

DaCe: Data Centric parallel programming for heterogeneous platforms

Tiziano De Matteis
Vrije Universiteit Amsterdam

Who am I?

Tiziano De Matteis

Assistant Professor at VU Amsterdam (joined March 2023)



Research Interests

I am interested in High-Performance Computing, Energy Efficiency, Systems.

Current focus: abstractions, methods, and systems to overcome the end of Moore's law.

Previously:

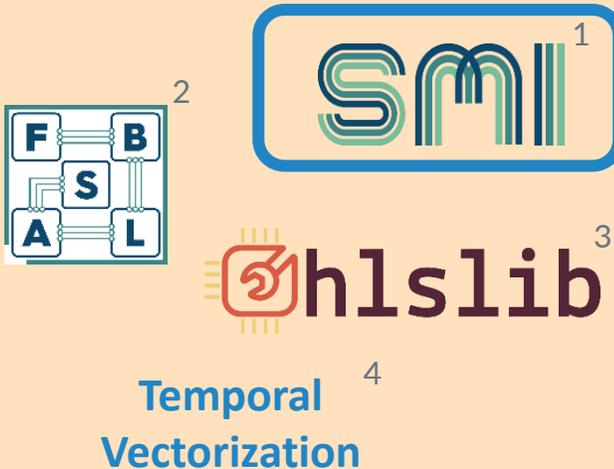
- ▷ PhD University of Pisa (on Parallel Streaming Analytics)
- ▷ PostDoc ETH Zurich (on FPGA for HPC)

Contributions to the FPGA-HPC community

Focus on methods and tools to productively program (with HLS) FPGAs for High-Performance Computing

Libraries/Tools

To increase programming productivity



Programming Model

To provide performance portability in Heter. Architectures (not only FPGA!)



Applications

To show the potentials of reconfigurable hardware in applications



¹ T. De Matteis, J. de Fine Licht, J. Beránek, T. Hoefler. “Streaming Message Interface: High-Performance Distributed Memory Programming on Reconfigurable Hardware”. SC’19

² T. De Matteis, J. de Fine Licht, T. Hoefler. “FBLAS: Streaming Linear Algebra on FPGA”. SC’20

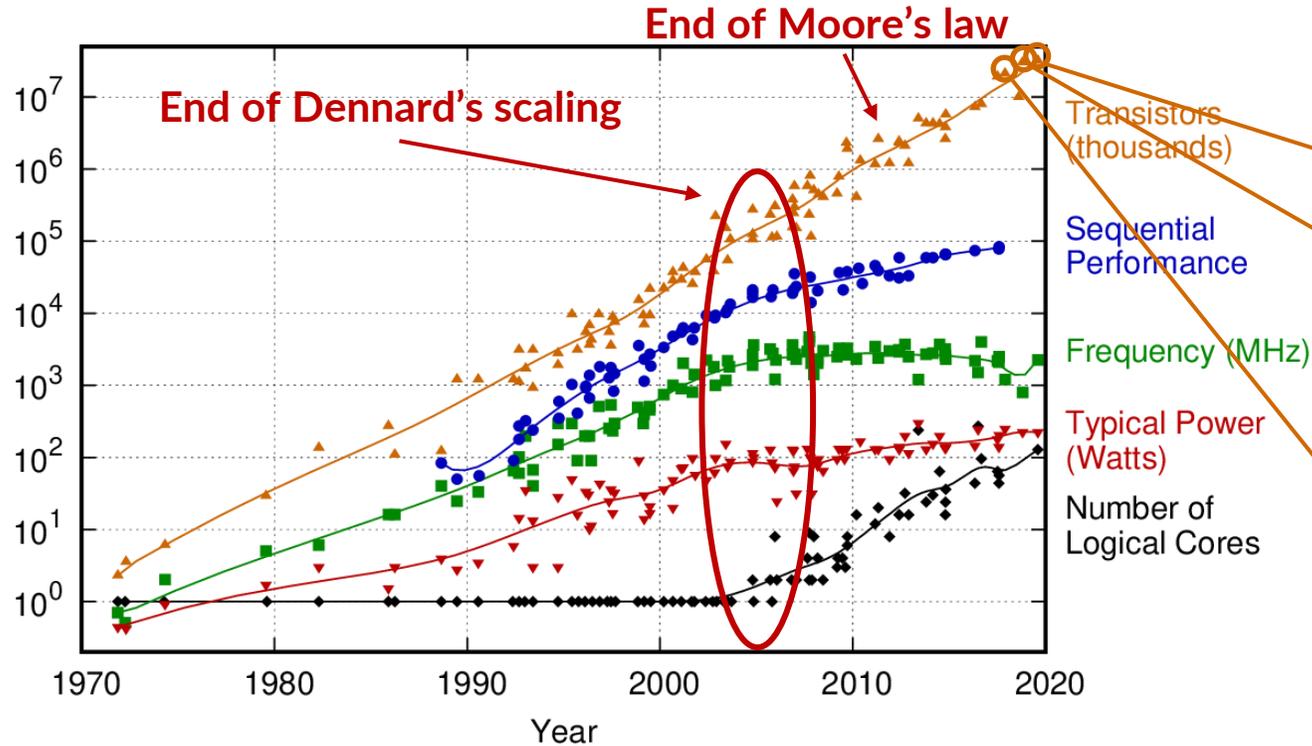
³ J. de Fine Licht, T. Hoefler. “hlslib: Software Engineering for Hardware Design”. H2RC’19

⁴ C. Johnsen, T. De Matteis, T. Ben-Nun, J. de Fine Licht, T. Hoefler de Fine Licht, T. Hoefler. “Temporal Vectorization: A Compiler Approach to Automatic Multi-Pumping”. ICCAD’22

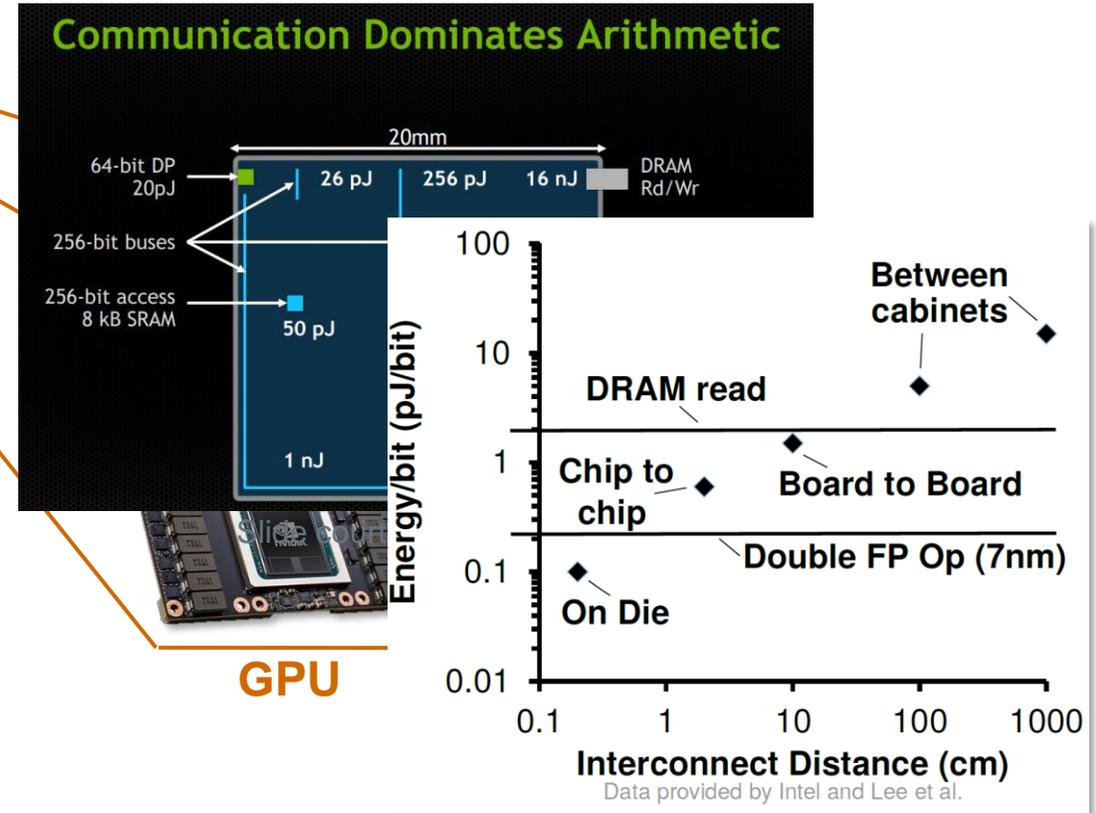
⁵ A. N. Ziogas, T. Schneider, T. Ben-Nun, A. Calotoiu, T. De Matteis, J. de Fine Licht, L. Lavarini, T. Hoefler. “Productivity, Portability, Performance: Data-Centric Python”. SC ‘21.

