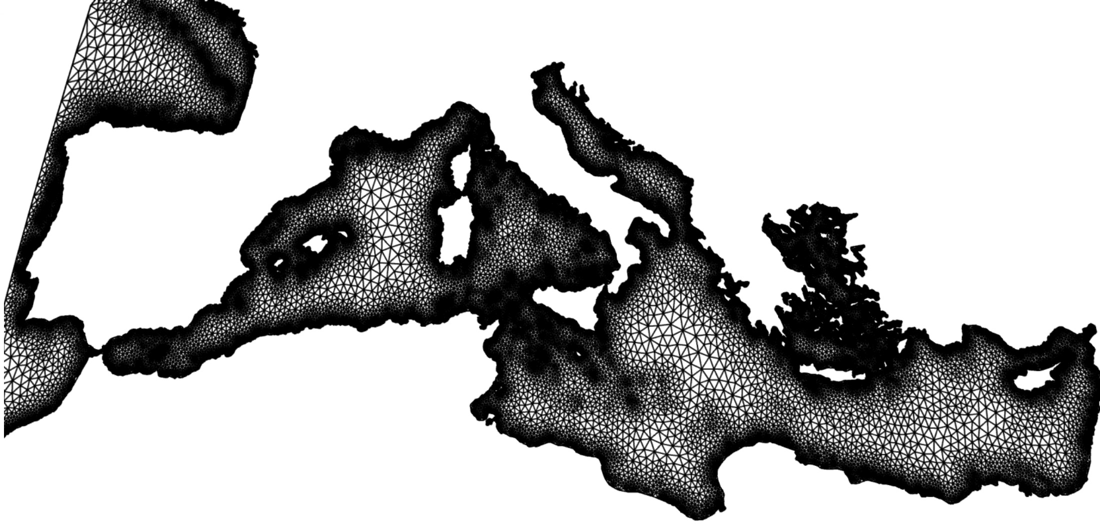

Modelling Ocean Dynamics with anisotropic adaptive mesh methods

Matthew Piggott
m.d.piggott@imperial.ac.uk

Grantham Institute for Climate Change &
Applied Modelling and Computation Group
Department of Earth Science and Engineering,
Imperial College London, UK

Acknowledgments

A large number of people are contributing to this work including groups at:

- Imperial College London
  - Applied Modelling and Computation Group
  - Grantham Institute for Climate Change
  - Department of Earth Science and Engineering
  - Department of Aeronautics
  - Department of Computing
  - High Performance Computing Service
- University of Oxford
  - Department of Physics - AOPP
- National Oceanography Centre, Southampton
  - Ocean Modelling and Forecasting
  - Ocean Observing and Climate
- National Oceanography Centre, Liverpool (POL)
- British Antarctic Survey
- Fujitsu Laboratories of Europe
- Daresbury Laboratory
- Florida State University
- and others

NATURAL ENVIRONMENT RESEARCH COUNCIL

Aims of this work

- To develop a flexible environmental fluid dynamics simulation framework
- built upon finite element based discretisation methods
- and unstructured meshes which may adapt, such that the mesh becomes an aid to simulation quality and efficiency rather than a constraint
- applicable to a wide range of problems with cross-fertilisation of ideas and algorithms
- and allowing comparisons between different numerical approaches

- ... deliver an open source tool that people will be able, and want, to use
  (implying software engineering, code robustness, V&V procedures, manual, tutorial problems, GUIs, etc)

Talk outline

- Motivation for multi-scale approach achieved via flexible meshes
- Brief description of the numerical methods used and mesh adaptivity
- Some idealised examples from classical CFD and GFD
- Very recent applications to lab, process and basin scale problems, trying to quantify the relative merits of some of our approaches
- Move to real world applications
- Some of the reapplication areas of the technology

Motivation: Resolution is critical, both to define adequately the domain of interest and also to resolve the relevant dynamics within it

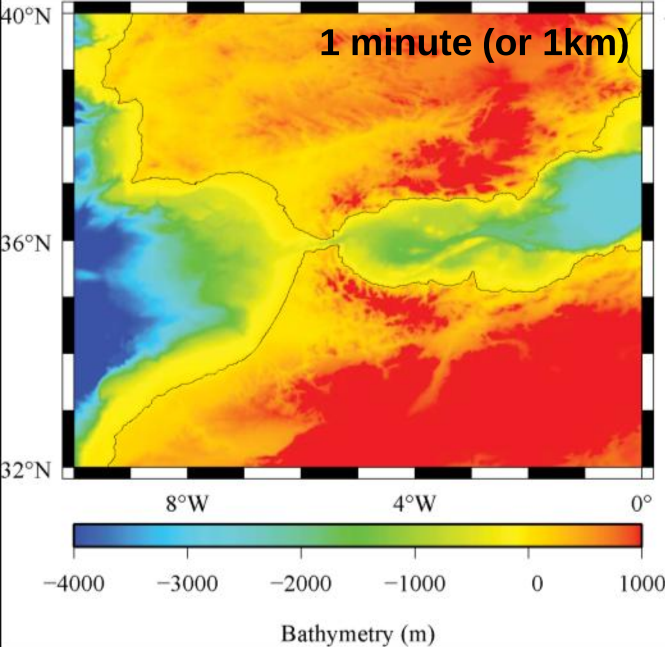

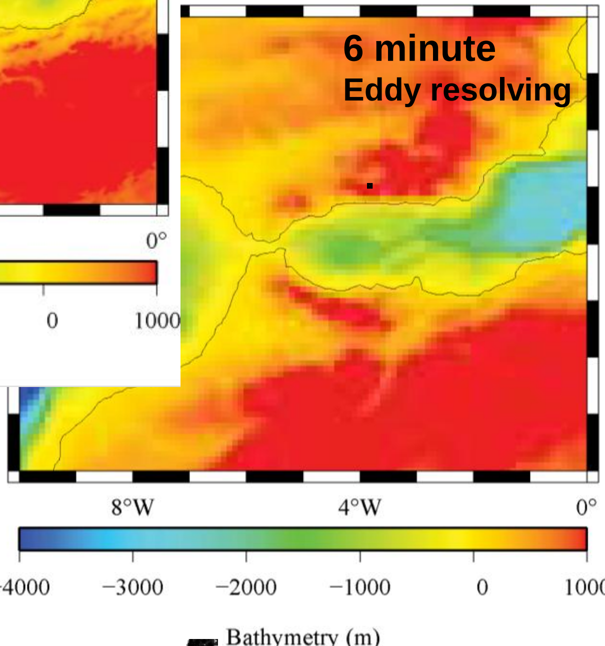

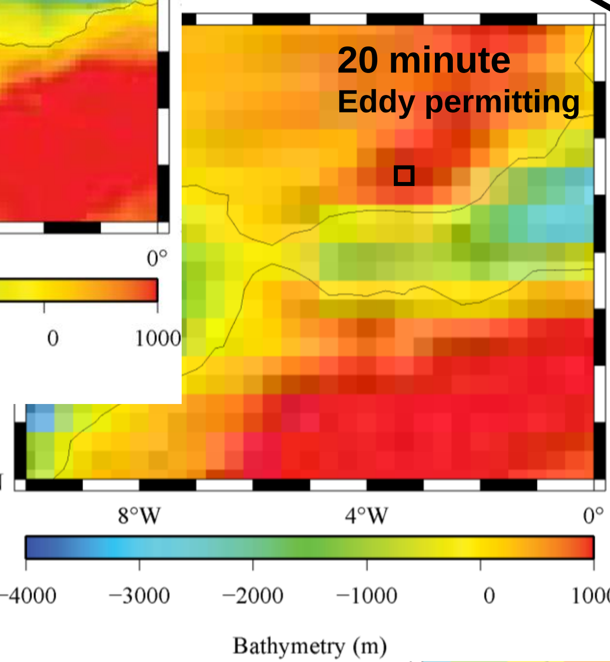

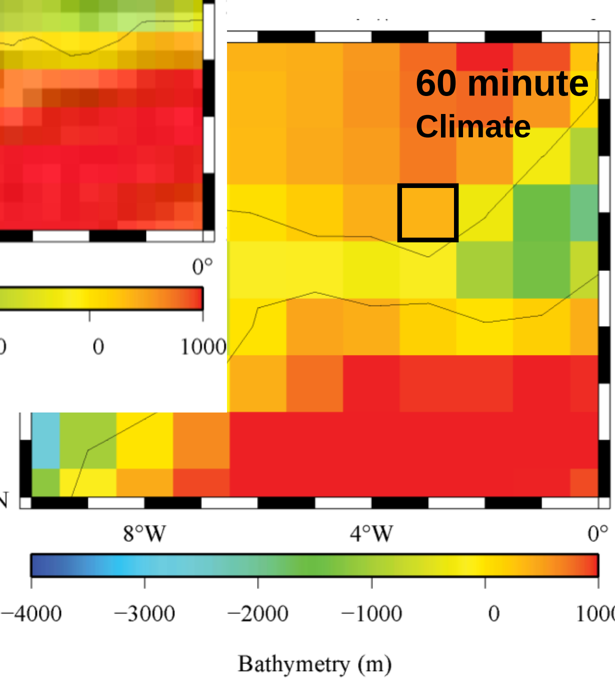

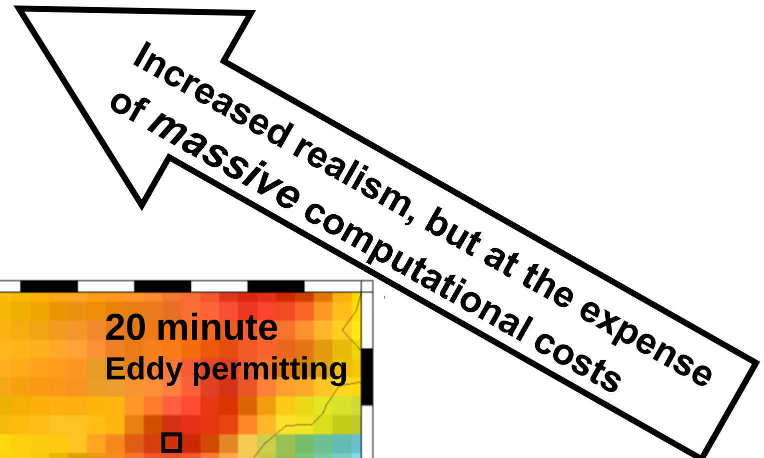

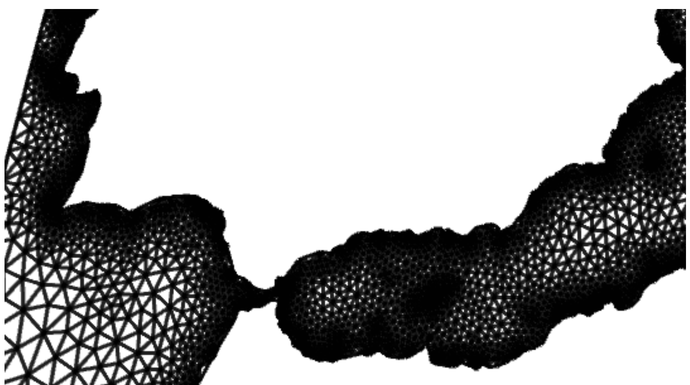

An unstructured (and adaptive) mesh is potentially able to give an optimal representation of coasts, bathymetry and the dynamics, without the need for high resolution everywhere

- Hence our interest in the flexibility for variable resolution that unstructured meshes provide
- However this flexibility does comes at a cost of higher CPU costs
- The aim is to mitigate this through the use of static or dynamic mesh adaptivity to focus computational resource where it is most required
- This can be due to numerical errors, or simply what is of relevance to the simulation being conducted, or the region you are interested in

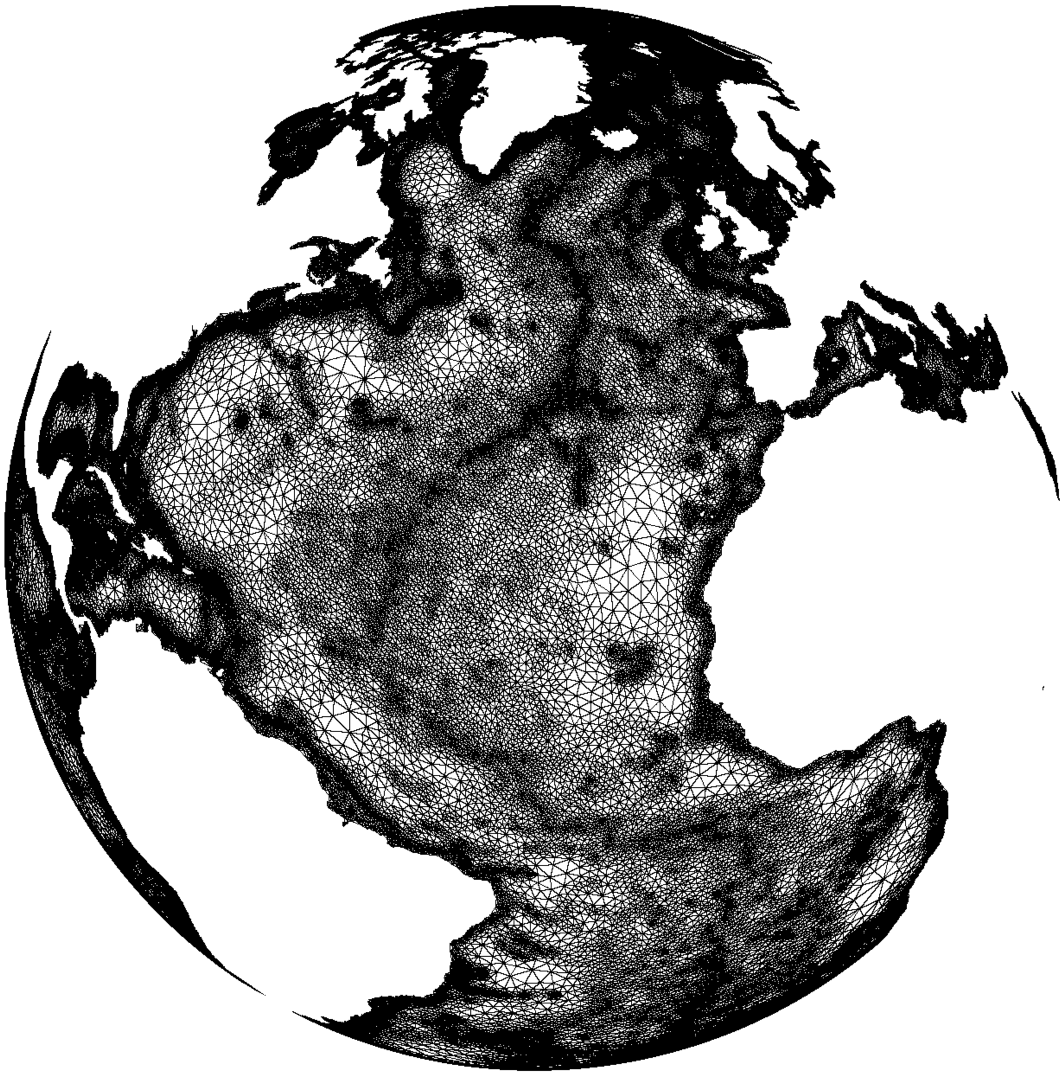

Note: mesh optimised methods are used to provide a generation procedure for coastlines and bathymetry (Terreno package; Gorman et al., 2006,7,8)

Structured vs unstructured meshes

Better representation of geometry, also the ability to the subdivide elements (hence smoothly varying resolution) without leaving ‘hanging nodes’ that you get with AMR or nesting

