

# GT4Py

A Python framework for performance  
portable weather and climate applications

Till Ehrenguber

# Vison

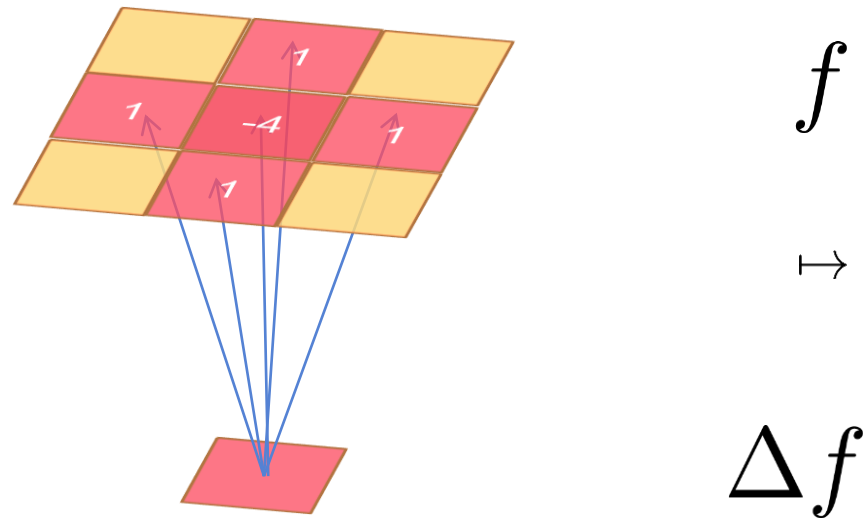
Make it **easy** to write **performant** weather  
and climate applications



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Prototypical example: 2D Laplacian



Mathematical formulation

$$\Delta f_{i,j} \approx -4f_{i,j} + f_{i-1,j} + f_{i+1,j} + f_{i,j-1} + f_{i,j+1}$$



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

**How do we achieve good performance?**



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# How do we achieve good performance?

⇒ increase data locality



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Traditional approach

```
for (int i = 0; i < n; ++i) {  
    for (int j = 0; j < n; ++j) {  
        g(i, j) = -4*f(i, j) + f(i-1, j) + f(i, j-1)  
                + f(i+1, j) + f(i, j+1);  
    }  
}
```



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Traditional approach

```
#pragma omp parallel for collapse(2)
for (int i = 0; i < n; i+=block_size) {
    for (int j = 0; j < n; j+=block_size) {
        for (int ii = i; ii < i+block_size; ++ii) {
            for (int jj = j; jj < j+block_size; ++jj) {
                g(ii, jj) = -4*f(ii, jj) + f(ii-1, jj) + f(ii, jj-1)
                               + f(ii+1, jj) + f(ii, jj+1);
            }
        }
    }
}
```



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Traditional approach: Optimization directives

| Model             | Number of directives |
|-------------------|----------------------|
| FVM (Fortran)     | 1365 x OMP PARALLEL  |
| ICON (Fortran)    | 1243 x ACC PARALLEL  |
| GFS-FV3 (Fortran) | 333 x OMP PARALLEL   |



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre



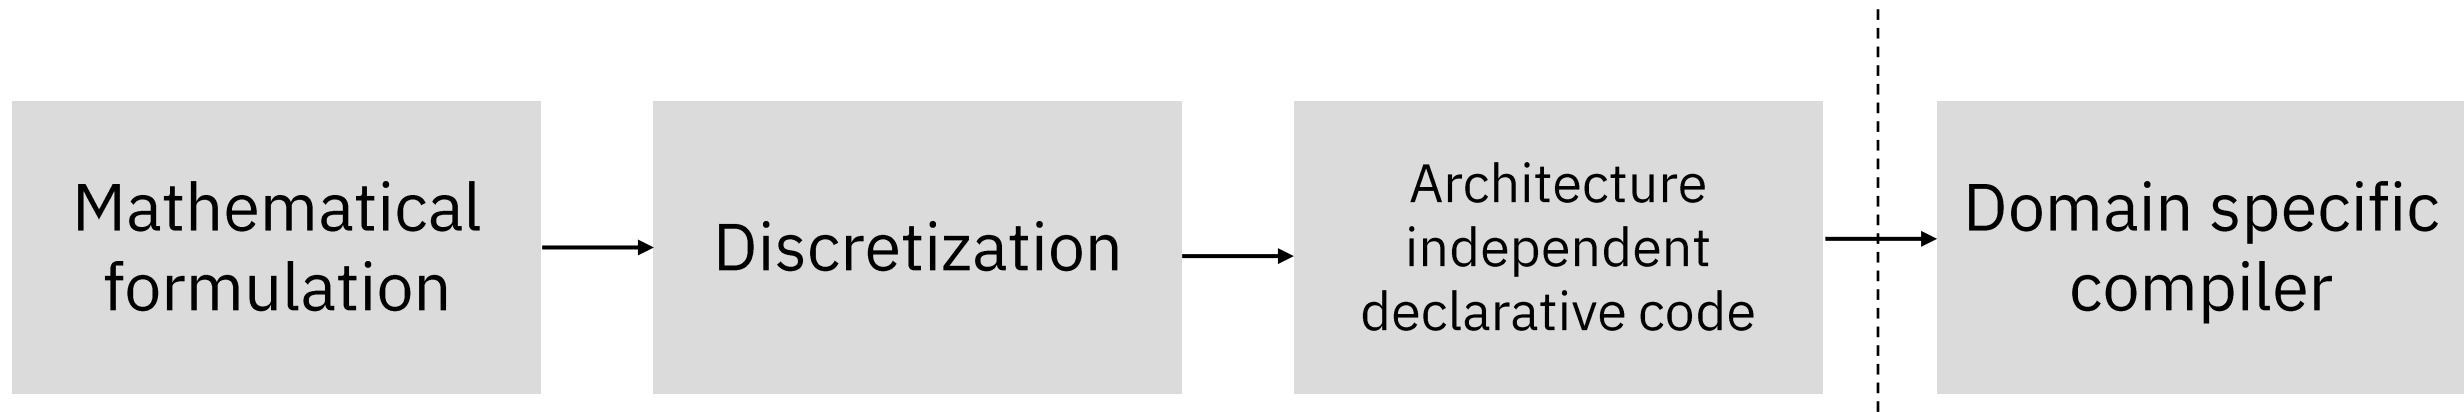
**Optimization directives are not scalable!**