⁶ J. de Fine Licht, A. Kuster, T. De Matteis, T. Ben-Nun, D. Hofer, T. Hoefler. “StencilFlow: Mapping Large Stencil Programs to Distributed Spatial Computing Systems”. CGO ‘21

Changes in the Computing Landscape



Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
 New plot and data collected for 2010-2017 by K. Rupp



GPU

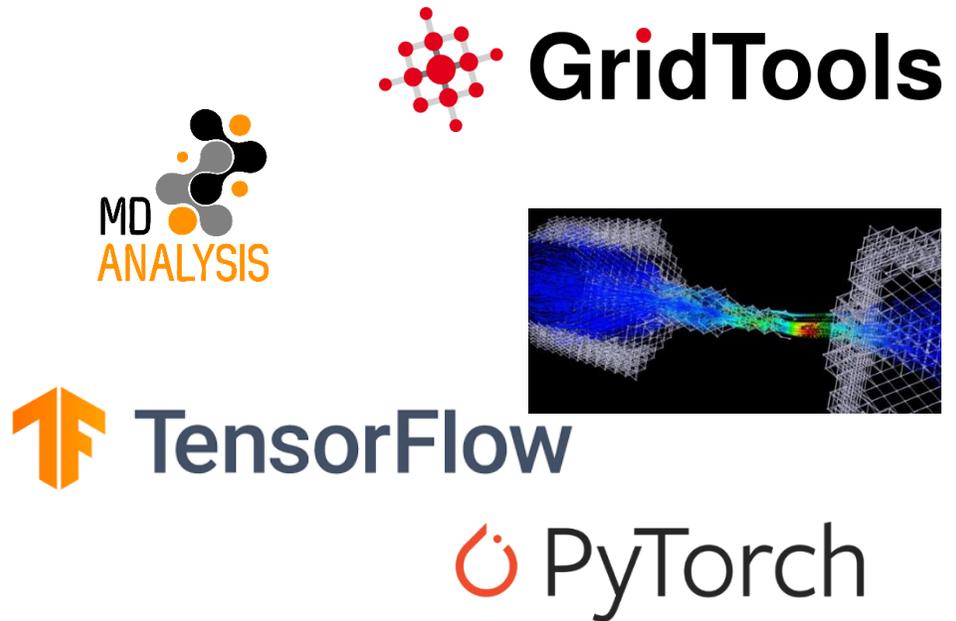
The holy grail of Performance
 Portability!

High-Performance Optimization -> Data
 Movement Reduction

Changes in the Software Landscape – The Scientific Languages



Vast software ecosystem



Productivity
Performance?
Portability?

DaCe Overview

Domain Scientist

Problem Formulation

$$\frac{\partial u}{\partial t} - \alpha \nabla^2 u = 0$$

Python

DSLs

PyTorch

C

...



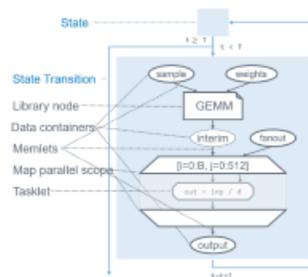
Scientific Frontend



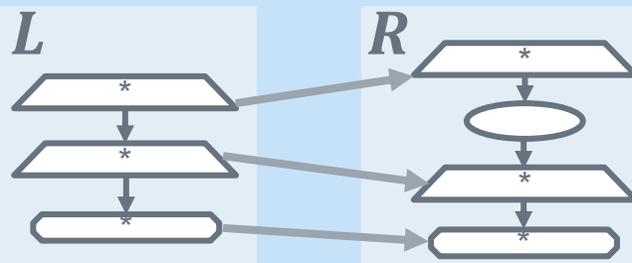
Performance Engineer

Data-Centric Intermediate Representation

1. Separate data containers from computation
2. Coarsening: multi-level view of data movement
3. Data movement as a first-class citizen (dependencies \rightarrow program order)
4. Control dependencies only when dataflow is not implied



Data-Centric Intermediate Representation (SDFG)



Graph Transformations



Transformed Dataflow



Performance Results



System

Hardware Information

Compiler

Runtime

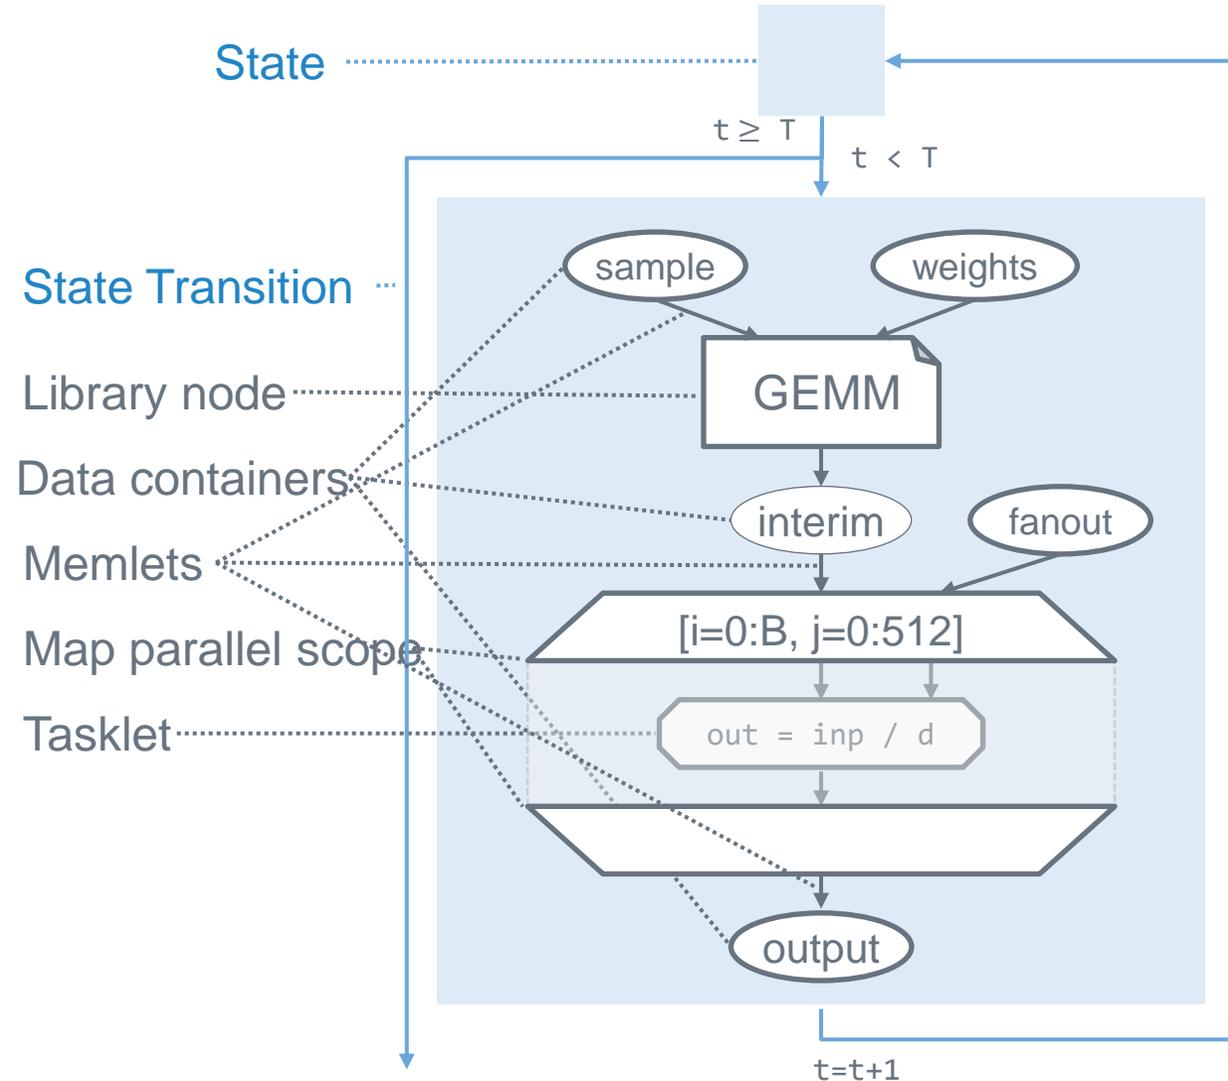
CPU Binary

GPU Binary

FPGA Modules

Data-Centric Intermediate Representation

1. Separate data containers from computation
2. Coarsening: multi-level view of data movement
3. Data movement as a first-class citizen (dependencies \rightarrow program order)
4. Control dependencies *only when dataflow is not implied*



From source code to SDFG



Code

```
glob_b = ...
class ClassA:
    def __init__(self, arr):
        self.q = arr

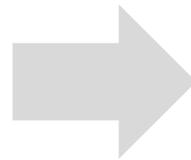
    @dace.method
    def __call__(self, a):
        return a * self.q + glob_b
```