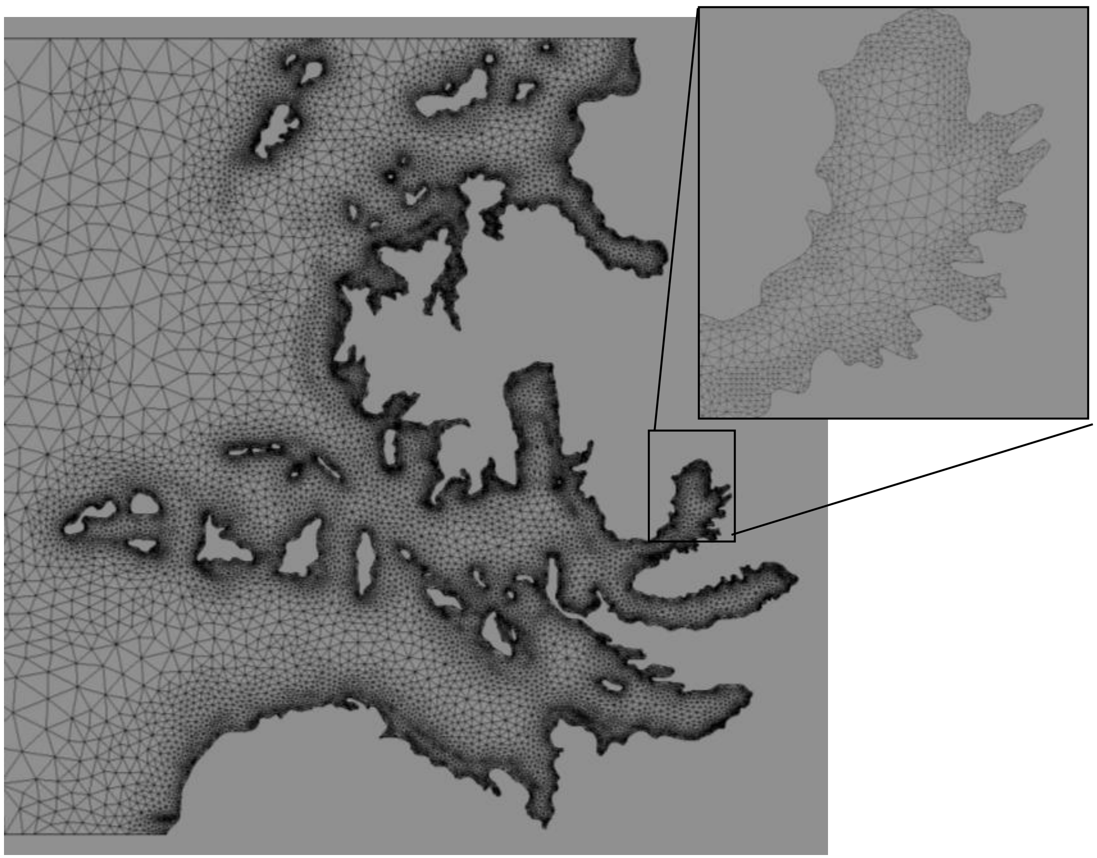

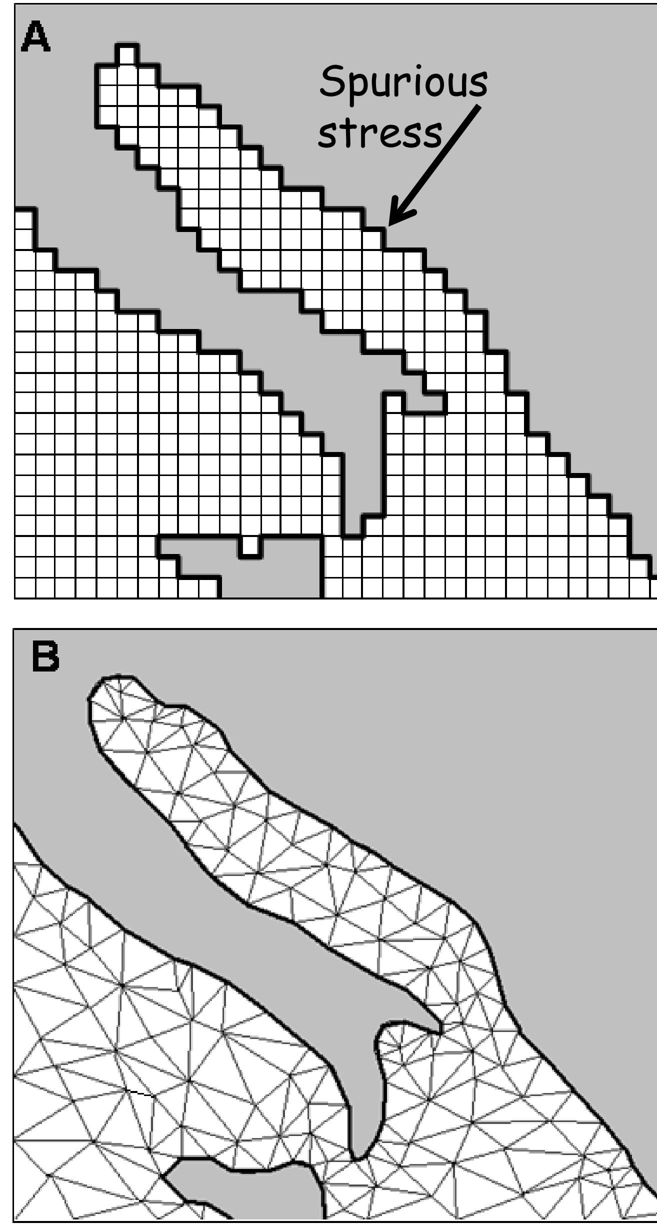

Unstructured meshes are an attractive choice for representing complex domains and a coupled range of scales without the need for grid nesting

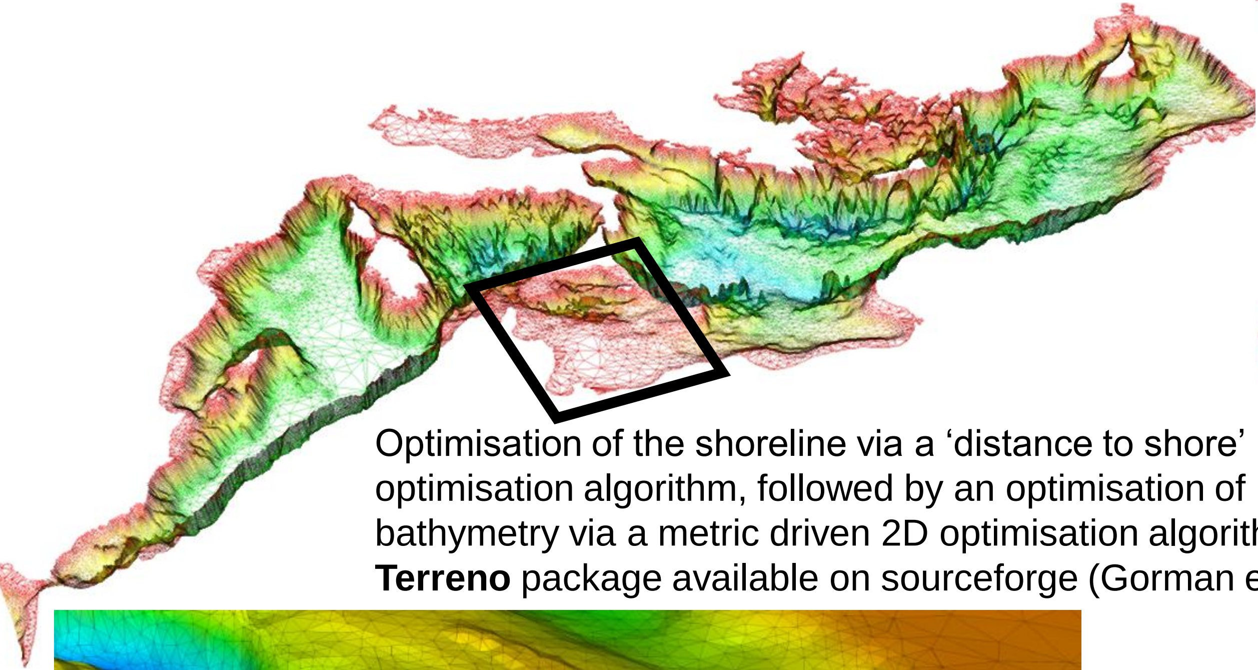

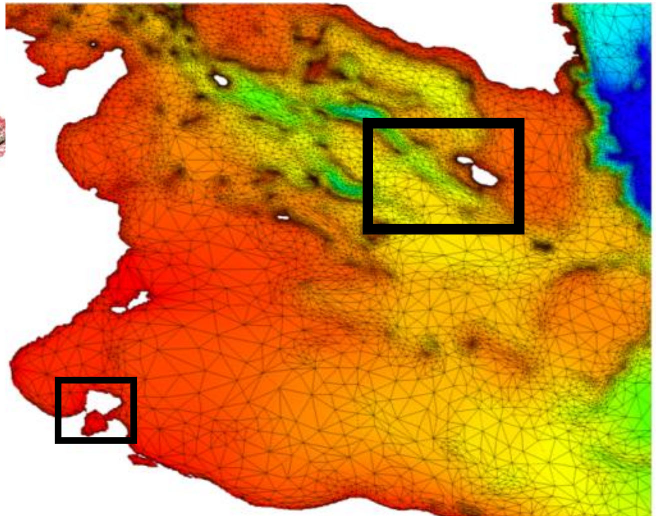

Optimisation of the shoreline via a ‘distance to shore’ optimisation algorithm, followed by an optimisation of bathymetry via a metric driven 2D optimisation algorithm. **Terreno** package available on sourceforge (Gorman et al.)

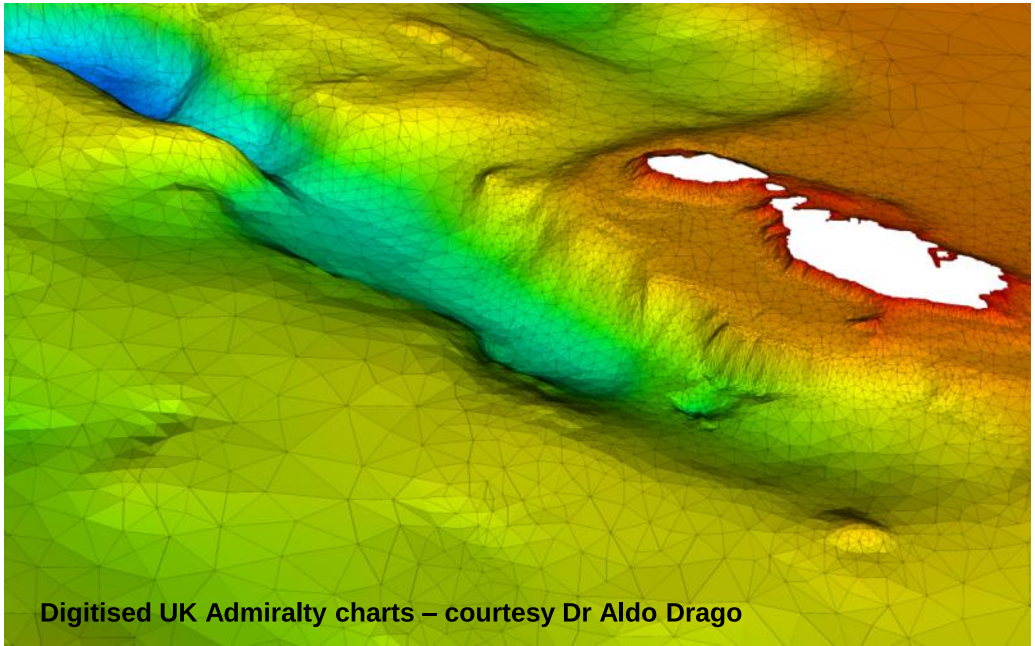

Digitised UK Admiralty charts – courtesy Dr Aldo Drago

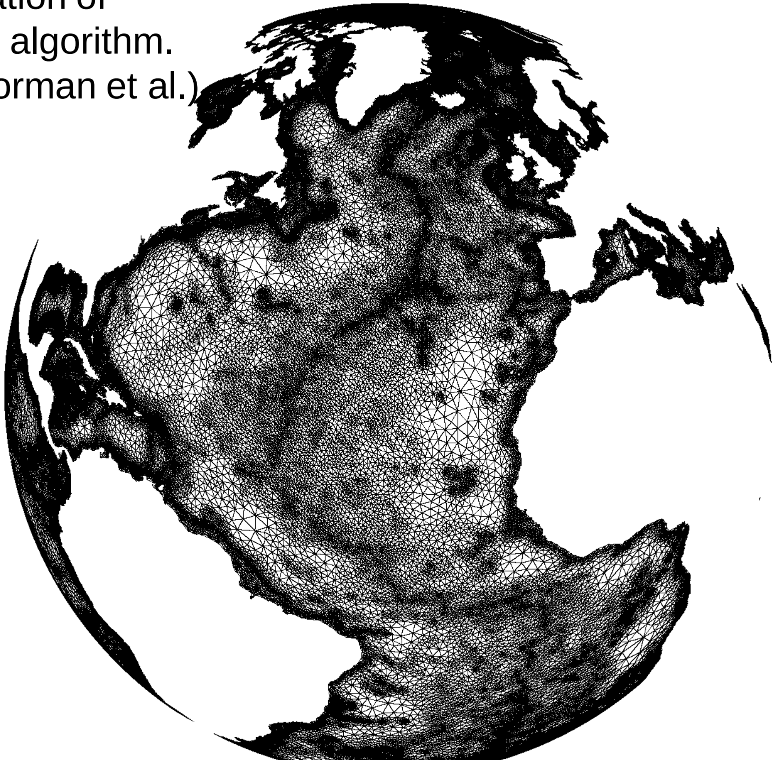

Mesh generation: First step is an accurate shoreline

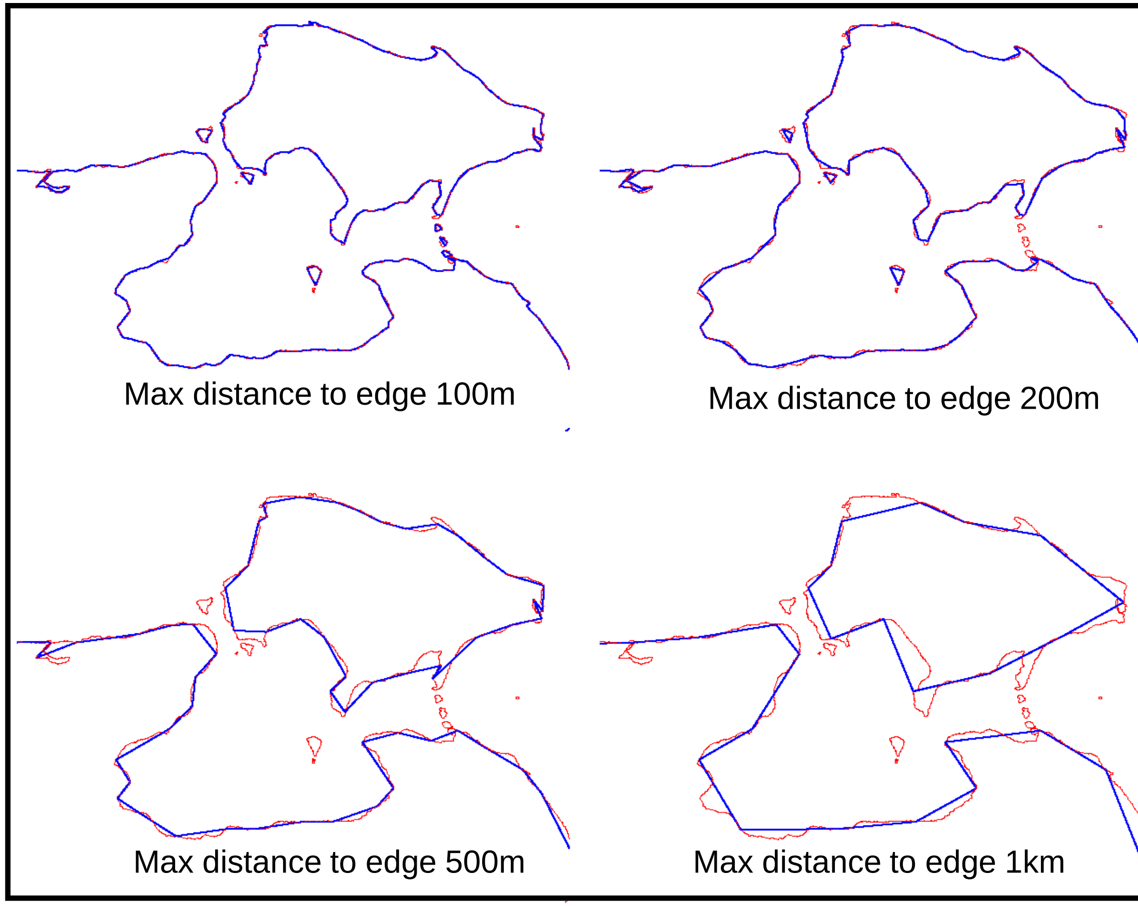

Given either the bathymetry zero contour or some high resolution coastline dataset, this initial ‘raw’ polyline data is optimized using a distance to edge and minimum edge length criteria.

The algorithm iterates so that previously ‘deactivated’ coastal points can be reactivated if they result in a better approximation to the original data.

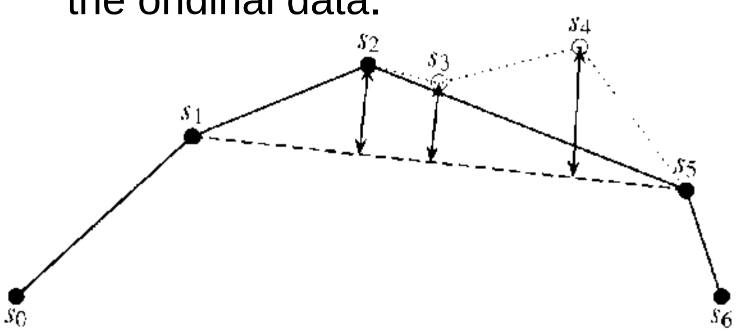

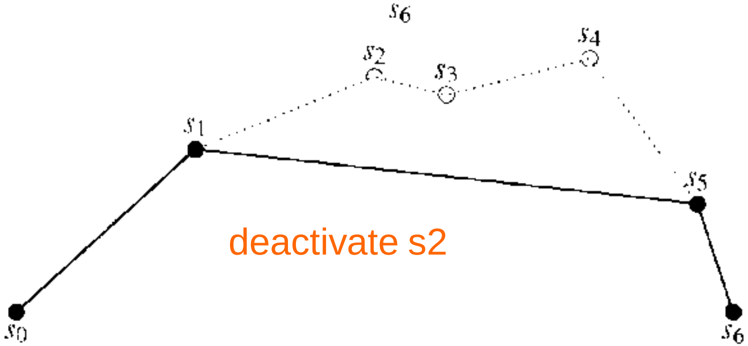

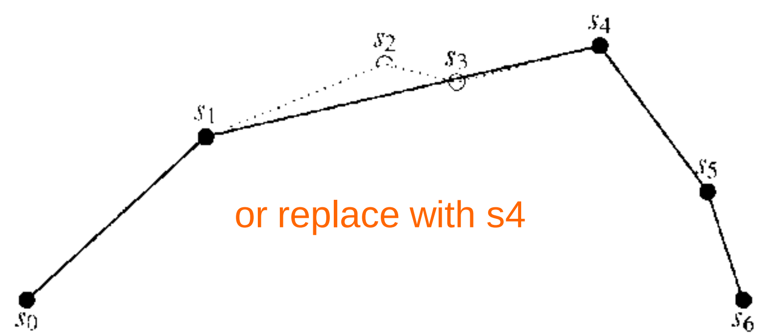

An example with same bounding box, but absolute and relative error measures used. Notice that the second case naturally preserves the opening at Gibraltar

