




Science and  
Technology  
Facilities Council

Hartree Centre

# Performance Portability for Existing Weather & Climate Models using PSyclone Application to the NEMO Ocean Model

Chris Dearden and **Andy Porter**, STFC Hartree Centre  
Wayne Gaudin, NVIDIA



# Acknowledgements



PSyclone is developed by a growing number of people:

Aidan Chalk, Rupert Ford, Andy Porter, Sergi Siso, **STFC Hartree Centre**  
Andy Coughtrie, Iva Kavcic and Chris Maynard, **UK Met Office**  
Joerg Henrichs, **Australian Bureau of Meteorology**

The work described here has received direct support from:



Hartree Centre

# Overview

1. The Problem: Performance, Portability and Productivity
2. The PSyclone tool
3. PSyclone: revolution versus evolution
4. Results of Applying PSyclone to NEMO
5. Next Steps

# The Problem

- 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

- 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 (arguably)
  - Work round limitations/bugs
  - Optimisations encoded as a 'recipe' rather than baked into the scientific source code
  - Different recipes for different computer architectures
  - Enables scriptable, whole-code optimisation
  - Support for profiling, debugging and kernel extraction

# PSyclone: Two Modes of Operation

## Revolution

Process code [written in a DSL](#)

Currently [two Domains](#) supported:

- [LFRic](#) - Mixed finite elements, mesh unstructured in horizontal, structured in vertical, embedded in Fortran
- [GOcean](#) - DSL for 2D, finite difference, stretched, structured grid, embedded in Fortran

## Evolution

Process [existing code](#) that follows strict coding conventions

Recognise certain code structures and construct higher-level Internal Representation

Transformations applied to this IR

In development for [NEMO](#) (plus associated models, e.g. SI3, MEDUSA). Also applied to [ROMS](#).

# PSyclone: Two Modes of Operation

## Revolution

Process code **written in a DSL**

Currently **two Domains** supported:

- **LFRic** - Mixed finite elements, mesh unstructured in horizontal, structured in vertical, embedded in Fortran
- **GOcean** - DSL for 2D, finite difference, stretched, structured grid, embedded in Fortran

## Evolution

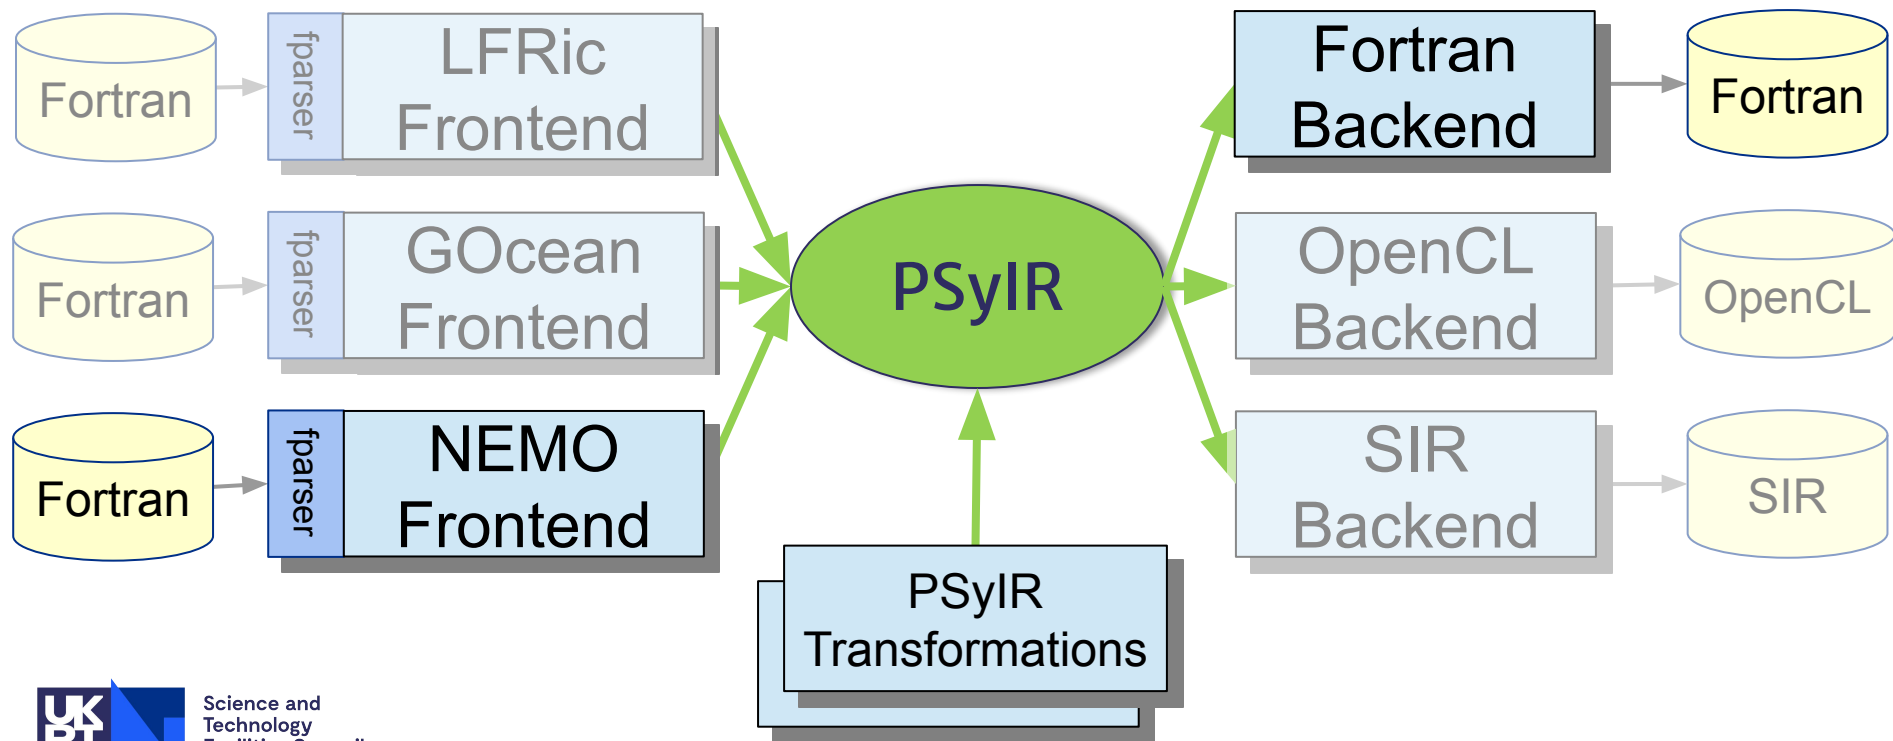
Process **existing code** that follows strict coding conventions

