NEMO: Future Look

Jason Holt, Adrian New, Hedong Liu, James Harle - NOC
John Siddorn, Richard Wood, Mike Bell – Met Office
Stephen Pickles, Mike Ashworth - STFC

Science & Technology
Facilities Council

Next Generation Ocean Dynamical Core Roadmap Project

Key science drivers for NEMO

- To improve the representation of diapycnal mixing
  - reduce numerical and improve physical mixing
  - address through ALE and advanced advection
- To accurately model the mesoscale and effectively parameterise the submesoscale
  - Refining resolution
  - Scale dependent subgrid-scale parameterisations (e.g. ocean-shelf)
- Maintain efficient lower resolution models (e.g. for ensembles, ecology etc)
- To improve the representation of coastal, shelf seas and key small-scale topographic features
  - in basin and global scale models
  - in regional (1km's) and local models (100m's)
  - through better process representation, resolution and parameterization,
  - towards multi-scale modelling approaches.
- To develop an easy/rapid relocation capability

Characteristics of a good ocean model:

- Have good discrete dissipation properties: i.e. minimal numerical diffusion;
- Give conservative solutions for mass, tracer, energy, momentum, PV and perhaps Enstrophy
- Have good discrete wave dispersion properties: e.g.. numerical modes are controlled
- **Be computationally efficient (for large and small models) and across a range of parallel architectures**
- Fit coastline at least as well as quadrilaterals, **ideally as well as triangles**
- Have flexible vertical coordinates/methods (**including ALE**)
- Have **multi-scale capability and geometric flexibility**
- Be accurate in 'realistic' and idealised test cases
- Be portable, adaptable and usable as a community model
- Be supplied with a set of tools that allows efficient handling of the increasingly large amounts of model output, including facilities for runtime calculation of derived statistics such as time and spatial means.

C-grid on quads (eg NEMO) meets many of these…. possible issues in bold

Scales: the motivation for finer resolution

Log10(km)

