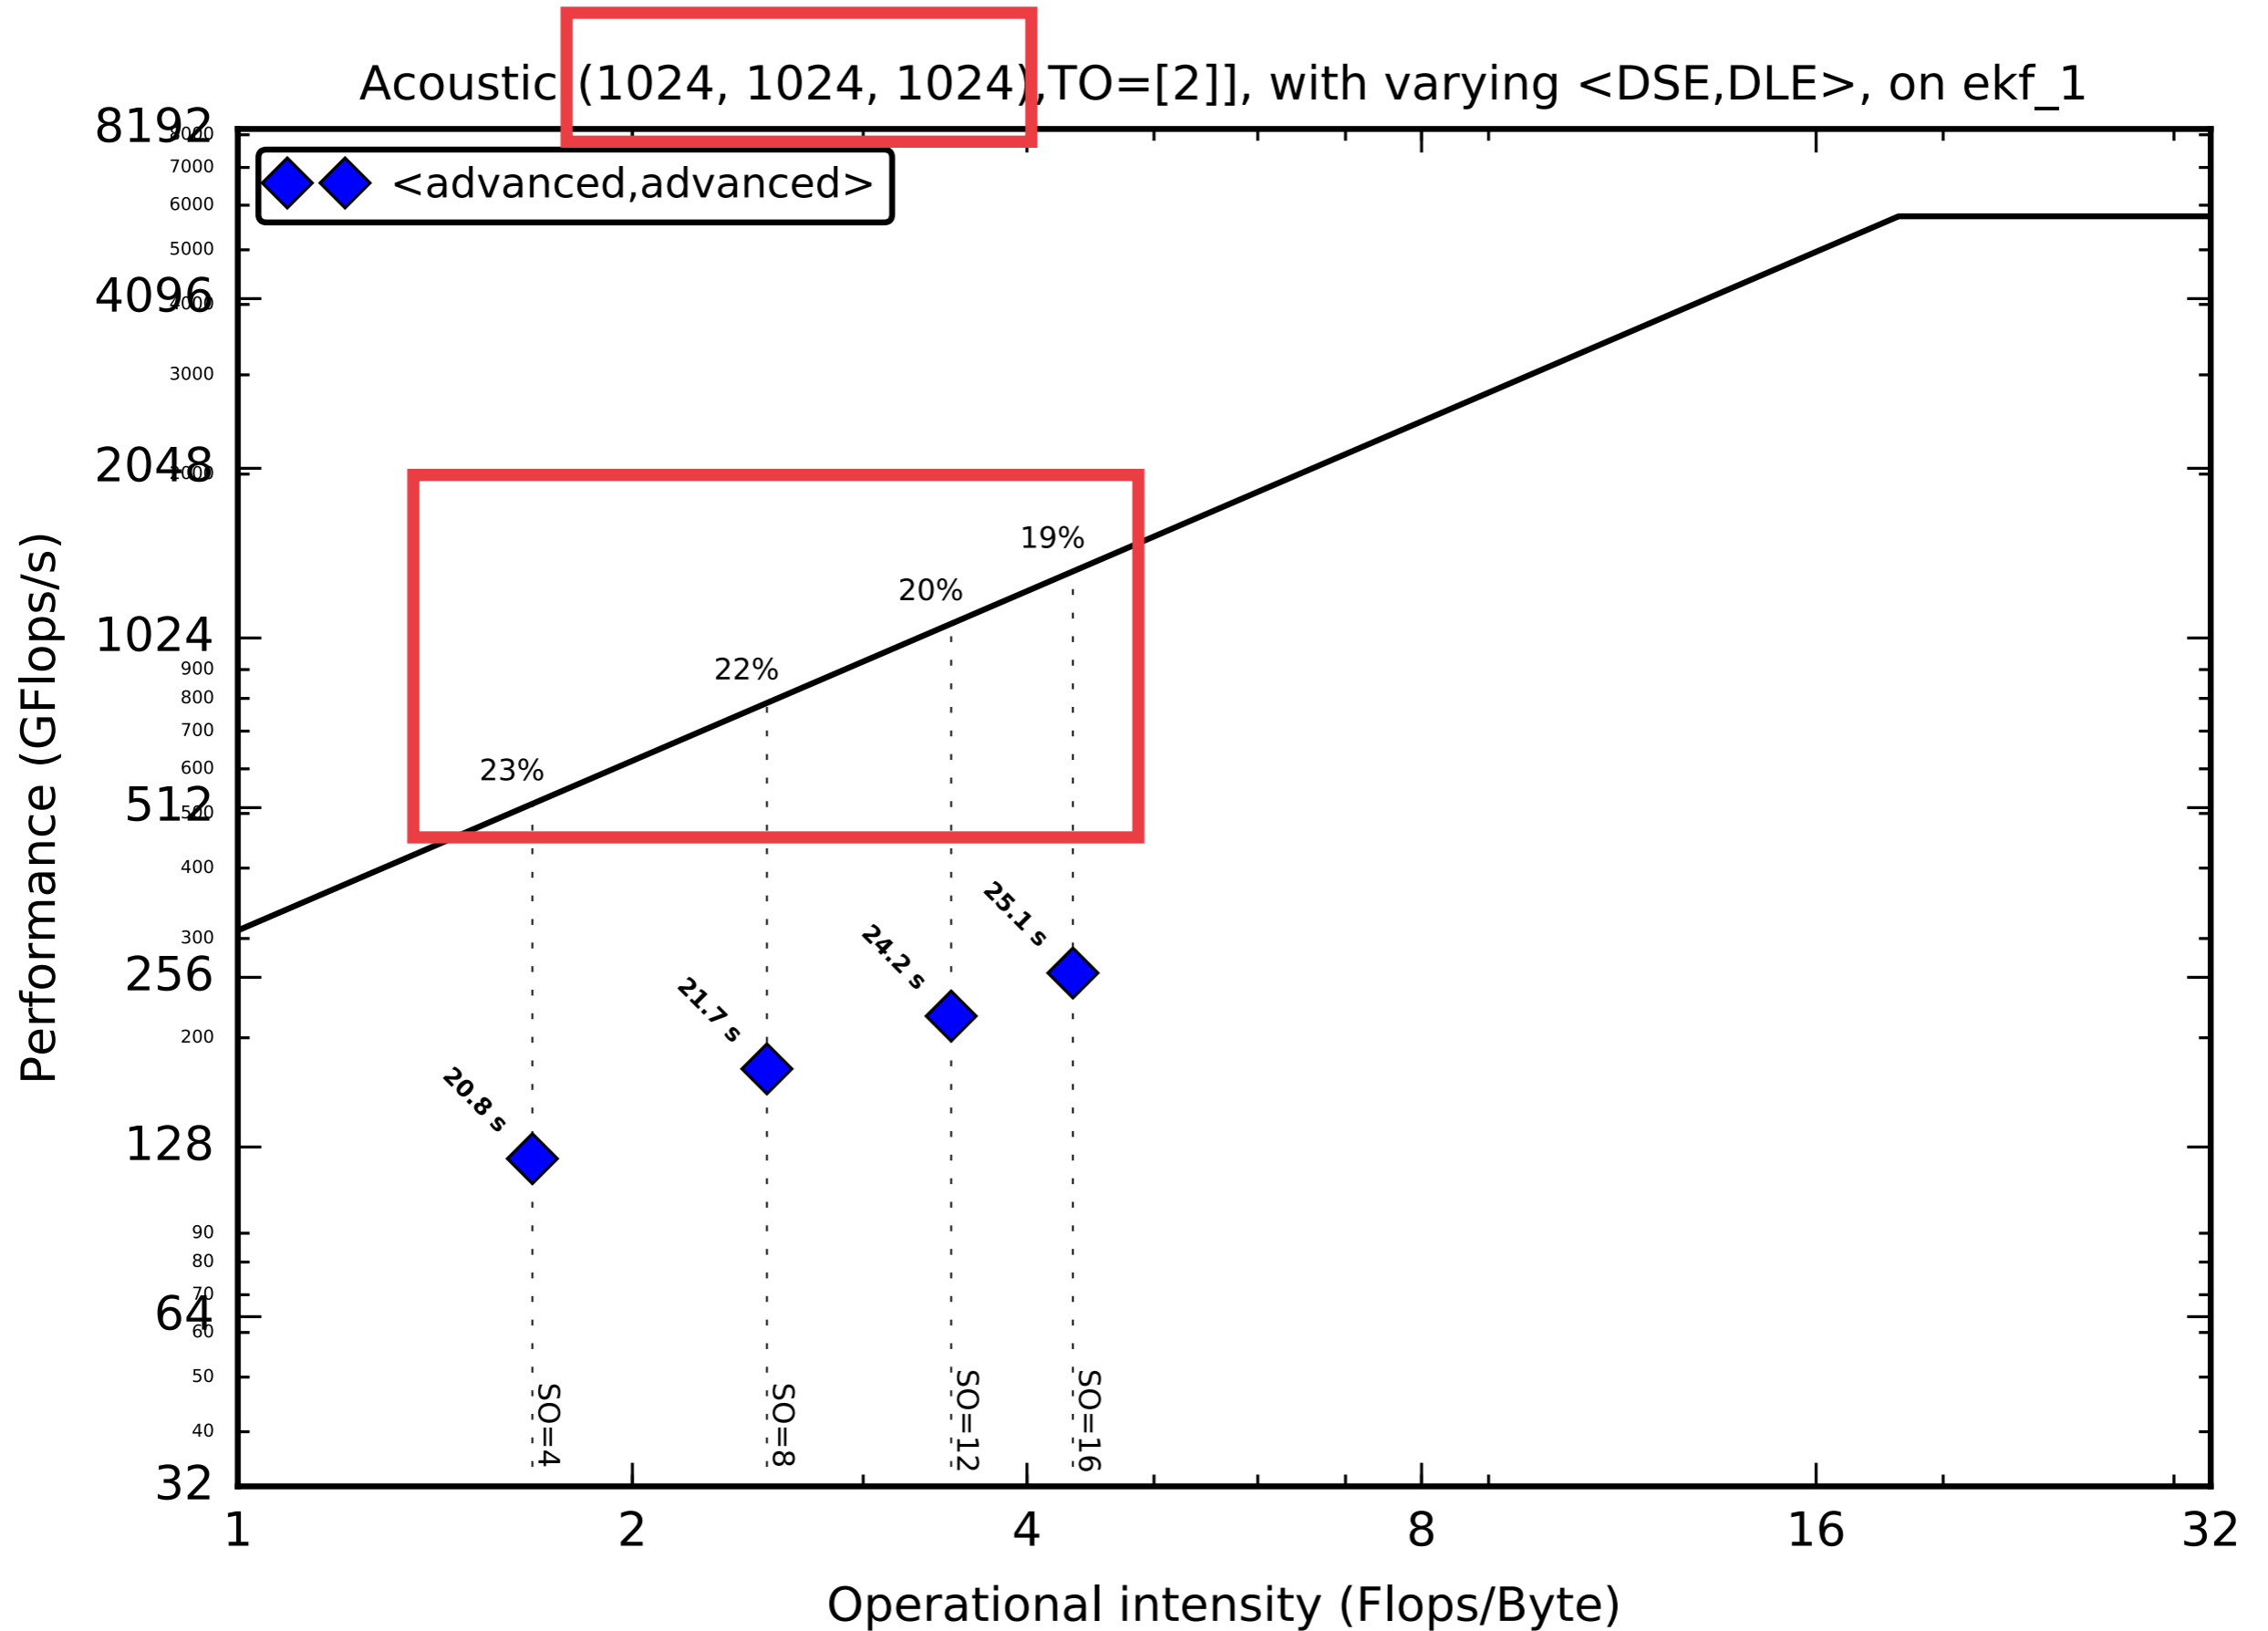
Acoustic on Xeon Phi (64 threads, in MCDRAM)



Acoustic on Xeon Phi (64 threads, in MCDRAM)



Acoustic on Xeon Phi (64 threads, needs DRAM)



Acoustic on Xeon Phi (64 threads, needs DRAM)