From source code to SDFG



```
glob_b = ...  
class ClassA:  
    def __init__(self, arr):  
        self.q = arr  
  
    @dace.method  
    def __call__(self, a):  
        return a * self.q + glob_b
```



```
@dace.program  
def ClassA__call__(a, __g_self_q, __g_glob_b):  
    return a * __g_self_q + __g_glob_b
```

Closure
Resolution

Constant
Folding

Conditional/Dead
Code Elimination

Call Tree
Evaluation

Array
Consolidation

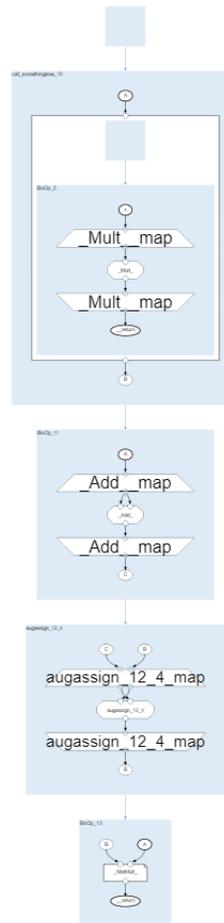


From source code to SDFG

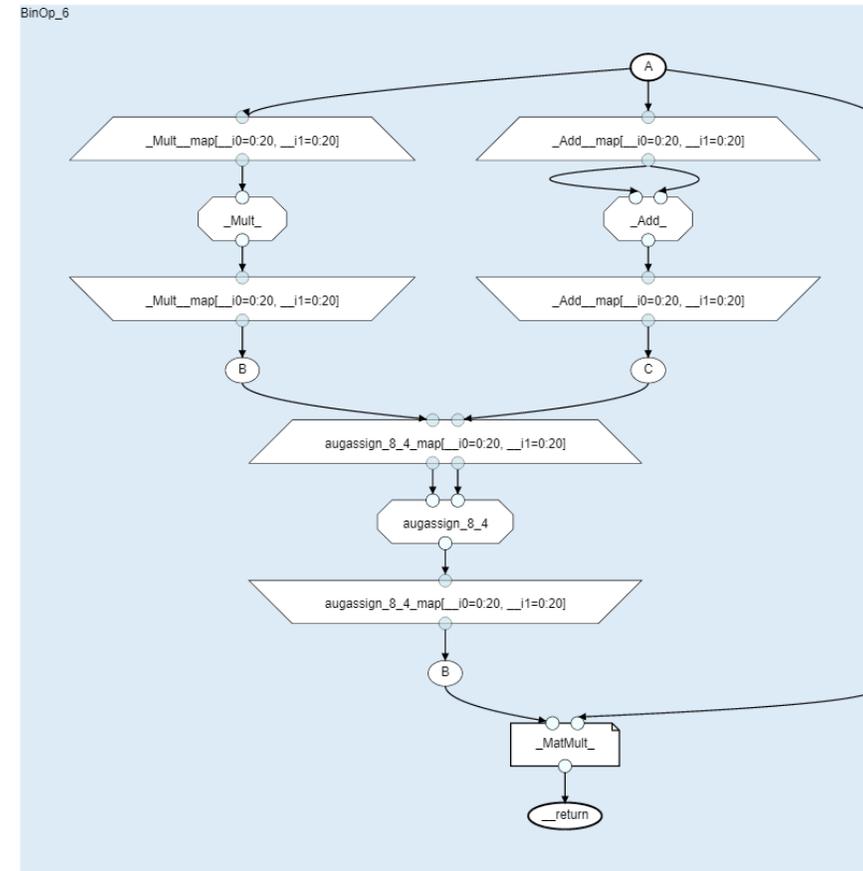
Code

Preprocess

Simplify



Control flow
coarsening passes



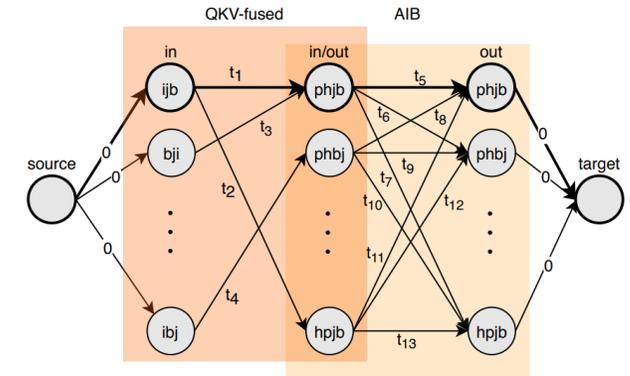
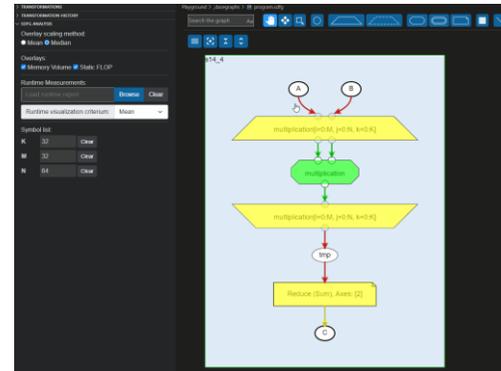
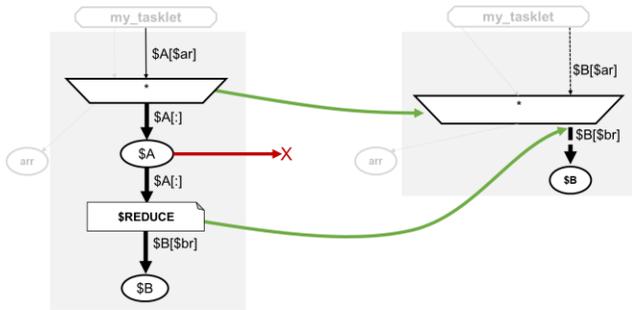
From source code to SDFG

Code

Preprocess

Simplify

Optimize



Graph Rewriting Transformations

Interactive Transformation and Instrumentation

Local and Global Tuning Interface

```
class ClassA:
    def __init__(self, arr):
        self.q = arr

    @dace.method(auto_optimize=True, device=dace.DeviceType.GPU)
    def __call__(self, a):
        return a * self.q + glob_b
```

Visual Studio Code Integration

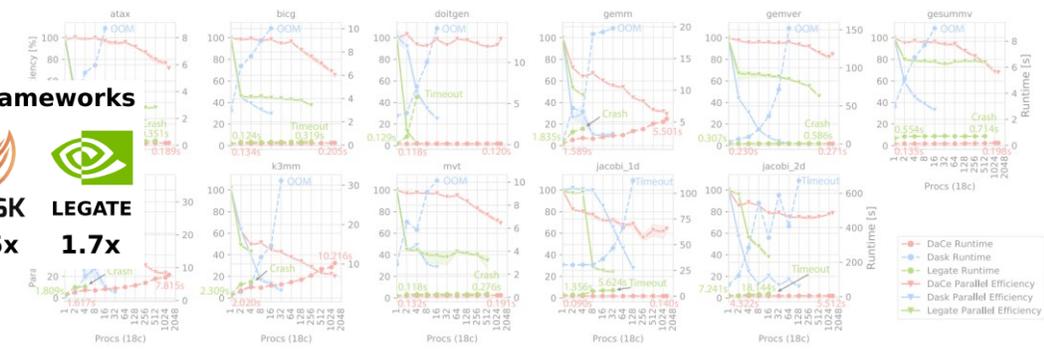
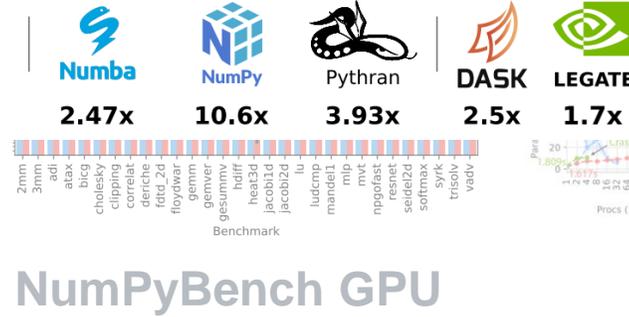
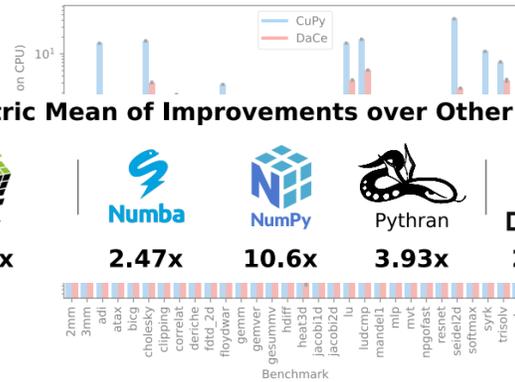
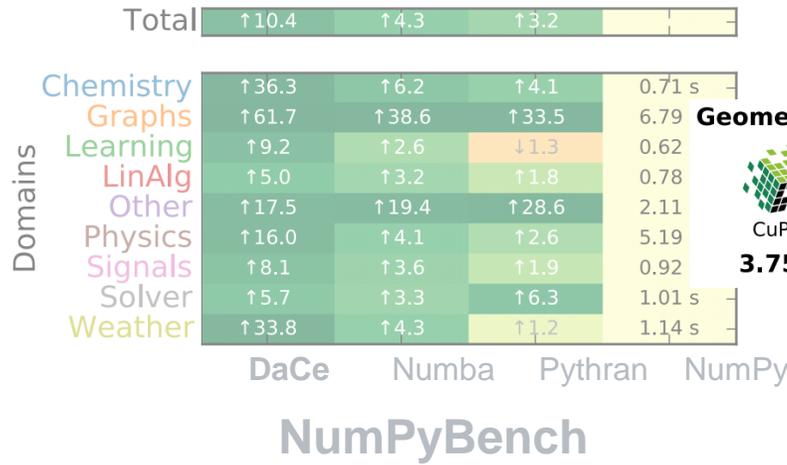
The screenshot displays the Visual Studio Code interface for Dace SDFG optimization. It is divided into three main sections:

- Left Sidebar (TRANSFORMATIONS):** Lists various optimization techniques such as Selection, Viewport, FPGATransformState, MapTiling, StripMining, MapDimShuffle, ReduceExpansion, MapReduceFusion, GPUTransformMap, GPUTransformLocalStorage, MapTilingWithOverlap, MapExpansion, OuterProductOperation, and Global. It also includes a TRANSFORMATION HISTORY section showing no previously applied transformations.
- Central Code Editor (gemm.py):** Contains the Python code for a Gemv (General Matrix-Vector) multiplication. The code defines a function `gemm` that takes matrices `A`, `B`, and `C` as input and performs a dot product using Dace's `map` and `sum` operations. The code includes comments and variable declarations for `M`, `N`, and `K`.
- Right Sidebar (SDFG gemm):** Shows the generated SDFG (Static Data Flow Graph) for the `gemm` function. The graph is a vertical flow starting with inputs `A` and `B`, followed by a `multiplication` node, a `tmp` node, and a `Reduce (Sum), Axes: [2]` node, leading to the output `C`. The right sidebar also includes configuration options for the SDFG, such as `arg_names`, `constants_prop`, `exit_code`, `global_code`, `init_code`, `instrument`, `openmp_sections`, and `symbols`.

DaCe Performance



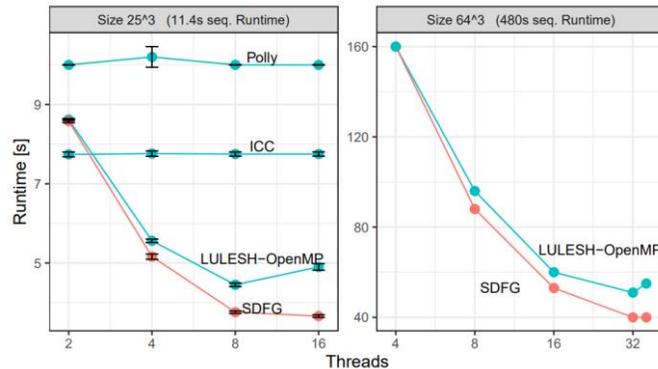
<https://github.com/spcl/npbench>



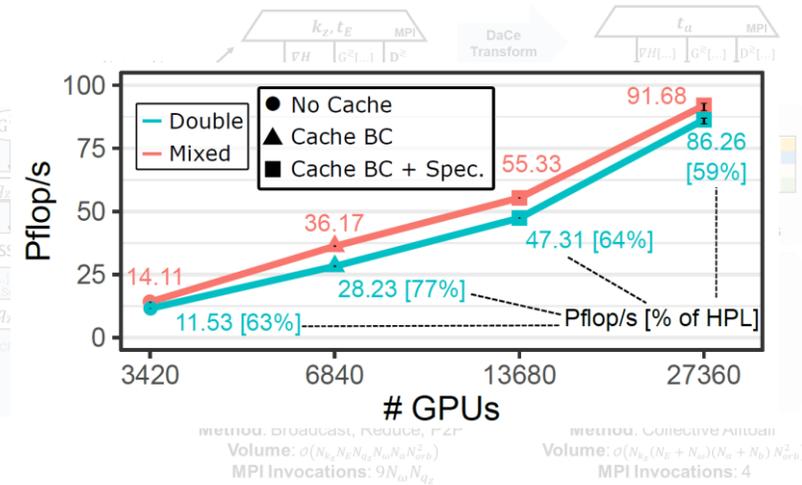
NumPyBench

NumPyBench GPU

NumPyBench Distributed



Unstructured Hydrodynamics (LULESH)



Quantum Transport Simulation

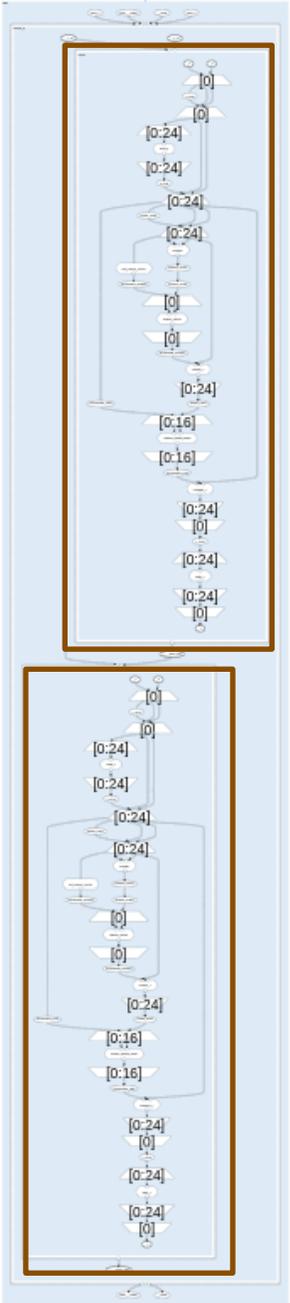
A. N. Ziogas, T. Schneider, T. Ben-Nun, A. Calotoiu, T. De Matteis, J. de Fine Licht, L. Lavarini, T. Hoefler. "Productivity, Portability, Performance: Data-Centric Python". SC '21.
 A. Calotoiu, T. Ben-Nun, G. Kwasniewski, J. de Fine Licht, T. Schneider, P. Schaad, T. Hoefler. "Lifting C semantics for Dataflow Optimization". ICS '22.
 A. Ziogas, et al. A data-centric approach to extreme-scale ab initio dissipative quantum transport simulations". SC'19, **Gordon Bell Award**

DaCe and FPGA

Let's use DaCe to compute $y = A^T Ax$ (ATAx)

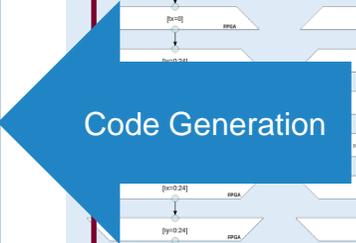
```
import dace

M, N = 24, 24
@dace.program
def atax(A: dace.float32[M, N],
        x:dace.float32[N]):
    return (A @ x)@A
```



```
void _MatMult_gemv_sdfg_1_0_3(const float* _A, const float* _x, float* _y,
int M, int N) {
    #pragma HLS INLINE
    for (int _o0 = 0; _o0 < M; _o0 += 1) {
        #pragma HLS PIPELINE II=1
        #pragma HLS LOOP_FLATTEN
        float out = 0
        ...
    }
}
```