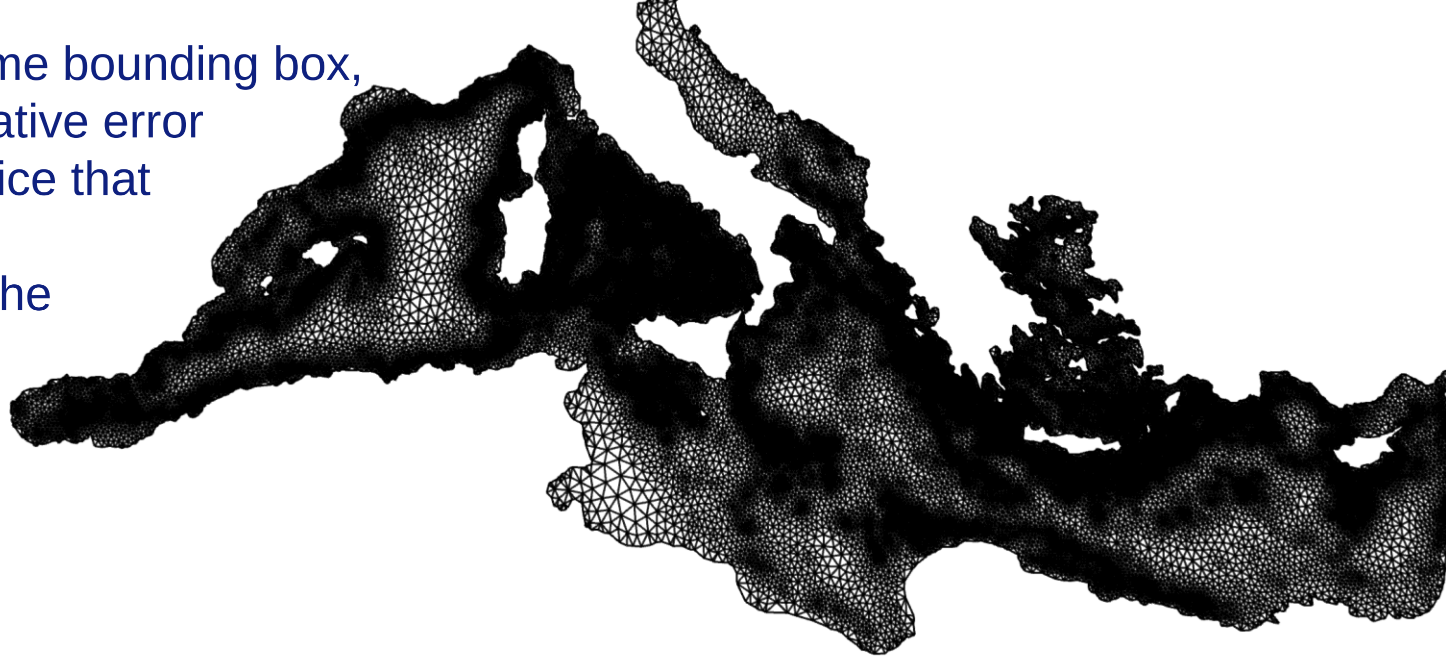

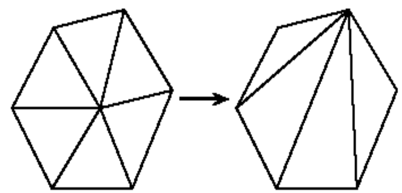

(a) Vertex collapse

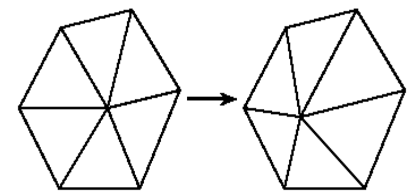

(b) Vertex repositioning

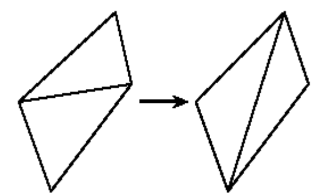

(c) Edge swap

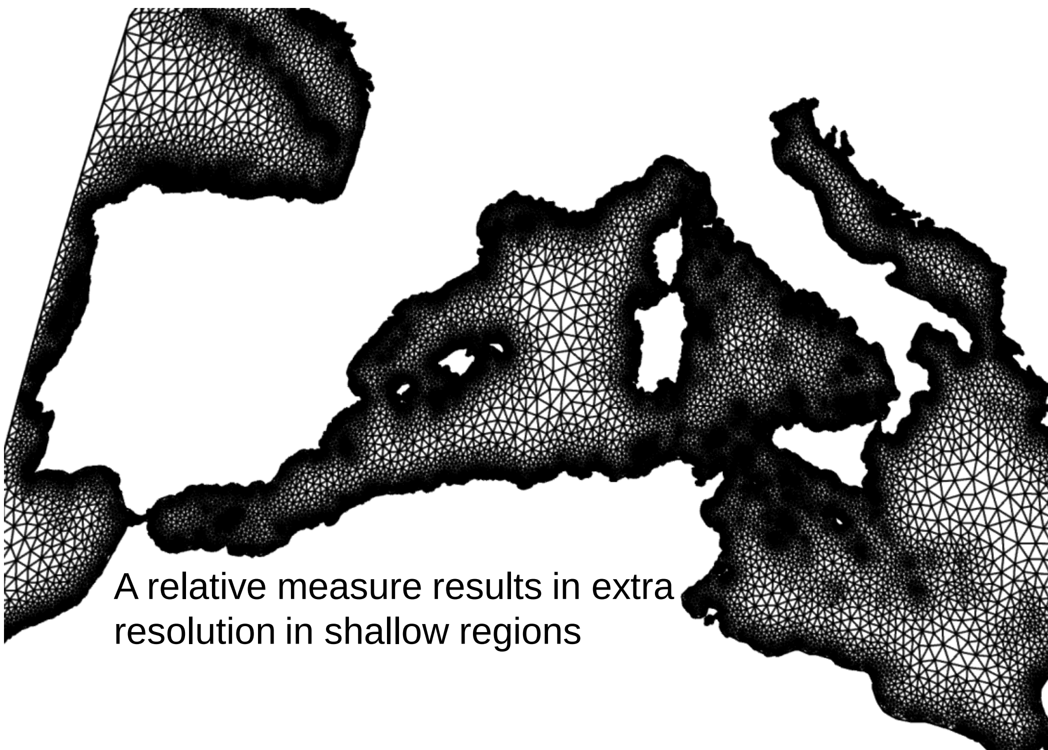

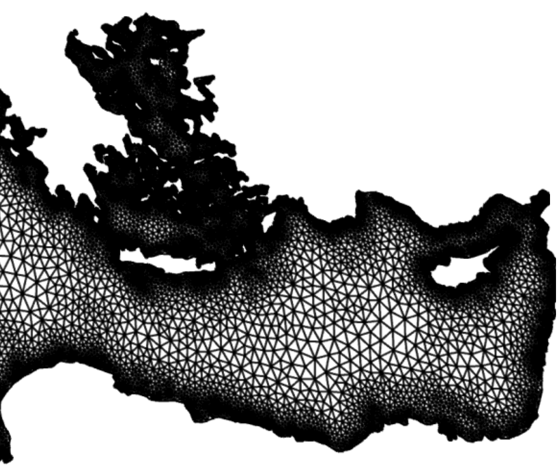

A relative measure results in extra resolution in shallow regions

Astronomically forced M2 tide in the North Atlantic coupled to the North and Baltic Seas (Mitchell et al.)

A high resolution mesh is required to resolve the channels connecting the Baltic and North Seas but a much coarser mesh is acceptable in the North Atlantic

More than two orders of magnitude difference between coarsest and finest resolutions

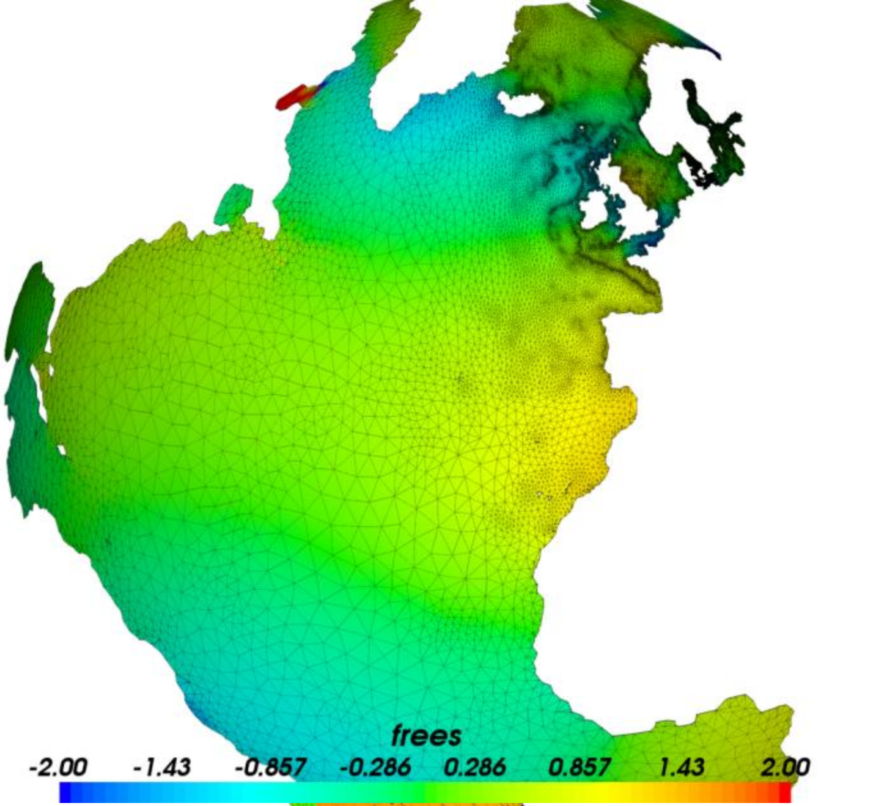

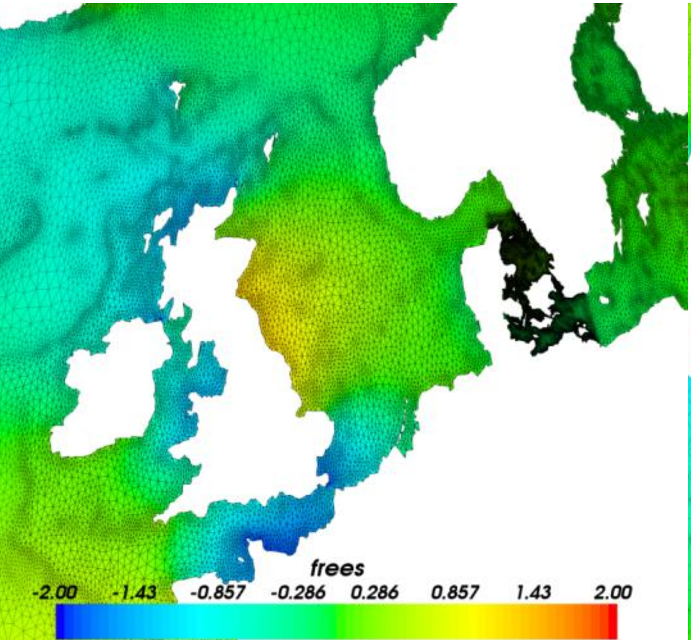

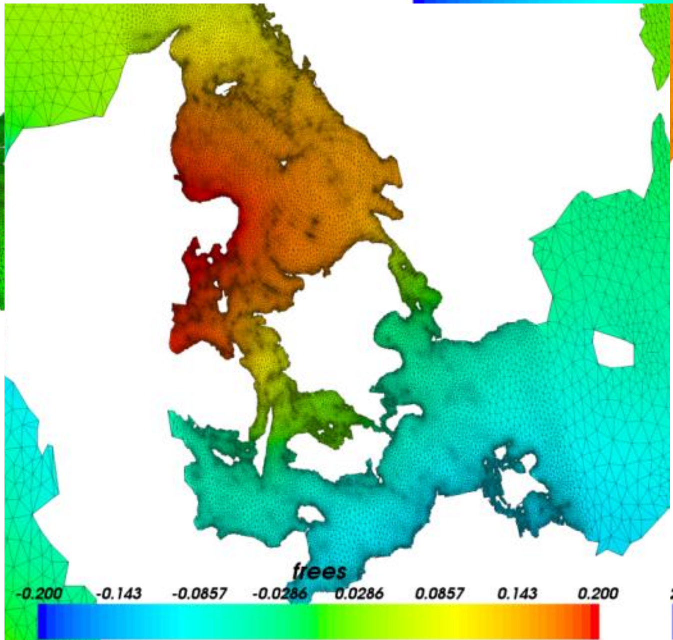

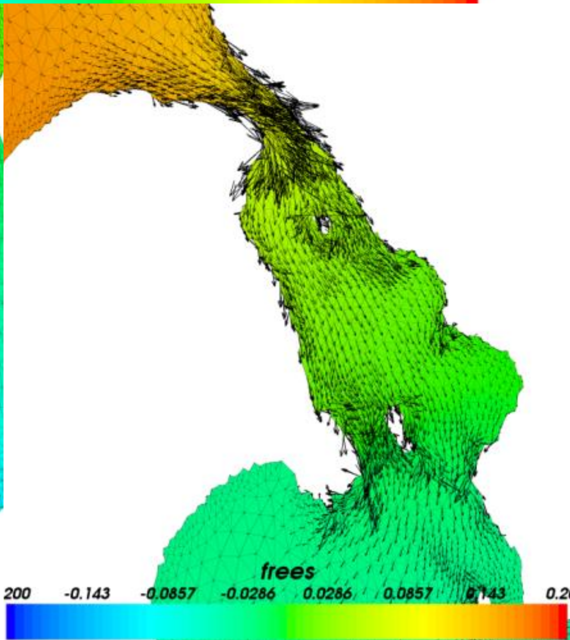

Numerical technology overview

- 3D Navier-Stokes flow solver with a range of CV-FEM discretisation methods
- Mesh adaptivity performs topological operations on the mesh to optimise size and shape of elements
- Load-balanced parallelisation, so far tested up to 4096 cores
- Open source community model development approach
- User-friendly graphical user interface, quick and flexible problem setup
- Automated quality assurance (validation & verification)

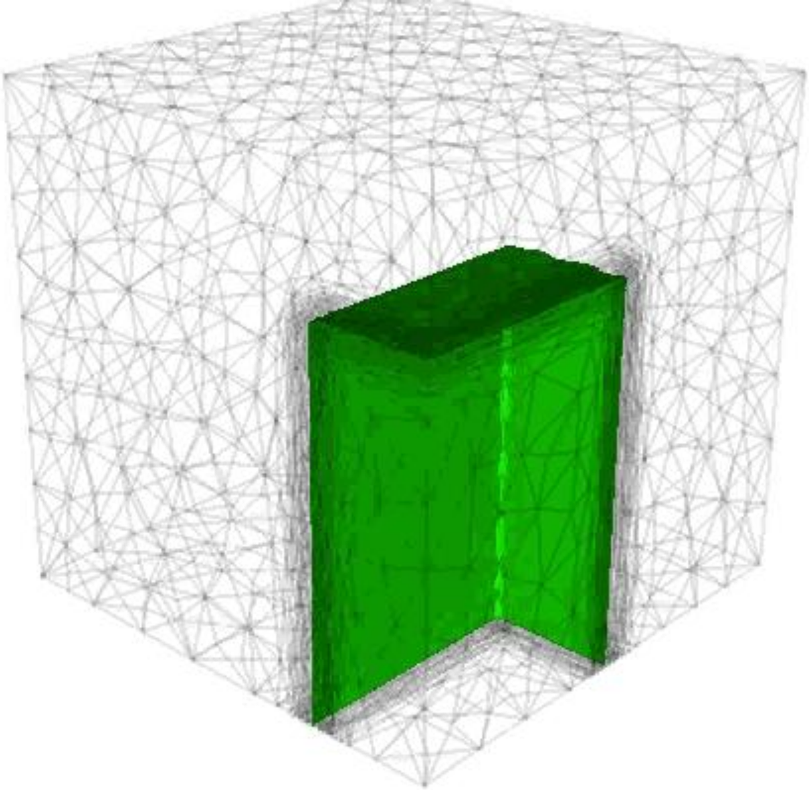

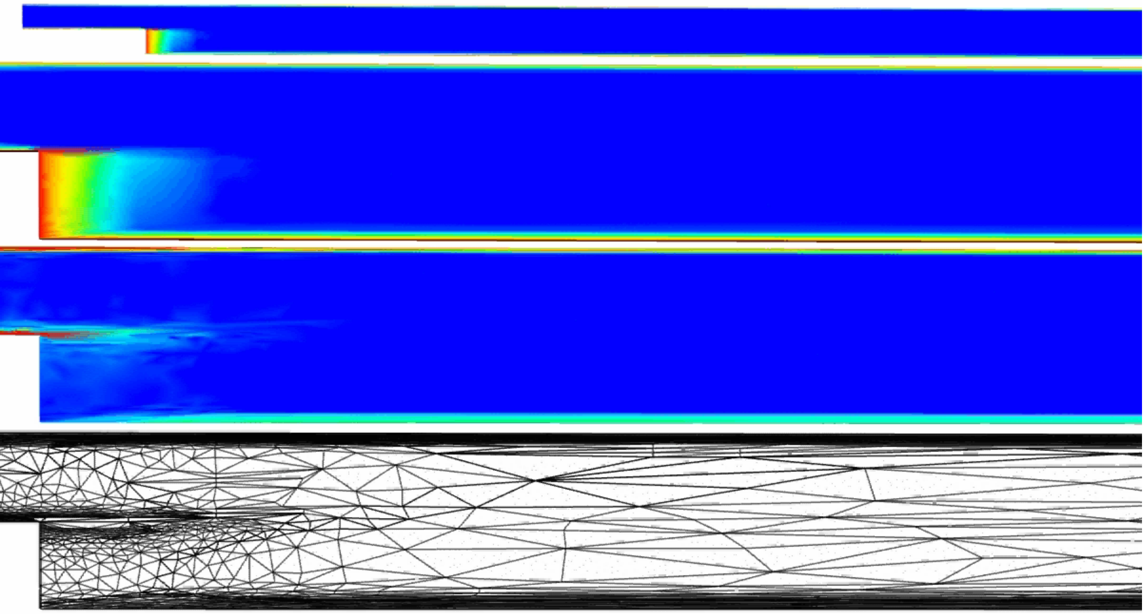

CFD tests used for code verification. Above: collapsing water column, volume fraction shown. Left: backward facing step, tracer, vorticity and mesh shown, number of nodes used shown below.

