Science and Technology Facilities Council

Recent Advances in PSyclone … the PSyIR

Rupert Ford, Andrew Porter, Sergi Siso, STFC Hartree Centre
Iva Kavcic, Chris Maynard, UK Met Office
Joerg Henrichs, Australian Bureau of Meteorology

Hartree Centre
Science & Technology Facilities Council

Met Office

Australian Government
Bureau of Meteorology

6th ENES HPC Workshop Tuesday 26th May 2020

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

PSyclone

PSyclone 1.9.0
BSD 3-clause
https://github.com/stfc/PSyclone
https://psyclone.readthedocs.io
> pip install psyclone

- 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

Levels of Abstraction

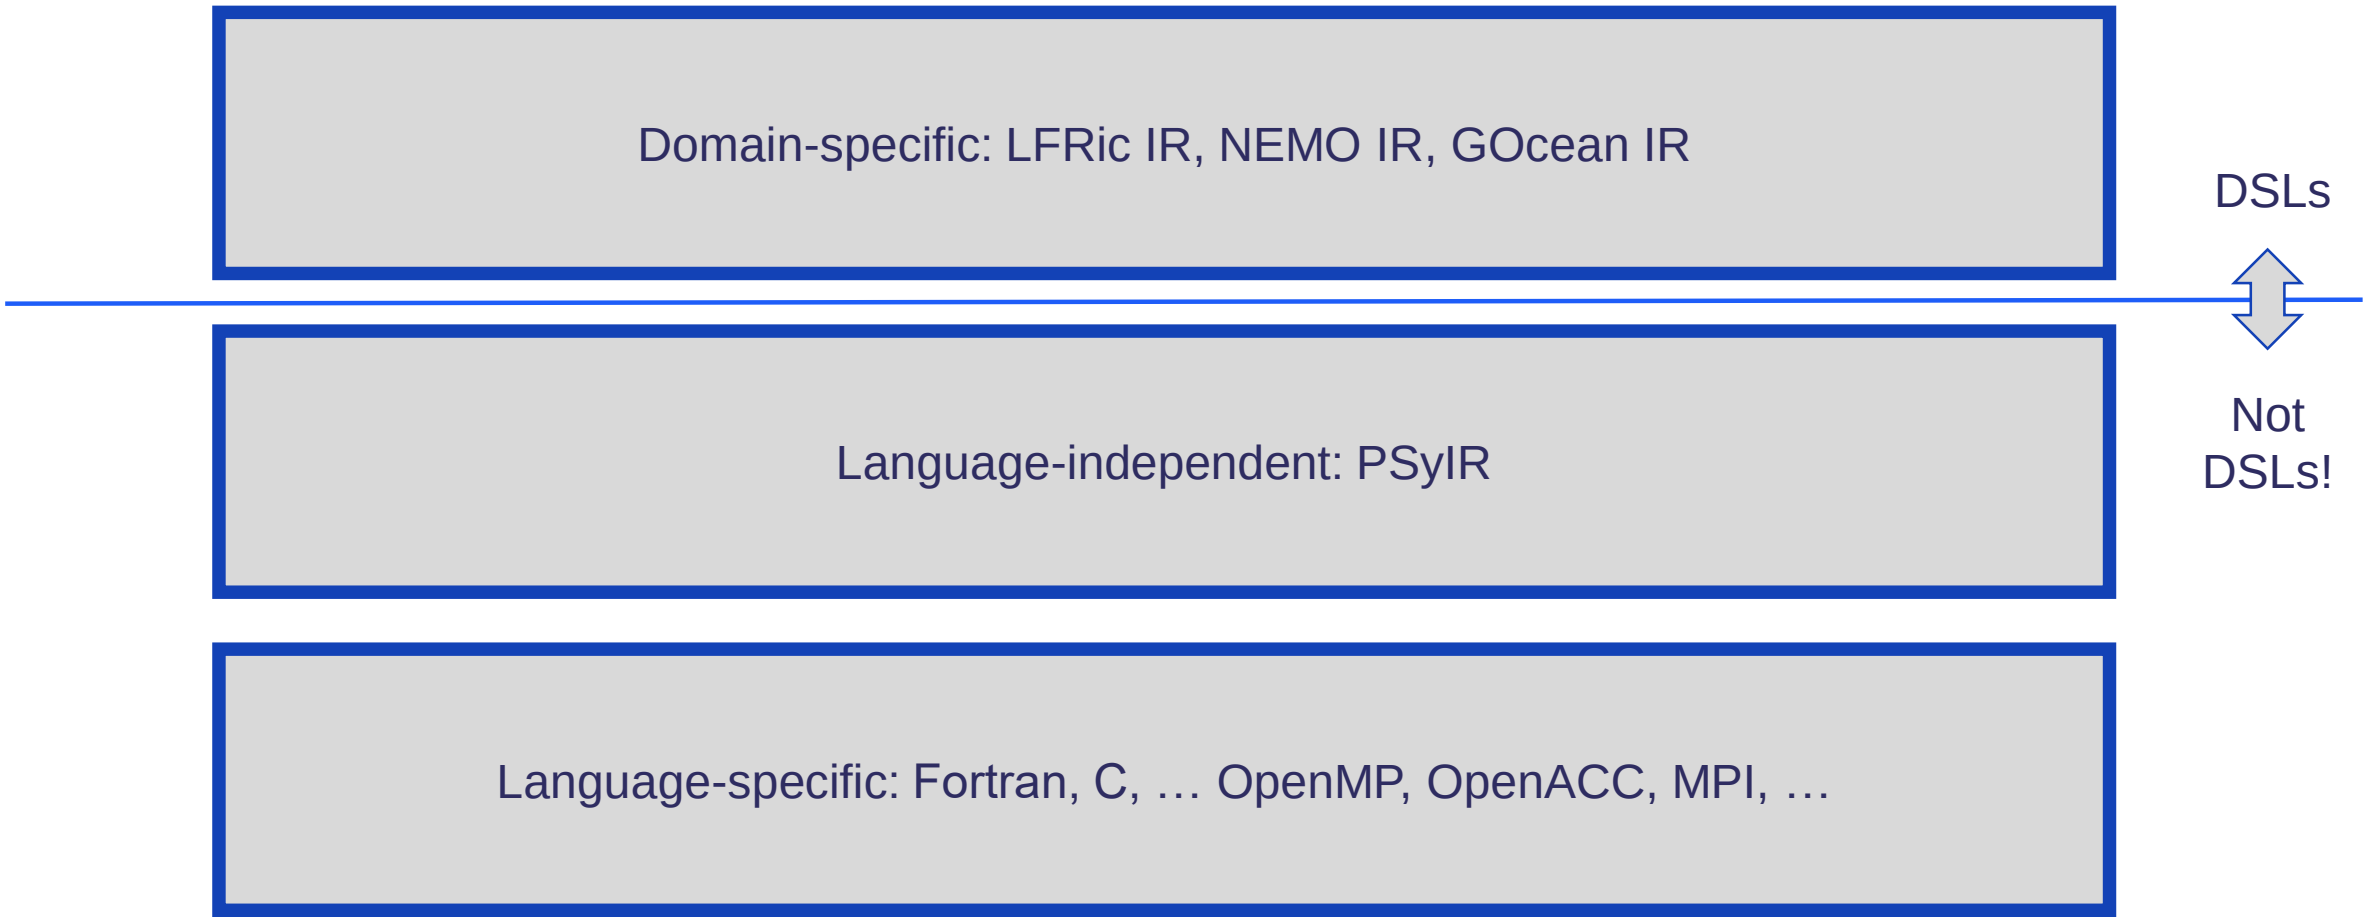

Fparser

- 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

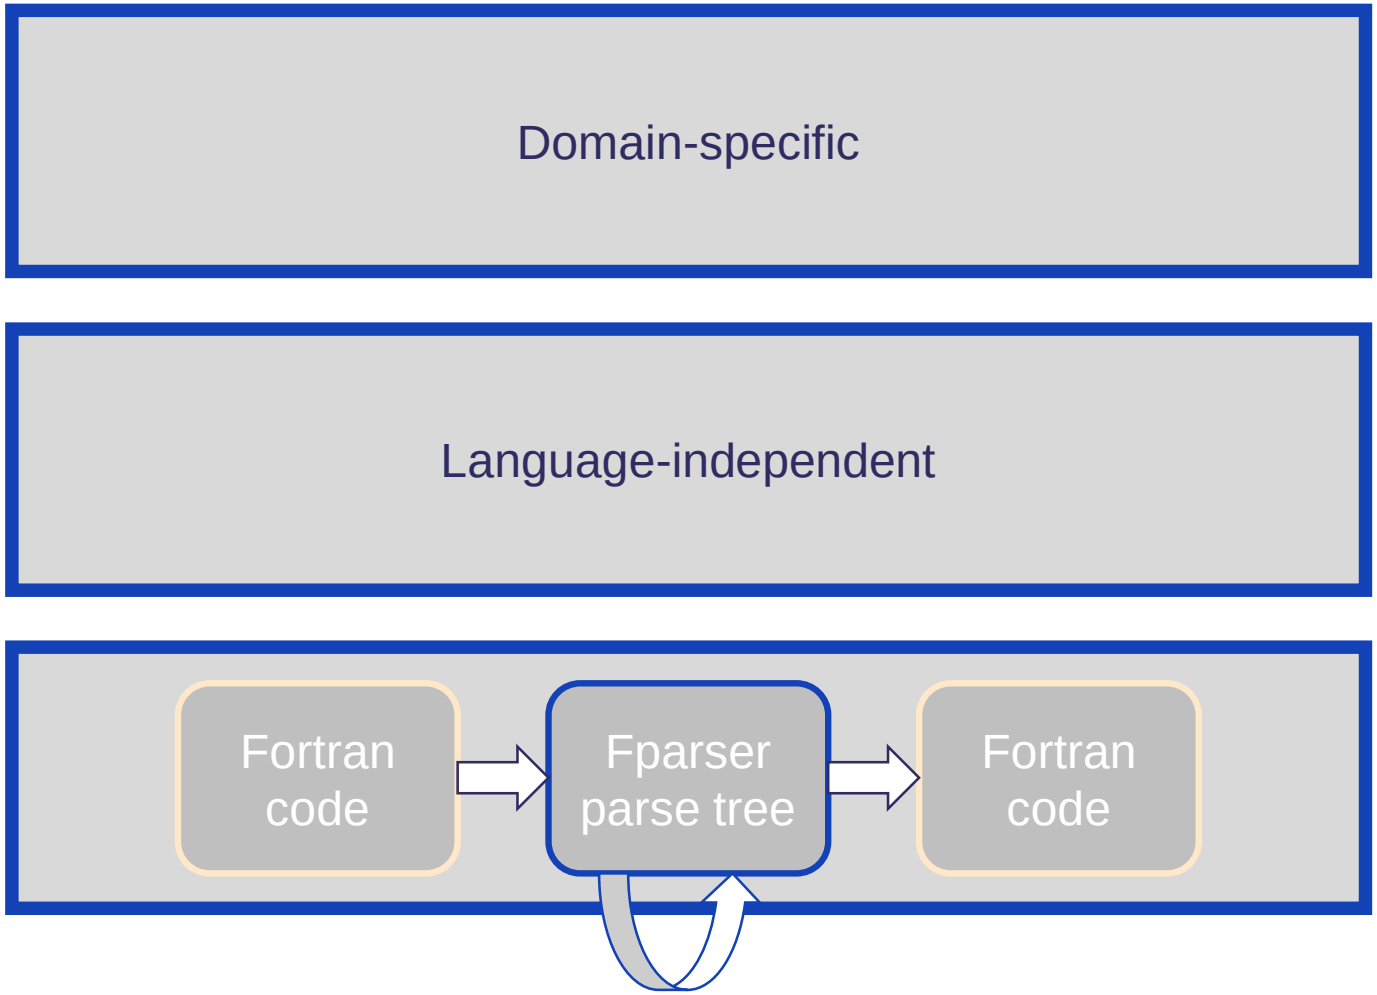

Science and Technology Facilities Council

LFRic : Separation of Concerns

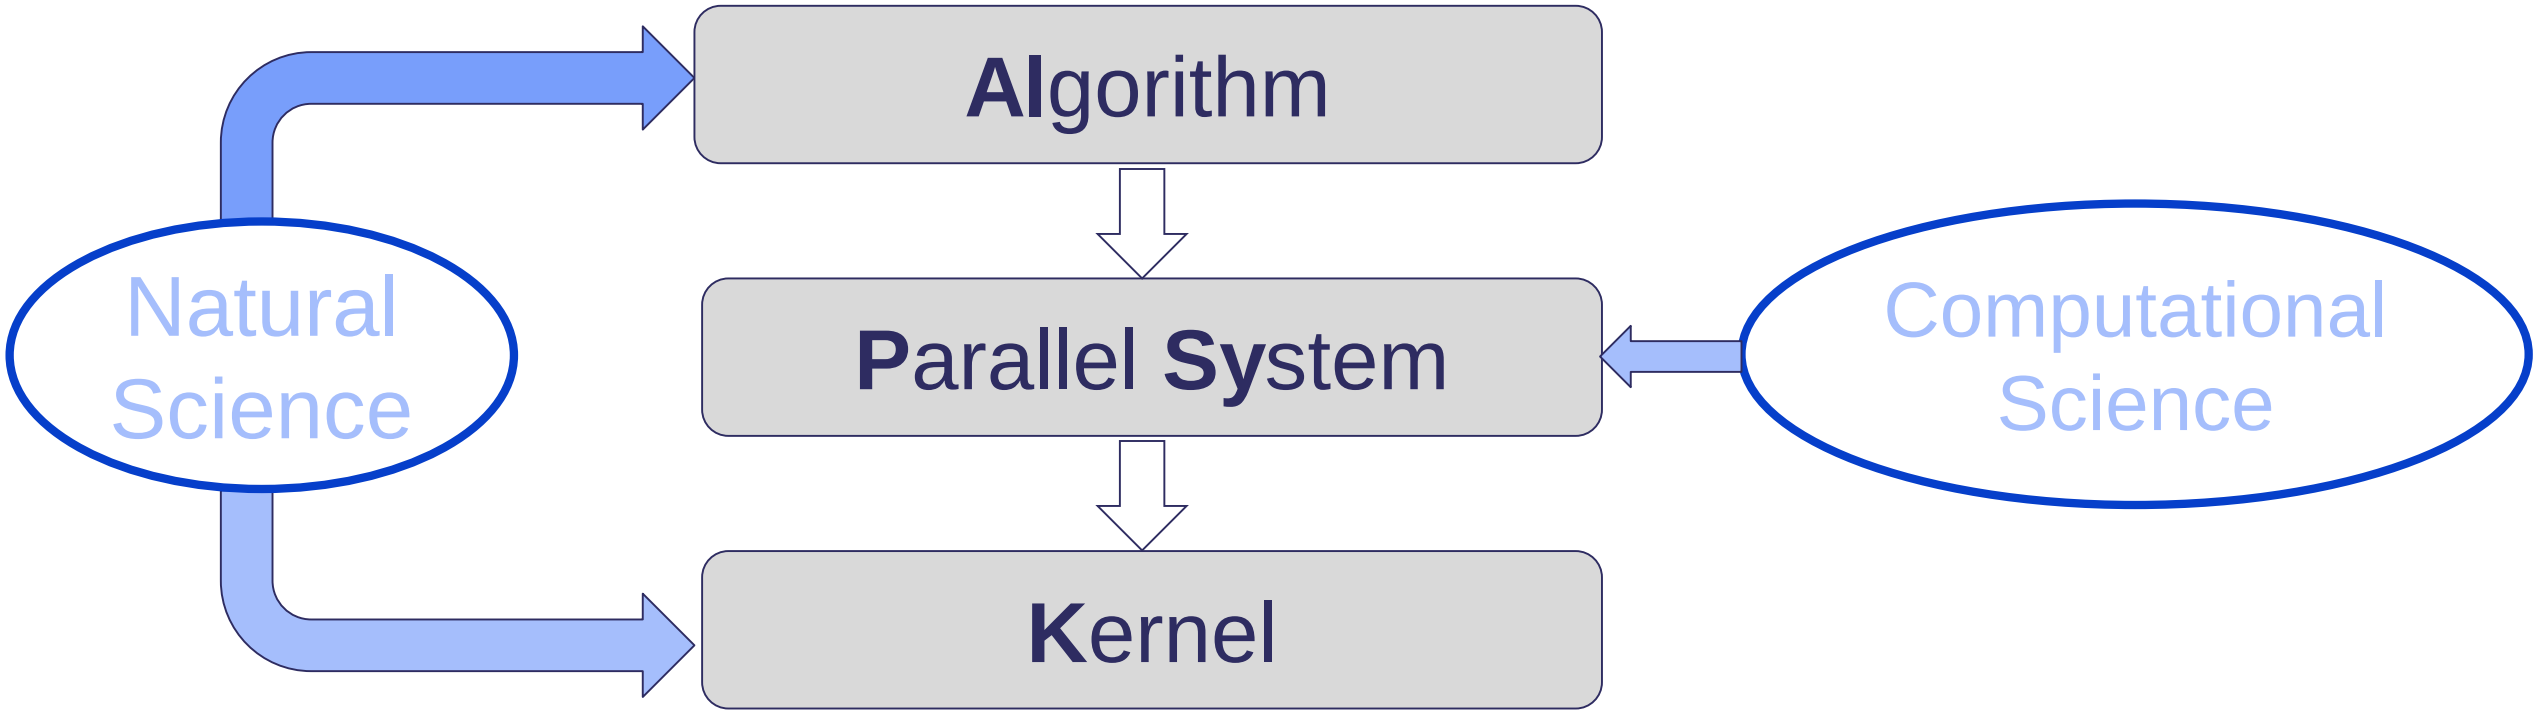

Science and Technology Facilities Council

PSyKAl : Separate the Natural Science from the Computational Science (performance)

