

Developing DSLs in ESiWACE2

STFC, MeteoSwiss, DKRZ, ECMWF, UNIMAN, ICHEC, CMCC,
METO, UREAD

ESiWACE2 Virtual Workshop on Emerging Technologies for
Weather and Climate Modelling, 30th June 2020

ESiWACE2

- EU H2020 Centre of Excellence in Weather and Climate
- WP2 “**Establish, evaluate and watch** new technologies for the community”
 - **Establish** DSLs for the community
 - **Evaluate** model concurrency
 - **Evaluate** containers
 - **Watch** future trends
- This talk presents progress with DSLs

Motivation

- 3P's : Performance, Portability and Productivity
 - Maintainable high performance software
 - Single-source science code
 - Performance portability
- Complex parallel code + Complex parallel architectures + Complex compilers = Complex optimisation space => unlikely to be a single solution
- Single-source optimised code is unlikely to be possible
- So ... separate science specification/code from code optimisation

Approach

- 2 DSL implementations : DAWN and PSyclone
- Comparison via benchmarks
- Adaptation/extension of DSLs (for NEMO, ICON and IFS)
- Demonstration on pre-exascale machines (LFRic, NEMO, ICON, IFS)
- Interoperability
- Performance evaluation

DAWN in a slide

DAWN
MIT License
<https://github.com/MeteoSwiss-APN/dawn>

- Compiler toolchain to enable generation of high-level DSLs for geophysical fluid dynamics models
- Multiple supported Frontends, current options either embedded in Python or C++
- Generates C++ code, accelerated using OpenMP or CUDA
- Support for structured meshes very mature, support for unstructured meshes currently being developed
- Intended as a tool for domain scientists
- Automatic optimisations
- Scientific code fully decoupled from performance considerations

PSyclone in a slide

- A domain-specific compiler for embedded DSL(s)
 - Configurable: FD/FV NEMO, GOcean, FE LFRic
 - Currently Fortran -> Fortran/OpenCL
 - Supports distributed and shared memory parallelism
 - Supports code generation and code transformation
- A tool for use by HPC experts
 - Hard to beat a human (debatable)
 - Work round limitations/bugs
 - Optimisations encoded as a 'recipe' rather than baked into the scientific source code
 - Different recipes for different architectures