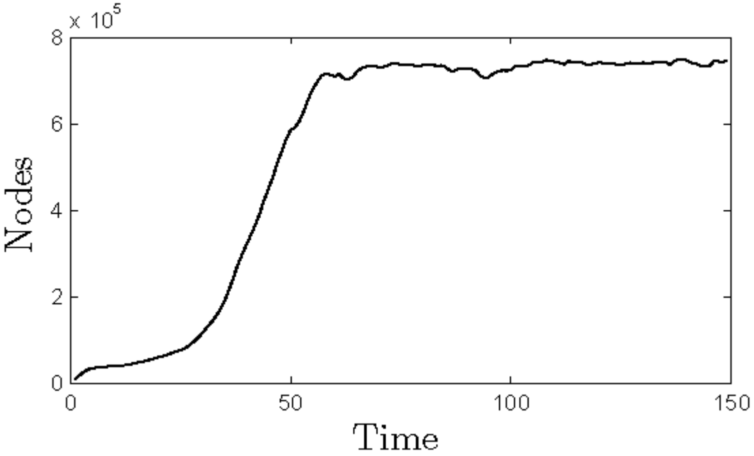

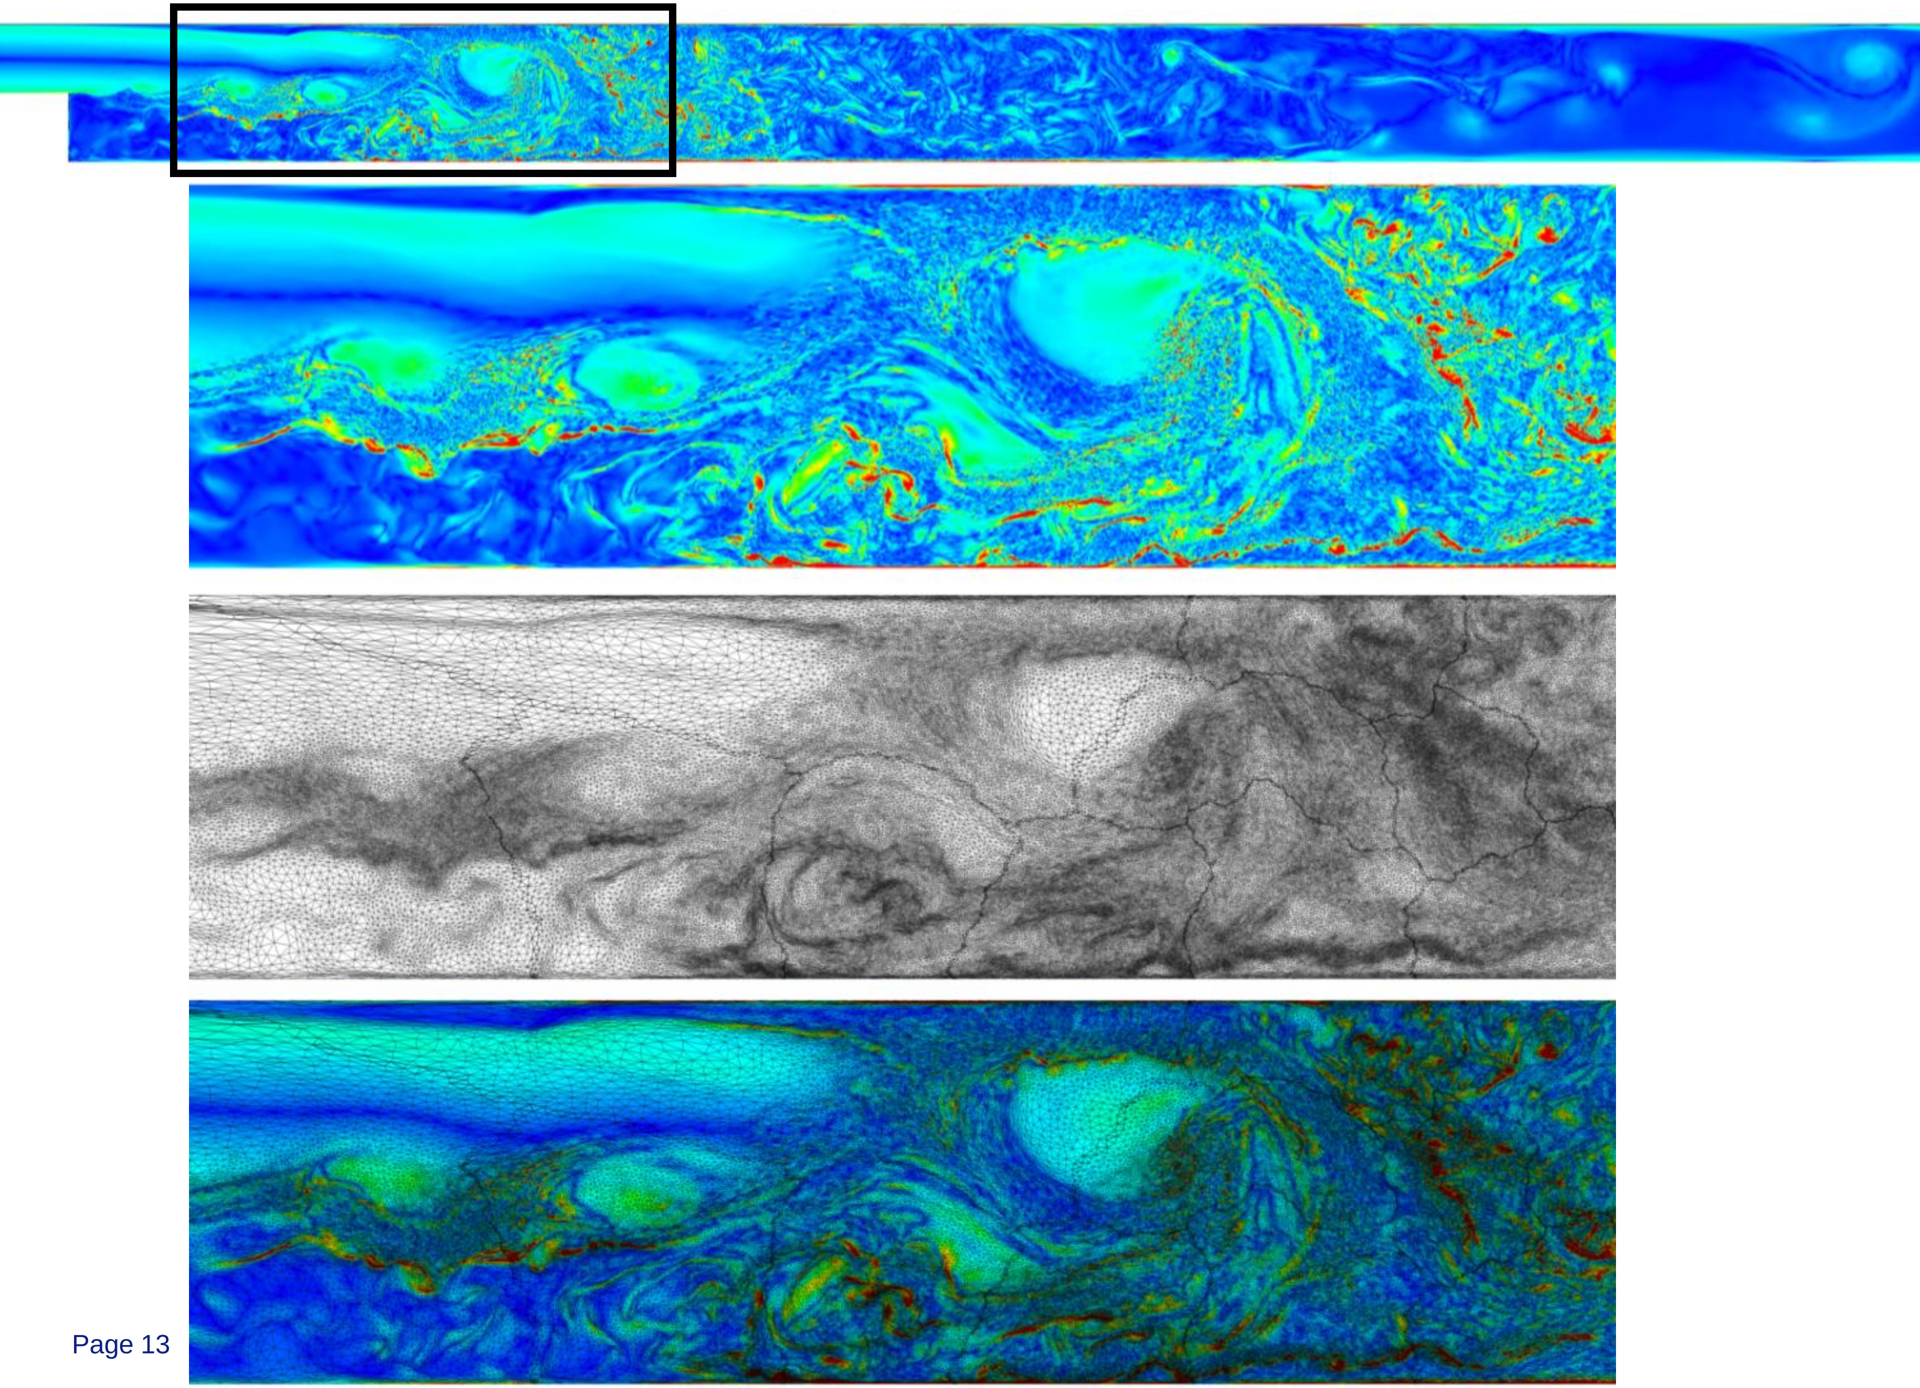

Typical discretisation methods used in the examples presented here

- Finite elements in space with either P1-P1, P0-P1 or P1dg-P2 momentum-pressure element pairs
- Option for auxiliary solvers to aid stability when buoyancy/rotation dominate
- Time-stepping: Crank-Nicolson or backward Euler here
- Projection method for pressure
- Two Picard iterations for nonlinearity
- PETSc library used for linear solves using CG or GMRES; AMG preconditioner for pressure solve
- CG+SUPG, flux-limited CV or DG with slope limiters used for scalar equations

- Compressible multi-phase the most general equation set but mainly non-hydrostatic Boussinesq considered here, some non-Boussinesq multi-material examples
- Stand alone QG and SW codes under development using same code framework

Anisotropic mesh optimisation via a Hessian based 'error' indicator

- Motivated initially by fact that discretisation errors can be linked to interpolation errors
- Linear interpolation error in 1D -> error is proportional to the mesh spacing squared multiplied by the second derivative of the exact solution
- In multi-dimensions it is bounded over an element by $v^T|H|v$ where H is the Hessian matrix and v is an arbitrary vector in the element (e.g. take the edges of the element)
- We seek an adapted or optimised mesh which equidistributes this error, i.e. a mesh where all edge lengths satisfy $v^T|H|v=\varepsilon$ – a user defined tolerance
- Equivalently all elements should be equilateral and of unit size when considered in a metric (or computational) space defined by the metric $M=|H|/\varepsilon$

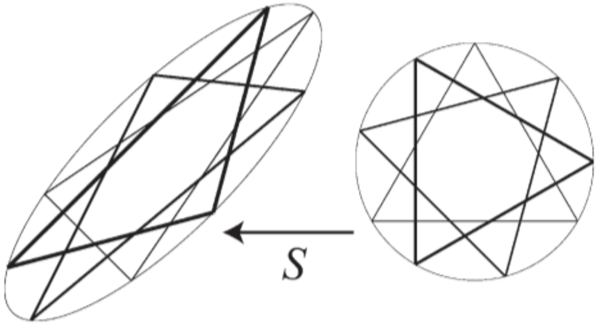

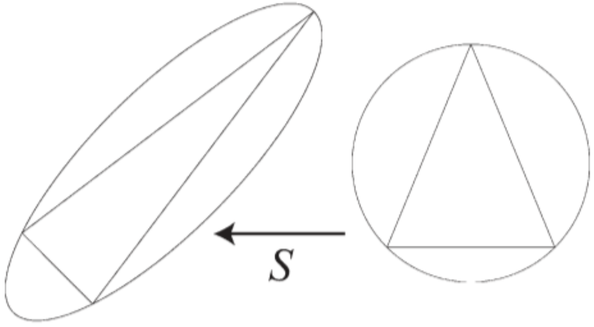

Optimal elements in physical space are those which are the realisation of equilateral elements under a coordinate transform obtained from the Hessian (directional stretching)

Formulating an optimisation problem (3D)

- Form an optimisation functional, in 3D we use:

$$F_e = \frac{1}{2}\sum_{l\in\mathcal{L}_e}(r_l - 1)^2 + \left(\frac{\alpha}{\rho_e} - 1\right)^2$$

- First term ensures good edge lengths, second term helps with element shape
- Evaluate with respect to a metric $M = \frac{|H|}{\epsilon}$ where H is the Hessian matrix, motivated by linear interpolation error
- Minimise the functional through local mesh operations
- The total metric is actually formed by combining metrics for multiple solution fields, incorporating user-defined weightings, and constraints on max/min edge lengths

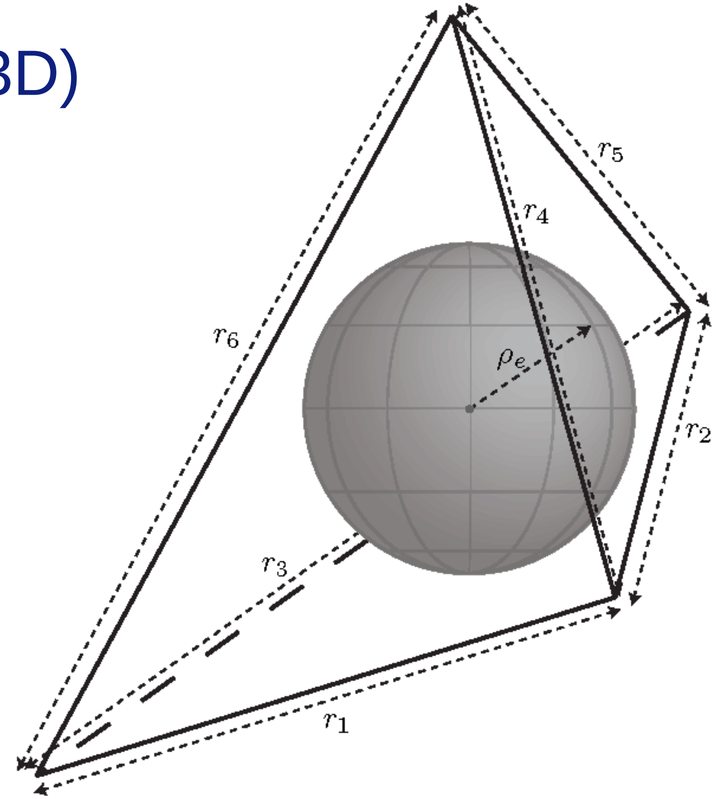

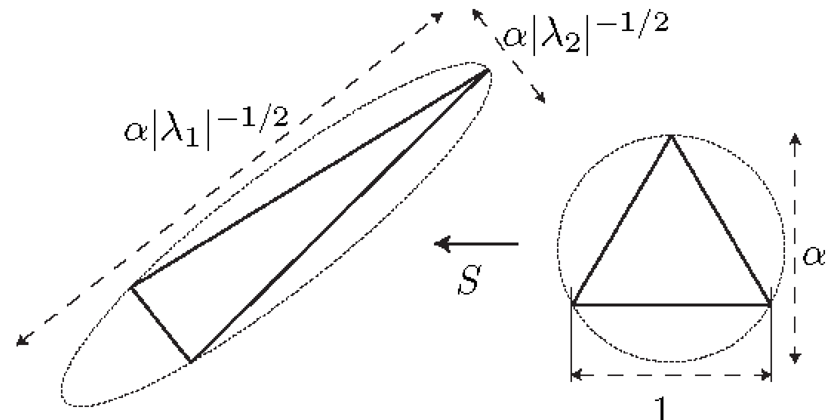

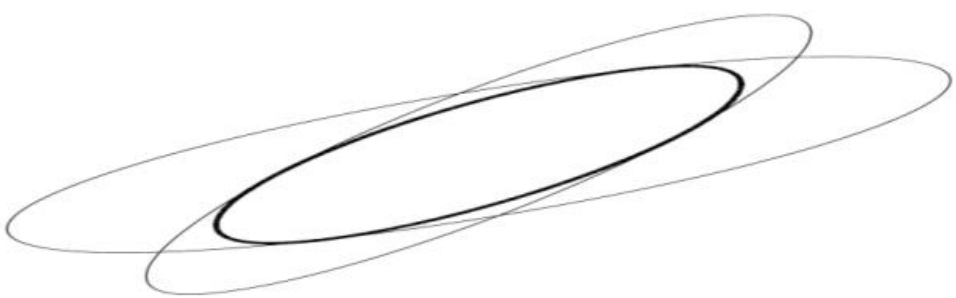

Following rotation to coordinate axis, edge lengths still satisfy:

$$\Delta x^2 \frac{\partial^2 u}{\partial x^2} = \Delta y^2 \frac{\partial^2 u}{\partial y^2} = \epsilon$$

Operations used by the mesh optimisation algorithm (3D)