LFRic DSL PSy layer

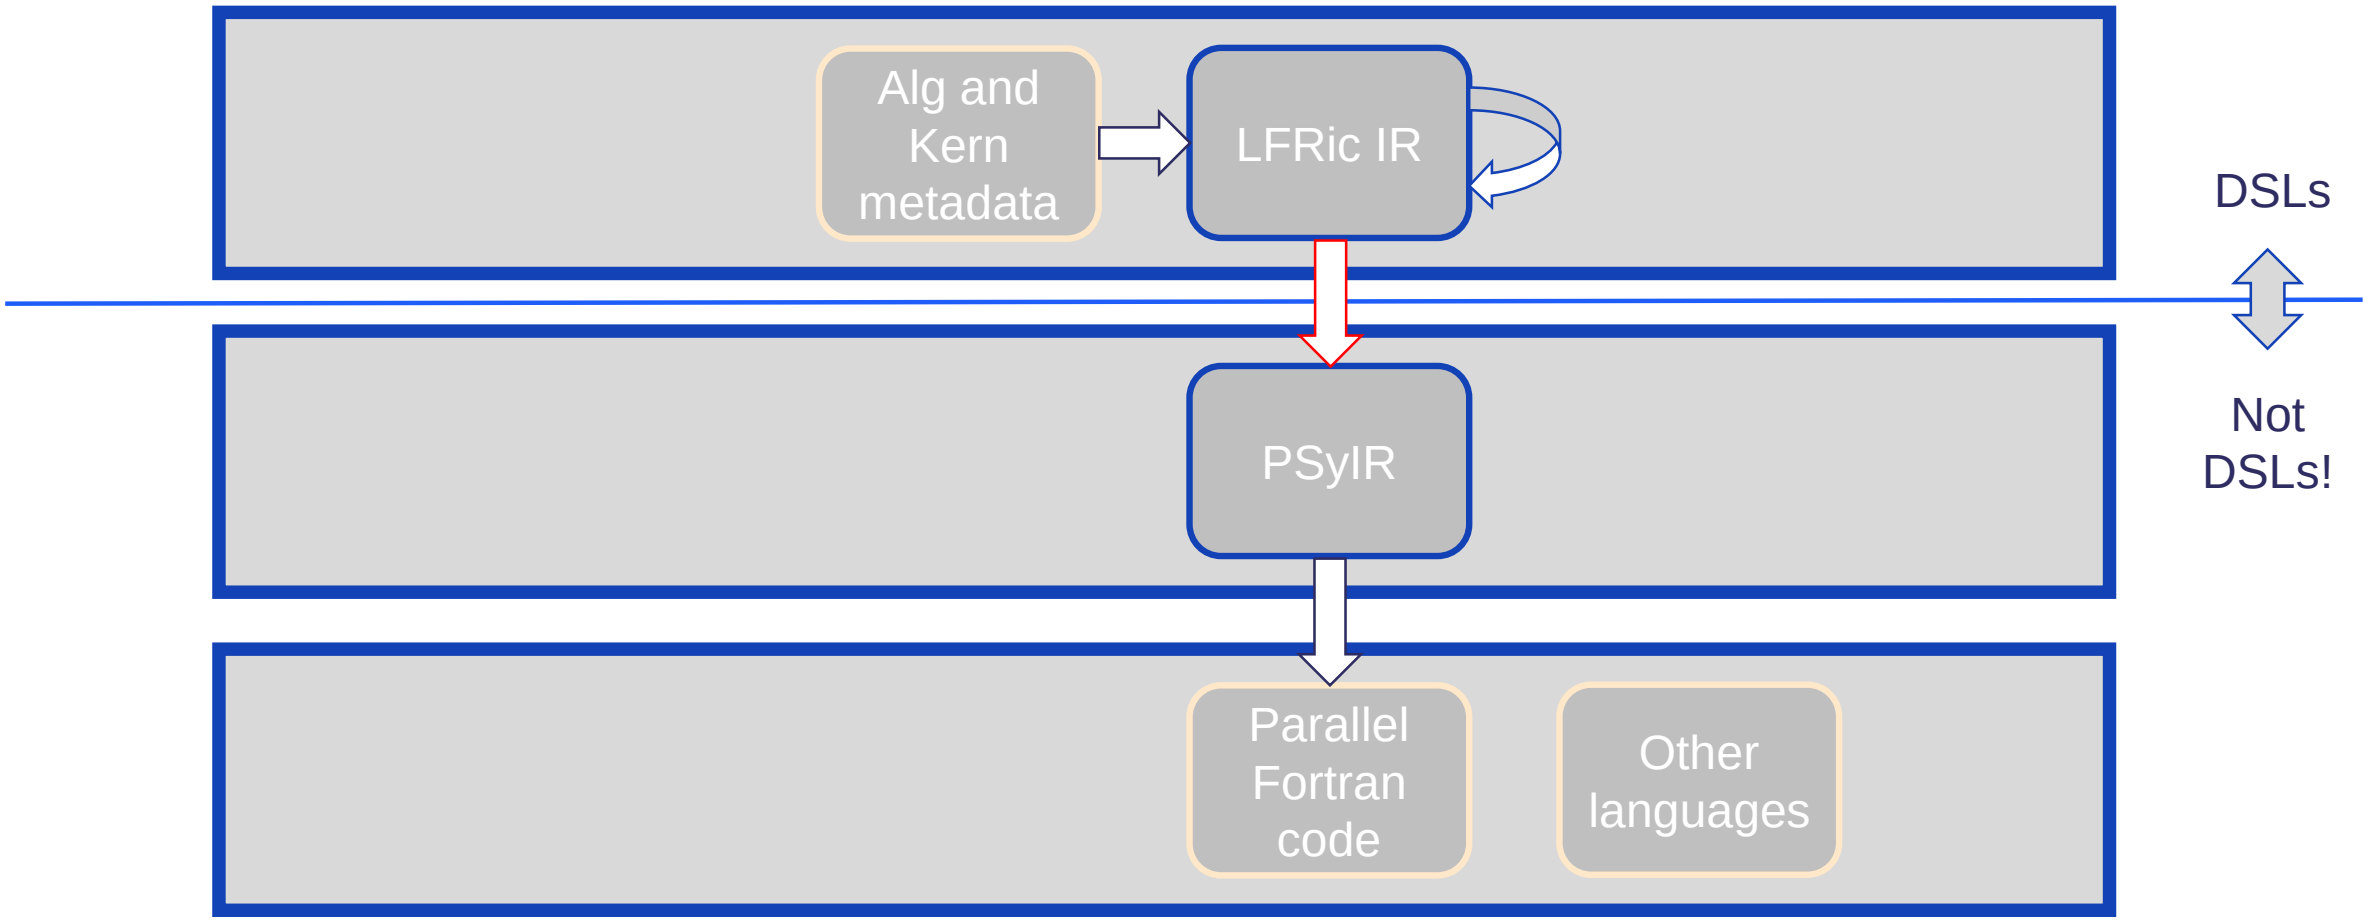

Science and Technology Facilities Council

Met Office

LFRic DSL PSy + Kern layer

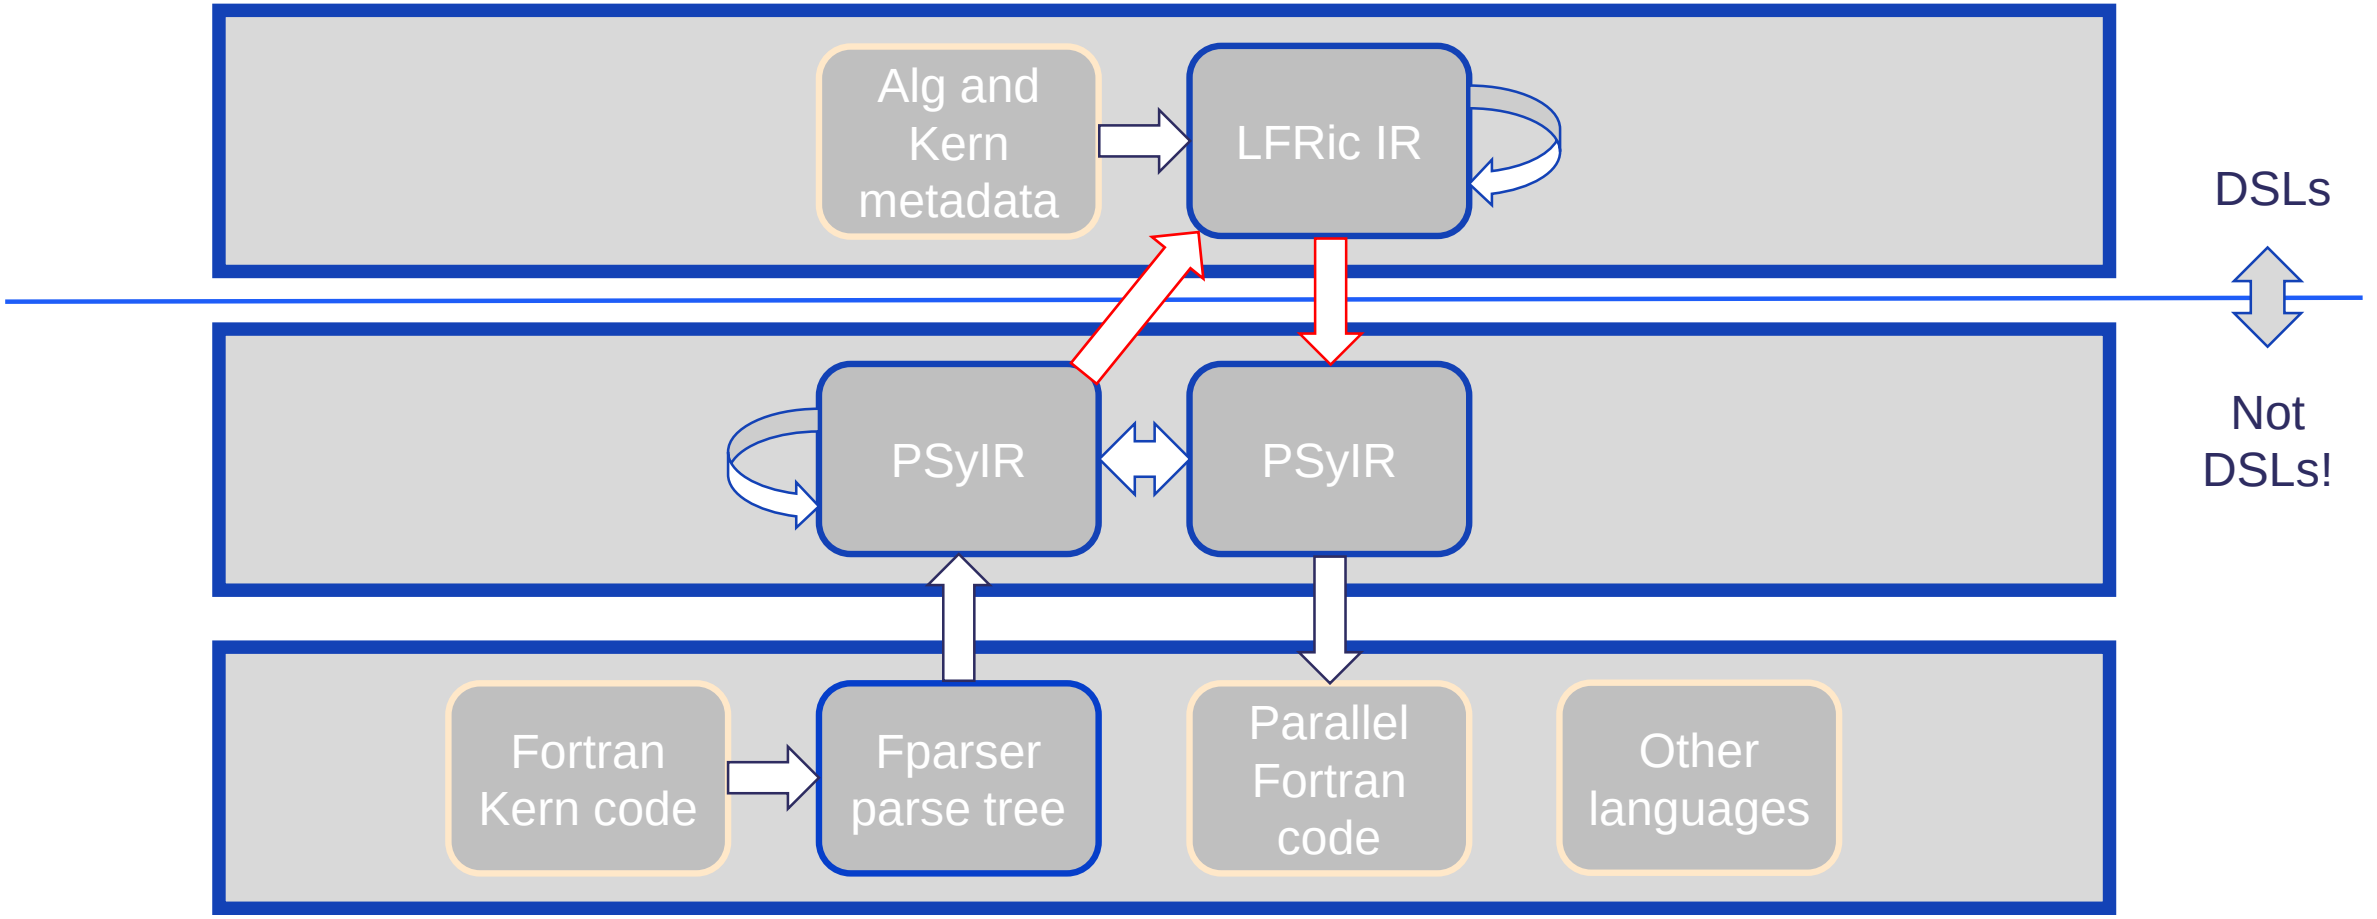

Science and Technology Facilities Council

Met Office