PSyclone 1.9.0

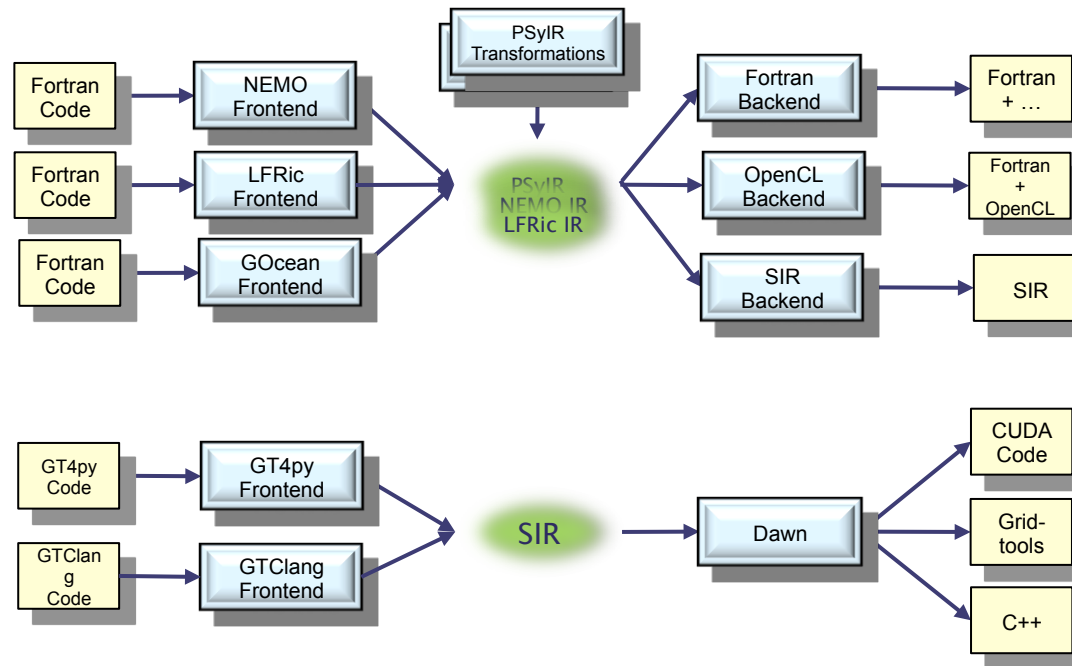
BSD 3-clause

<https://github.com/stfc/PSyclone>

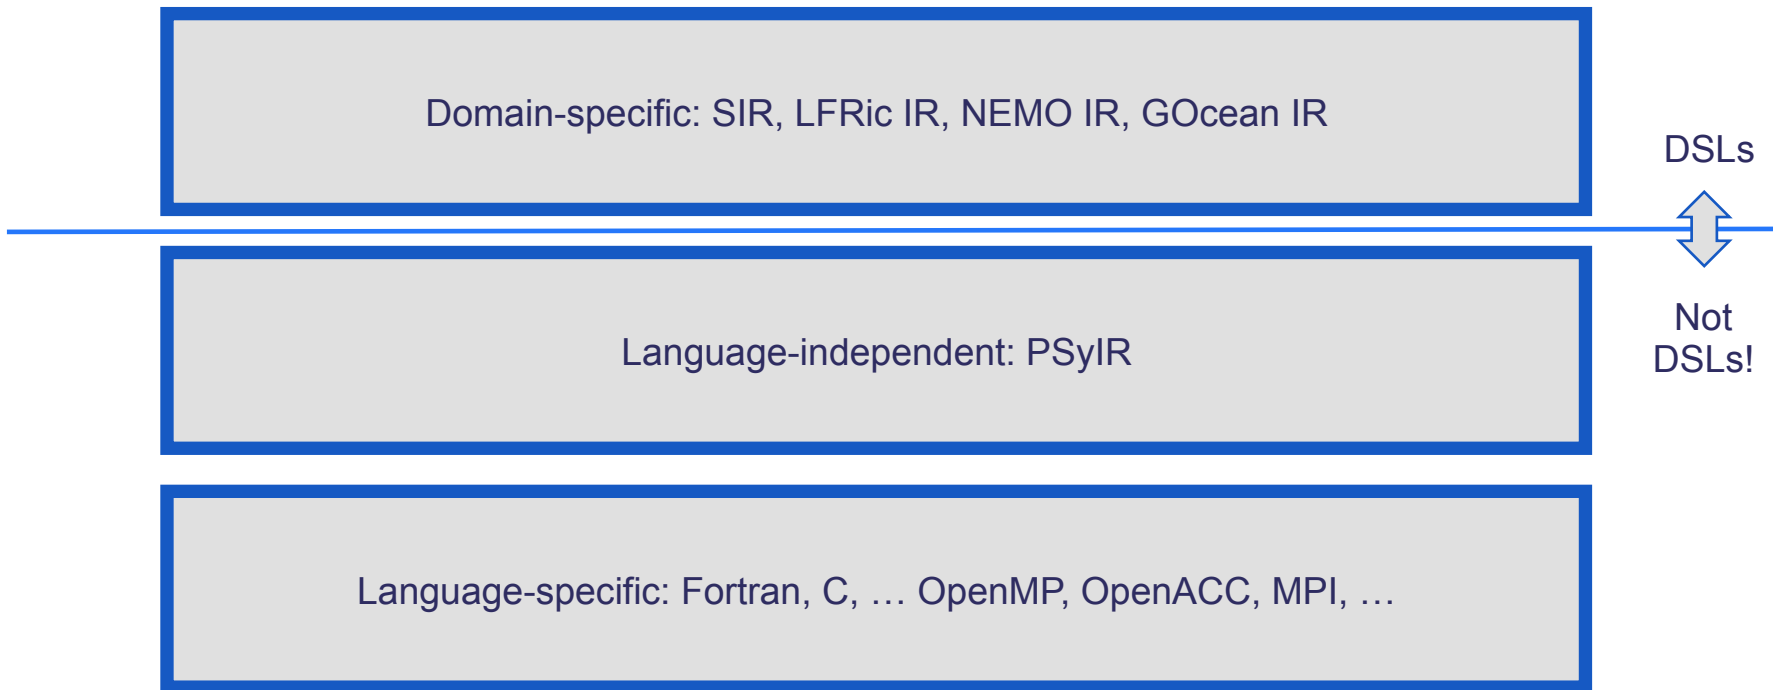
<https://psyclone.readthedocs.io>

```
> pip install psyclone
```