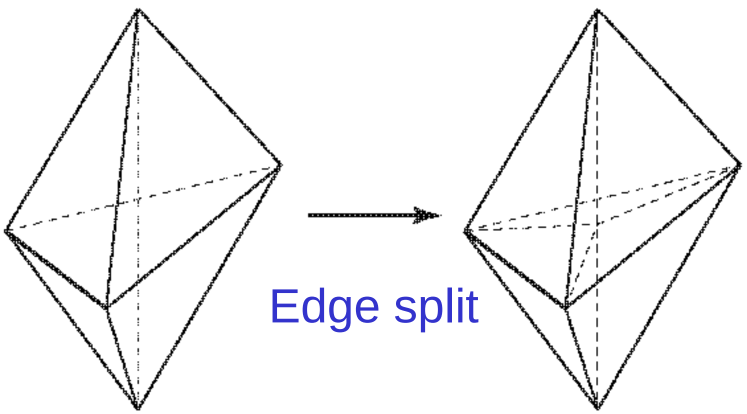

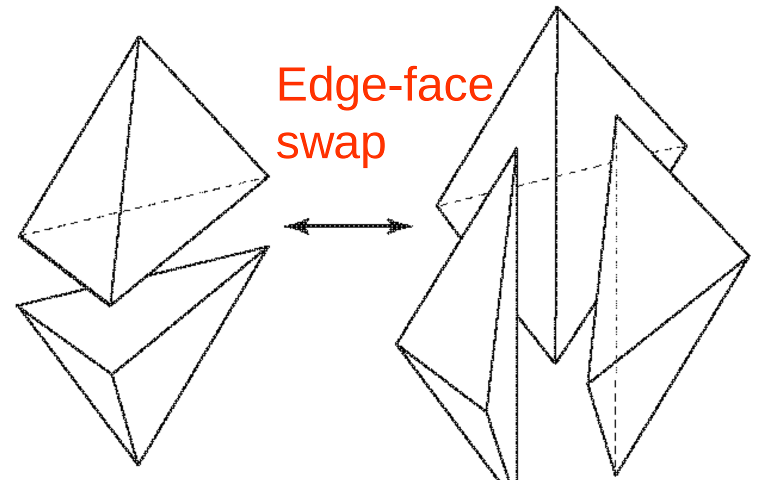

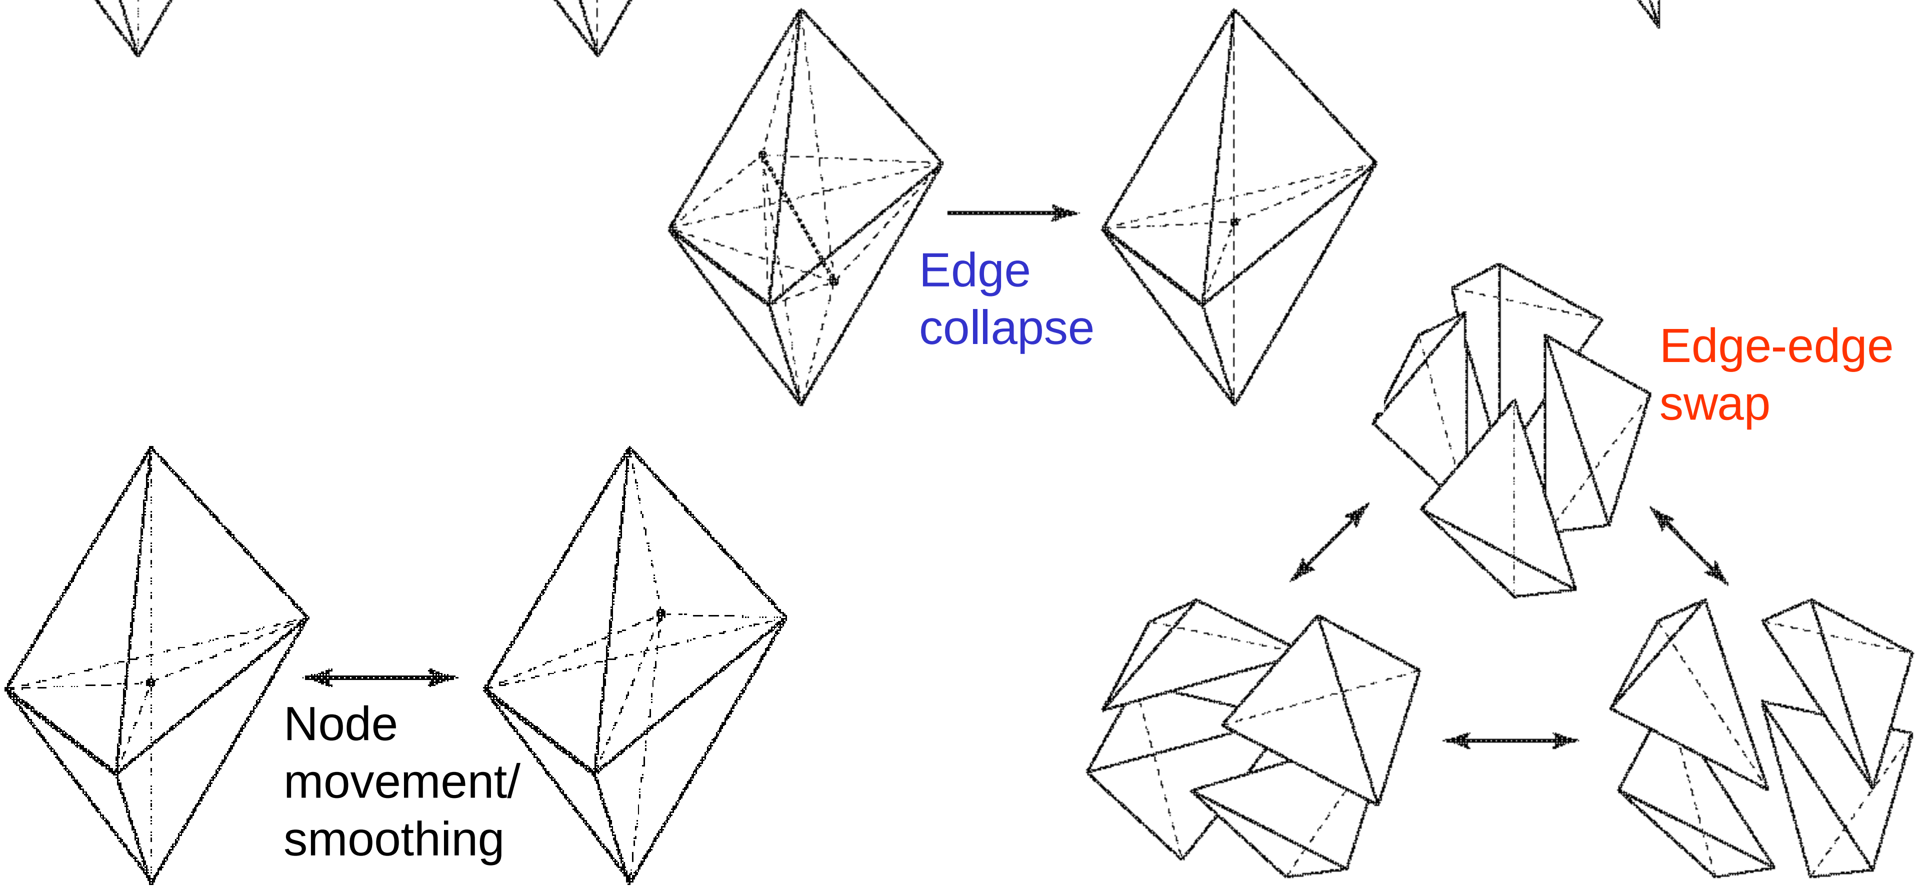

Two-dimensional mesh optimisation

For 2D applications we use the approach of Lipnikov et al. and their libmba library

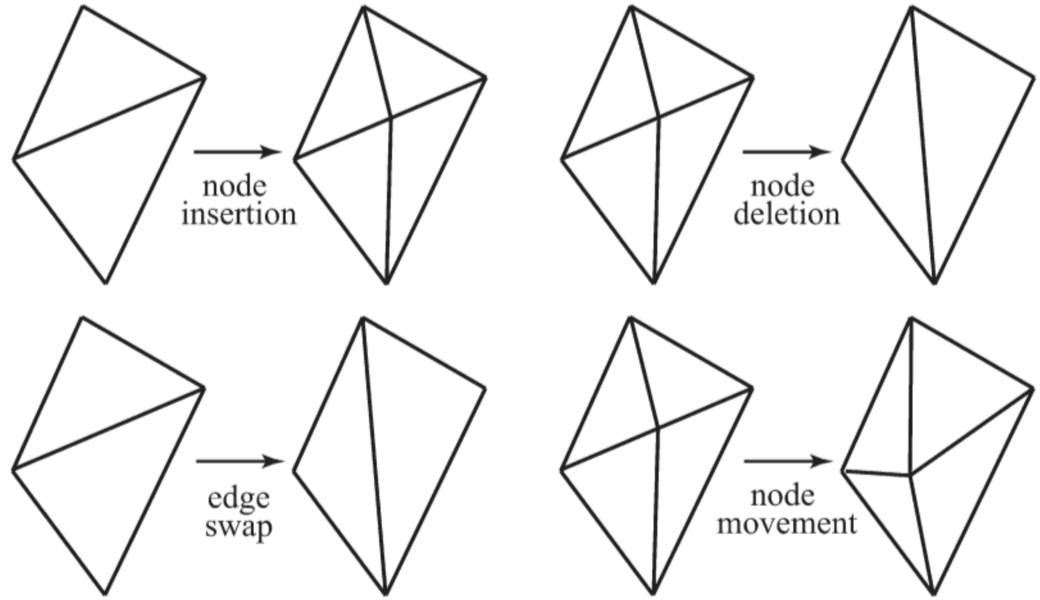

E.g. An example of the effect of an edge swap on the element quality functional for a pair of elements in 2D. The 2D functional is optimal for a value of unity here

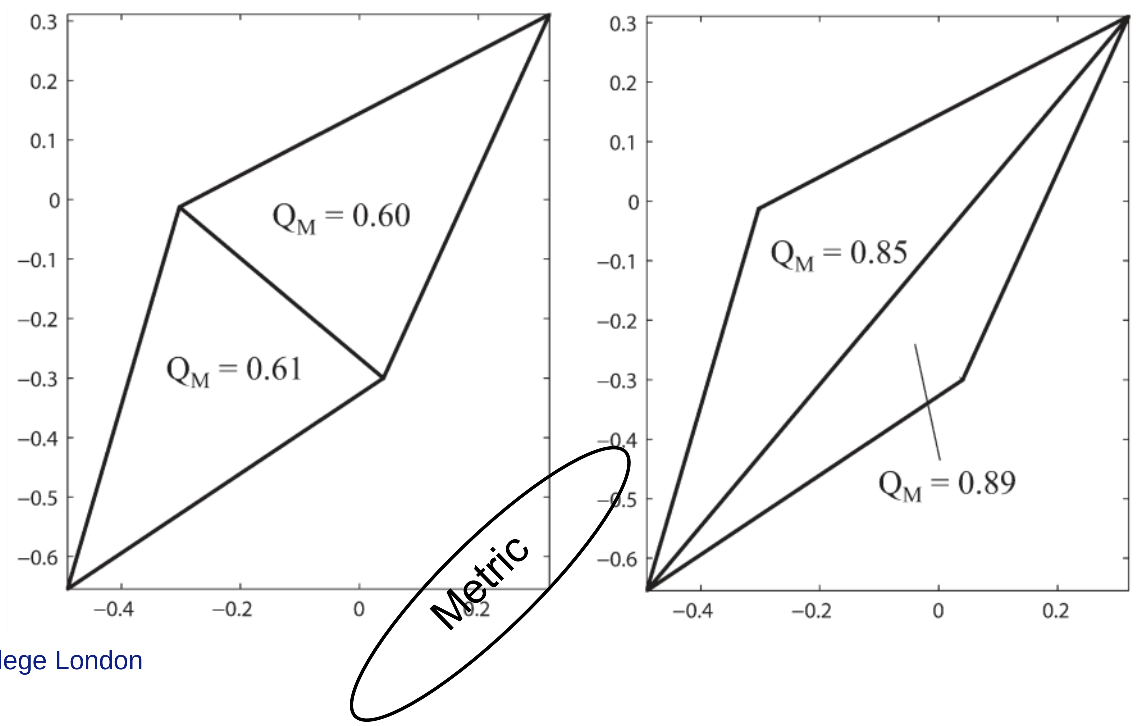

Dynamic balancing of the computational load between processors as the mesh adapts in parallel