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

LFRic DSL Alg layer

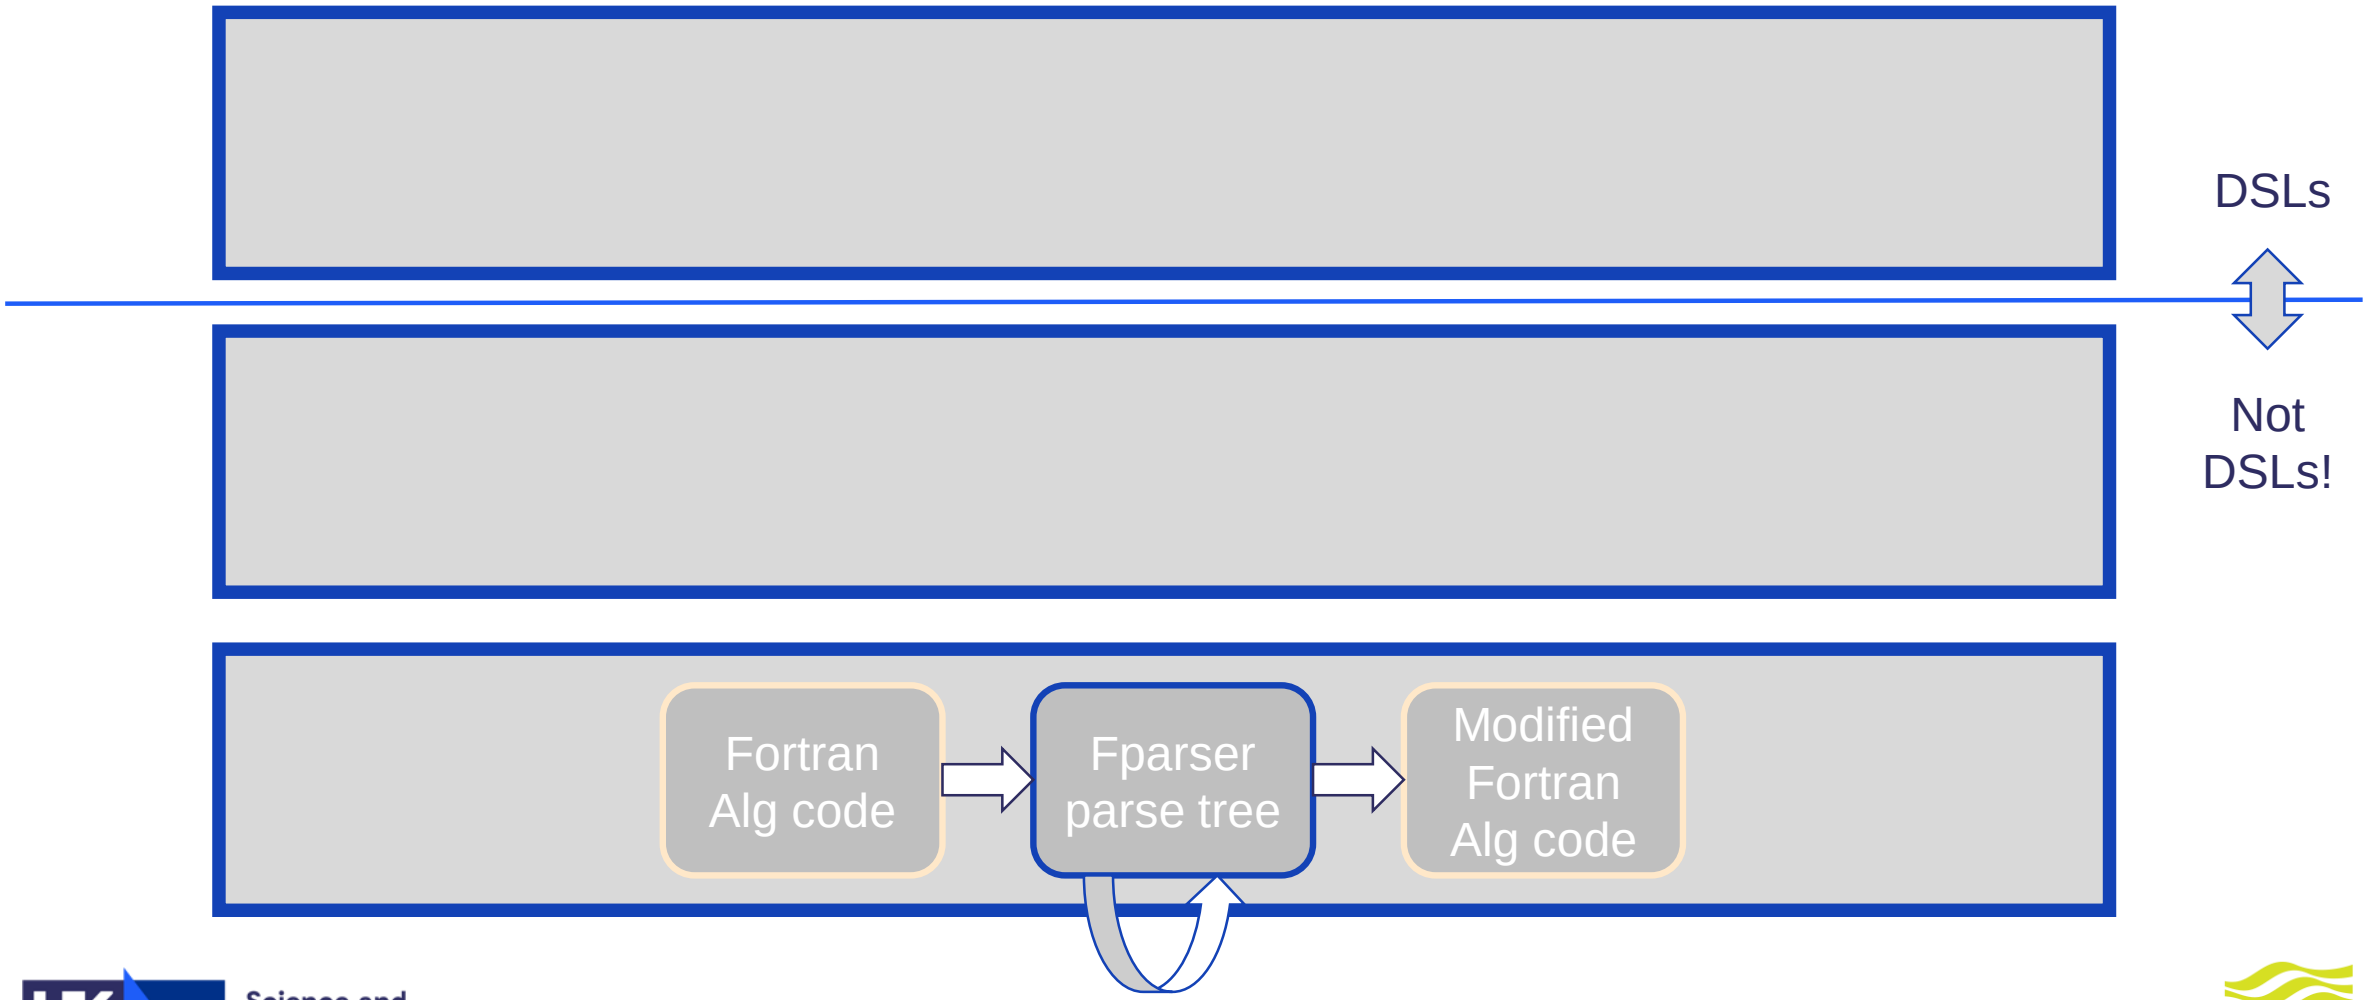

Science and Technology Facilities Council

Met Office

NEMO DSL

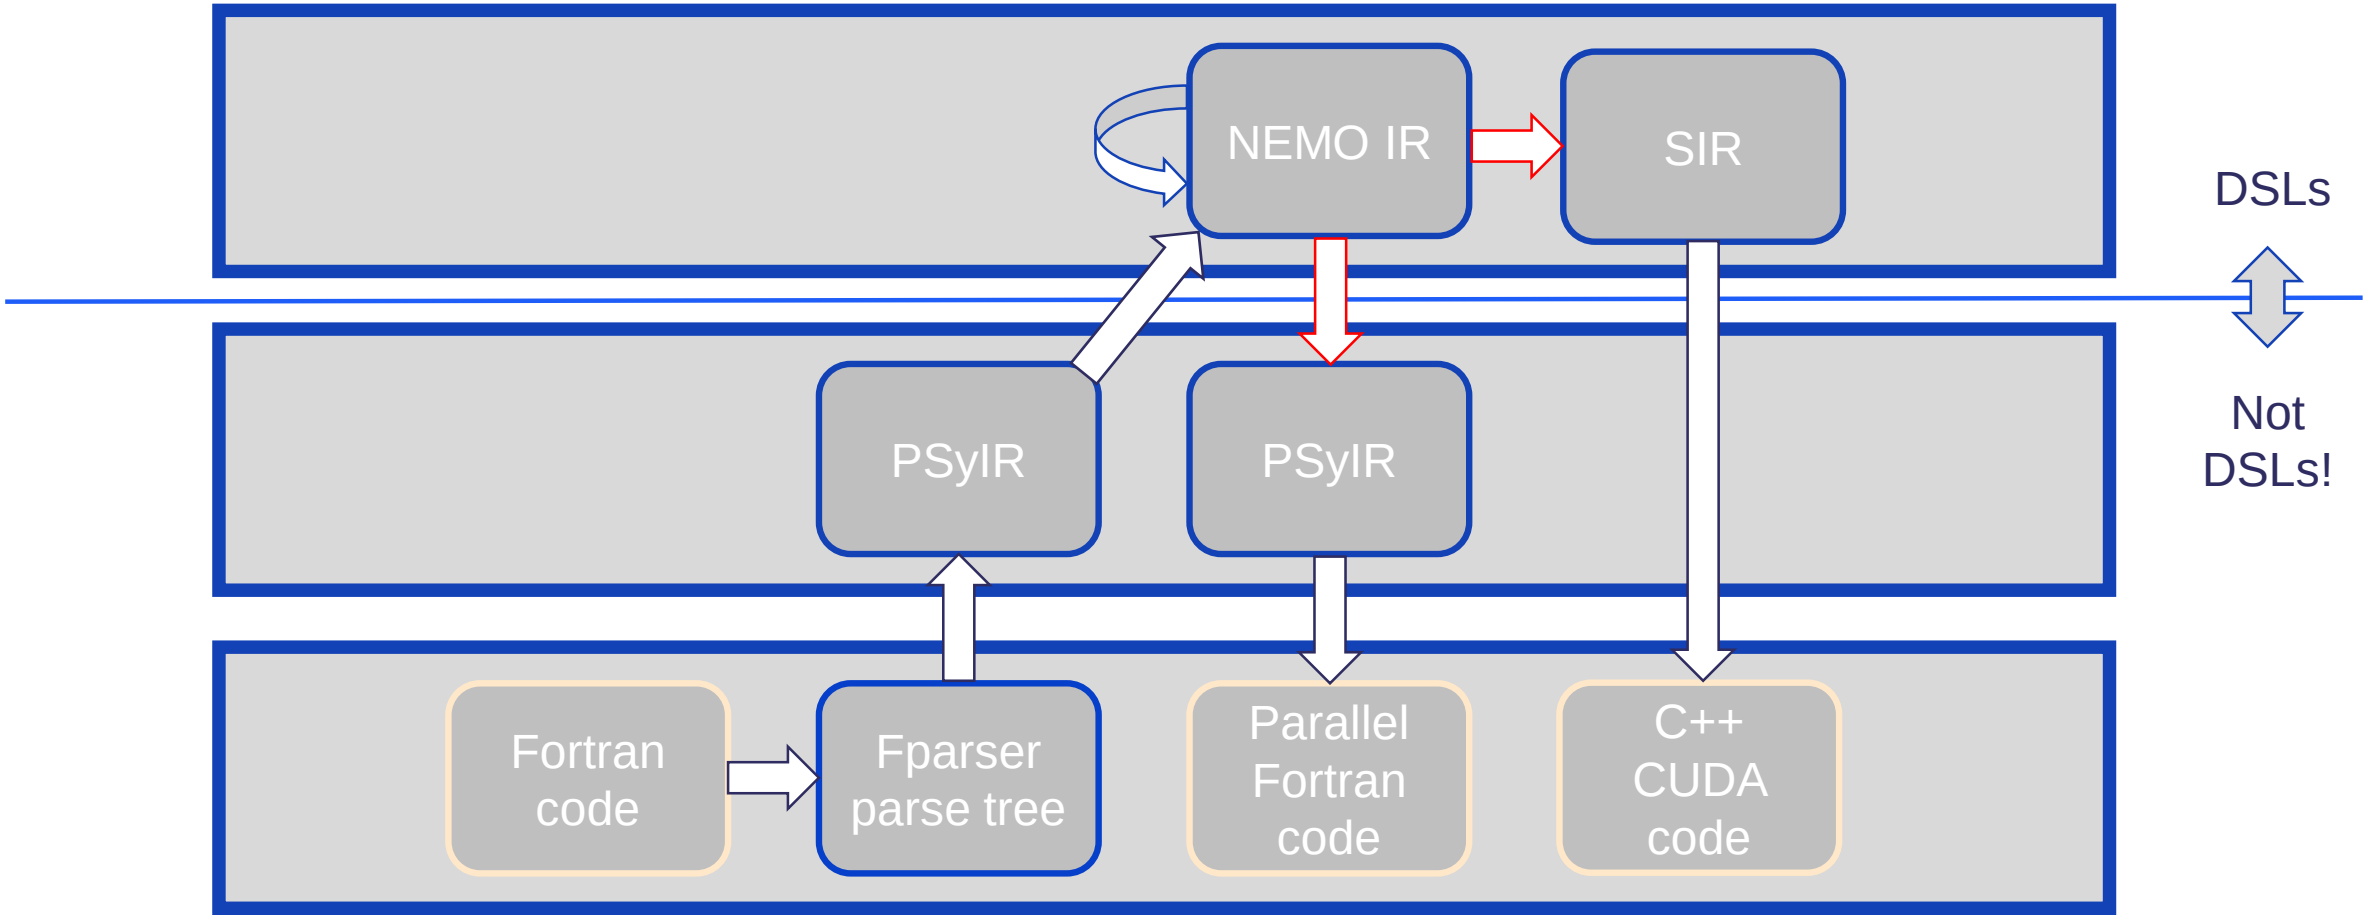

Science and Technology Facilities Council