DAWN and PSystem



Levels of Abstraction



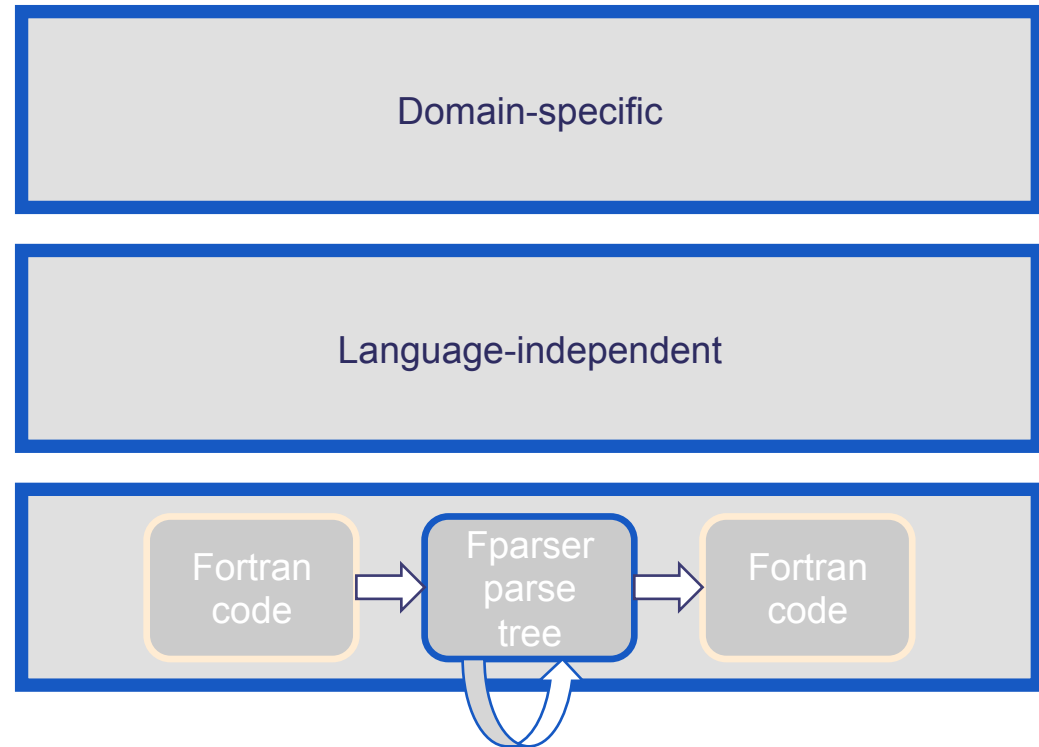
Fparser – shameless plug

- Fortran parser release 0.0.11
- Use fparser2!
- Supports Fortran 2003 + some 2008
- Written in Python
- Open source BSD3 licence
- Developed on Github
- Can fully parse UM, LFRic and NEMO source
- Work-in-progress to parse IFS source
- Used by PSyclone, Stylist, Loki