//450+ lines of code + Make files

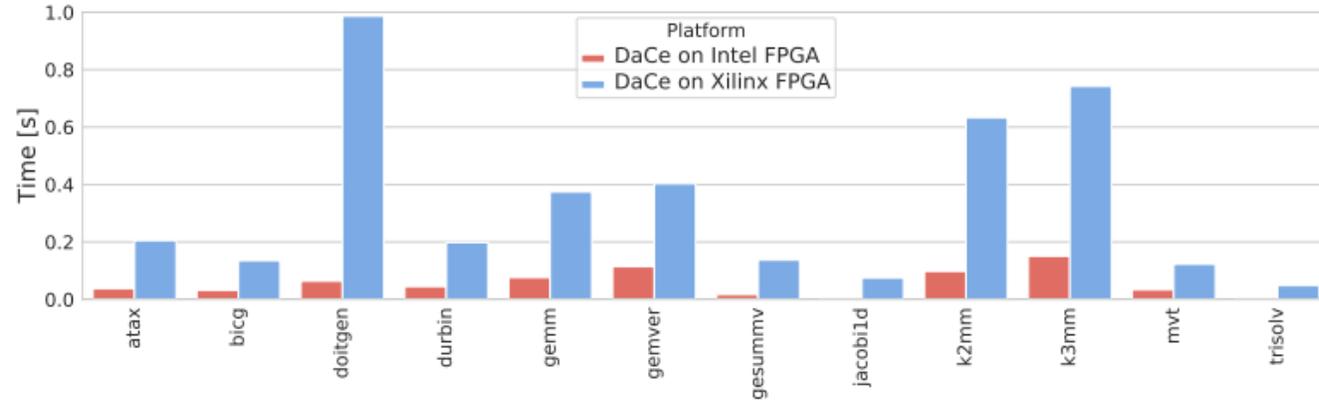


DaCe and FPGA

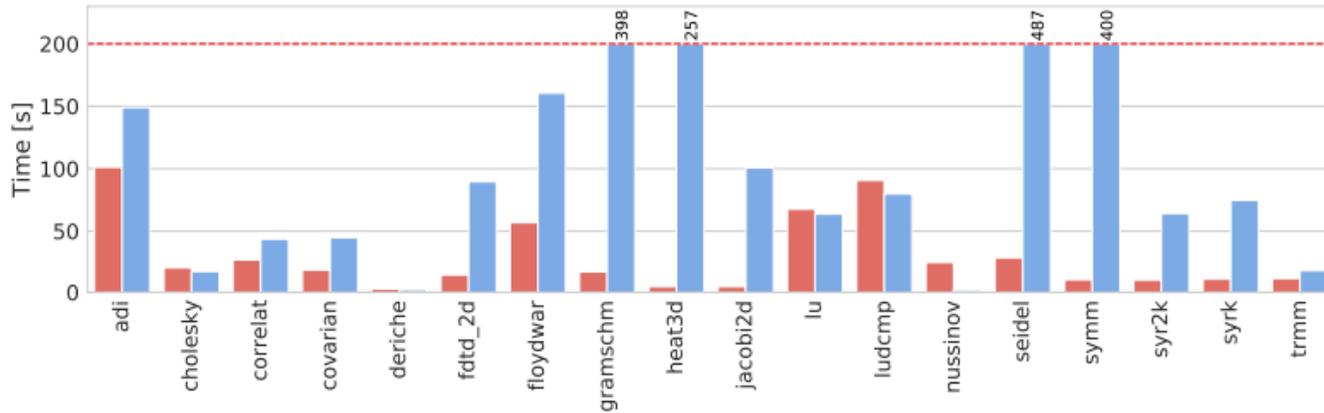
Let's use DaCe to compute $y = A^T Ax$ (ATAx)

```
import dace

M, N = 24, 24
@dace.program
def atax(A: dace.float32[M, N],
        x: dace.float32[N, 1])
    return (A @ x)@A
```



```
void _MatMult_gemv_sdfg_1_0_3(const float*_A, const float*_x, float* out,
int M, int N) {
    #pragma HLS INLINE
    for (int _o0 = 0; _o0 < M; _o0 += 1)
        #pragma HLS PIPELINE II=1
        #pragma HLS LOOP_FLATTEN
        float out = 0
    ...
}
```



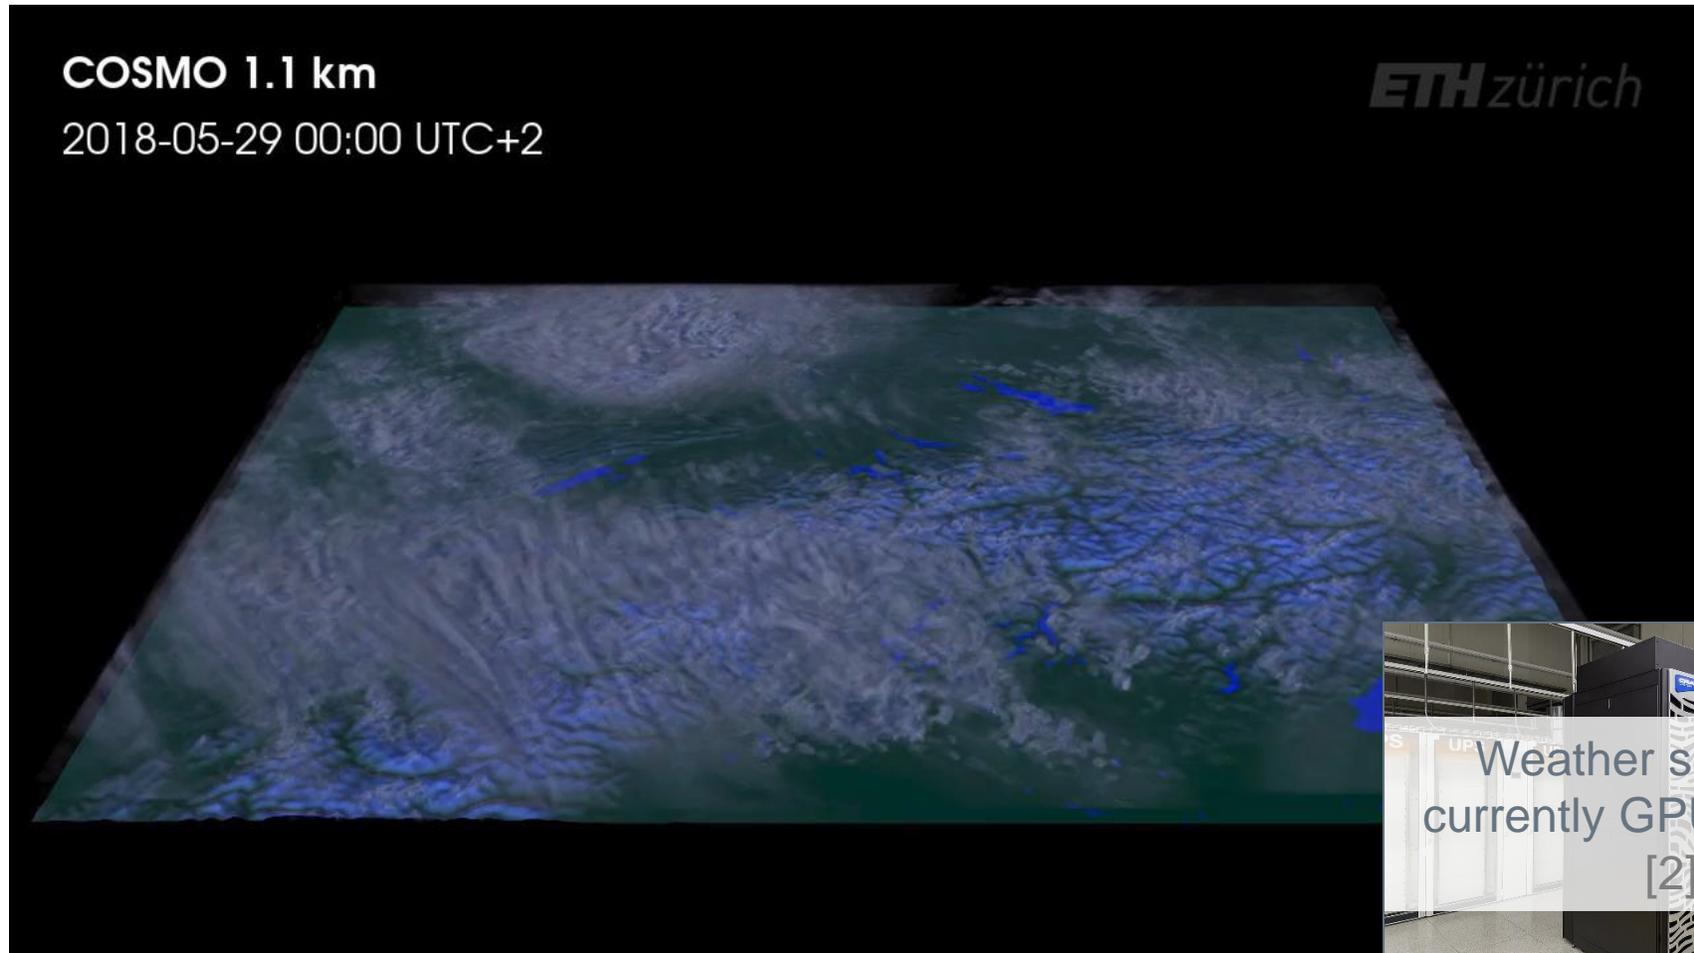
Polybench kernels code generated from Numpy for both Xilinx and Intel