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

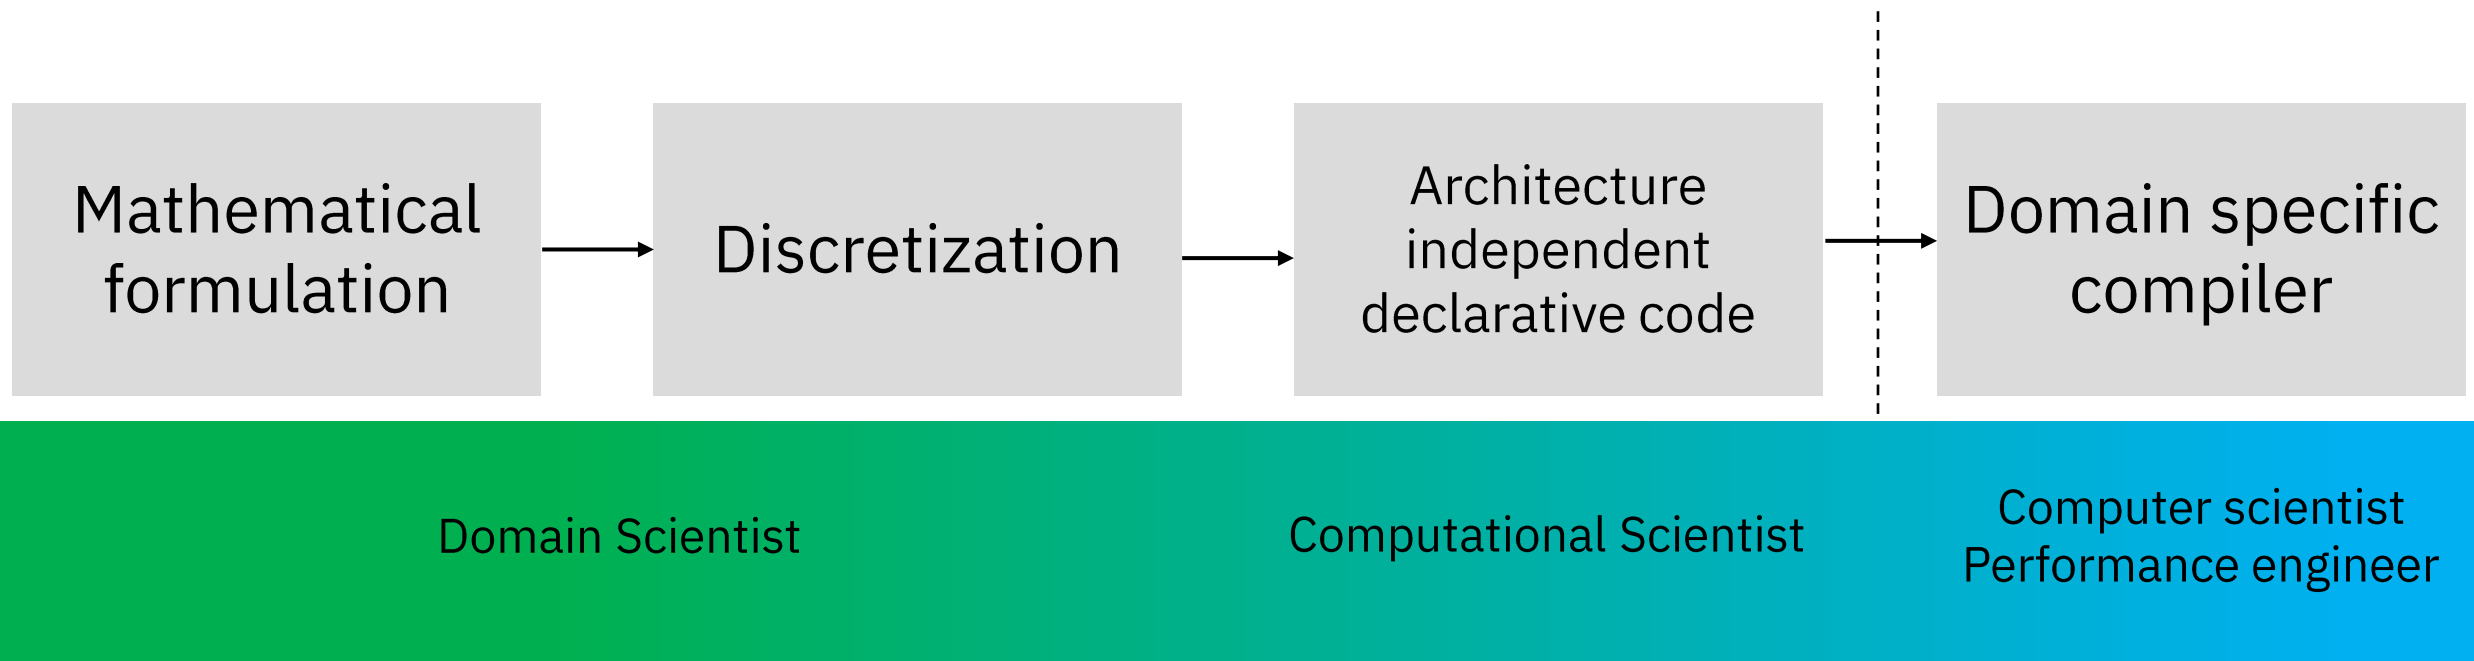
# Separation of concerns



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

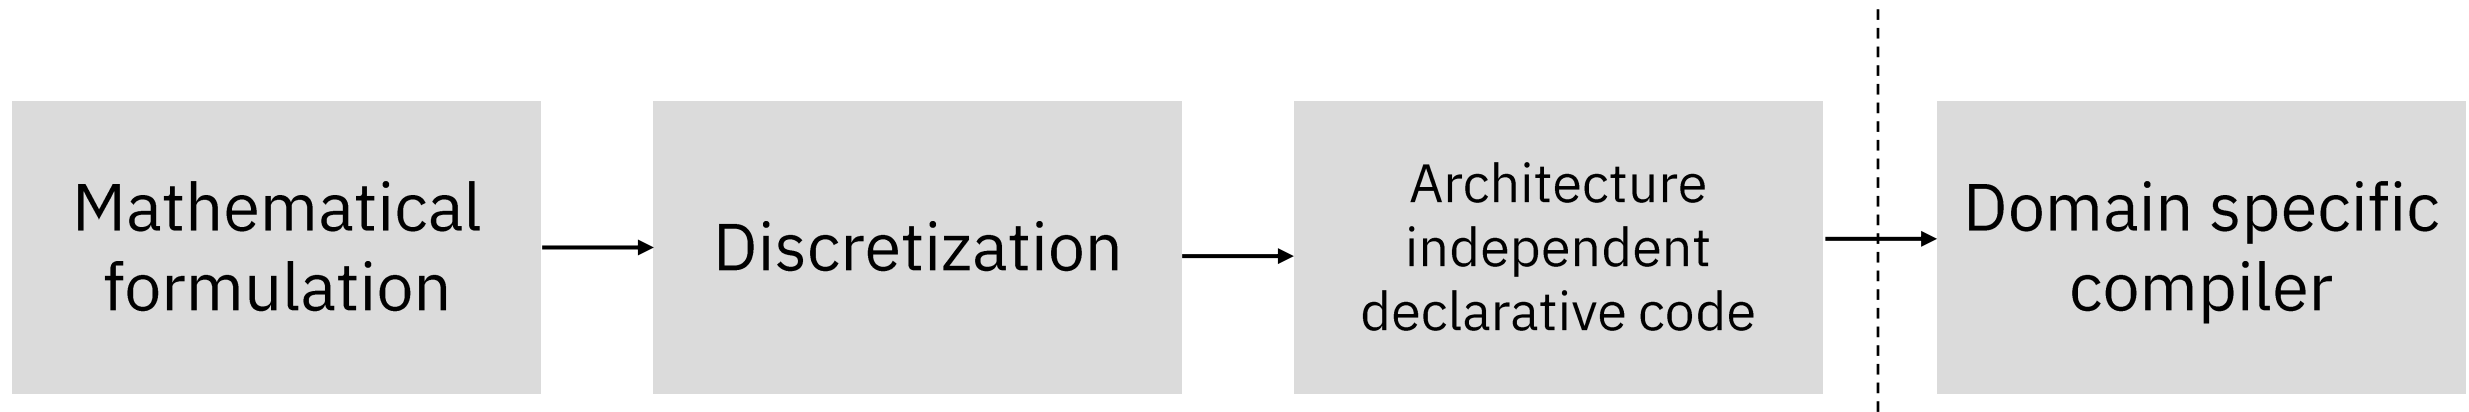
# Separation of concerns



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