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

NEMO GPU Status

- Full unmodified code: ORCA1 SI3 GO8 configuration
- Version ~NEMO 4.0
- 227 files, 110,000 lines of code
- PSyclone 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

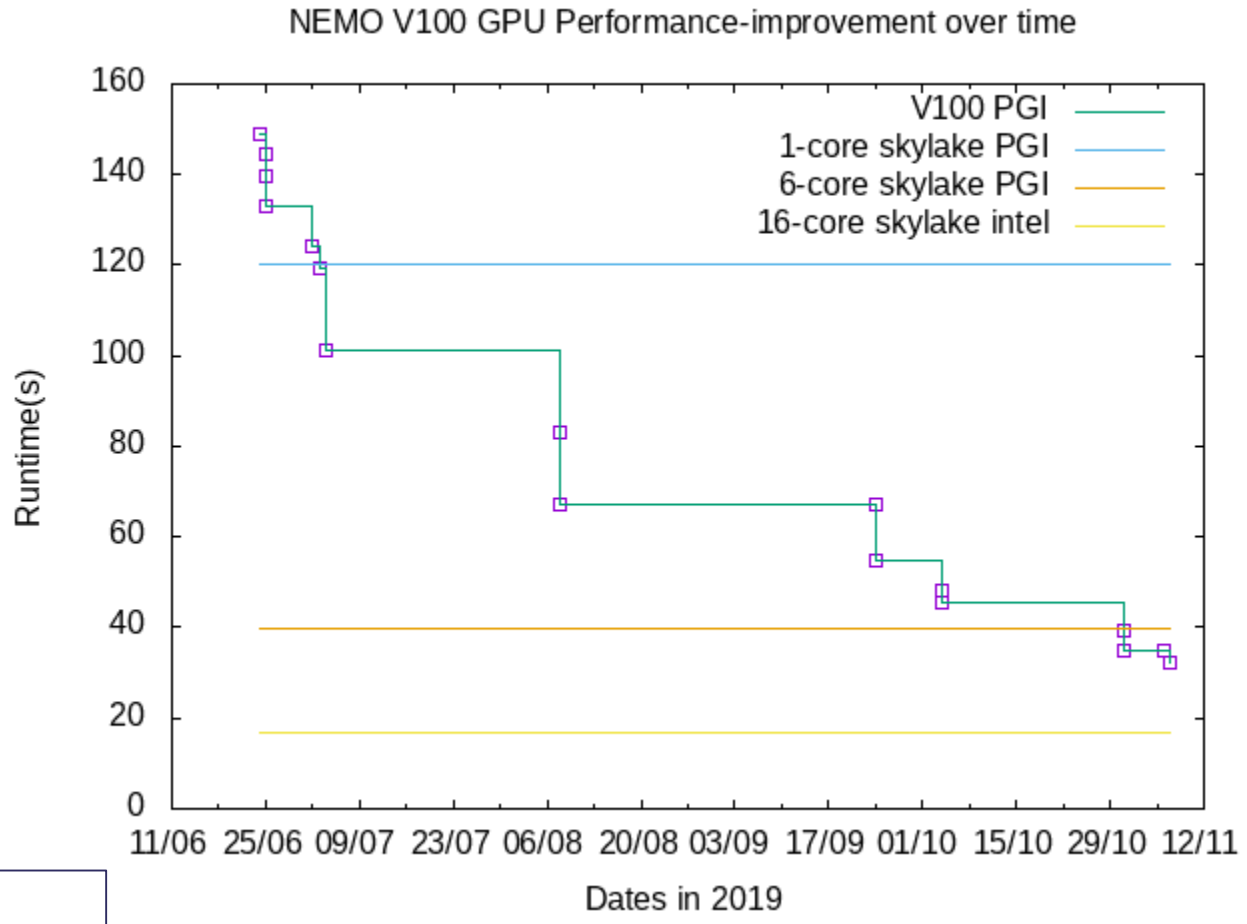

