



## HPC WG June 2019

Work in progress or planned



# Contents

- Purposes of this review (we've done a stock-take recently!)
- Inter-node communications
- Mixed precision calculation
- Horizontal tiling and OpenMP directives
- Loop fusion and other single node optimisations
- Adaptation for GPUs
- Exploration of DSL approaches (Escape-2, STFC)

# Purposes of review

- The purpose of the group is co-ordination
- It takes concentration to work well together
  
- Things have evolved quite a bit since the stock-take
- The code base in 2019 is evolving rapidly (e.g. variable renaming as part of IMMERSE project by Dave Storkey & Andrew Coward)
- There are several separate EC projects (EsiWACE2, Escape-2, EuroExa, IMMERSE, IS-ENES2, PRACE)
  
- As group leader I'm not sure I'm up-to-date.
- Quite likely that some others aren't either
- Opportunity to get a consistent picture

# Inter-node communications

## Work in 2018

- The last stock-take (by Seb & Eric) itemised a lot of progress
- Including major improvements to north-pole fold bottleneck
- This all went into version 4.0

## Work in 2019

- Multi-grid-point halos **underpin major opportunities**
  - >1 step between exchanges in sea-ice and barotropic solvers (using halos of 3, 5 or 7 points)
  - Huge reduction in places where halo exchange is needed
  - Horizontal tiling and loop fusion depend on it
- Italo & Seb have started work and will co-ordinate; Italo has MUSCL working with 2-point grid halo
- Seb points out that **consequences for restart & diagnostic files & users** is a nettle that needs to be grasped (work for NEMO System Team?)
- Introduction of MPI3 collective neighbours (EasiNESS, CMCC)

# Mixed precision calculations

- Oriol has visited Gurvan and Seb. Outcomes:
  - Oriol's results seem to match "expectations" based on physical / numerical insights
  - It is easy to include a uniform precision default option
  - There is no reason not to include the code in next release
- A large number of declaration statements must be changed
  - Oriol will agree starting point for his implementation with DS & AC
  - ECMWF have a new recruit starting 1<sup>st</sup> October who will contribute

# Horizontal tiling & OpenMP

- Eric Maisonnave has shown that horizontal tiling should reduce the cost of 3D subroutines by about a factor of 2
- Horizontal tiling is much less complicated than 3D tiling
- It can be implemented across all 3D routines so is better than 3D tiles
- Gurvan has proposed that 2D / 3D loops be replaced by loop macros
- It makes the code slightly simpler and tiling easier to implement
- Andrew Coward has written a perl script to implement these macros
- Tiling will be driven at the level of stp.
- We expect to be able to implement OpenMP directives efficiently and effectively at this level
- Some experimentation will be needed
- This work requires multi-point halos (can't get far without them)
- Storage of diagnostics needs to be clarified (Seb to discuss with Yann whether can XIOS do it)
- Daley Calvert can work on this with assistance from Maff Glover, Gurvan and AC
- Erwan is working on a proof of concept for coarse grain OpenMP to avoid problems with MPI

# Loop fusion

- Italo has proposed several ideas for re-organising calculations
- He's also suggested that some code segments be re-factored into macros
- Andrew Coward has volunteered to take responsibility for reviewing these proposals

# Adaptation for GPUs

- Wayne Gaudian & Dave Norton of Nvidia led a hackathon at the Met Office in Dec 2018 to kick-start work
- A mini-app calculating horizontal pressure gradients (and later isopycnal diffusion) was ported to Isambard GPUs
- “Good” performance was obtained using 2-3 approaches with OpenACC directives
- Mirek Andrejczuk has confirmed that Jeremy Appleyard’s NEMO3.5 port to GPUs gives “good” performance
- In most routine the directives are not very intrusive
- It’s possible to devise simple rules on where to put them
- Andy Porter, Mirek & Wayne are working together on an automated implementation using Psyclone (next slide)
- There are alternative strategies for implementation possible (e.g. OpenMP or OpenACC; hard-code in NEMO or use Psyclone pre-processor)



# Exploration of DSL tools

- Domain Specific Languages (DSLs) (an awesome idea)
- Basic ideas:
  - adaptation of code for a given HPC is usually done by a **small number** of “transformations” of it
  - Our codes can be written in an **intermediate language** (HIR) which will allow these transformations to be done
  - A suite of transformations for finite volume codes has already been developed
- ESCAPE-2 includes work to translate NEMO mini-apps into an HIR and test performance (Italo, BSC, Nils Wedi)
- Andrew Porter has developed a Psychone front-end for NEMO; it converts NEMO into the Psychone IR (so no manual re-writing necessary)
- it can then be translated into a standard HIR; or have OpenACC directives added to it
- The NEMO Psychone front-end could be used as a NEMO **style checker**