ng
y

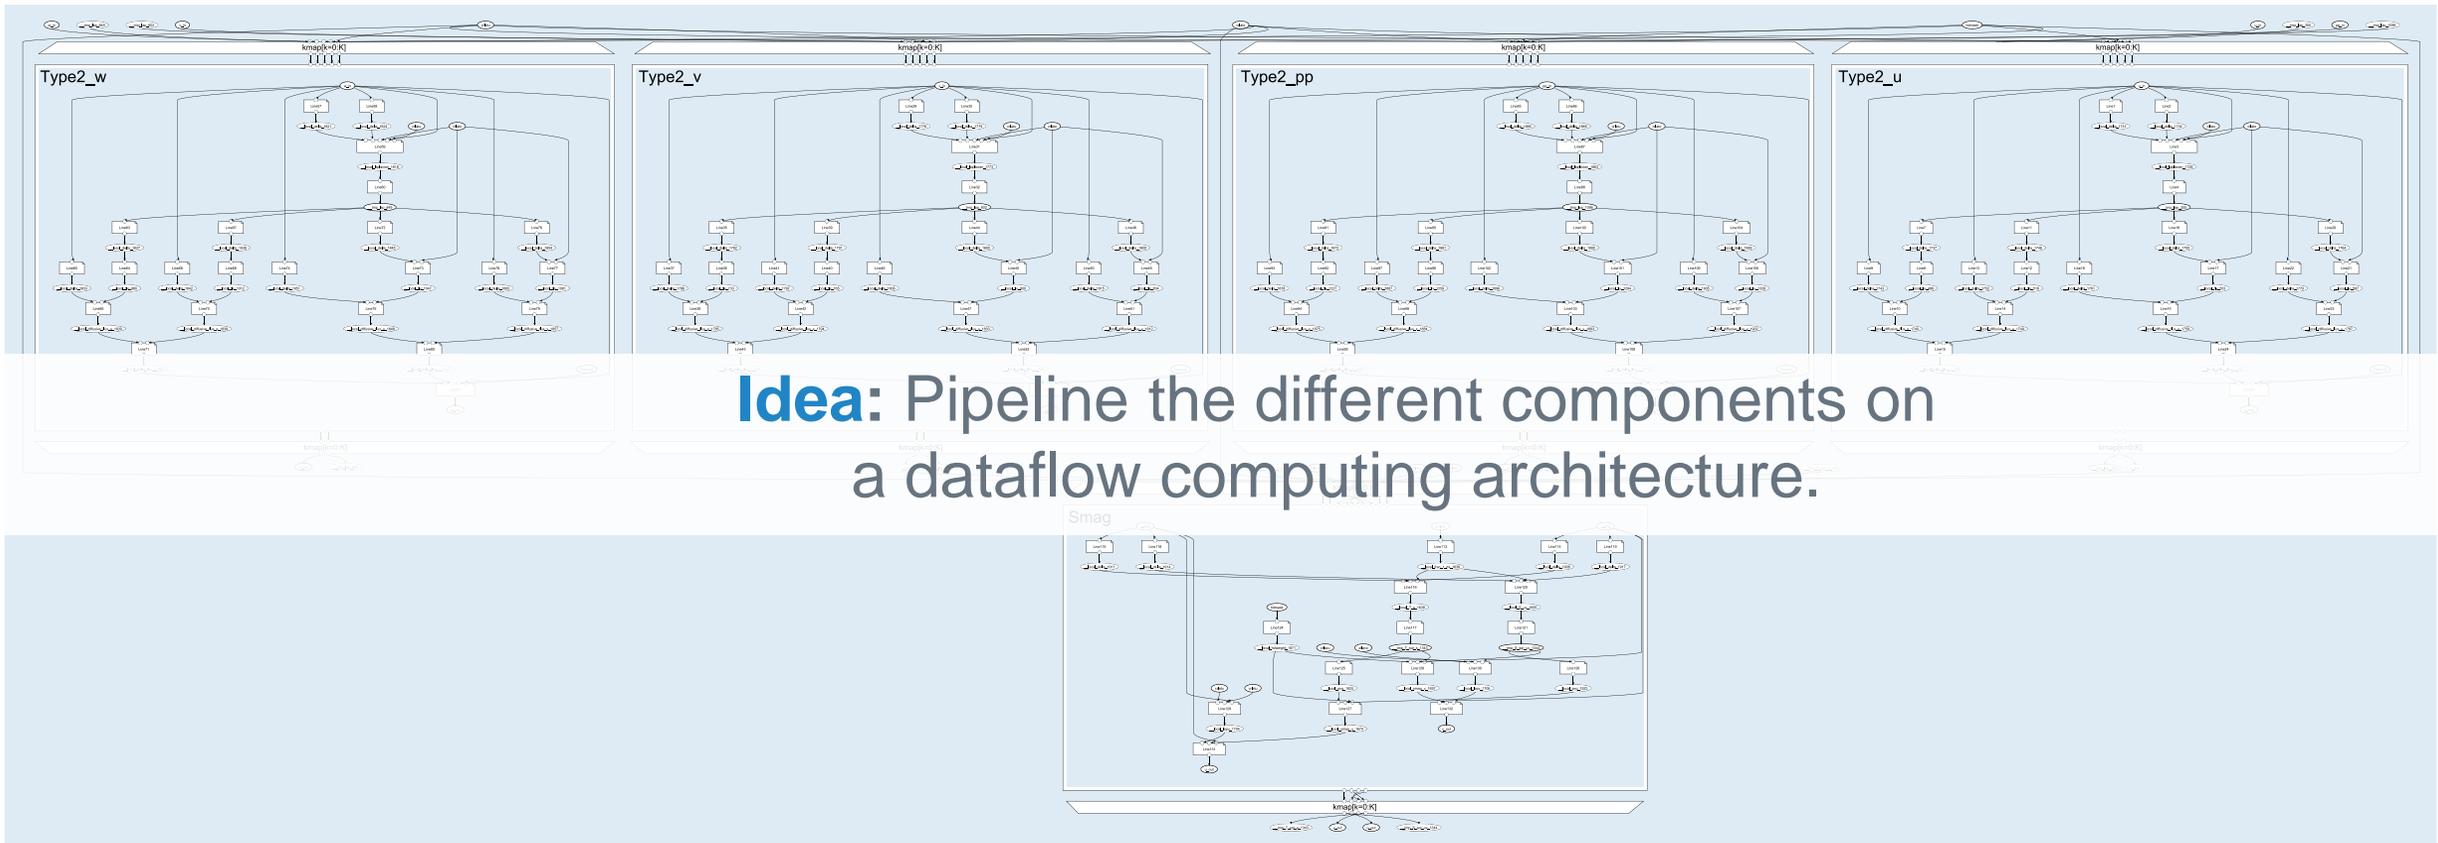
Weather simulation at MeteoSwiss

[1]