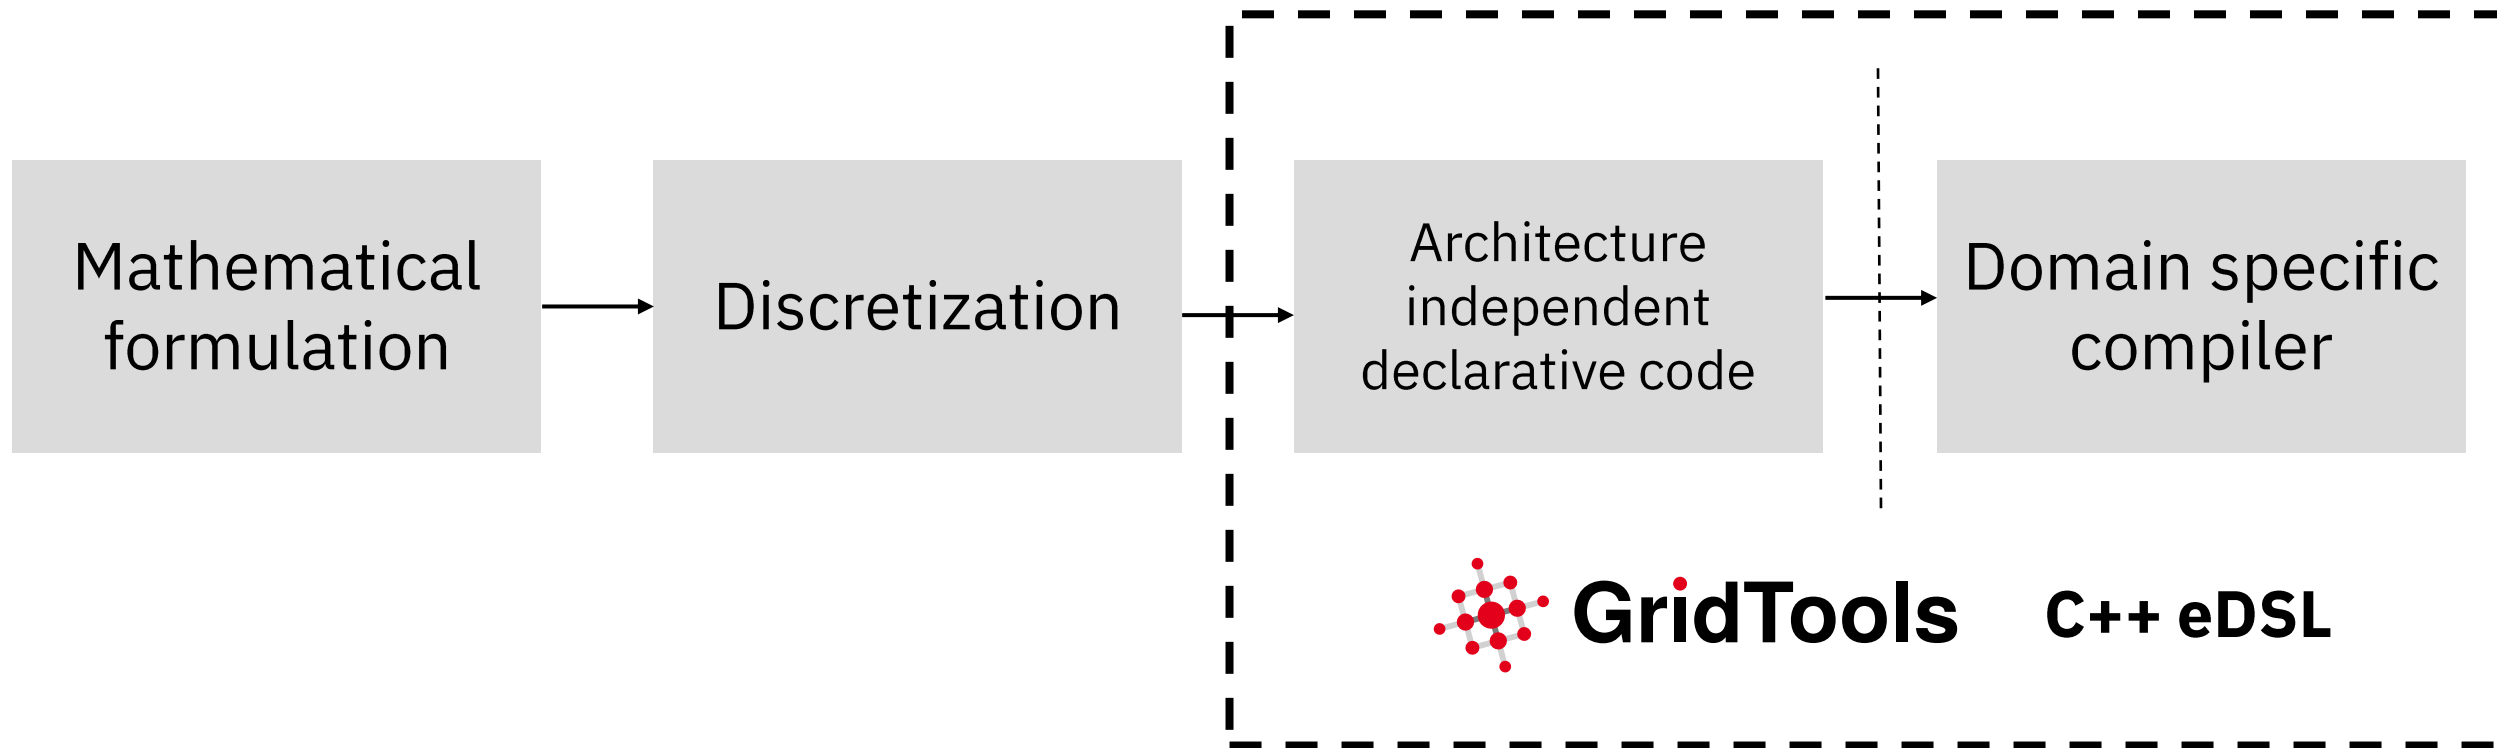
# Separation of concerns



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# GridTools C++ eDSL

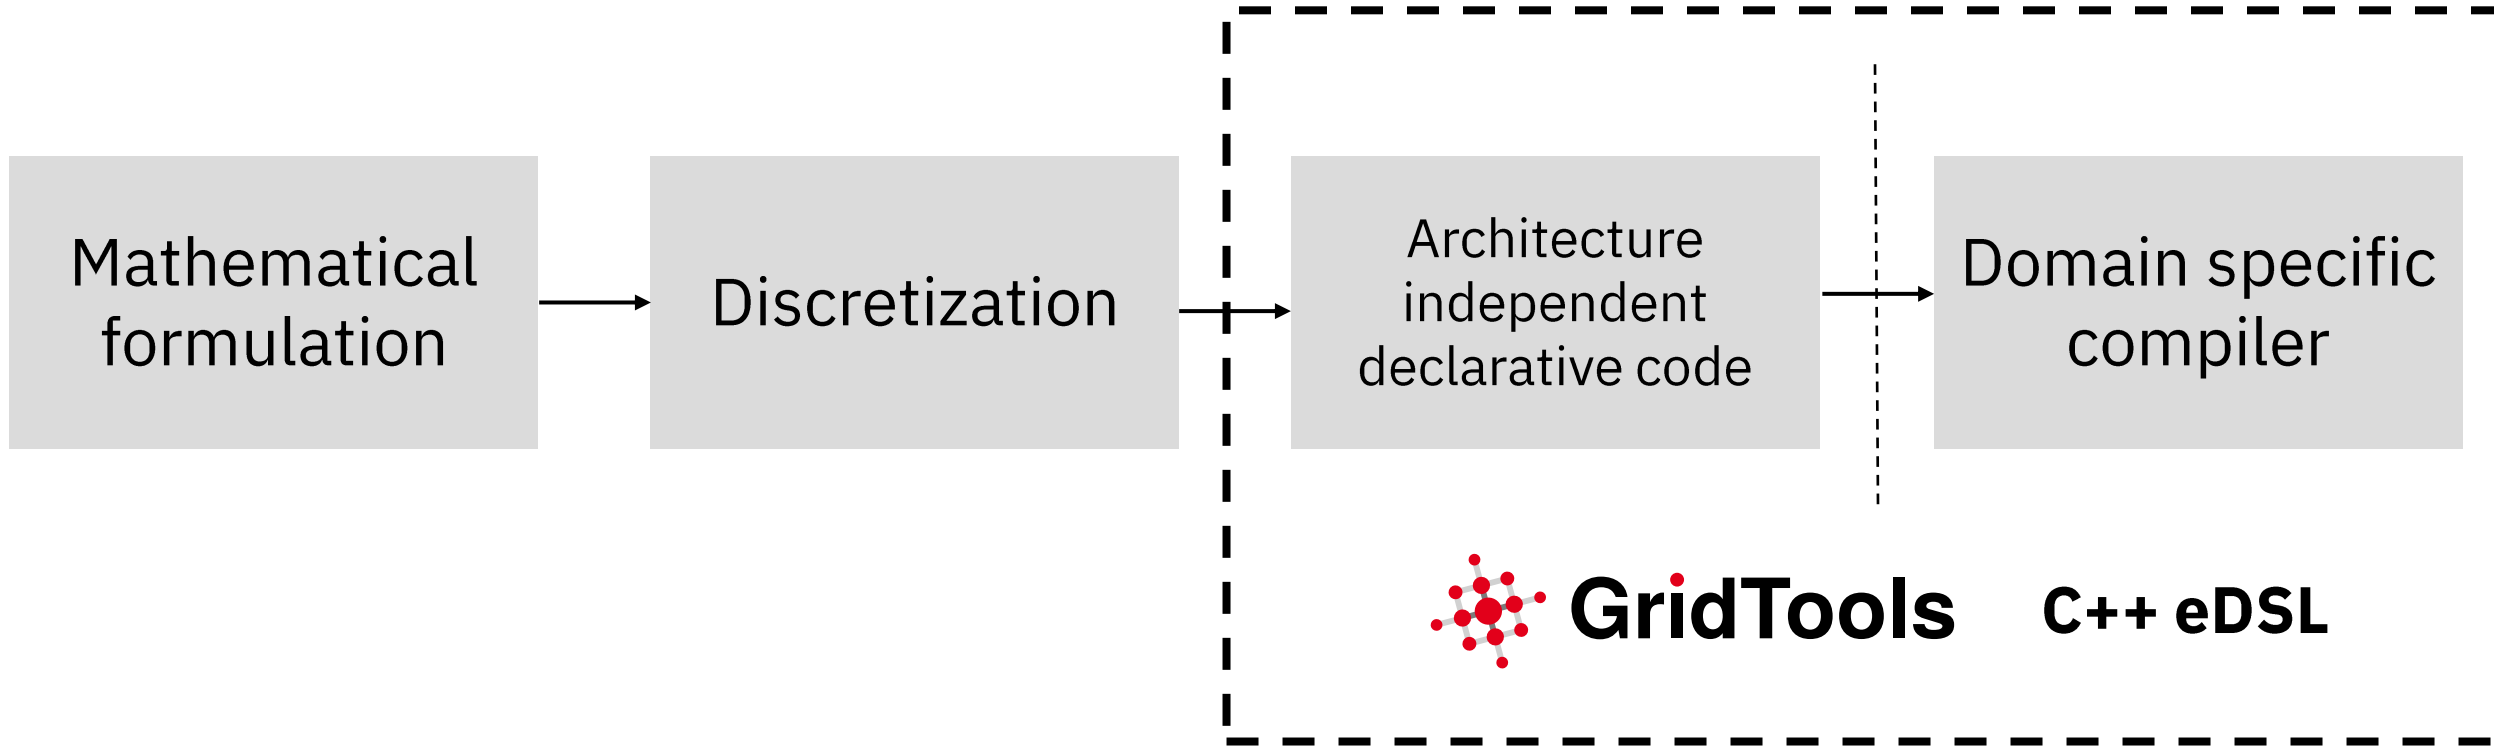


**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# GridTools C++ eDSL

Operational at MeteoSwiss since 2019 (COSMO)