Cumulative distributions

1st Rossby Rad
- •Eddies
- •Fronts
- •Coastal upwelling
- •Internal Tides

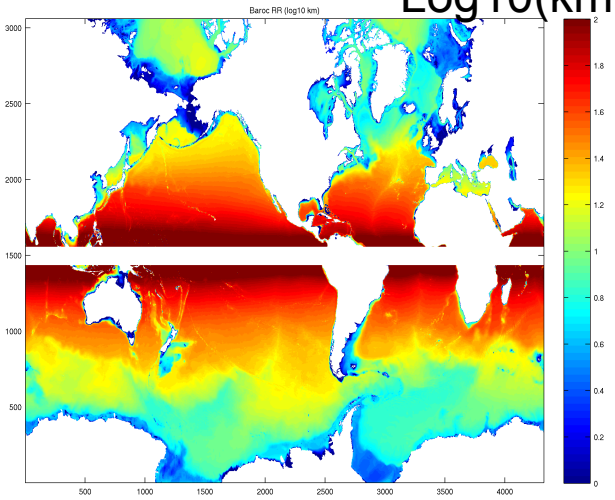

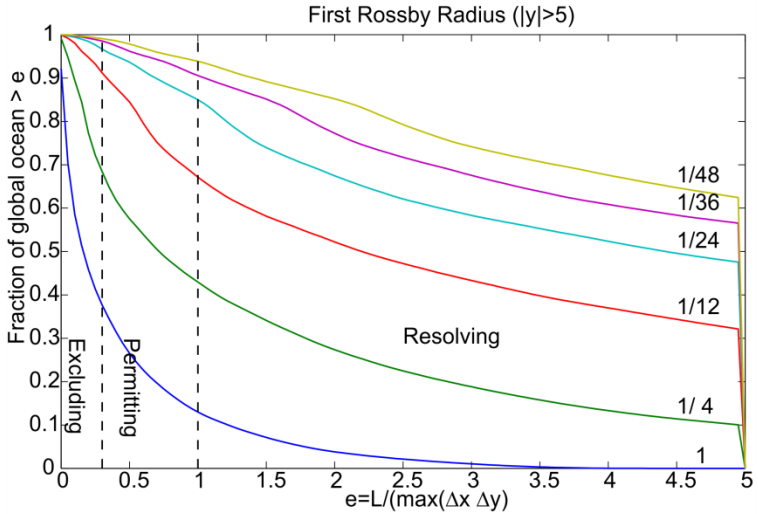

H/(dH/dx)
Scale of topographically steered barotropic currents

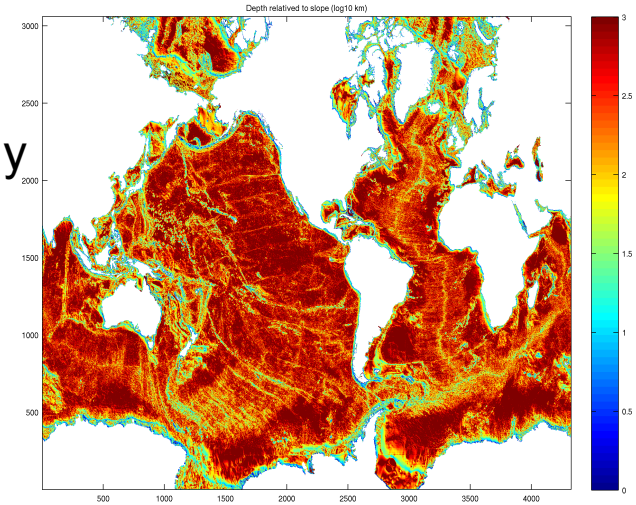

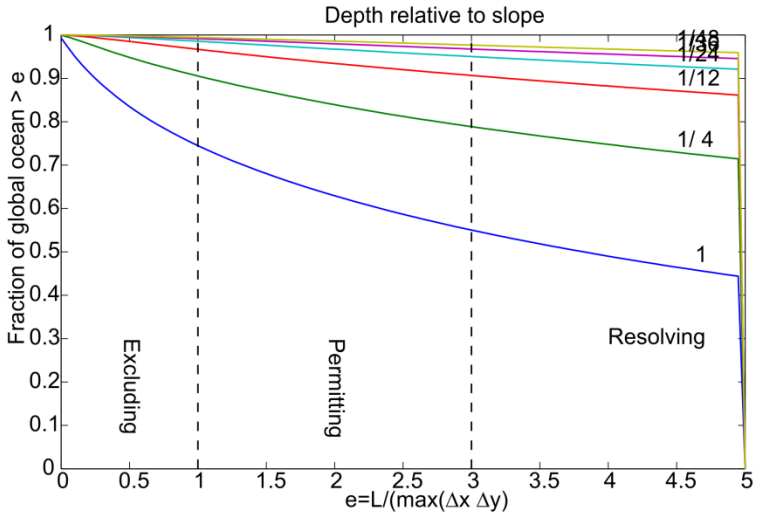

ΔXΔY based on scaling ORCA083

Regionally specific resolution dependence

Global – Basin – Ocean Margin

Ocean Margin - Shelf - Coastal

1º

1/4º

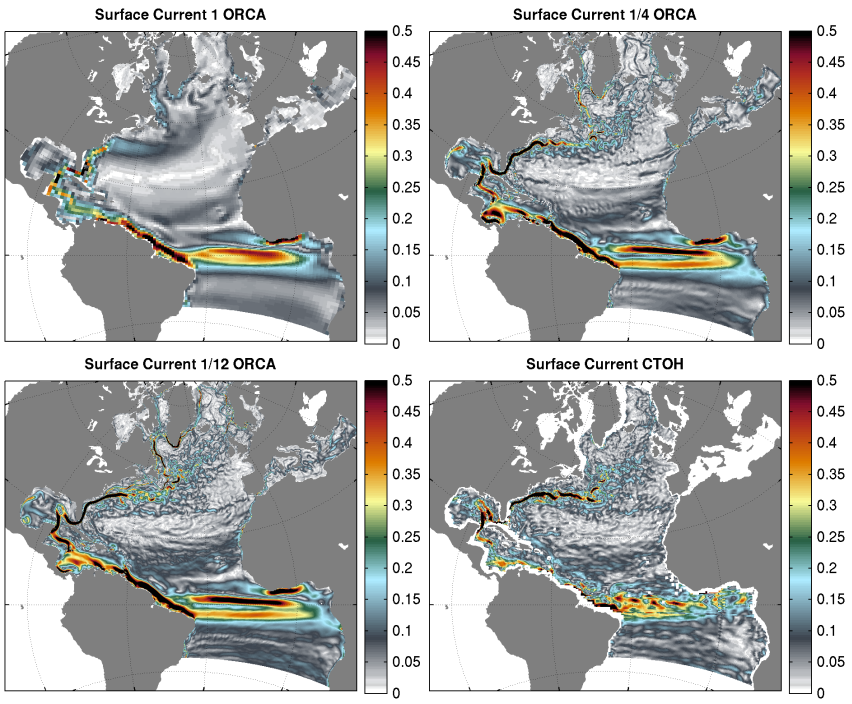

1/12º

NEMO

Obs

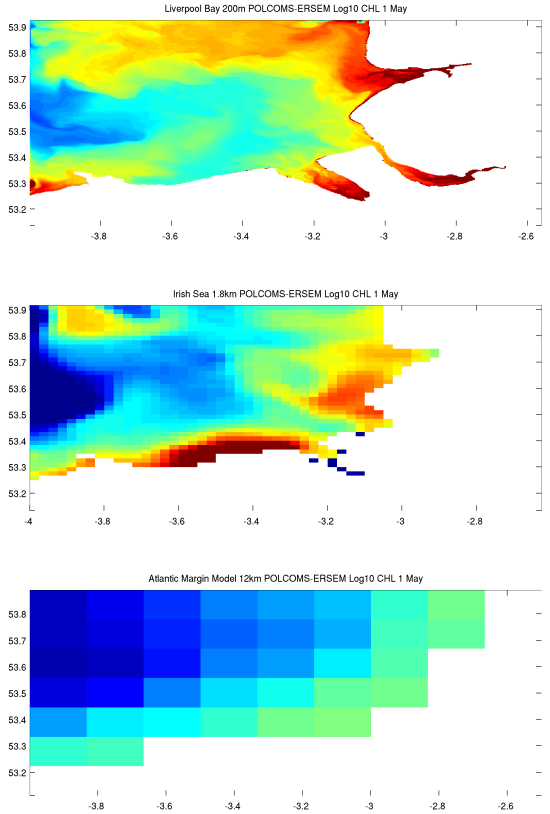

180m

1.8km

12km

POLCOMS-ERSEM

All grid cells were not created equal – the driver for multi-scale modelling (more later)

Process representation and geographic detail: related through non-linear interactions

Computational Challenge

•Increases in computer performance are now occurring primarily through increased parallelism (cores/chip)
•Utilising this growth will be a challenge
•Memory sizes and band widths per core are steadily decreasing
•Need much greater parallelism within the model:
    • eg mixed MPI-OpenMP approach
•Can this be achieved within the current code base?

Incrementally optimising the code without radically restructuring it will seriously hamper the ability of the community to develop it