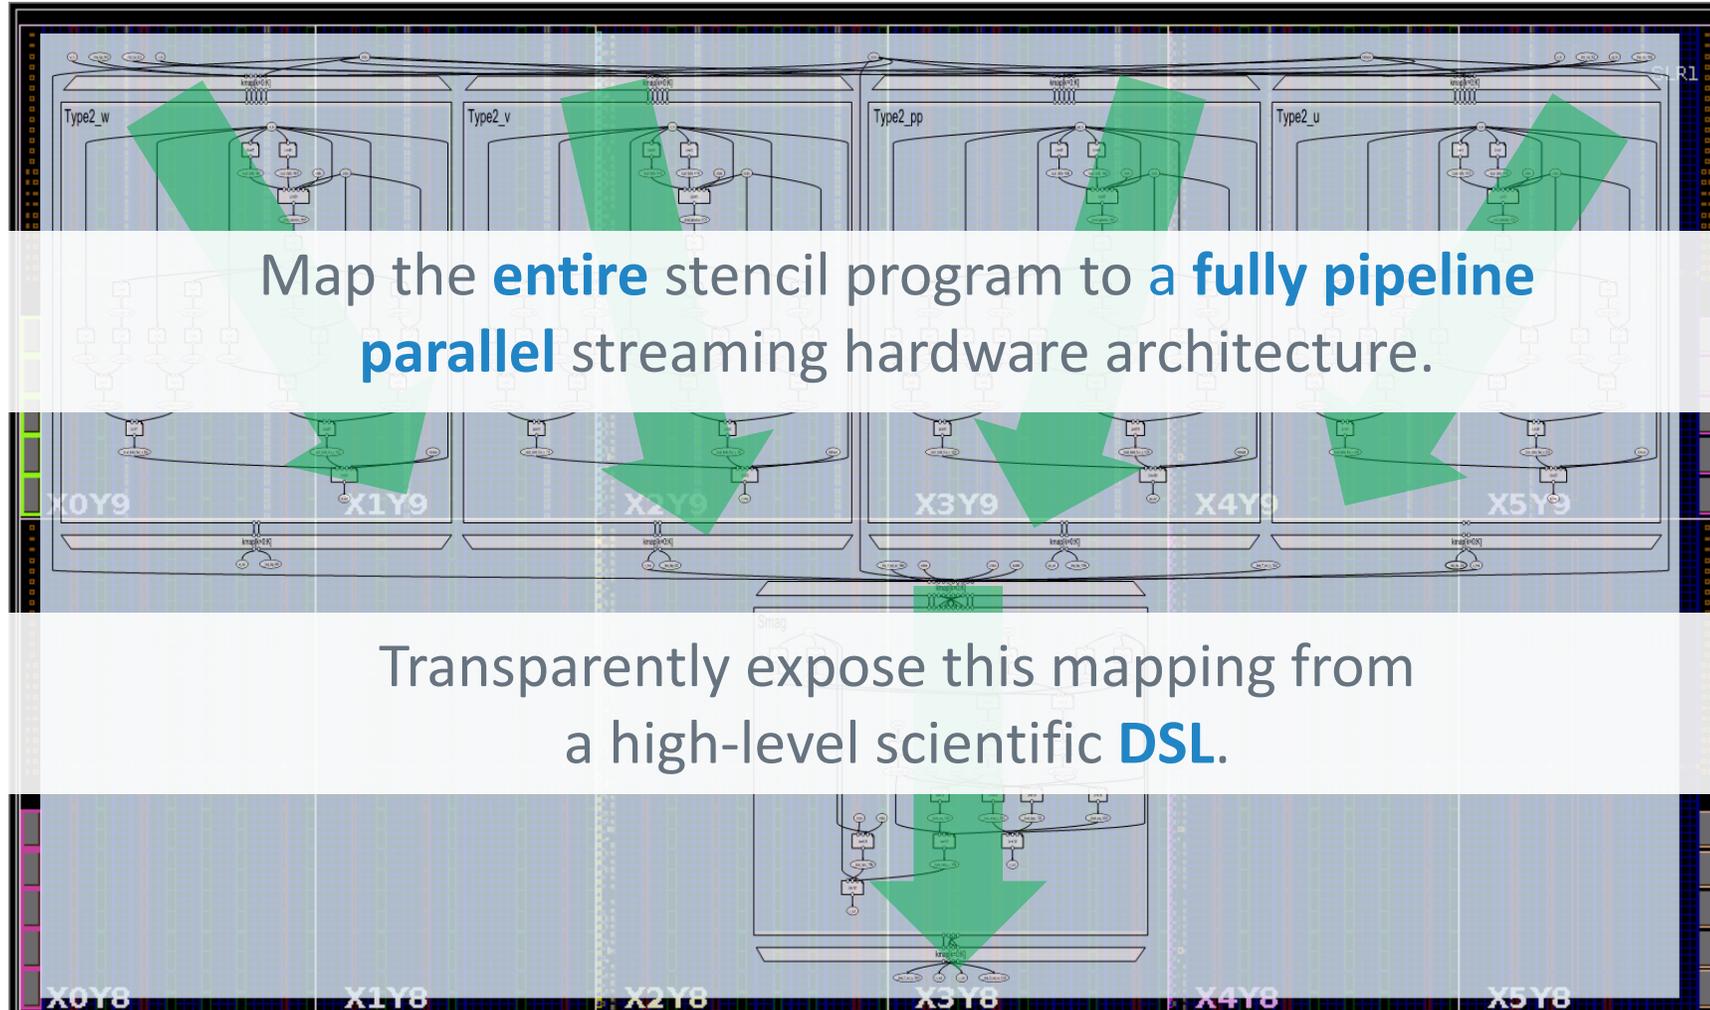
[1] Institute for Atmospheric and Climate Science and Computer Graphics Laboratory, ETH Zürich [<https://vimeo.com/389292423>]
[2] Swiss National Supercomputing Center (CSCS) [<https://www.cscs.ch/computers/arolla-tsa-meteoswiss>]

Weather programs

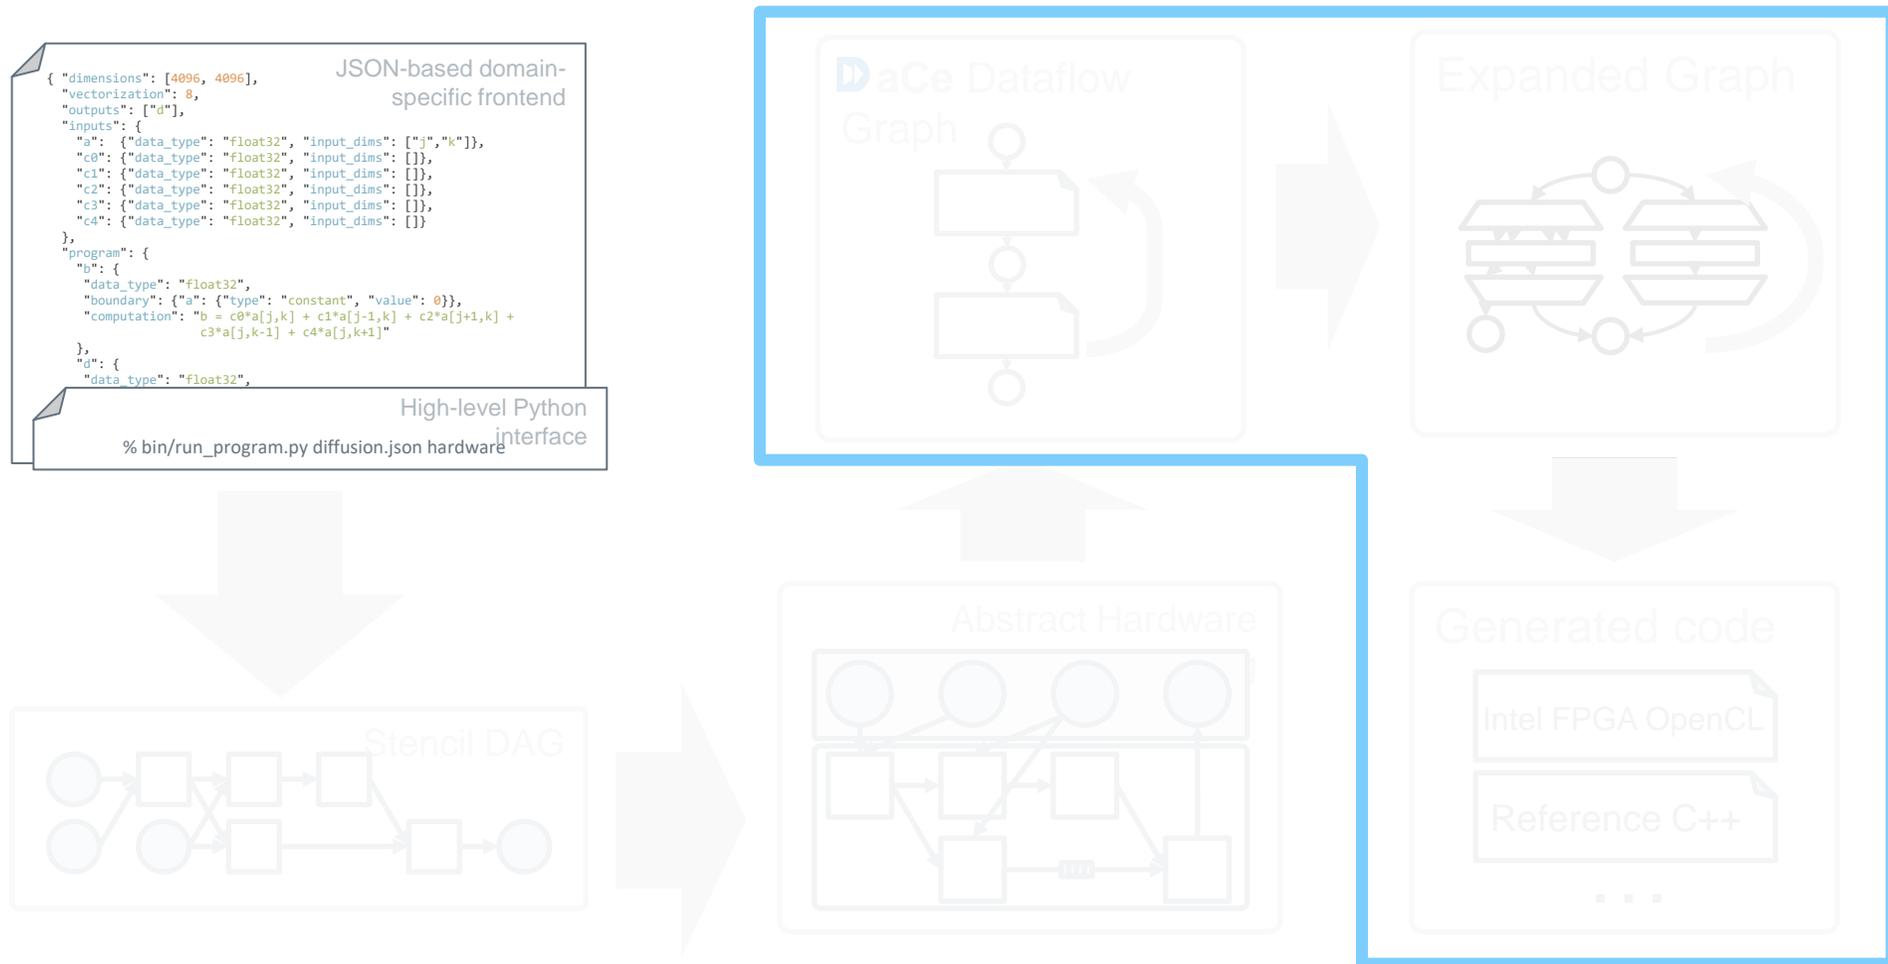


Idea: Pipeline the different components on a dataflow computing architecture.