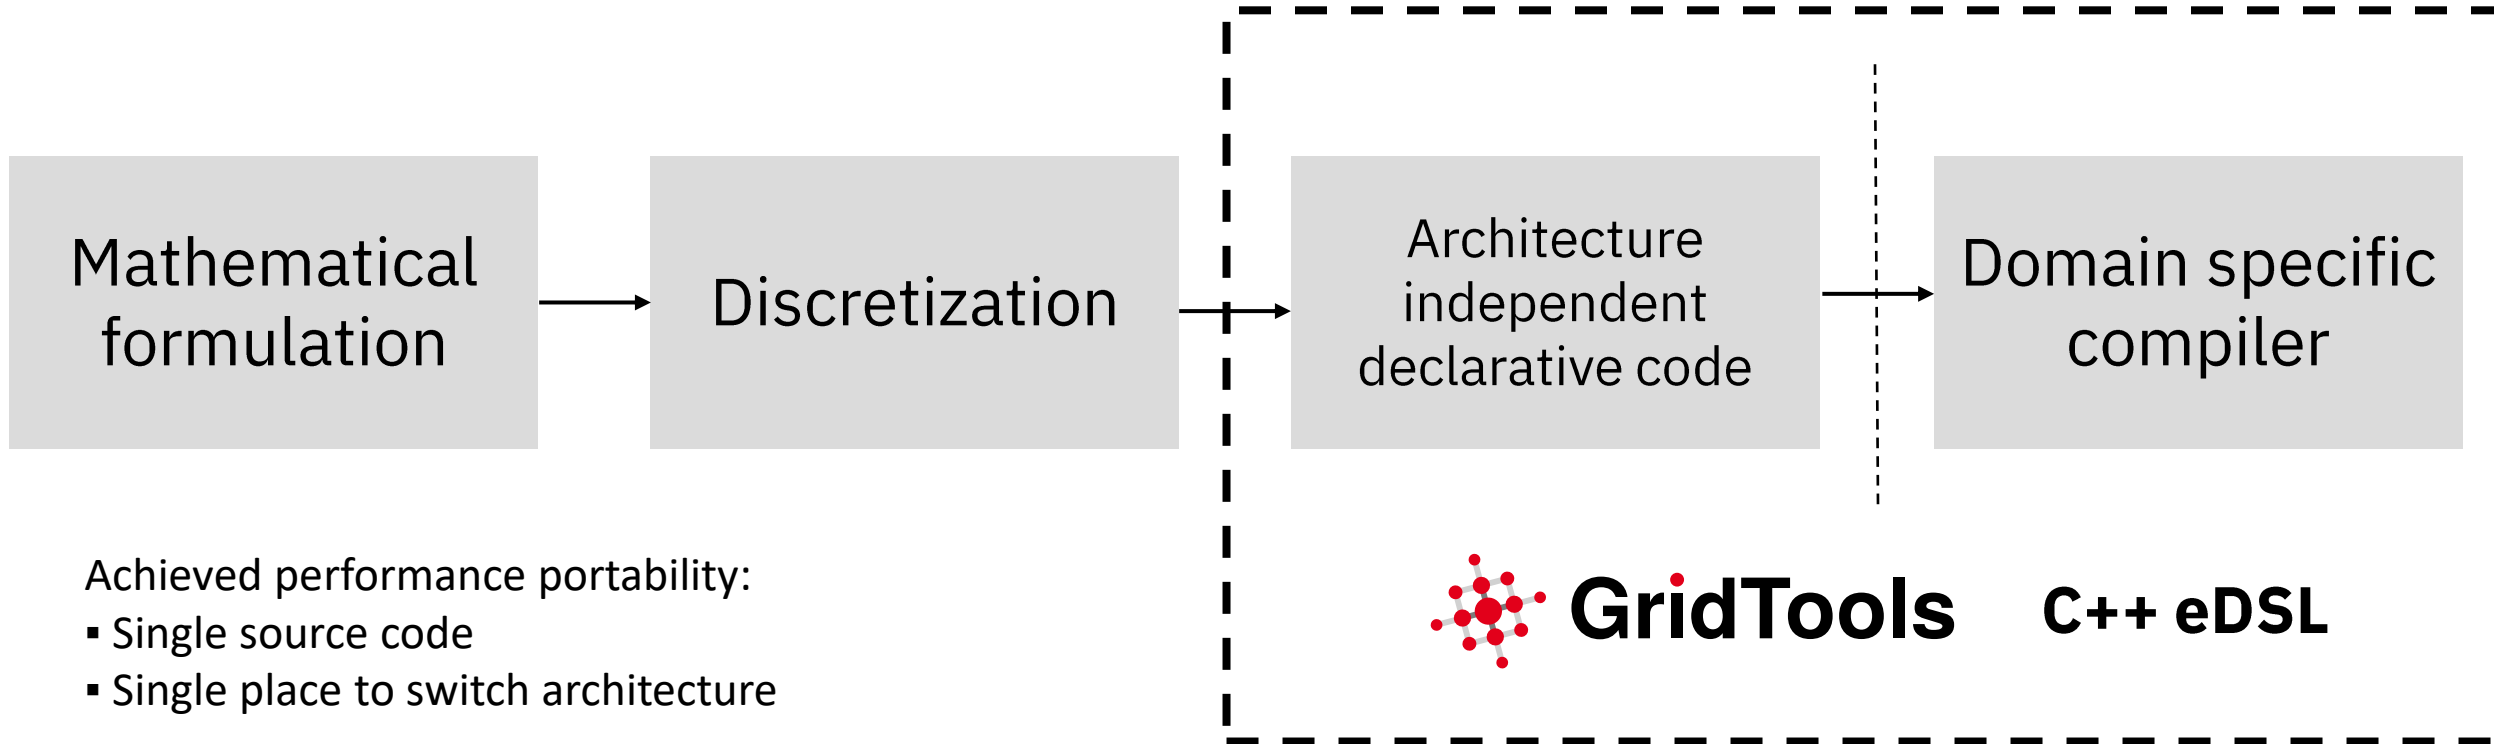


**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# GridTools C++ eDSL

Operational at MeteoSwiss since 2019 (COSMO)