Cost of adaptivity here approx 10% of simulation, cost of data migration minimal

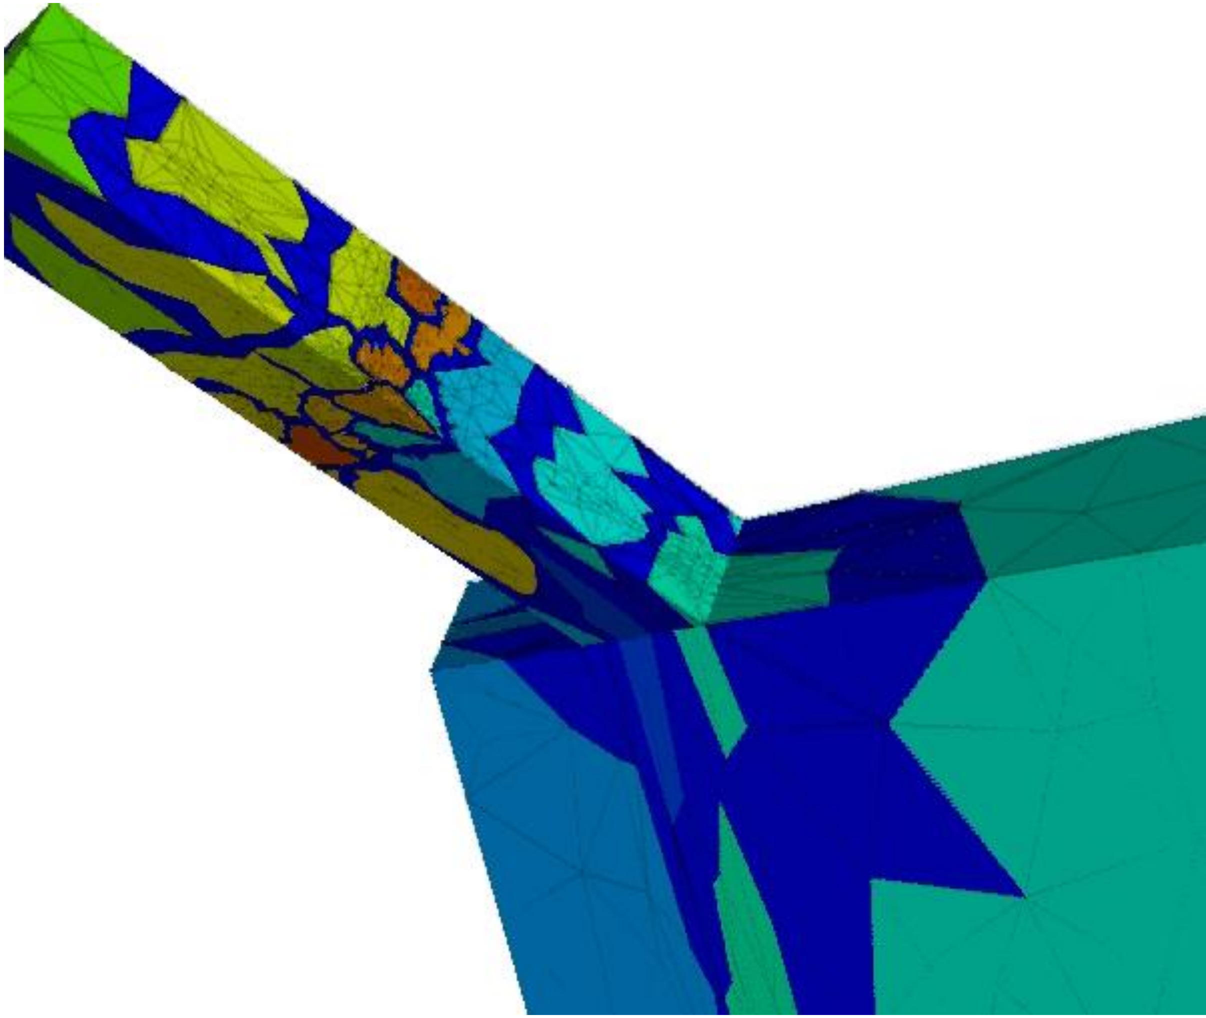

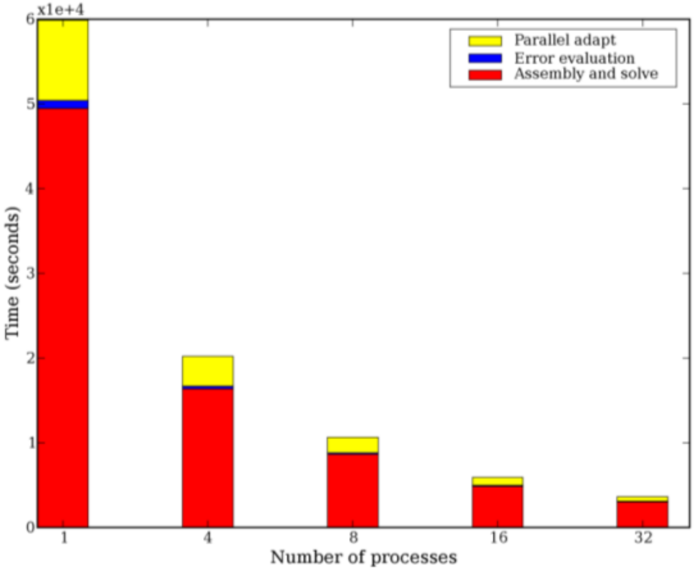

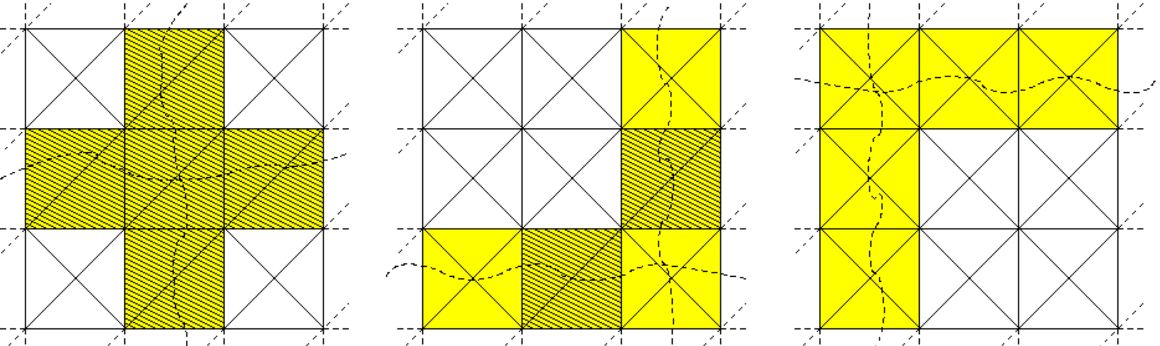

Currently extending these ideas for periodic domains

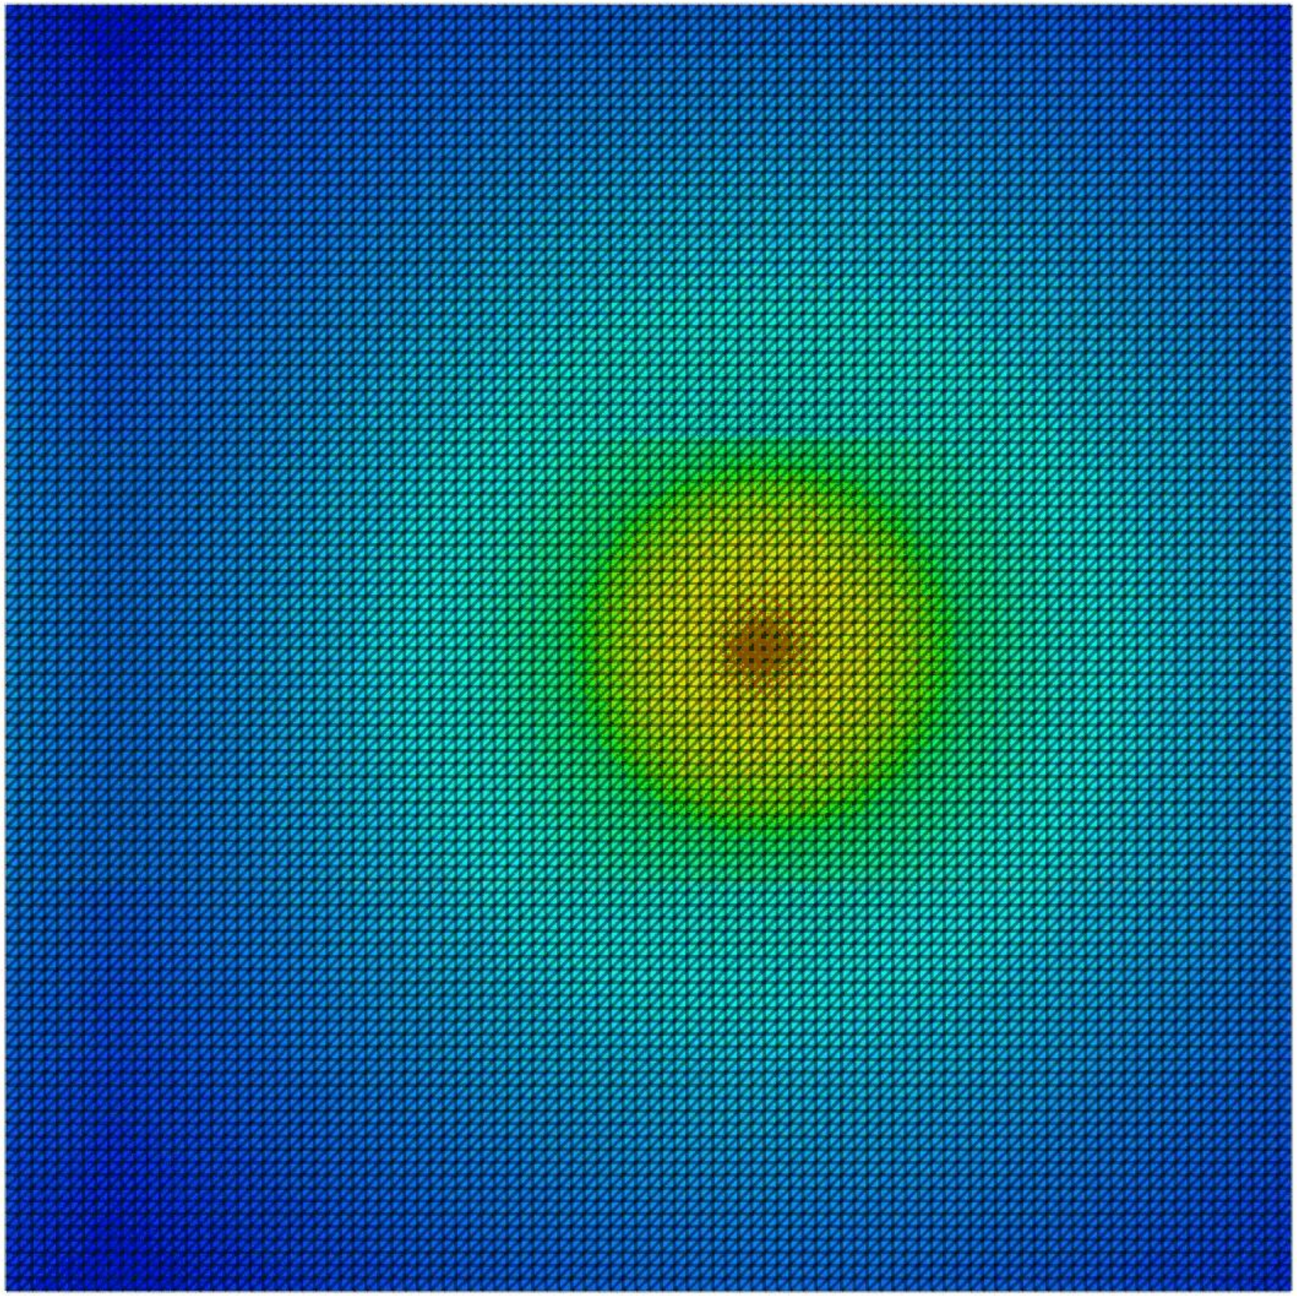

Conservative Interpolation: crucial for many applications, with adaptivity, coupled models

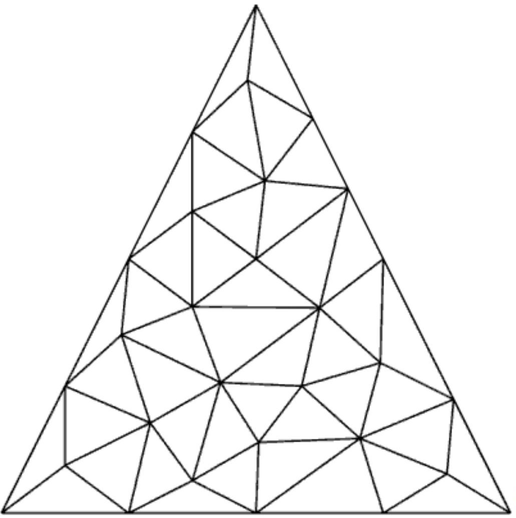

‘donor’ mesh

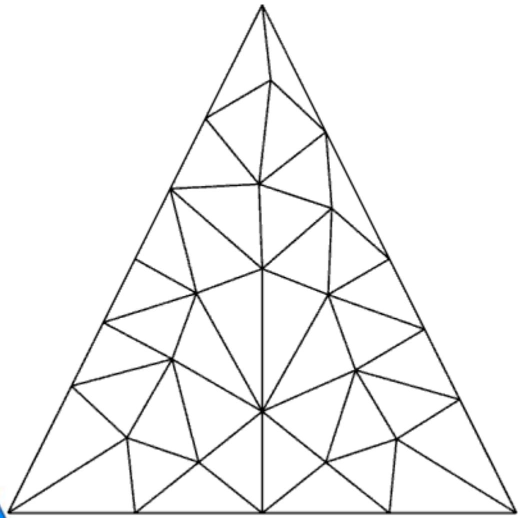

‘target’ mesh

‘super mesh’ and mapping from donor mesh

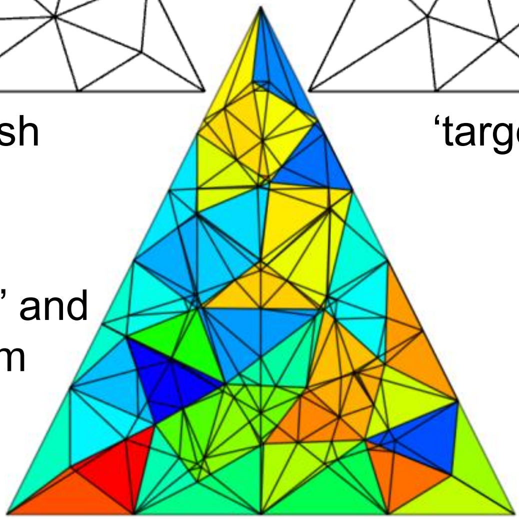

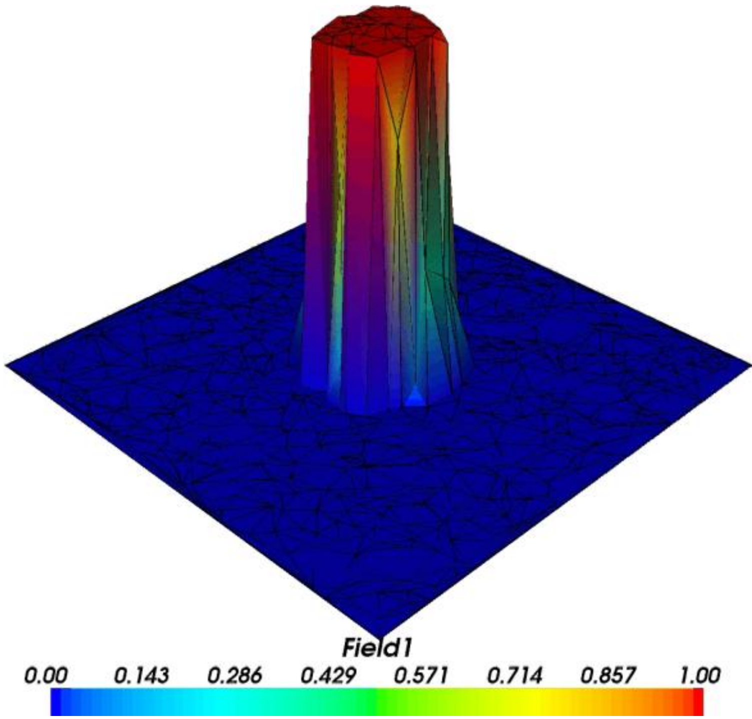

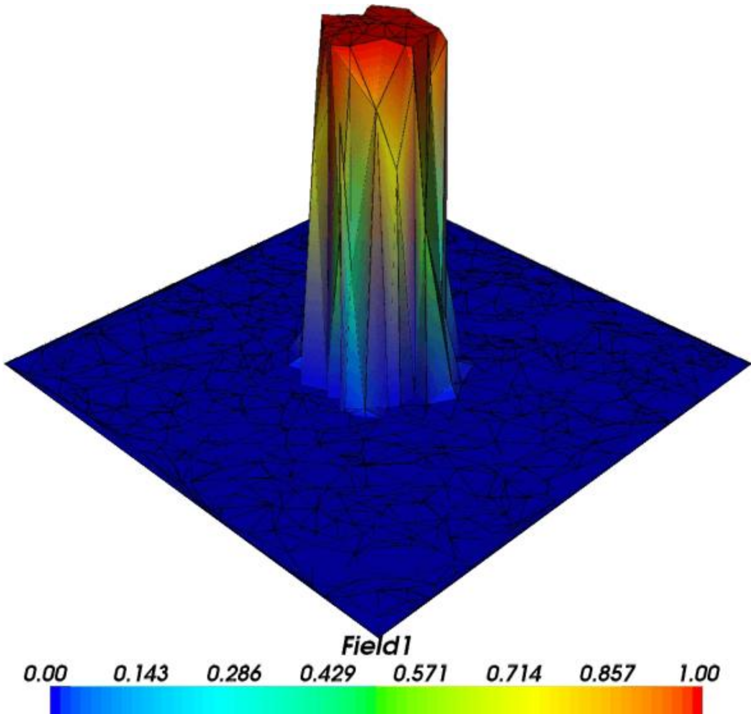

Following construction of the supermesh we are able to perform projection operations to achieve conservative, bounded interpolation schemes.

2D: Farrell et al., 2009

3D: Farrell, Maddison, 2009

Mesh movement

The use of mesh movement has many advantages, but for robustness should be used in combination with mesh optimisation (hr)

'Isopycnal' like capability in the vertical; mesh optimisation does less, and can be used less frequently; less interpolation; better load-balancing.

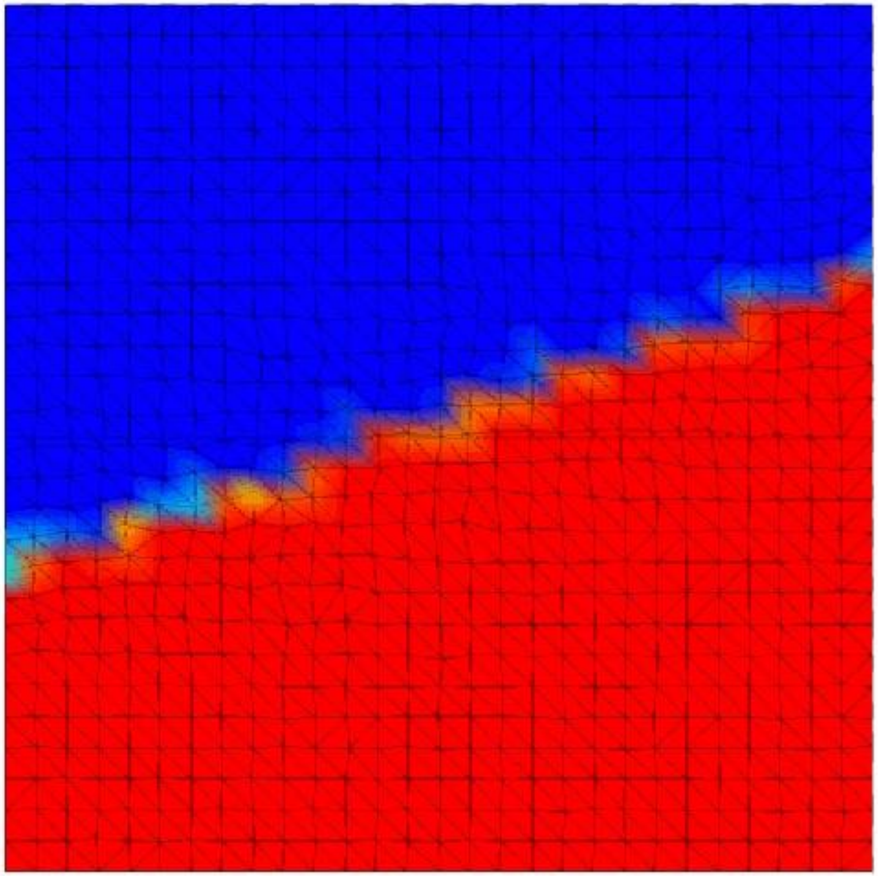

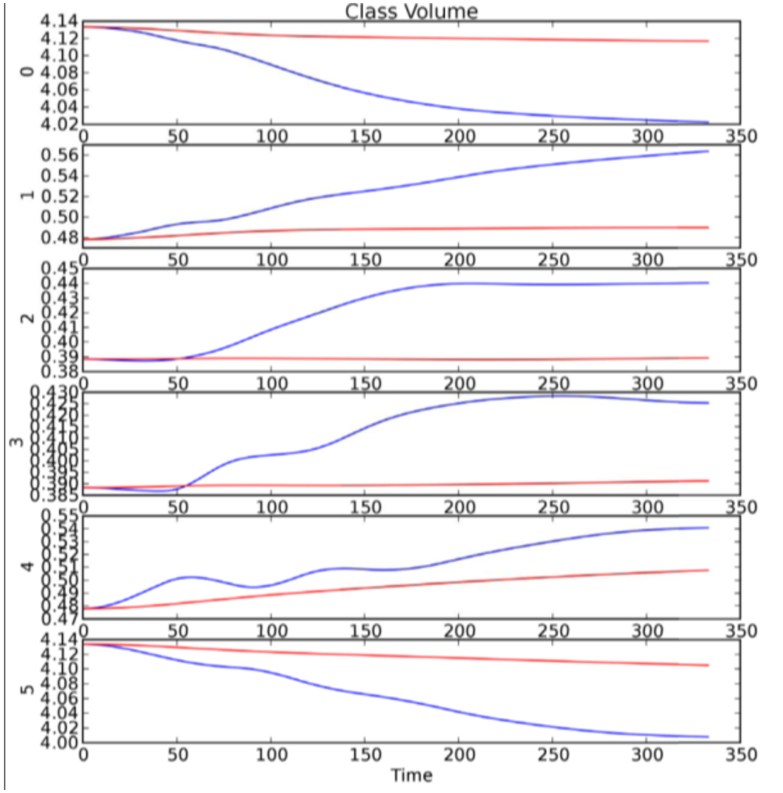

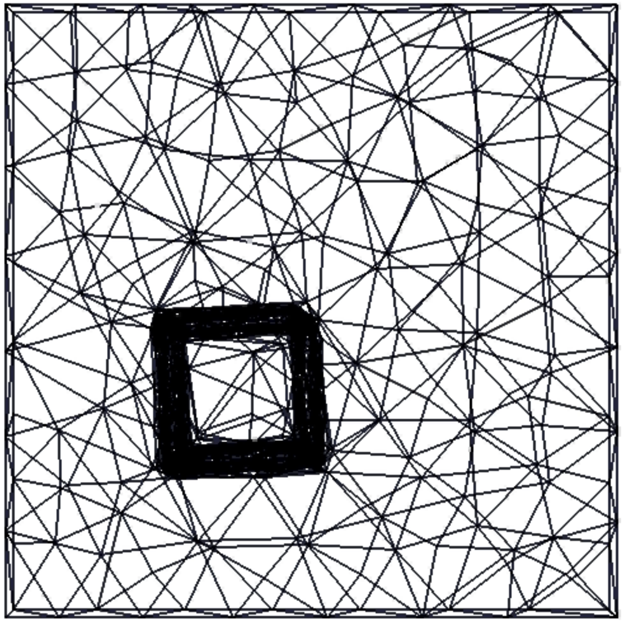

Use mesh optimisation to avoid tangling

Mesh movement

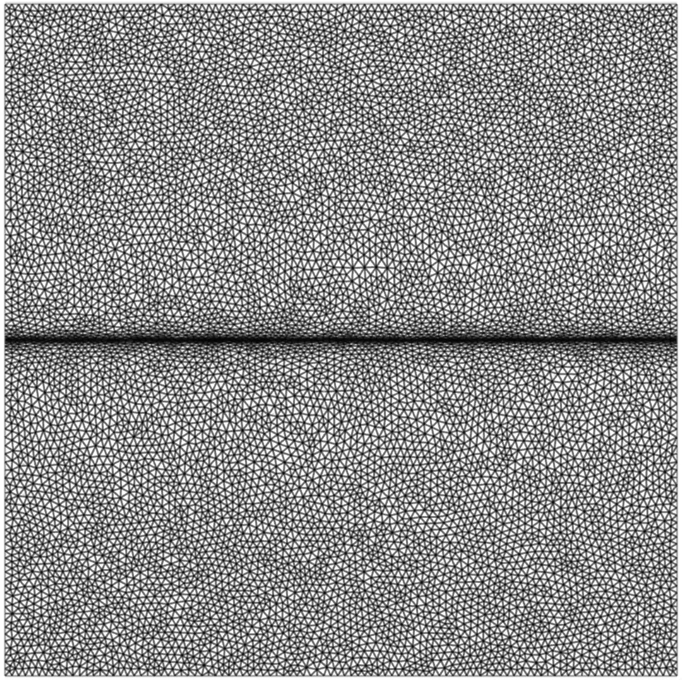

Solve a system of equations for node locations of the form

$$\frac{\partial x}{\partial \tau} = \nabla' \cdot (w \nabla' x)$$

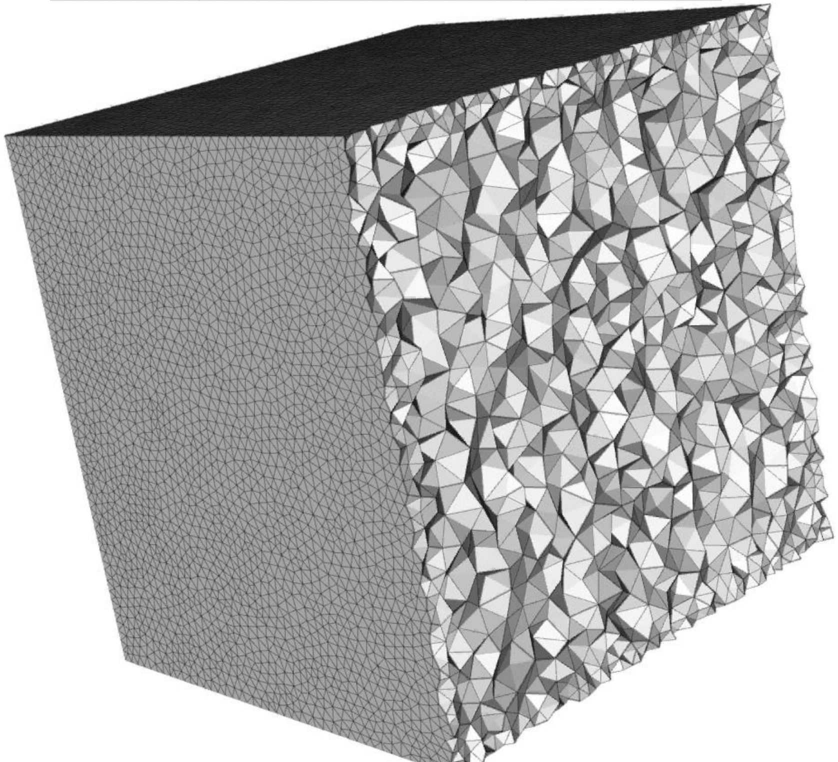

where w is a monitor function which is large in regions requiring finer resolution, e.g. Ceniceros & Hou 2000.