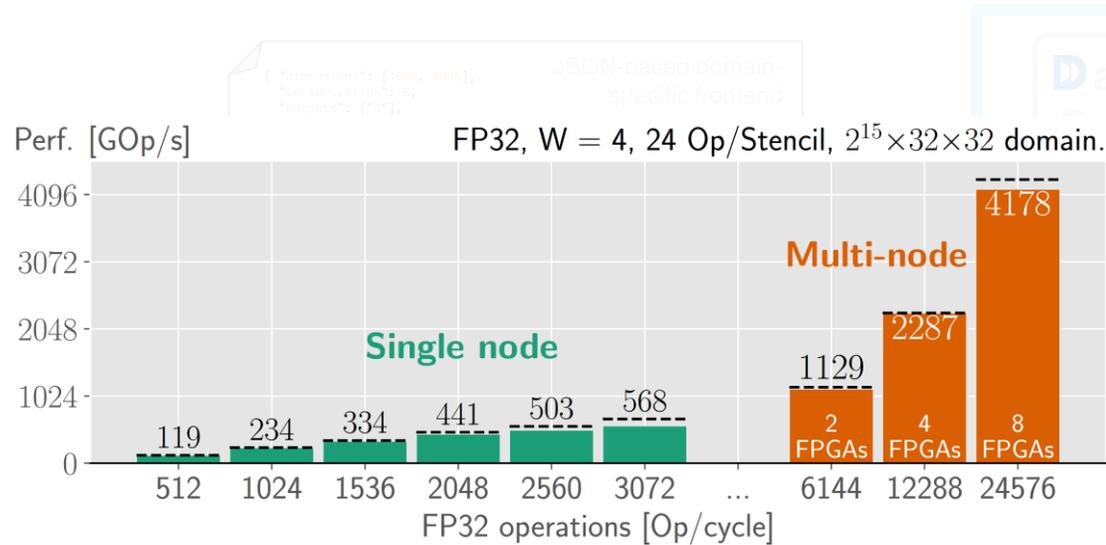
StencilFlow – FPGAs for Climate Modeling



Stencilflow



Stencilflow



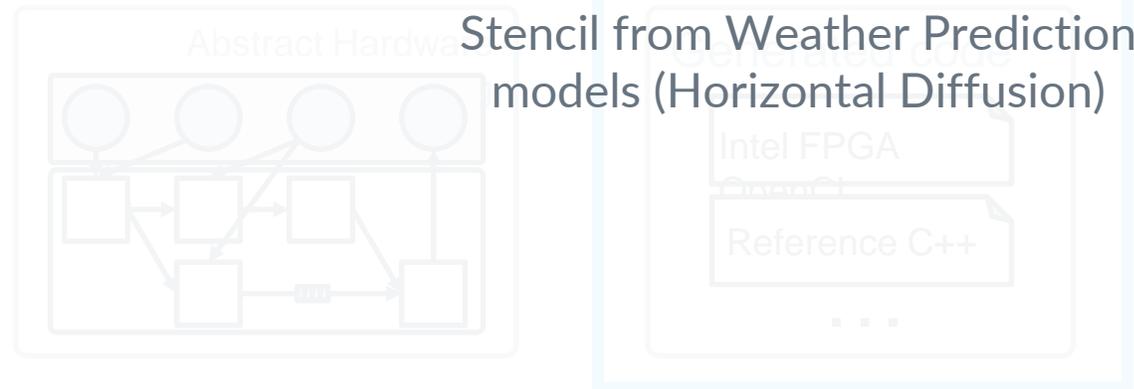
	Runtime	Performance	Peak BW.	%Roof.
Stratix 10	1,178 μ s	145 GOp/s	77 GB/s	52%
Stratix 10*	332 μ s	513 GOp/s	∞ GB/s	—
Xeon 12C	5,270 μ s	32 GOp/s	68 GB/s	13%
P100	810 μ s	210 GOp/s	732 GB/s	8%
V100	201 μ s	849 GOp/s	900 GB/s	26%

*Without memory bandwidth constraints.

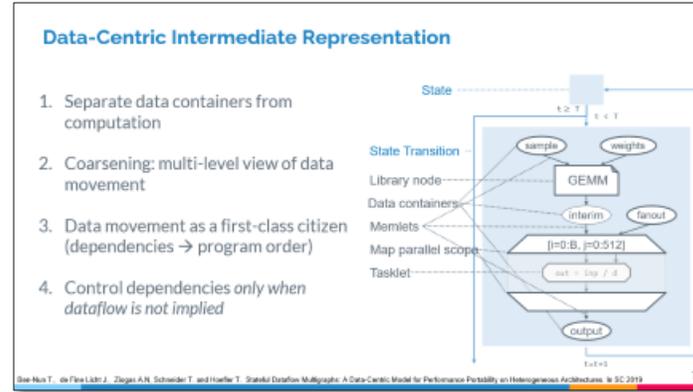
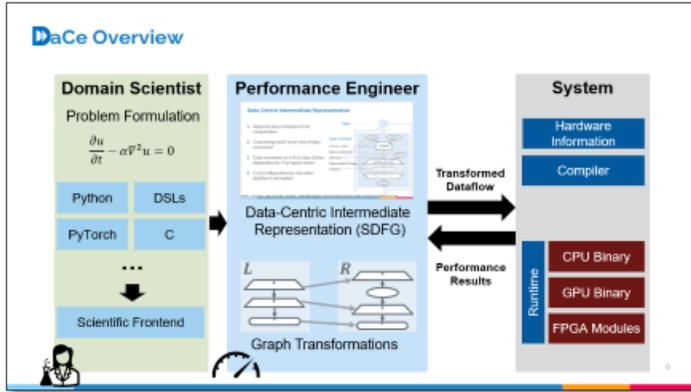
Iterative stencils (e.g., Jacobi, Diffusion, ...) on multiple FPGA devices using **SMIL**



Stencil from Weather Prediction models (Horizontal Diffusion)



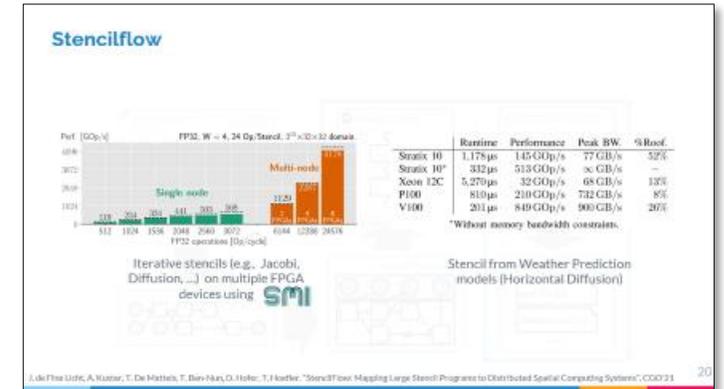
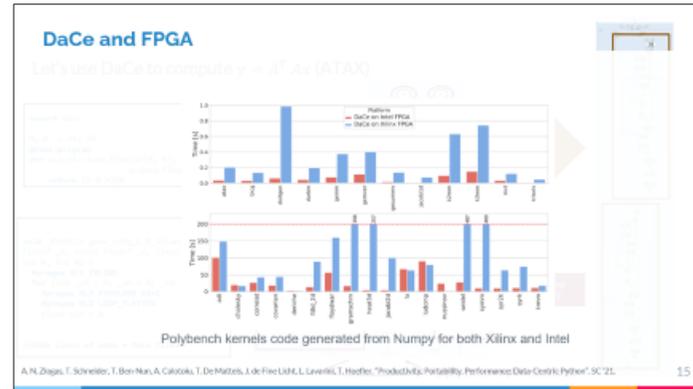
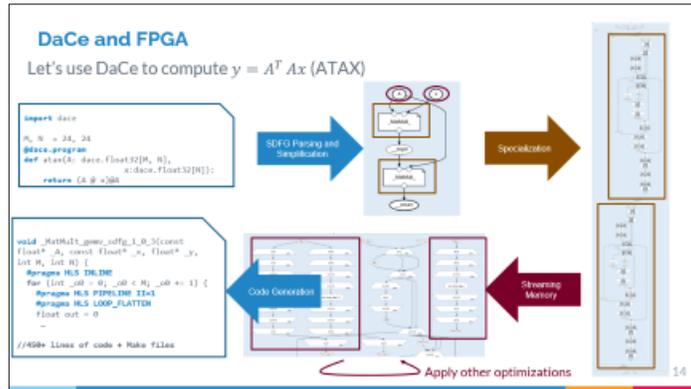
Conclusion



DaCe is in active development.
If you want to know more/contribute:



github.com/spcl/dace



t.de.matteis@vu.nl