•Need multi-disciplinary teams to address science and computational issues

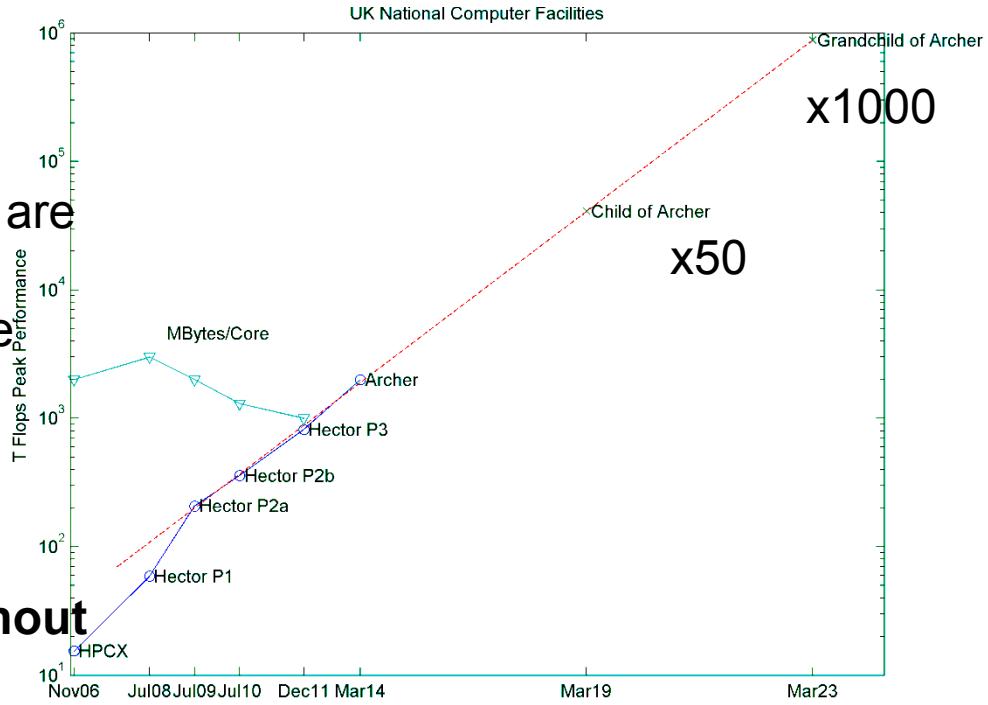

The Risk: If we do nothing?

- Increasingly under-utilise the core's per chip to gain memory bandwidth
- E.g. On Archer next year we may need to under populate by 0.5 to run NEMO AMM optimally
  - This is acceptable to start with, e.g. if resource is increasing
- But cost per physical chip is rising
- This will be unsustainable as new computers are upgraded and core count per chip increases: the resource we have available to utilise (in this manner) will decrease... So the cost (€) per run will increase
- There maybe radical changes in computer architecture that will solve the 'power wall' – but very risky to assume this

However, there is an immense opportunity here if we rise to it.....

Can the ocean modelling community match Moore's law?

Options for the future

1. Development of the current NEMO code base

- A ‘science neutral’ rewrite to address emerging computational issues
  - E.g. Introduce mixed MPI/openMP
  - Ideally separate science and computer science layers
  - This would be a radical change requiring substantial effort
- Key science developments (as above)

2. Development of the next generation: NEMO-2

- Separate science and computer science layers
- New multi-scale capability: Finite Volume or Finite Element
  - This would be a radical change requiring substantial effort
- Could draw on the Gung-Ho project
  - Provide infrastructure, CSE Layer, elements of dynamics

Aligning an ocean model with Gung Ho

- Gung Ho is potentially a huge opportunity for ocean model development
  - Levers substantial extra effort
  - Aligns with advanced CFD modelling community
- The needs/challenges of atmospheric modelling differ from oceanographic
  - the pole problem in atmosphere
  - Need for multiscale and coast line fitting in ocean
- Still need to map Model Properties (previous list) against Approaches (FV, FD, FE, element shape) to decide on an optimal choice
  - All choices may not be available (e.g. may be limited to FE)
  - The limitations need to be weighed against the needs of the science drivers (e.g. FE without improved coastline or multiscale would be a bit perverse)
- The science driver of multiscale modelling is the big push here – is this strong enough?
- Next step: NOC/NERC has funding for a workshop in the Autumn to assess the Properties/Approaches matrix

Discussion points...

- What are the scientific drivers for and against finer resolution modelling (at a global and shelf scale)?
- At what point does a truly multiscale approach become worth the effort?
- Can we better utilise the current capability in NEMO
- Is there a the niche for NEMO in the world of near coastal modelling?

Better use of current NEMO capability

- To what extent could the NEMO curvilinear grid be used for local refinement in global (and regional) applications? (Straits, large estuaries,...)?
-
- Global & regional scales: to what extent unstructured grids are more interesting than the AGRIF nesting strategy?
-
- To what extent a generalized AGRIF approach (unlimited numbers of "*AGRIFED subdomains*") could be considered equivalent to an unstructured grid?
-
- The *generalized AGRIF* strategy: technically possible or an unaffordable dream?
-

Operational oceanography & urgent crisis situations:

-
- When grid focusing is needed (examples: Air-France crash, Fukushima accident) which strategy is easier: reshape the unstructured numerical domain or insert an AGRIFed subdomain?