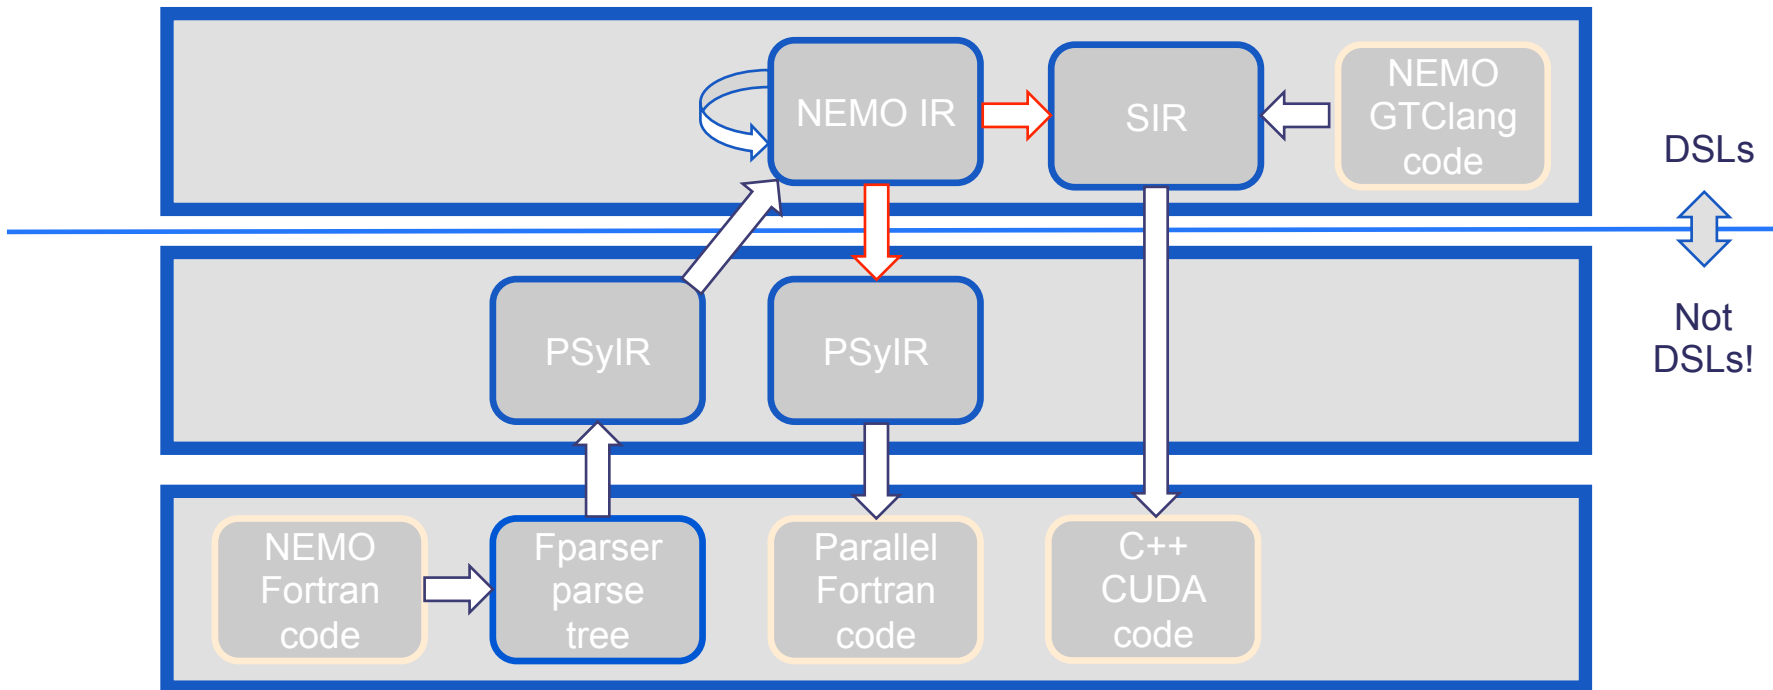
<https://github.com/stfc/fparser>

<https://fparser.readthedocs.io/>

> pip install fparser

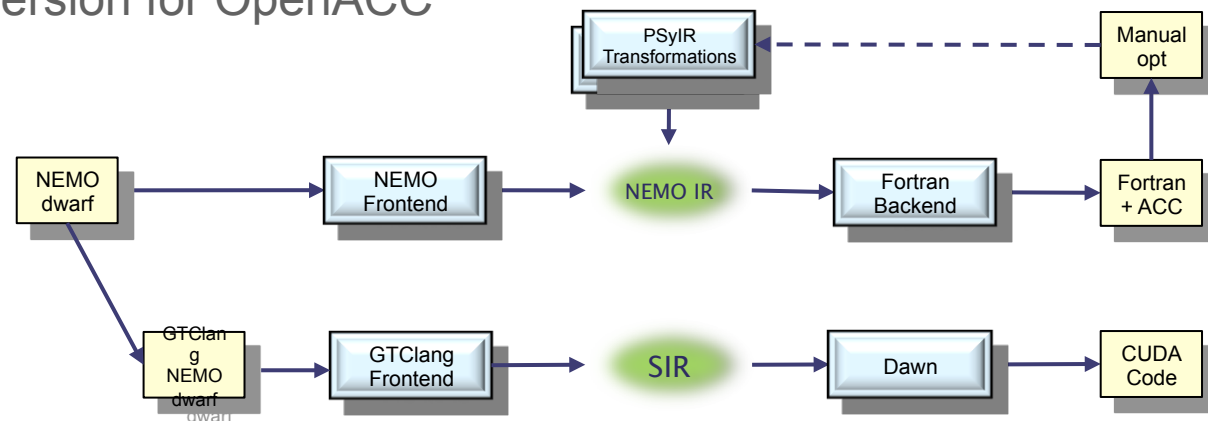


NEMO Example



Comparison via benchmarks

- NEMO dwarf (ESCAPE2)
- Tracer advection (most costly routine in NEMO)
- GTClang version in development
- PScyclone version for OpenACC [and OpenMP]
- Manually optimised version for OpenACC



Comparison via benchmarks


- Early version of NEMO dwarf (tracer advection benchmark) generated by PSystem run on GPU
- Performance analysed and hand-optimised
 - Better use of data regions
 - Better use of OpenACC Loop collapse
 - Loop fusion to remove temporary arrays
- Up to 21% performance improvement
- No significant further improvement expected

Adaptation/Extension : NEMO



- Script developed to add OpenACC to full NEMO code
- OpenACC transformation support extended to allow tuning (e.g. `independent and collapse`)
- Interface PSystem profiling API to NVIDIA profiling API
- Extend OpenACC transformation script to add profiling of any code outside OpenACC regions
- Configurable – recognised loop types (longitude, latitude, vertical levels, tracers) specified in config file
- Fparser extended to support non-ascii characters
- General restructuring

Performance : NEMO

- Full unmodified code: ORCA1 SI3 GO8 configuration
- Version ~NEMO 4.0
- 227 files, 110,000 lines of code
- PScyclone script: find largest valid regions containing loops – exclude static arrays, derived types, write statements, subroutine calls etc.
- Inserts 3,315 KERNELS directives
- General approach: should work on other configs and different versions of NEMO but not tested



ORCA1-SI3 Performance Evolution - time per step.



ESIWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988




esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER
AND CLIMATE IN EUROPE

Adaptation/Extension : ICON

- Dawn has been extended to cover more complicated stencils found in the ICON dynamical core

```
cellField div_c;
edgeField flux, n;
div_c =
reduce(CELL>EDGE,
flux*n)
```

- **larger neighbourhoods** / "wider" stencils
- weighted reductions (e.g. gradients)
- FD-type stencils




primal location	on_vertices()	on_edges()	on_cells()
vertex			
edge			
cell			



Performance: ICON

- Dawn: first version of a horizontal diffusion example using a Python frontend.