CFD validation: Comparisons between drag calculation in flow past a sphere at a range of Re numbers

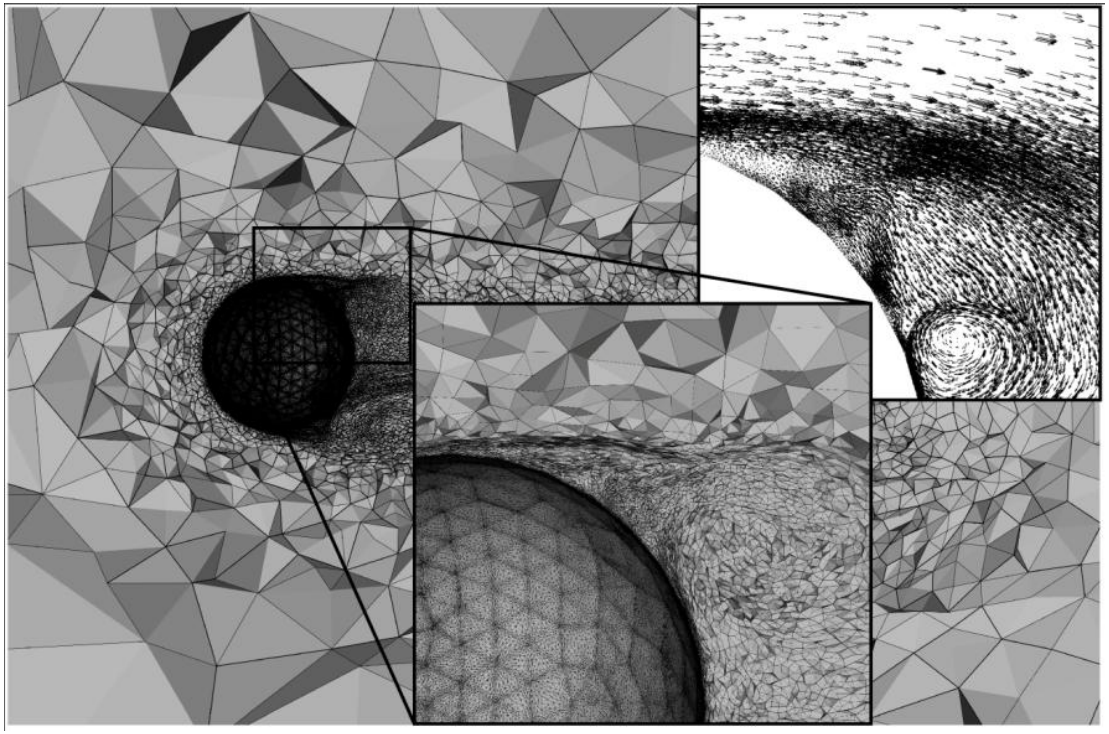

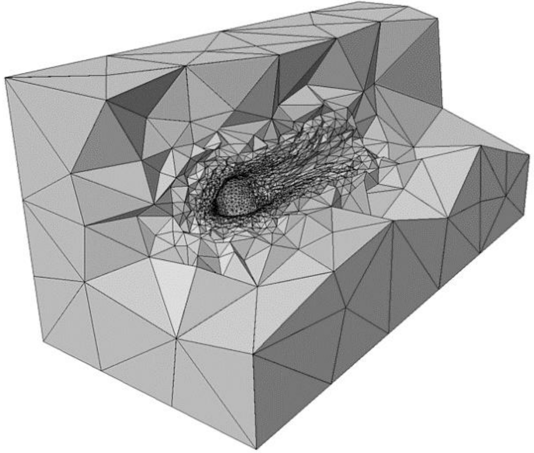

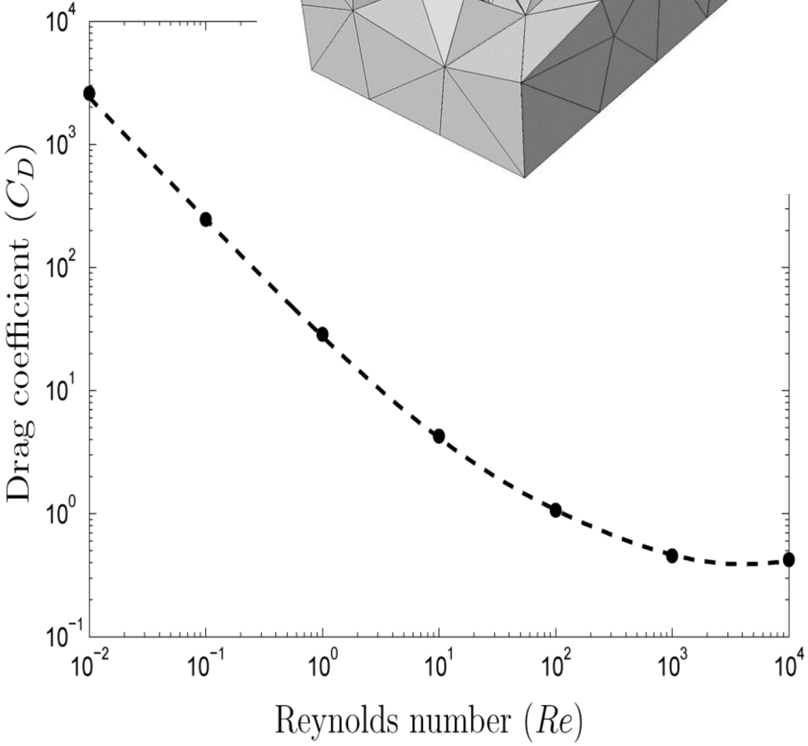

Computed drag coefficient compared against correlation (from Brown and Lawler, 2003) with lab data

$$C_D = \frac{24}{Re}\left(1 + 0.15 Re^{0.681}\right) + \frac{0.407}{1 + \frac{8710}{Re}}.$$

3D backward facing step parallel benchmark – performance optimisation

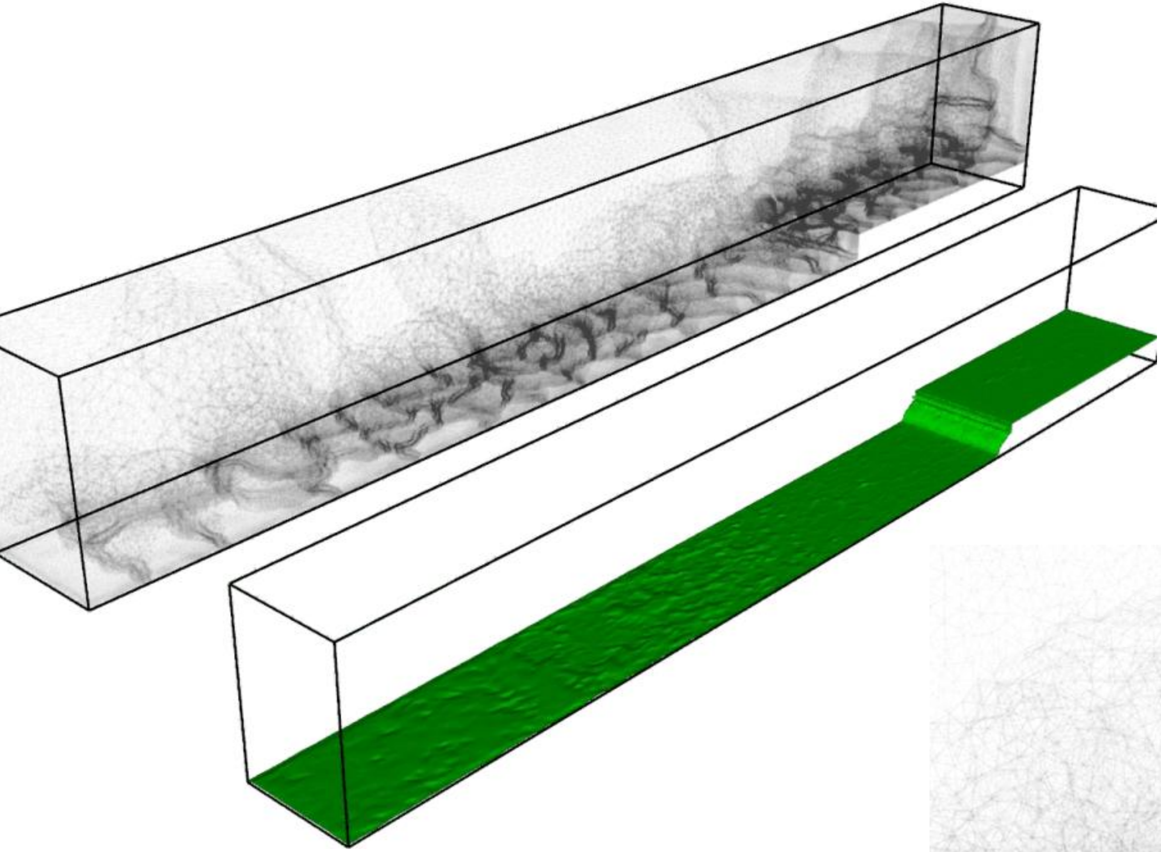

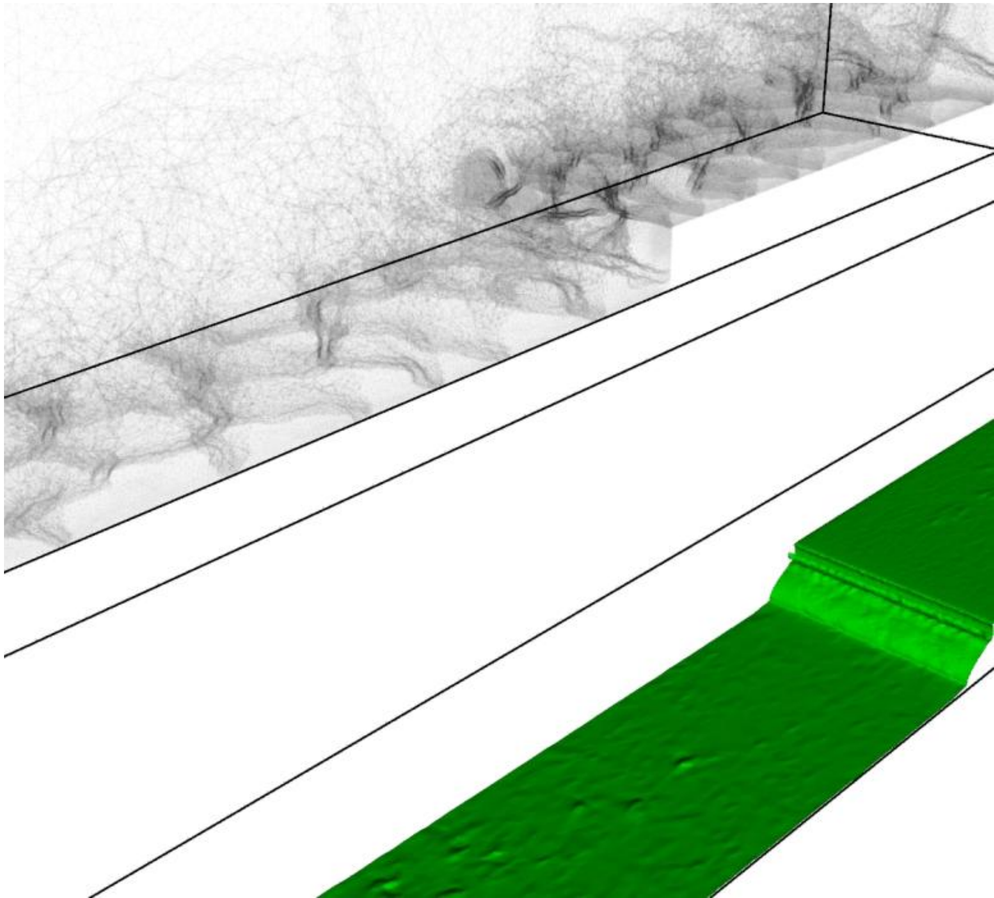

Here the adaptive mesh is focusing in and preserving the integrity of the thermocline in a lab scale simulation of internal wave generation, breaking and mixing over bathymetry