Recognise certain code structures and construct higher-level Internal Representation


Transformations applied to this IR

In development for **NEMO** (plus associated models, e.g. SI3, MEDUSA). Also applied to **ROMS**.


# PSyclone: Basic Architecture








- **Finite-difference** model using a tripolar, stretched latitude, longitude mesh ('ORCA')
- Three core components:
  - NEMO-OPA: ocean dynamics, thermodynamics
  - NEMO-SI3: sea-ice (thermo)dynamics, brine inclusions...
  - NEMO-TOP/PISCES: tracer transport and biogeochemistry
- Mesh rotated so that poles are over land
  - Can go to high resolution without the 'pole problem'
  - Relatively large (core of ~100K lines of Fortran)




# Results - ORCA1 ocean, OpenACC

Single GPU performance of ORCA1 NEMO-OCE since May 2020



Using 4.0.2 of  
NEMO & G08  
configuration  
provided by Mirek  
Andrejczuk, Met  
Office

# Results - ORCA1 ocean, GPU




# Results - ORCA12 ocean, GPU + MPI

Large-scale resources  
accessed through  
**ESiWACE2**

Run on up to **192 GPUs**  
on **Marconi** (V100) and  
**JUWELS Booster**  
(A100)

Scaling performance still  
under investigation



**CAUTION**




**PEOPLE  
WORKING**




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

# Ongoing - ORCA1 ocean+sea ice, GPU



# Results - ORCA12 ocean, OpenMP+MPI



Loops over **vertical levels** parallelised using OpenMP.  
More work to be done...

# Next Steps

- Other GPU hardware (OpenMP Offload)
- Optimise OpenMP CPU performance
- Data Movement
  - Explicit management instead of unified memory
- Investigate applicability to NEMOVAR and WAVEWATCH III
- Adjoint Generation

# Summary

- **PSyclone** is a tool for code-generation and transformation
- Aimed at the HPC expert
- Supports revolution (e.g. **LFRic**) and evolution (e.g. **NEMO**)
- Used with NEMO to add **OpenACC** or **OpenMP** directives
- Good single GPU performance obtained for NEMO
- Work in progress on multi-GPU performance
- Work in progress on application to **SI3**, **MEDUSA** and **NEMOVAR**
- Support for **OpenMP offload** in development





Science and  
Technology  
Facilities Council



# Thank you

User, Developer and Reference Guides are available:

[psyclone\[-dev,-ref\].readthedocs.io](https://psyclone[-dev,-ref].readthedocs.io)

For more information please contact:

[rupert.ford@stfc.ac.uk](mailto:rupert.ford@stfc.ac.uk)

[andrew.porter@stfc.ac.uk](mailto:andrew.porter@stfc.ac.uk)



Australian Government  
Bureau of Meteorology