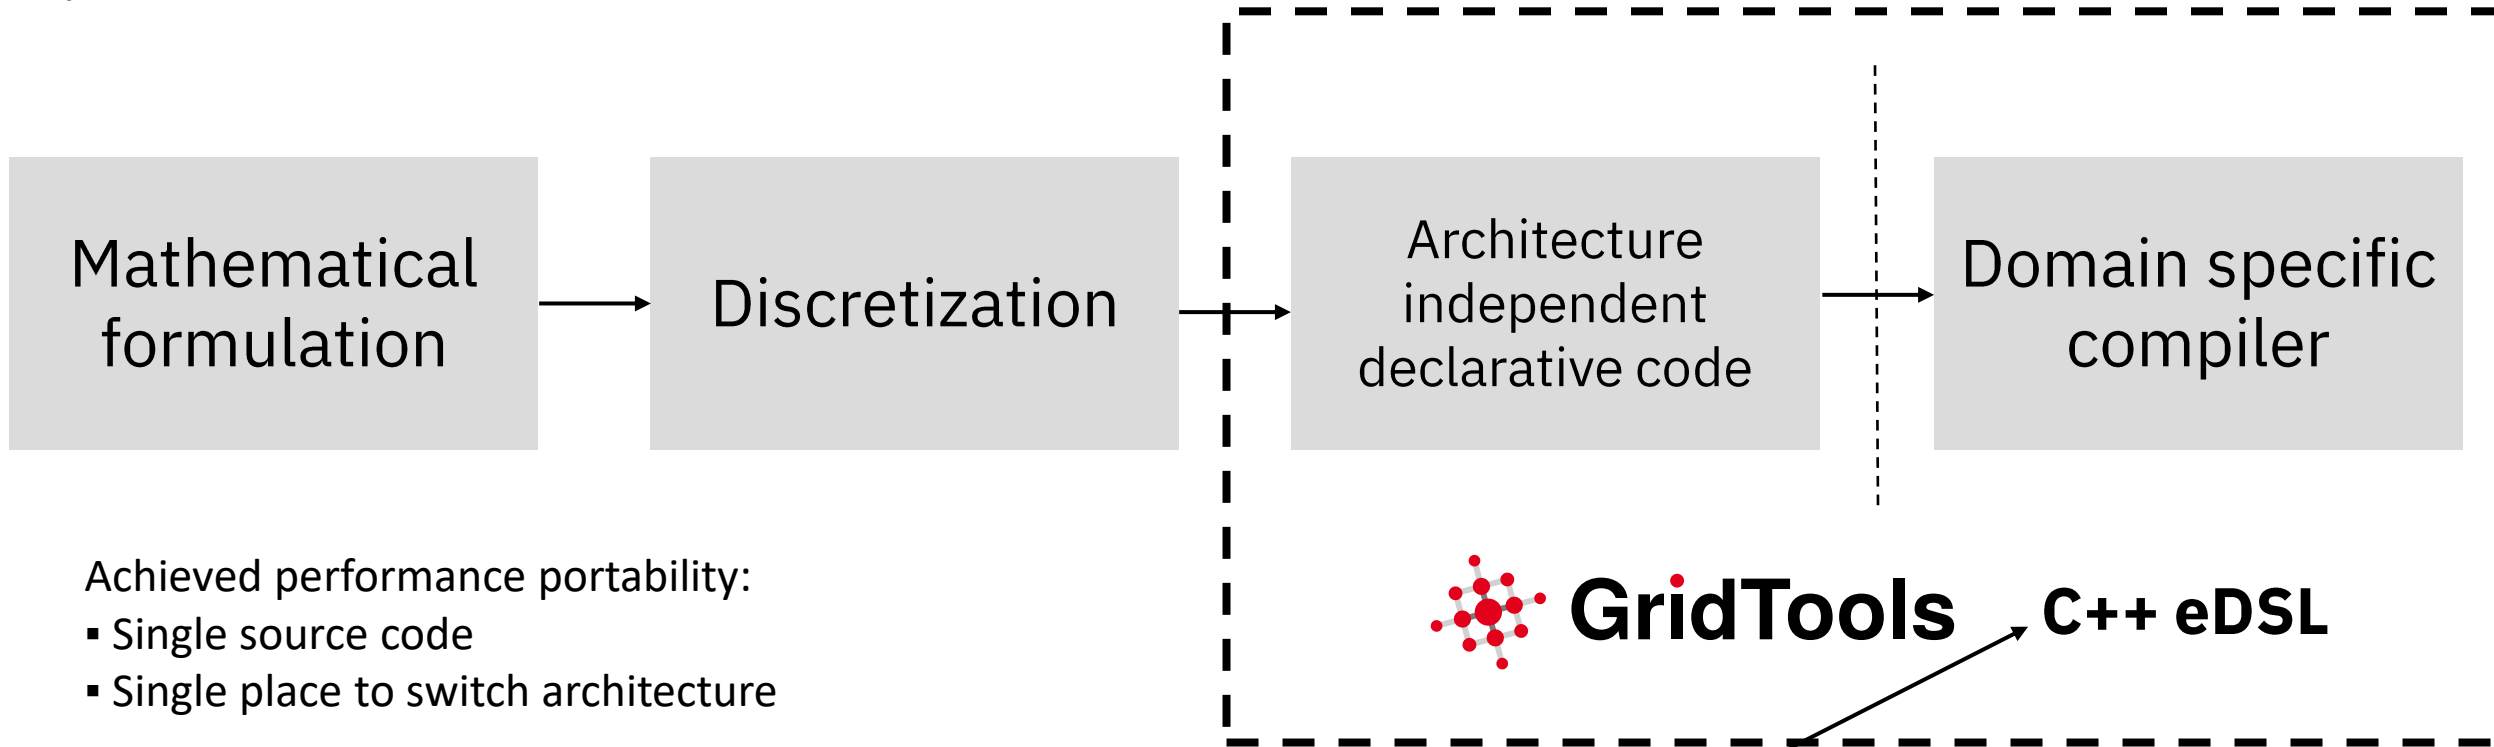


**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# GridTools C++ eDSL

Operational at MeteoSwiss since 2019 (COSMO)



Achieved performance portability:

- Single source code
- Single place to switch architecture

We did not manage to convince the domain scientist to use this language

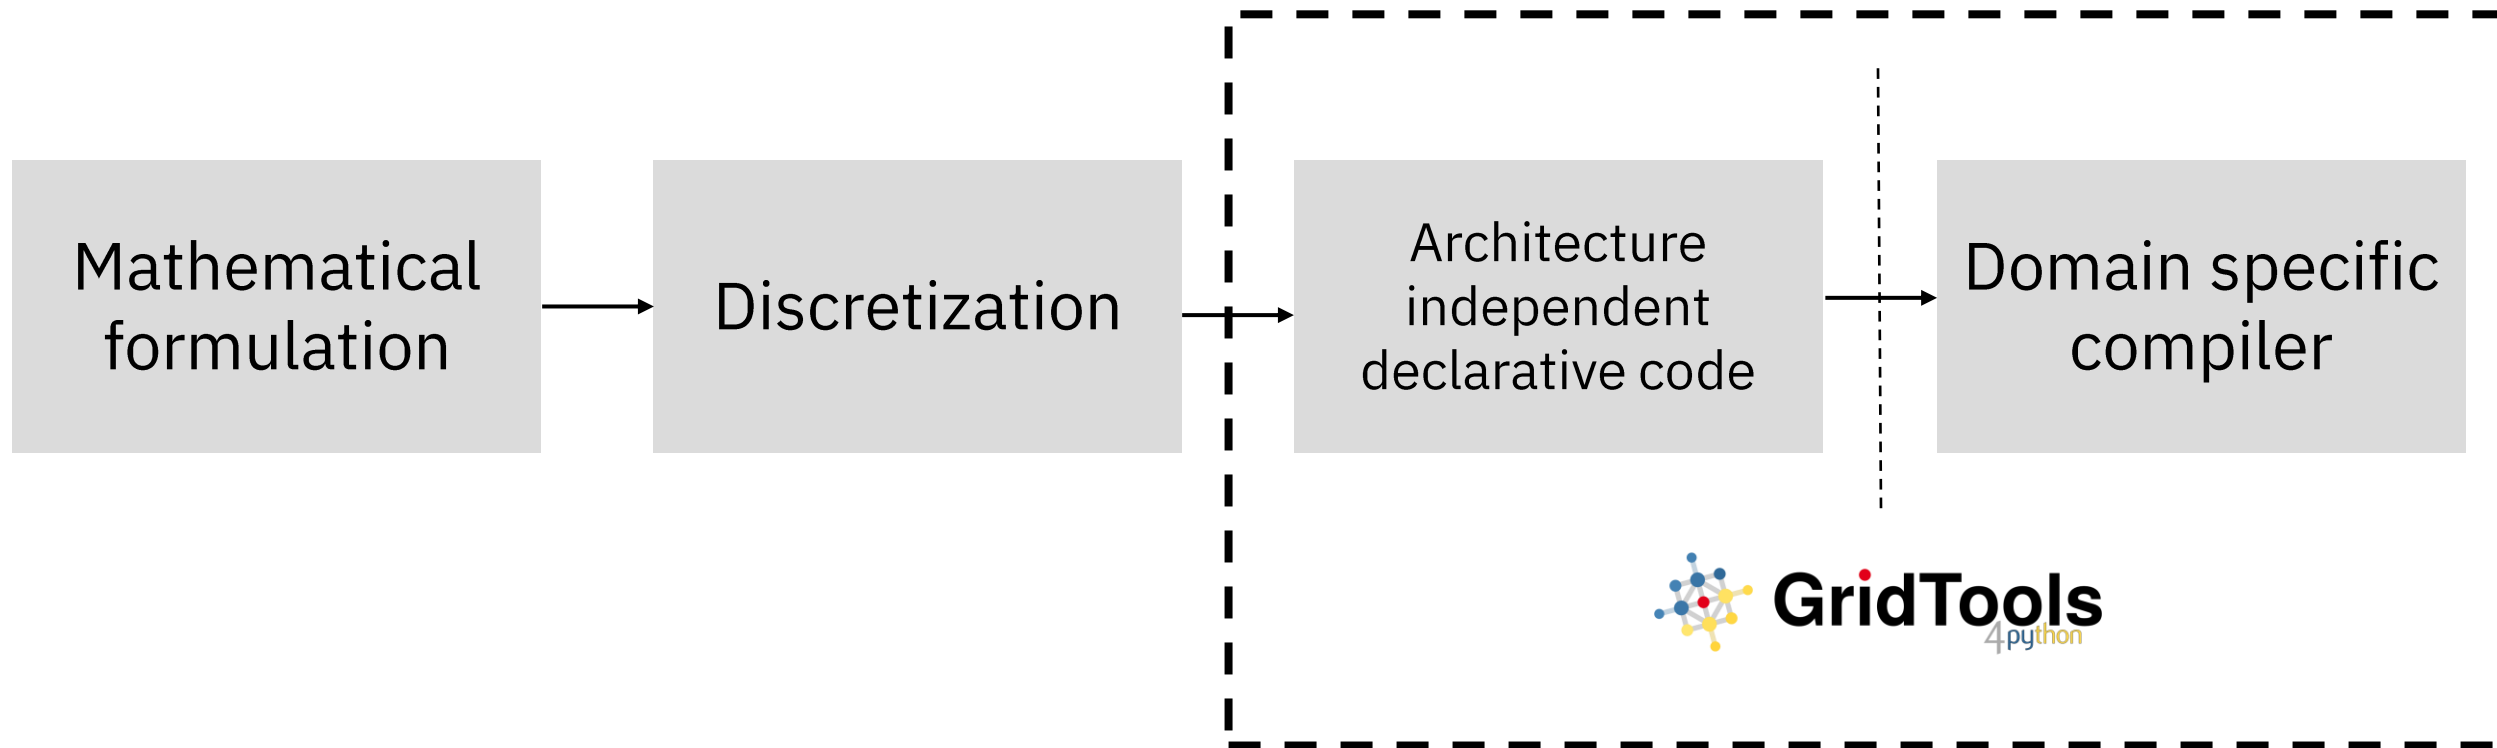


**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre



# GT4Py



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# GT4Py

- Domain scientists write high-level code specifying the algorithm and the mathematical operations
  - No need to pollute the code with high-performance optimizations techniques
  - Hardware-agnostic description of the computation → performance portability
- GT4Py works as an optimizing compiler with several backends
  - Code generation optimized for a specific architecture
  - Backend selects implementation strategy (parallelism, memory layout, etc.)
  - Backends can be added to provide efficient implementations for new platforms
  - Leverages knowledge of the typical computation patterns in the domain
- Smooth integration with Python scientific ecosystem provides access to a huge collection of Python libraries
  - data processing, analysis, visualization, etc.