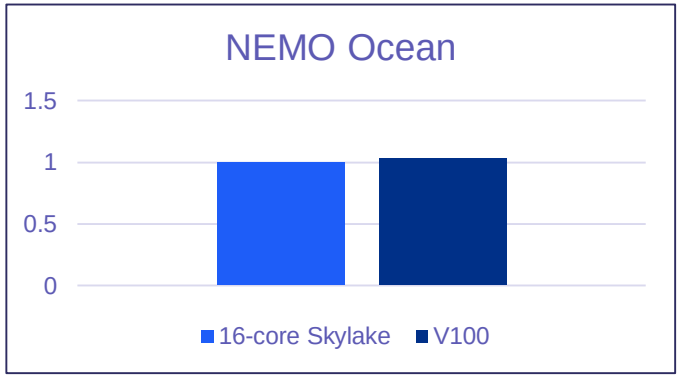

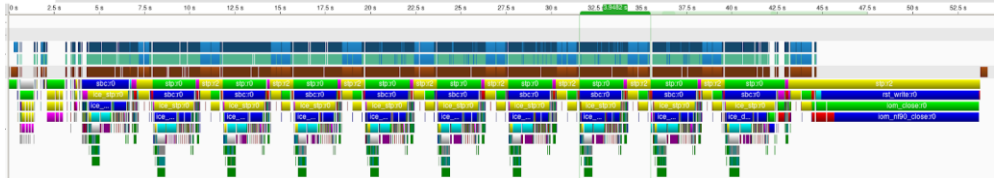

Met Office NVIDIA Hartree Centre Science & Technology Facilities Council

UKRI Science and Technology Facilities Council

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

NEMO: DSL Comparison and Interoperability

- Planning comparisons of the DSL approaches with UofM and CMCC

- Working towards translating the NEMO Dwarf to SIR
- Large amount of the code can be translated, but issues remain
- Generates 1,130 lines of SIR