Timings: Tesla P100, Single GPU, 1.3M Edges, Laplacian + Smagorinsky Diffusion



ESIWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988



esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER
AND CLIMATE IN EUROPE

```
for _ in forward:

    # fill sparse dimension vn vert using the loop concept
    for _ in neighbors[Edge > Cell > Vertex]:
        vn_vert = u_vert * primal_normal_x + v_vert * primal_normal_y

    # dvt_tang for smagorinsky
    dvt_tang = reduce(
        (u_vert * dual_normal_x) + (v_vert * dual_normal_y),
        Edge > Cell > Vertex,
        [-1.0, 1.0, 0.0, 0.0],
    )

    dvt_tang = dvt_tang * tangent_orientation

    # dvt_norm for smagorinsky
    dvt_norm = reduce(
        u_vert * dual_normal_x + v_vert * dual_normal_y,
        Edge > Cell > Vertex,
        [0.0, 0.0, -1.0, 1.0],
    )

    # compute smagorinsky
    kh_smag_1 = reduce(vn_vert, Edge > Cell > Vertex, [-1.0, 1.0, 0.0, 0.0])

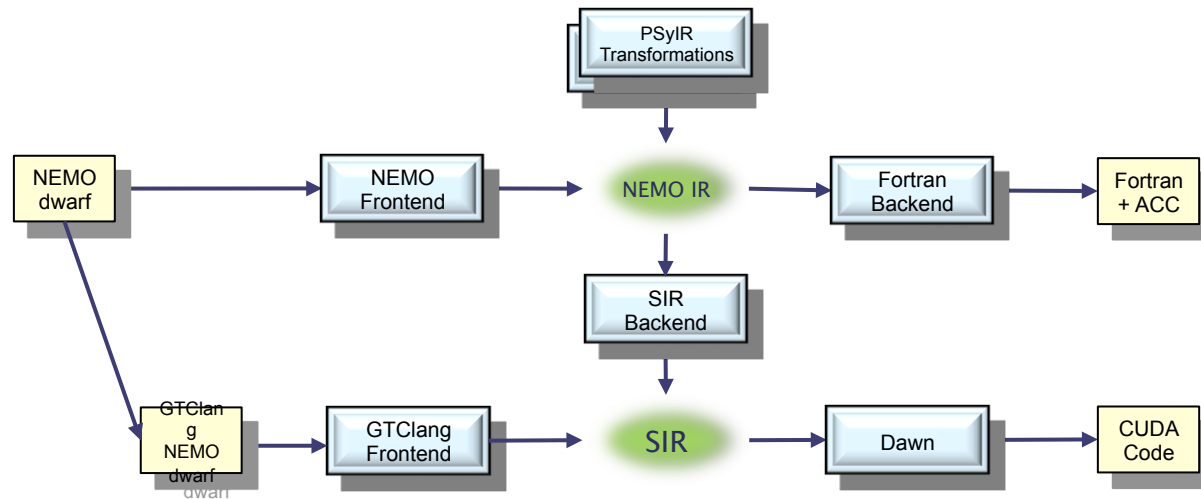
    kh_smag_1 = (kh_smag_1 * tangent_orientation * inv_primal_edge_length) + (
        dvt_norm * inv_vert_vert_length
    )

    ...
```


Interoperability

- Prototype NEMO IR to SIR backend written
- Allows PSystem front-end to use Dawn back-end
- Works for simple examples
- Working towards supporting NEMO dwarf
 - Intrinsic (sign, abs, ...)
 - Implicit loops ($a(:) = b(:) * c(:)$)
 - ...

Interoperability



Summary

- Introduced 2 DSL implementations : DAWN and PSyclone
- These separate the science from its parallelisation and optimisation
- Initial comparisons between DSLs will be via the NEMO dwarf benchmark
- Good progress has been made adapting DAWN and PSyclone for ICON and NEMO respectively
- Performance improvement shown for ICON horizontal diffusion example cf existing code.
- NEMO ocean-only performance currently equivalent to CPU, NEMO ocean + sea-ice performance improving.
- Interoperability of DSLs via their Intermediate Representations shows promise.



Thank you

ESIWACE2 has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 823988



esiwace
CENTRE OF EXCELLENCE IN SIMULATION OF WEATHER
AND CLIMATE IN EUROPE

GOcean DSL : OpenCL

- EuroExa project
- Translate Fortran kernels to OpenCL
- Bind Fortran to OpenCL via wrapper layer
- Tested on Nemolite2d Fortran benchmark
- Can automatically generate OpenCL version and run on an FPGA
- OpenCL allows running on CPU and GPU as well