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Laplacian in GT4Py

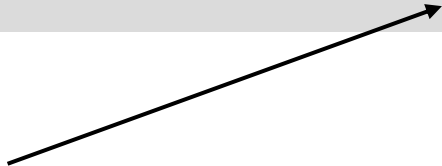
Mathematical formulation

$$\Delta f \approx -4f_{i,j} + f_{i-1,j} + f_{i+1,j} + f_{i,j-1} + f_{i,j+1}$$

GT4Py code

```
@field_operator
def laplacian(f: Field[[I, J], float]):
    return -4*f + f(I-1) + f(I+1) + f(J-1) + f(J+1)
```

Hardware-agnostic description  
of the computation



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Operators

## Field operator (@field\_operator)

Covering most patterns of explicit finite-difference and finite-volume discretizations, multiple field operations can be grouped together into a field operator.

```
@field_operator
def edge_average(
    vertex_field: Field[[Vertex], float]
) -> Field[[Edge], float]:
    return 0.5*(vertex_field(E2V[0])+vertex_field(E2V[1]))
```

Field operators are composable, allowing the description of high-level operators from basic building blocks.

```
laplap = laplacian(laplacian(f))
```

## Scan operator (@scan\_operator)

Scan operators are useful for expressing computations with dependencies across an entire dimension, which commonly occur in implicit solvers and physical parametrizations. The output from the previous level (i.e.,  $k+1$  or  $k-1$ , depending on the direction) is used by a scalar function to derive a new value for the current grid point, iteratively building up a complete field.

```
@scan_operator(axis=KDim, forward=True, init=0.0)
def simple_scan_operator(
    carry: float, current_value: float
) -> float:
    return carry + current_value
```

```
result = simple_scan_operator(inp_field)
```



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Hardware Portability

```
@field_operator  
def laplacian(f: Field[[I, J], float]):  
    ...
```




**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Hardware Portability

Single line change to switch to a different hardware architecture



```
@field_operator(backend=gtn_cpu)
def laplacian(f: Field[[I, J], float]):
    ...
```




**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Hardware Portability

Single line change to switch to a different hardware architecture



```
@field_operator(backend=gtn_gpu)
def laplacian(f: Field[[I, J], float]):
    ...
```

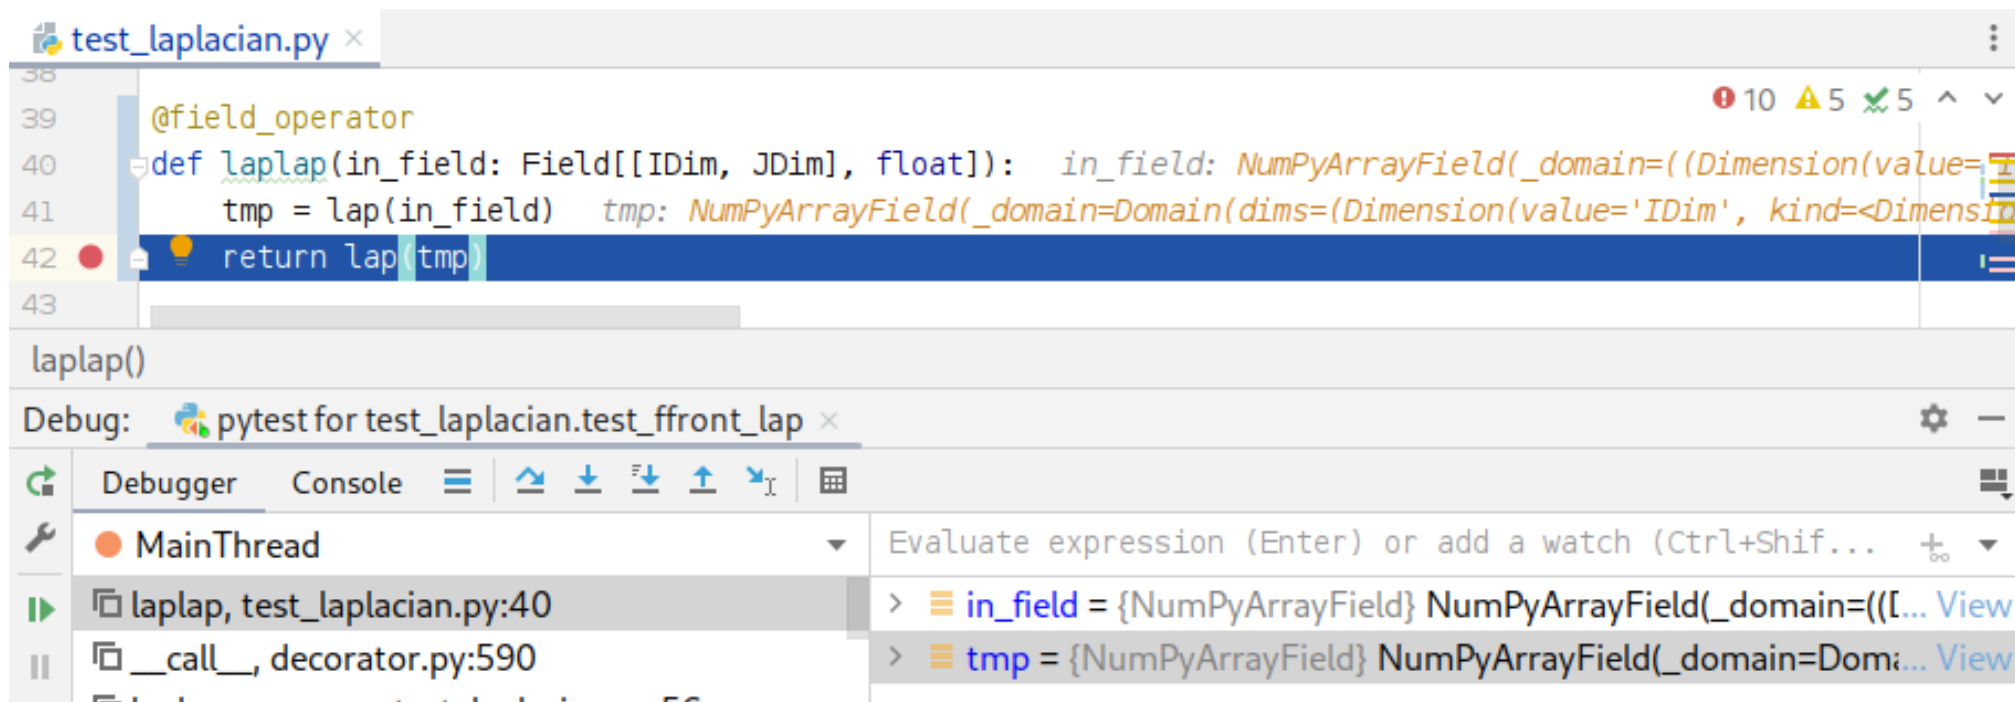


**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Debuggability

Intermediate results can be inspected using regular debugging tools (at the cost of performance).



The screenshot shows a Python IDE with a file named `test_laplacian.py` open. The code defines a function `laplap` decorated with `@field_operator`. The function signature is `def laplap(in_field: Field[[IDim, JDim], float])`. The function body contains two lines: `tmp = lap(in_field)` and `return lap(tmp)`. A red dot on line 42 indicates a breakpoint. Below the code editor, a debugger window is open, showing the current execution state. The debugger is running `pytest for test_laplacian.test_ffront_lap`. The current thread is `MainThread`. The call stack shows the function `laplap` at `test_laplacian.py:40` and the `__call__` method of the `decorator.py:590`. The debugger's watch window shows the current values of `in_field` and `tmp`, both of type `NumPyArrayField`.

```
test_laplacian.py x
38
39 @field_operator
40 def laplap(in_field: Field[[IDim, JDim], float]): in_field: NumPyArrayField(_domain=((Dimension(value=
41     tmp = lap(in_field) tmp: NumPyArrayField(_domain=Domain(dims=(Dimension(value='IDim', kind=<Dimension
42     return lap(tmp)
43

laplap()
Debug: pytest for test_laplacian.test_ffront_lap x
Debugger Console
MainThread Evaluate expression (Enter) or add a watch (Ctrl+Shif...
laplap, test_laplacian.py:40 > in_field = {NumPyArrayField} NumPyArrayField(_domain=(([... View
__call__, decorator.py:590 > tmp = {NumPyArrayField} NumPyArrayField(_domain=Dom...
```