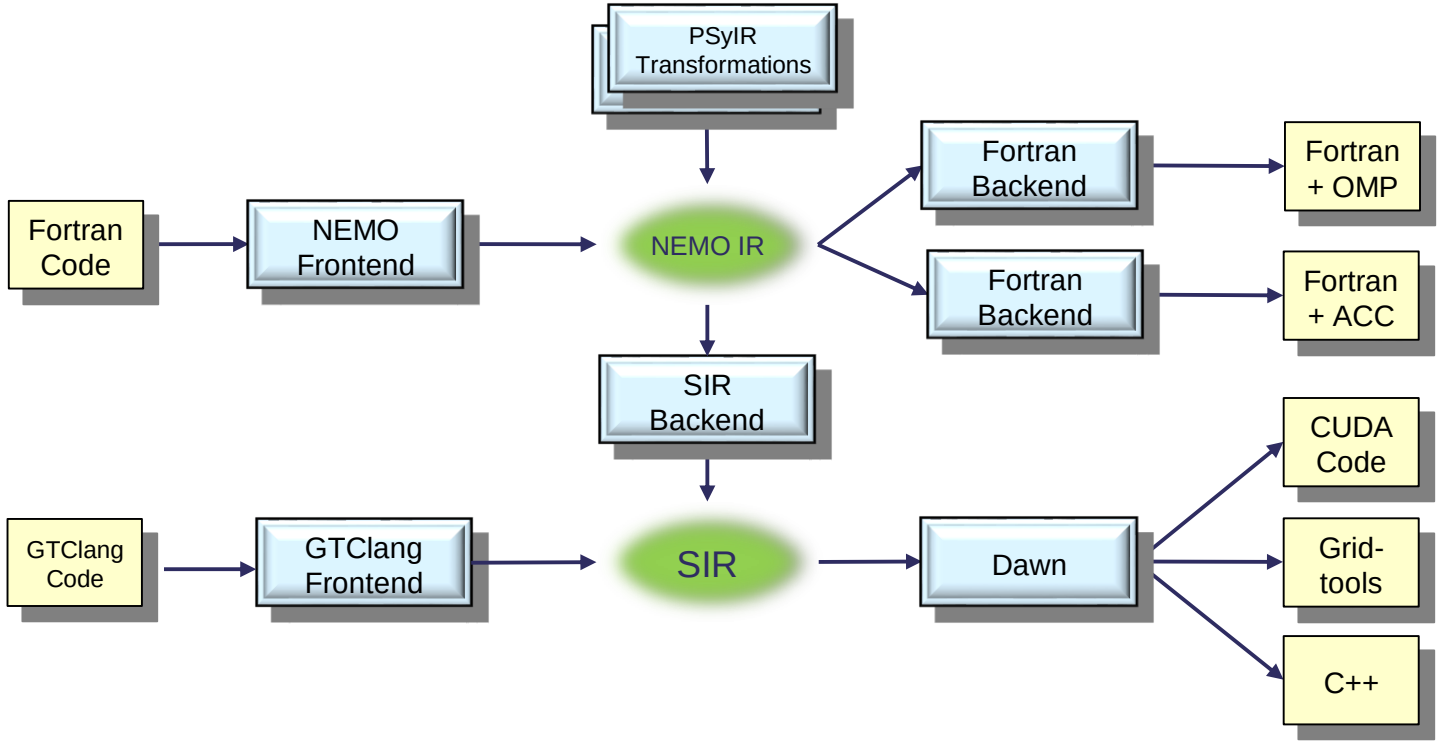

MeteoSwiss

Hartree Centre
Science & Technology Facilities Council

Science and Technology Facilities Council

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

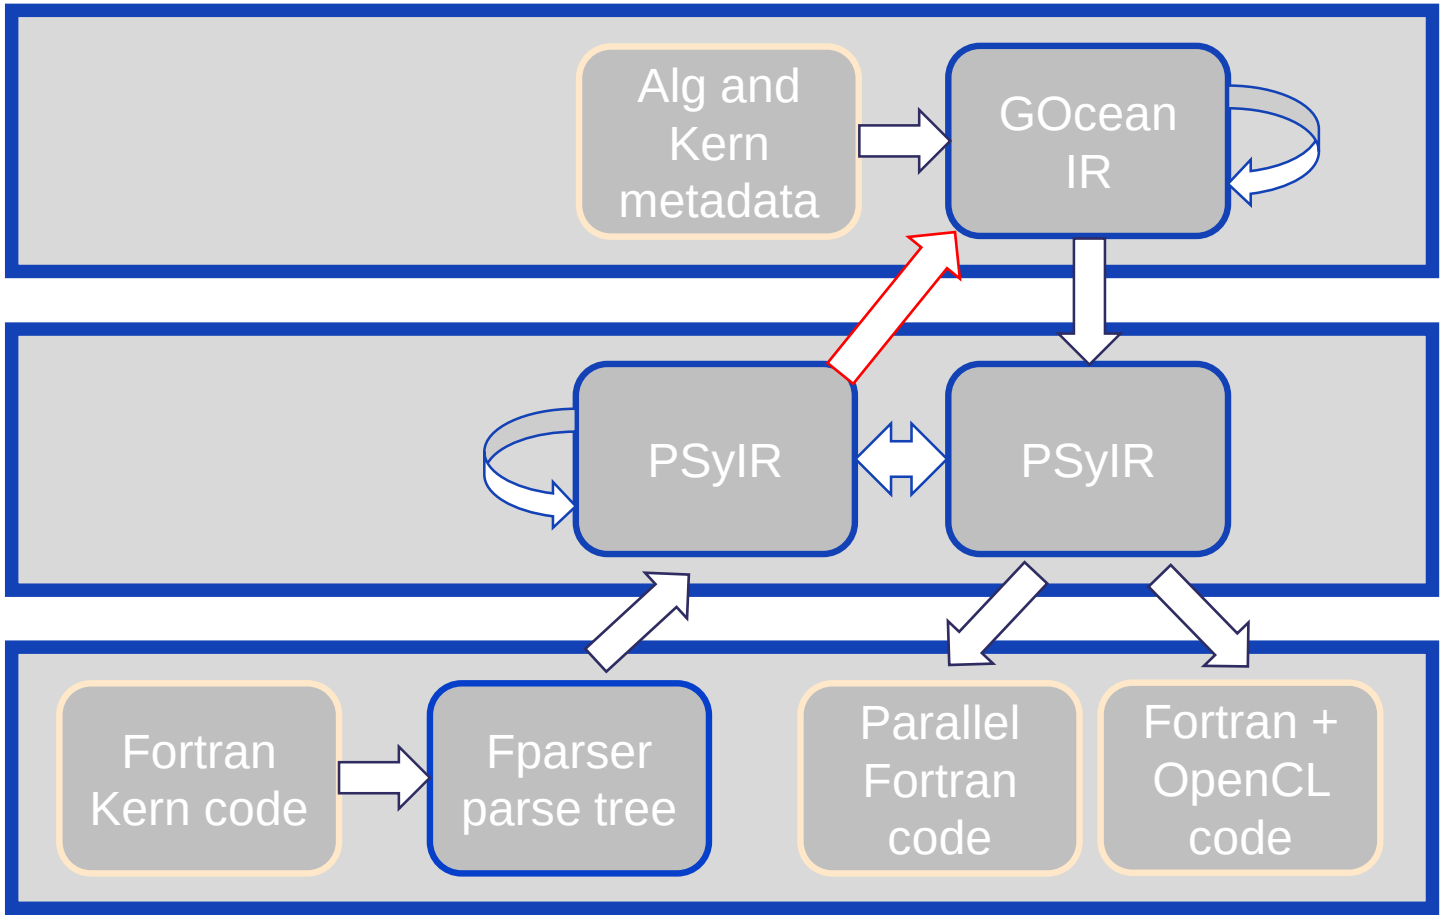

Science and Technology Facilities Council

EUROEXA has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 754337

Summary and Future work

PSyIR is a language independent representation
Specialise PSyIR for particular domain
Opens ability to translate to different languages

Continue PSyclone development
Comparison and interoperability between PSyclone and Gridtools
Collaboration with ECP – use their software stack e.g. Kokkos, SYCL, Legion
Excalibur funding – further develop NEMO for GPUs

Science and Technology Facilities Council

Thank you

Facebook: Science and Technology Facilities Council

Twitter:@STFC_matters

YouTube: Science and Technology Facilities Council