How practical might a global multi-scale approach be?

Number of grid cells to refine to global criteria compared to high resolution everywhere

$N=\sum(n/L_i)^2\ \Delta X\Delta Y$

$Lmin<L_i/n<max(\Delta X,\Delta Y).f$

Physical scles:
- Slope
- Baroc Rossby Rad
- Barot Rossby Rad
- Smallest of these

No account is made for the cost of multiscale approach

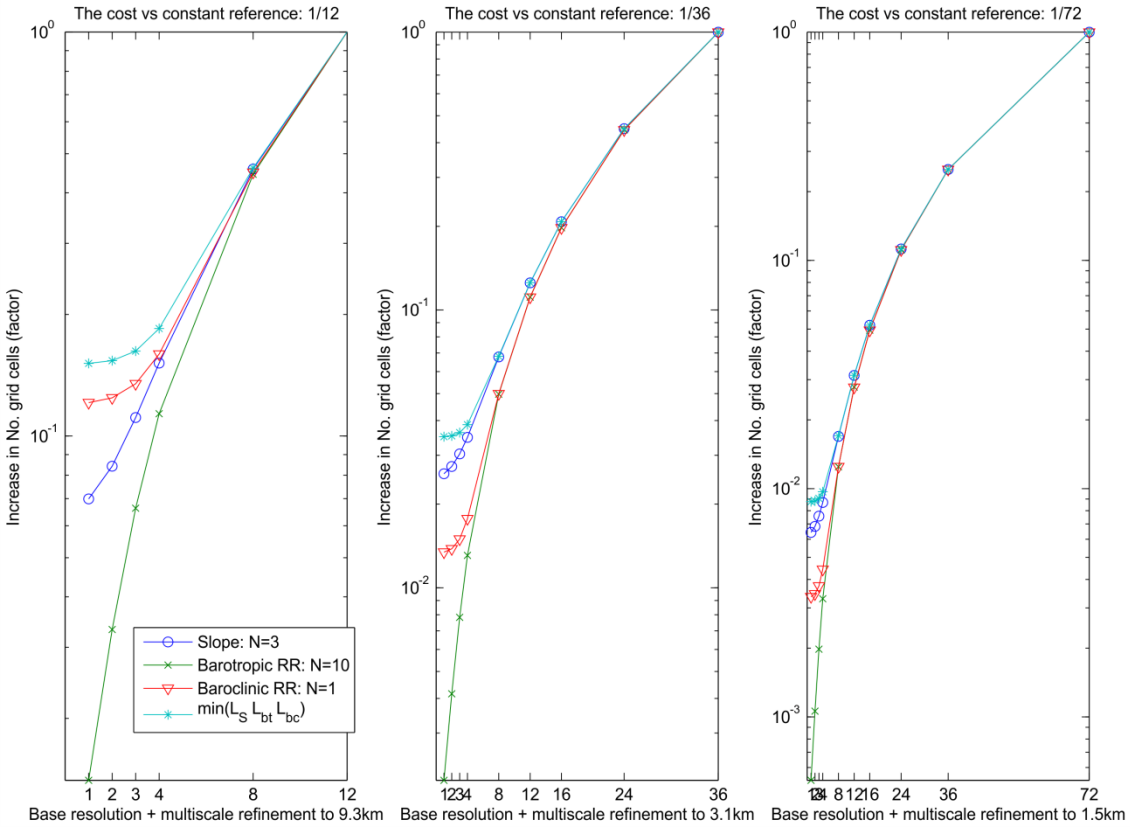

What niche for NEMO-Coast?

- Can we provide a suitable habitat to attract near-coastal modellers?
- There are many other models for the near coastal modeller to choose from:

Why NEMO?

- Ocean-shelf-coast nesting
  - single model system
- Computational efficiency
  - Aim at the longer integrations
  - Timesteps are the real challenge
  - Unstructured approaches are expensive
- Need aspects of the physics to be considere
  - Wetting/drying
  - Wave coupling
  - Momentum advection(?)
- Need funded projects