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre



# Interoperability

GT4Py can be integrated into existing Fortran (or C++) code

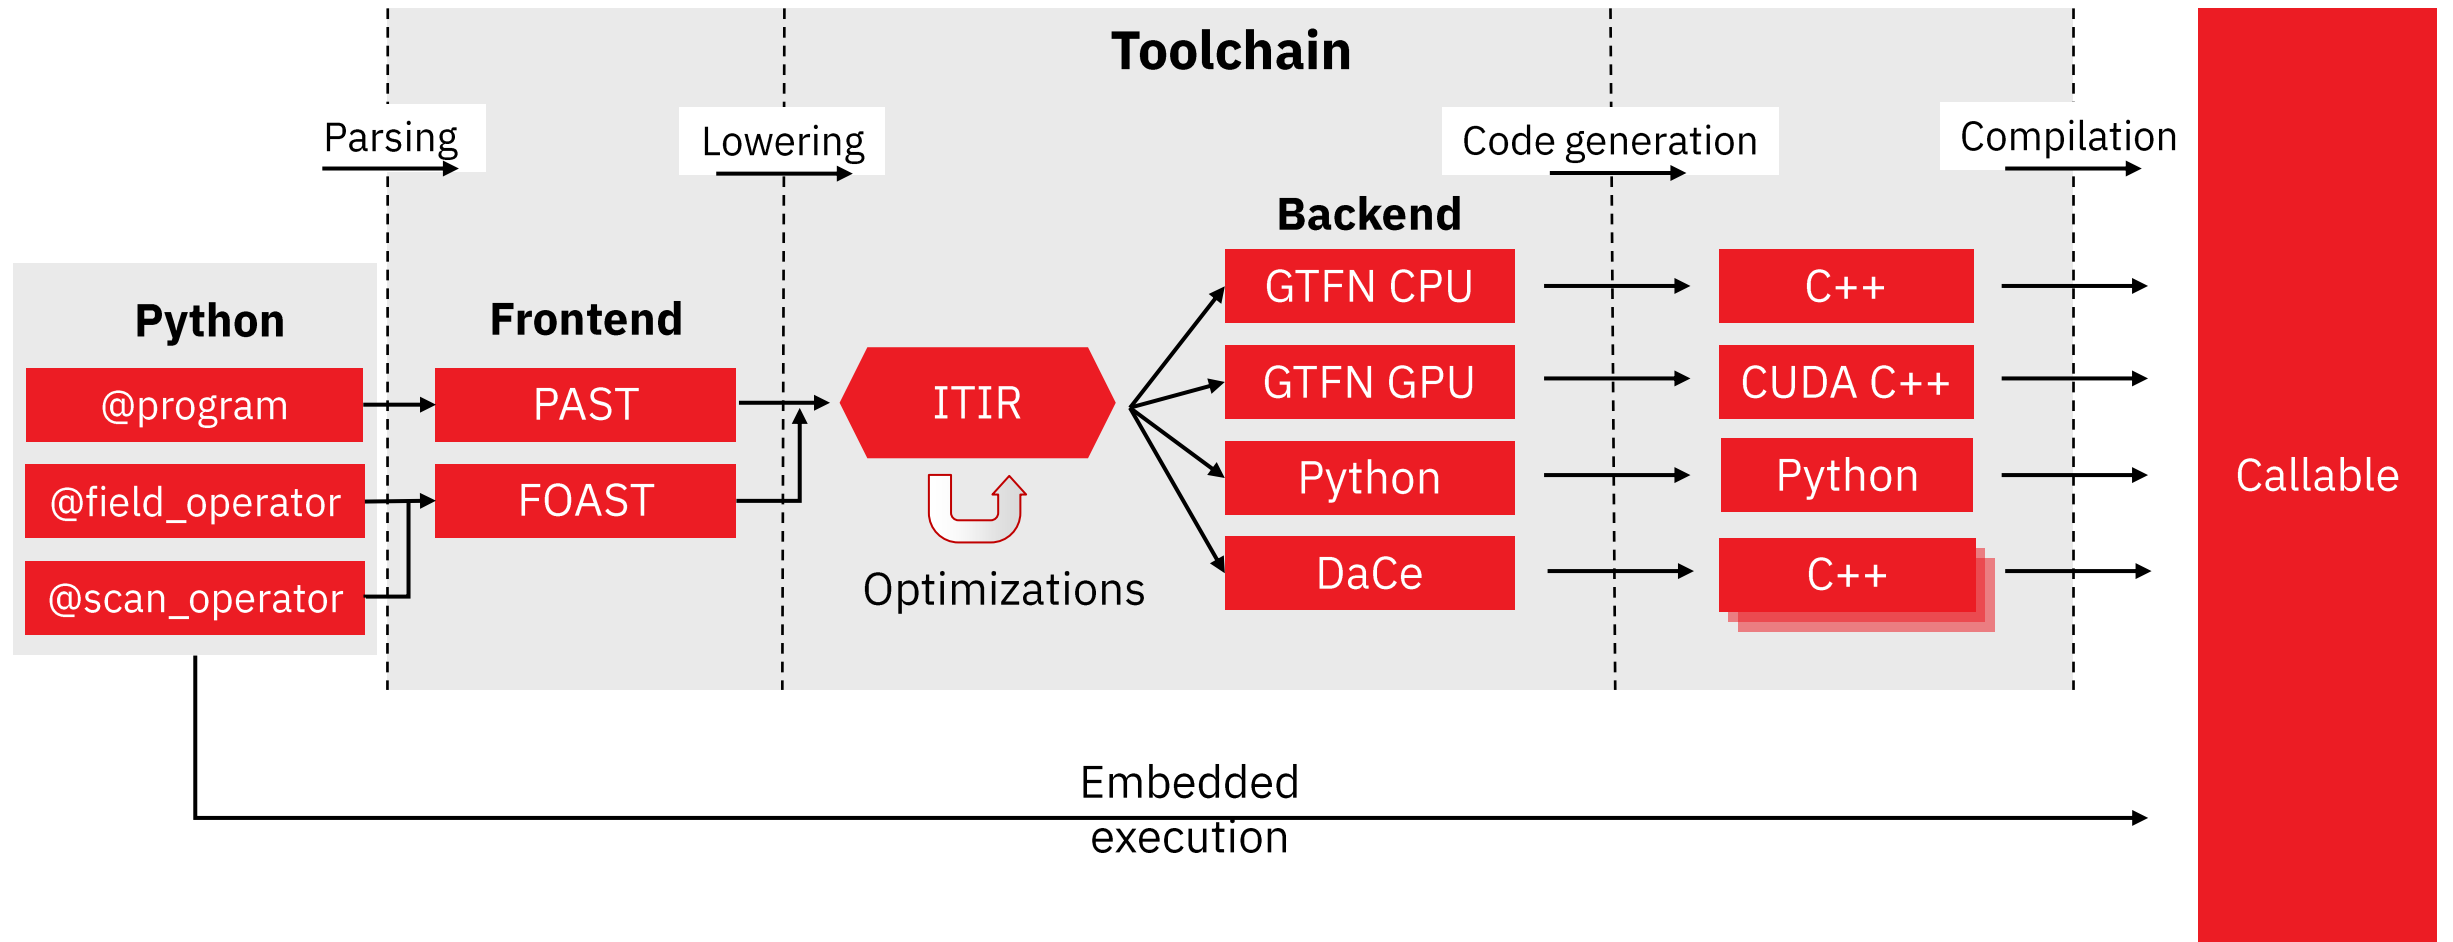
- Fortran driver, Python operators
- Incremental transition process towards Python
  - Approach taken in Exclaim project (ICON)
- Allows continuous validation during porting
- Some limitations persist
  - Memory layout fixed until everything ported (or costly reshuffling / H2D transfers)



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# GT4Py Toolchain



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Sample optimization: Constant folding

```
if_(True, true_branch, false_branch)
```

↓ transformation

```
true_branch
```

Pattern

```
if (
    isinstance(new_node.fun, itir.SymRef)
    and new_node.fun.id == "if_"
    and isinstance(new_node.args[0], itir.Literal)
):
    if new_node.args[0].value == "True":
        new_node = new_node.args[1]
    else:
        new_node = new_node.args[2]
```

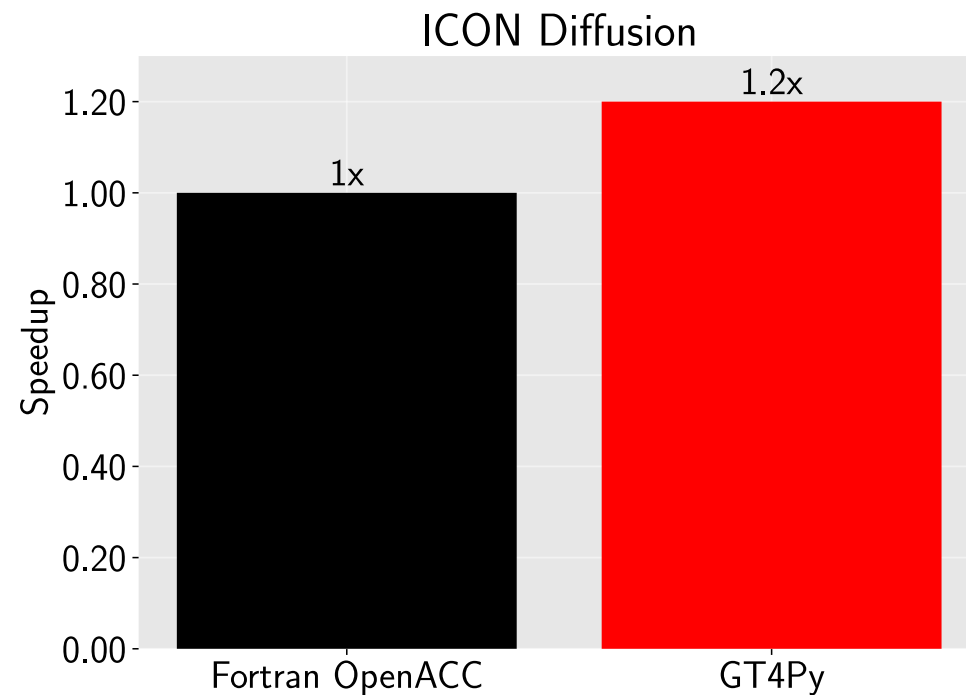
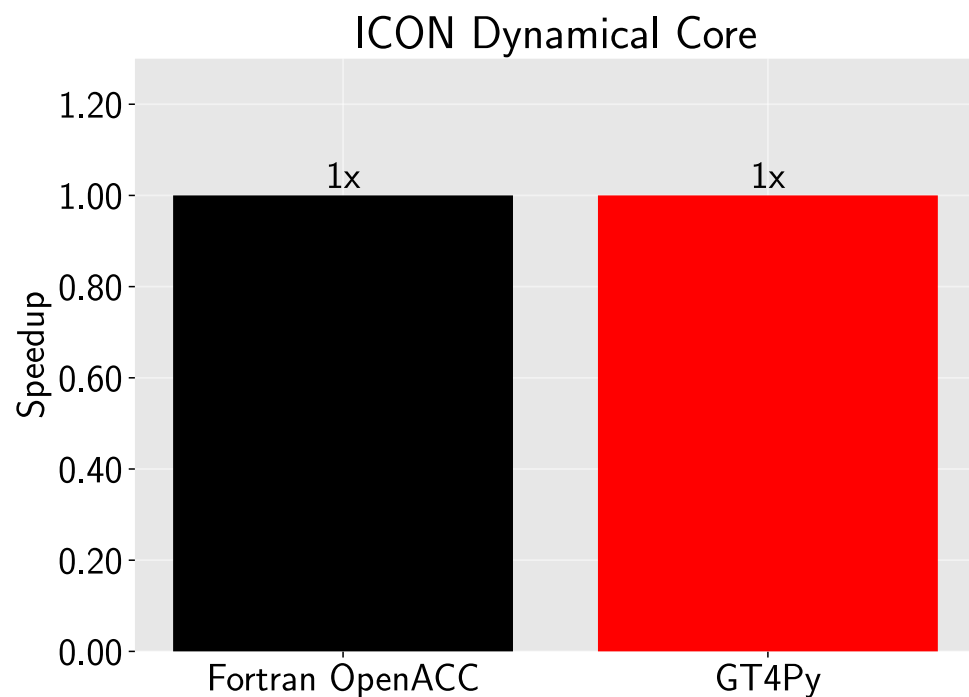
Optimization pass  
(excerpt)



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# Performance



*Hardware:*

NVIDIA® V100 GPUs 32GB  
2 x Intel® Xeon® Skylake 6134  
384 GB DDR4

*Software:*

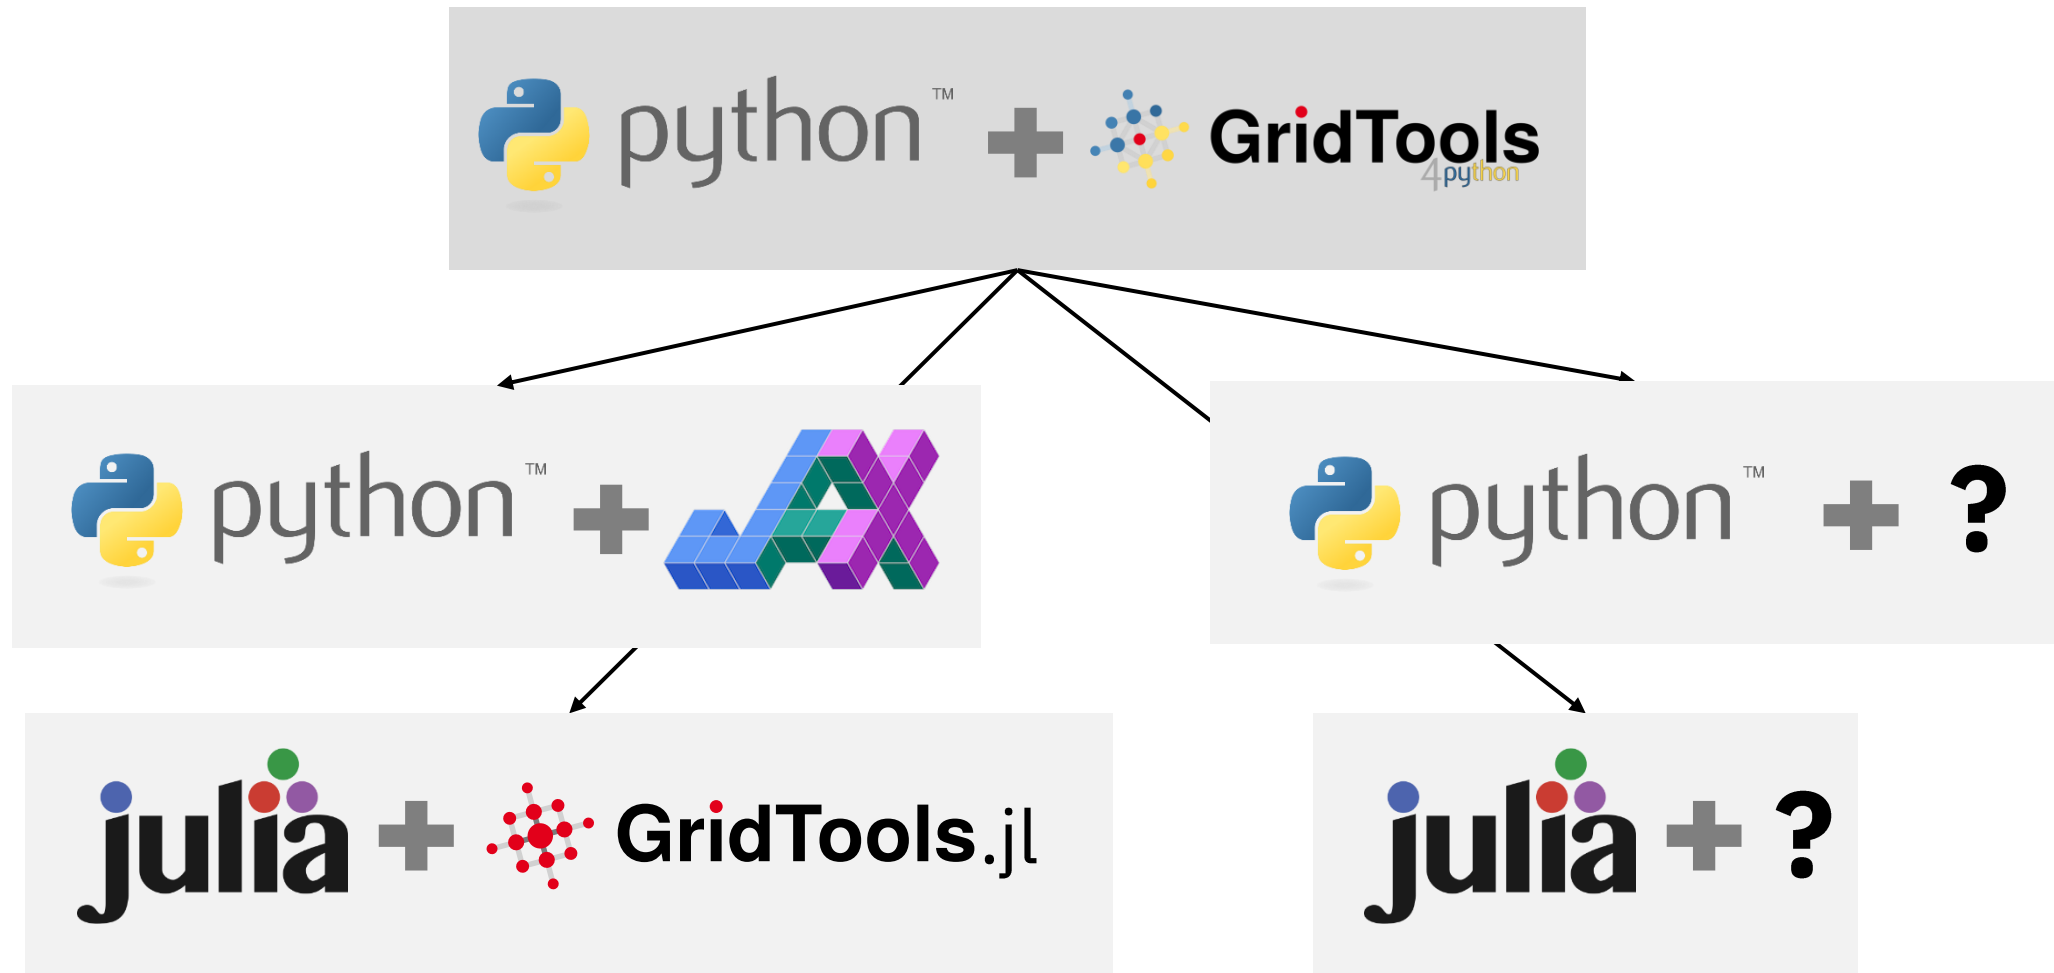
cuda/11.2.2  
nvhpc/21.2-cuda-11.2



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# GT4Py is not a dead end!



**CSCS**

Centro Svizzero di Calcolo Scientifico  
Swiss National Supercomputing Centre

# NWP & climate applications

- **ICON** atmospheric model dynamics and physics ported to declarative GT4Py (V2) by MeteoSwiss, EXCLAIM project at ETH Zurich  
→ See Daniel Hupp's poster
- **FVM** for IFS porting and further development with GT4Py (V1+V2) by ECMWF, CSCS, and PASC-funded project KILOS at ETH Zurich.  
→ See Christian Kühnlein's talk
- **Pace** is GT4Py (V1) implementation of the FV3GFS / SHiELD atmospheric model of NOAA and GFDL by Allen Institute for AI (AI2)  
→ See Oliver Elbert's talk  
→ See Christopher Kung's talk

# Thanks!



<https://github.com/GridTools/gt4py>

M. Bianco<sup>1</sup>, Nina Burgdorfer<sup>3</sup>, T. Ehrenguber<sup>1</sup>, N. Farabullini<sup>2</sup>, A. Gopal<sup>2</sup>,  
R. Häuselmann<sup>1</sup>, Daniel Hupp<sup>3</sup>, P. Kardos<sup>2</sup>, S. Kellerhals<sup>2</sup>, M. Luz<sup>2</sup>, C.  
Müller<sup>3</sup>, E. G. Paredes<sup>1</sup>, M. Roethlin<sup>3</sup>, H. Vogt<sup>1</sup>

<sup>1</sup> Swiss National Supercomputing Center, CSCS

<sup>2</sup> Center for Climate Systems Modeling, ETH Zurich

<sup>3</sup> Federal Office of Meteorology and Climatology, MeteoSwiss