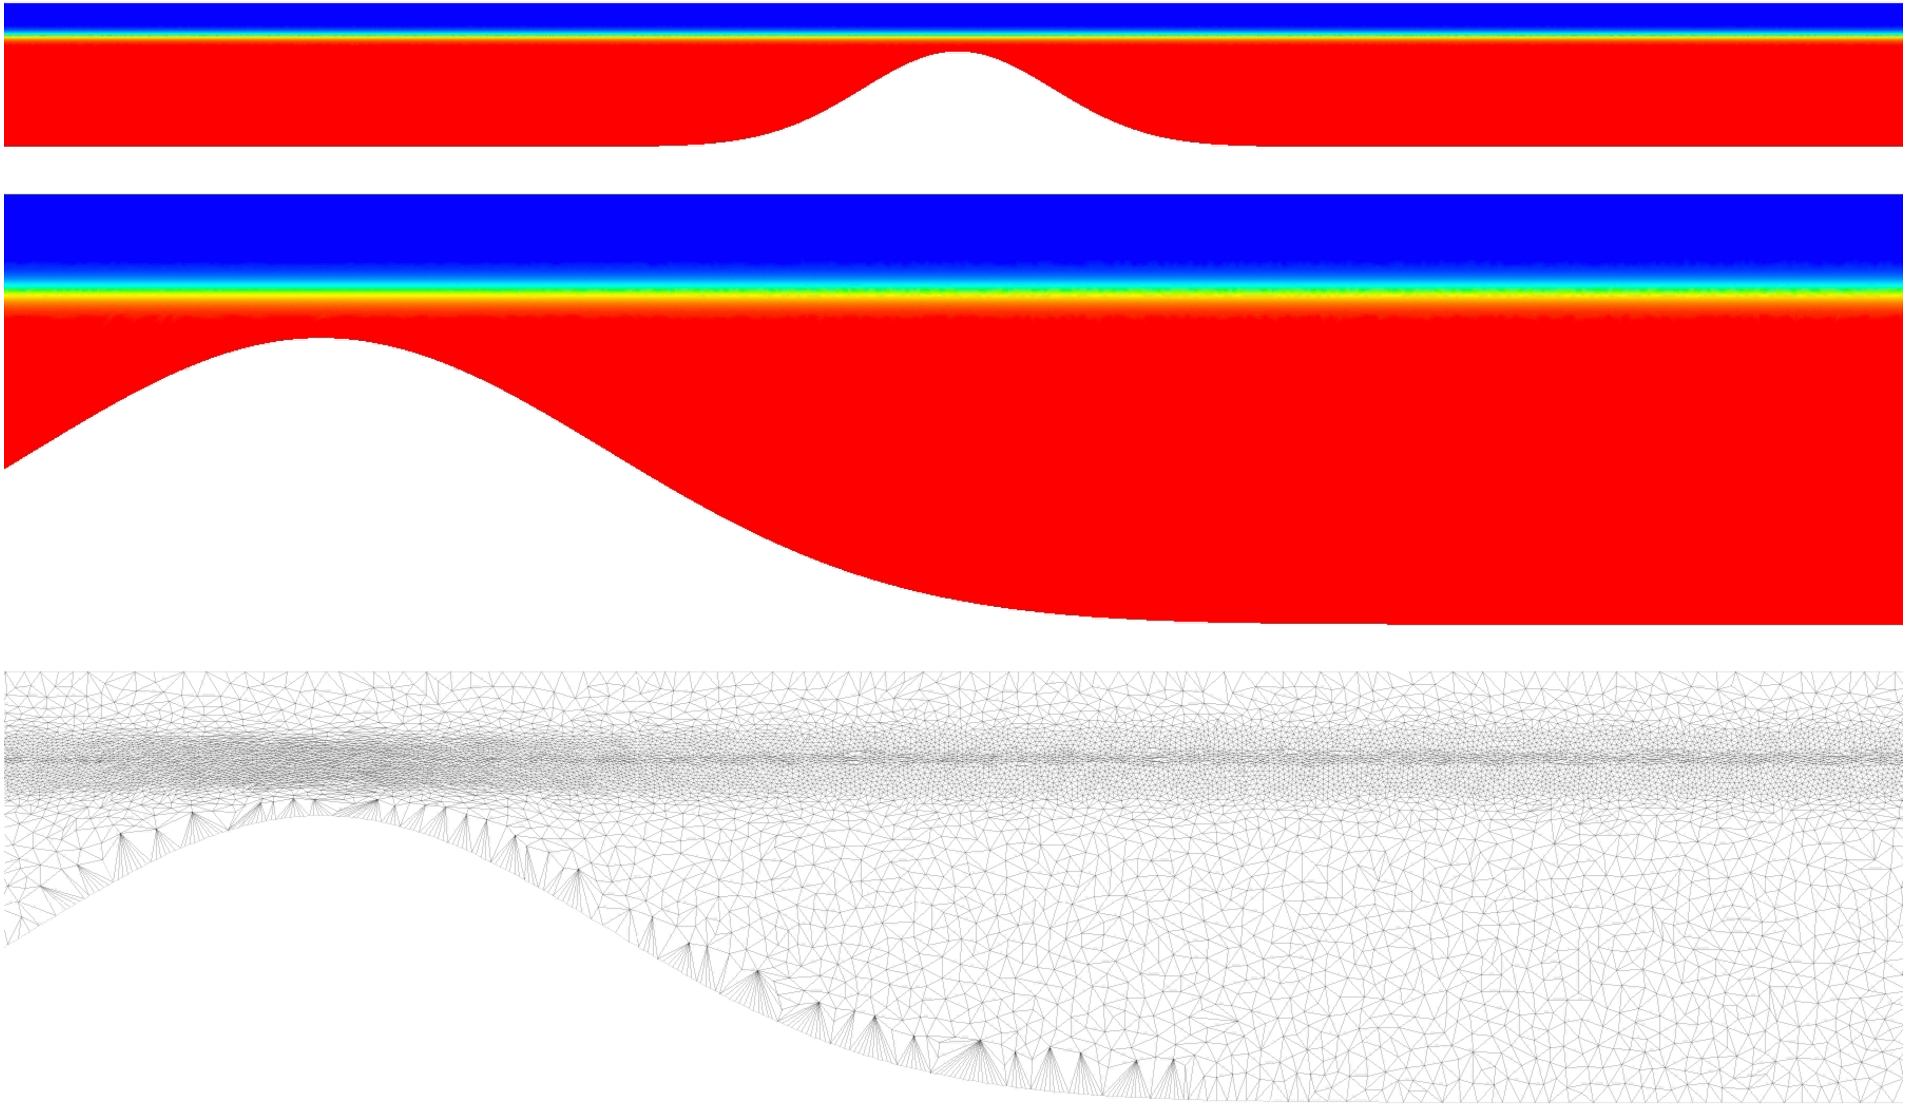

Validation against Laboratory experiments

The 3D differentially heated rotating annulus at two rotation rates. The bounds of the normalised temperature have been limited to aid visualisation at the end of the movie

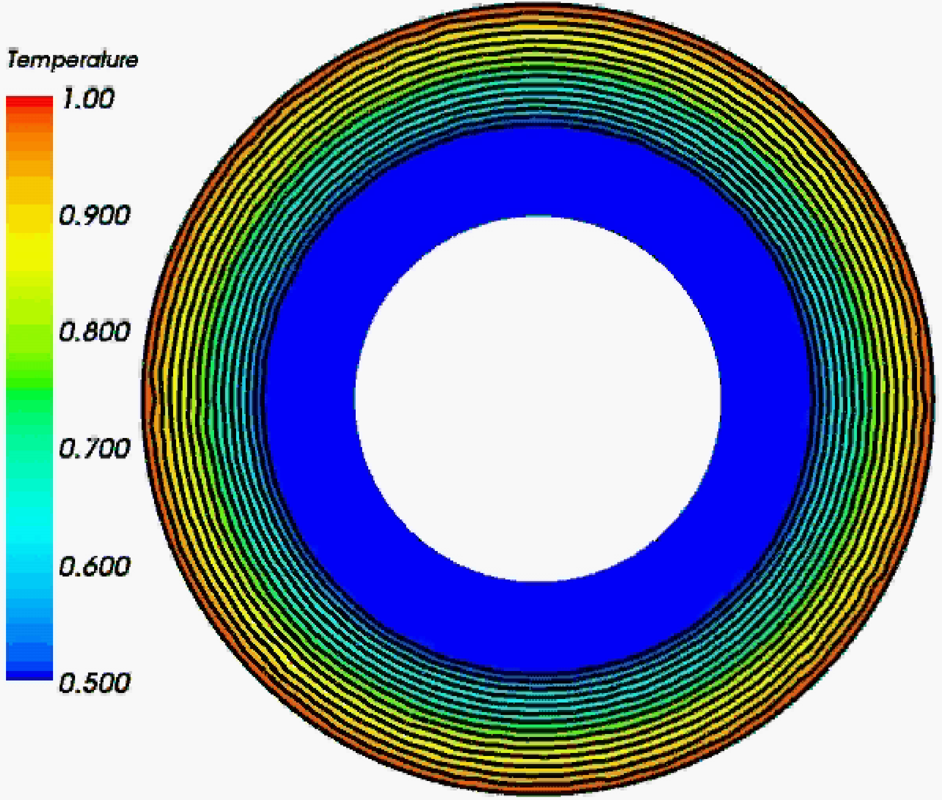

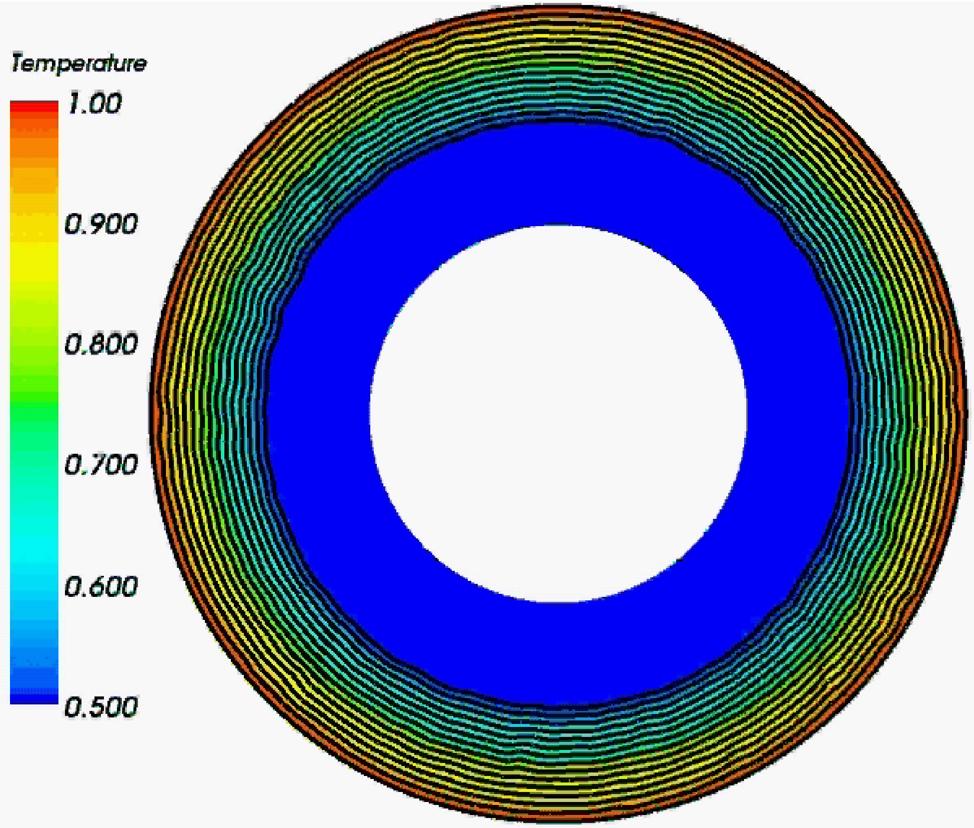

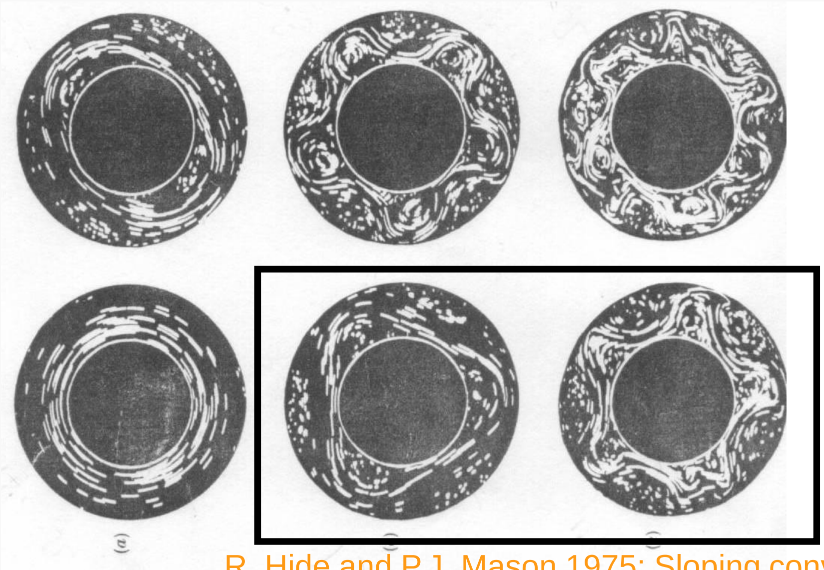

R. Hide and P.J. Mason 1975: Sloping convection in a rotating fluid. Adv. Phys. 24, 47-100

Two snapshots in time at the faster rotation rate

Recent quantitative comparisons with heat transport data from lab runs shows that it is hard to beat an optimised fixed mesh at lower rotation rates, but adaptivity does seem to be needed for higher rotation rates

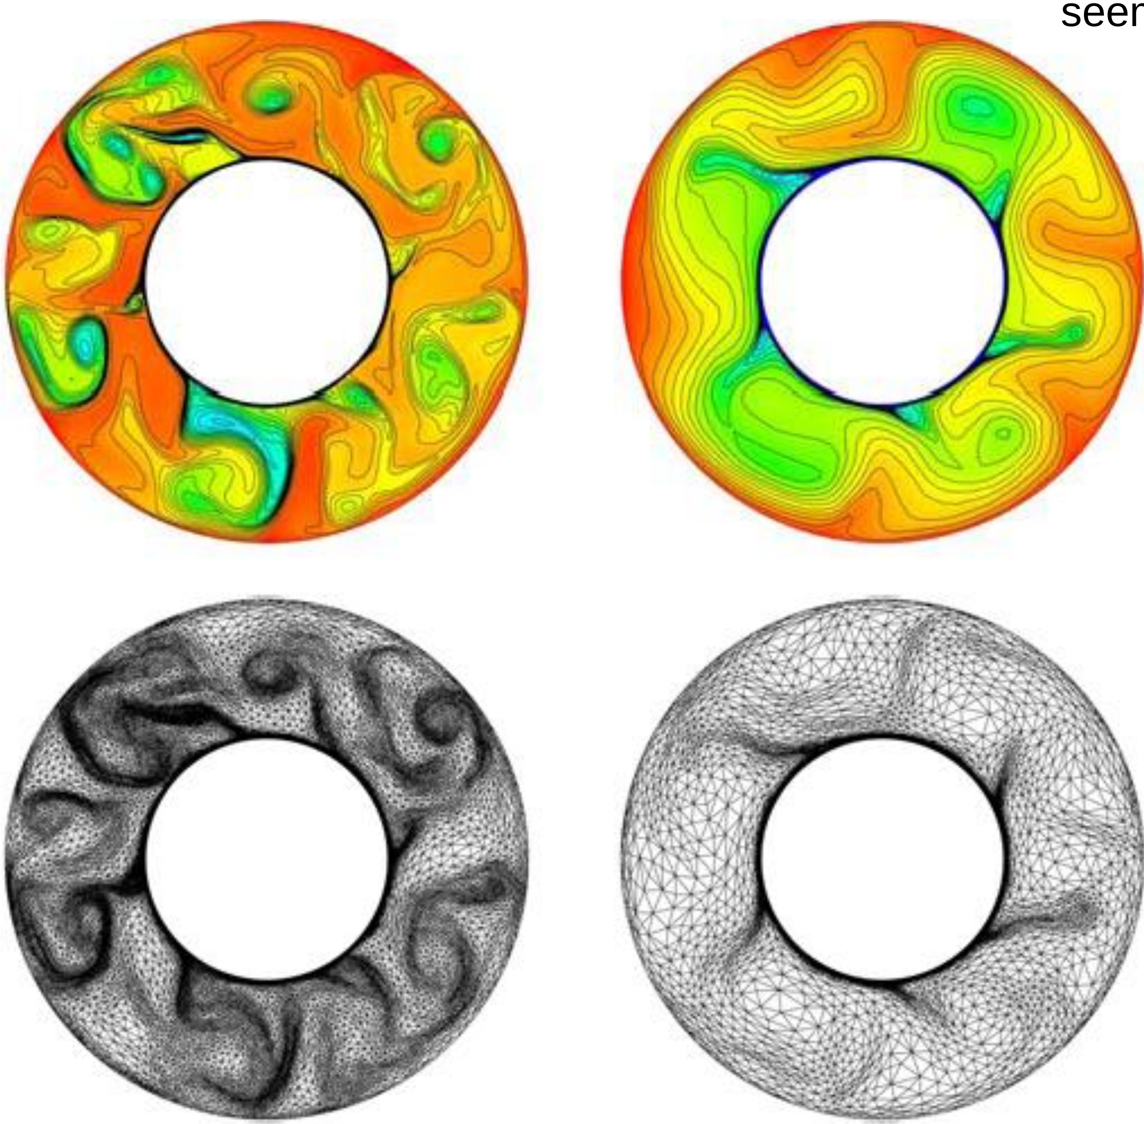

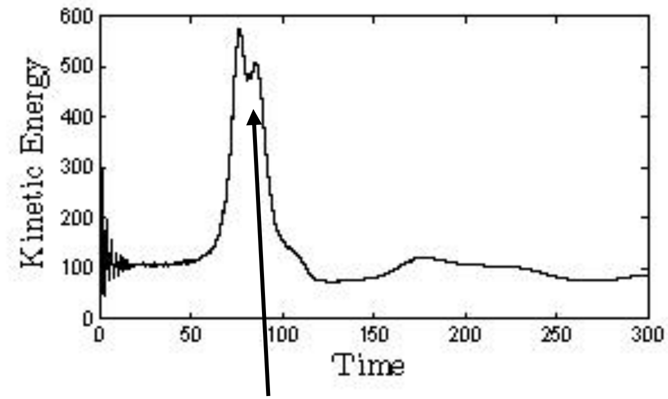

Here the maximum computational cost corresponds to the time of maximum kinetic energy

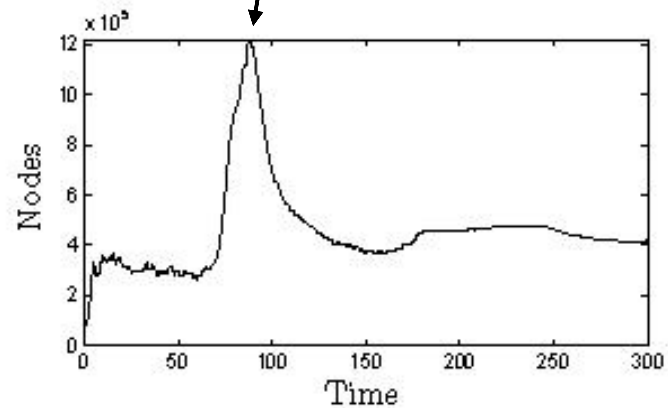

OODC: An isosurface of density coloured by speed (Garcia-Sagrado)

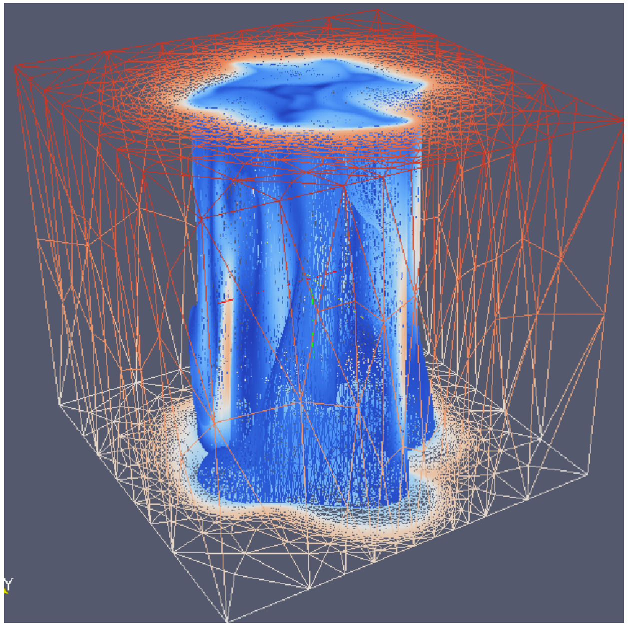

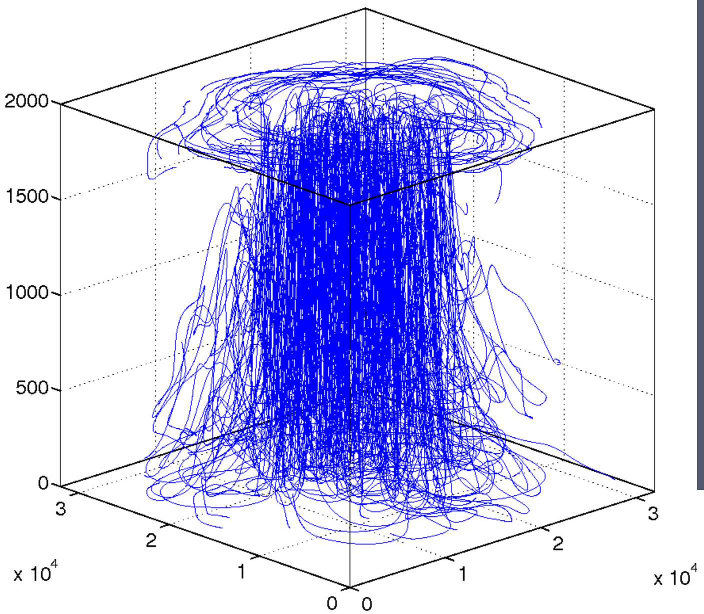

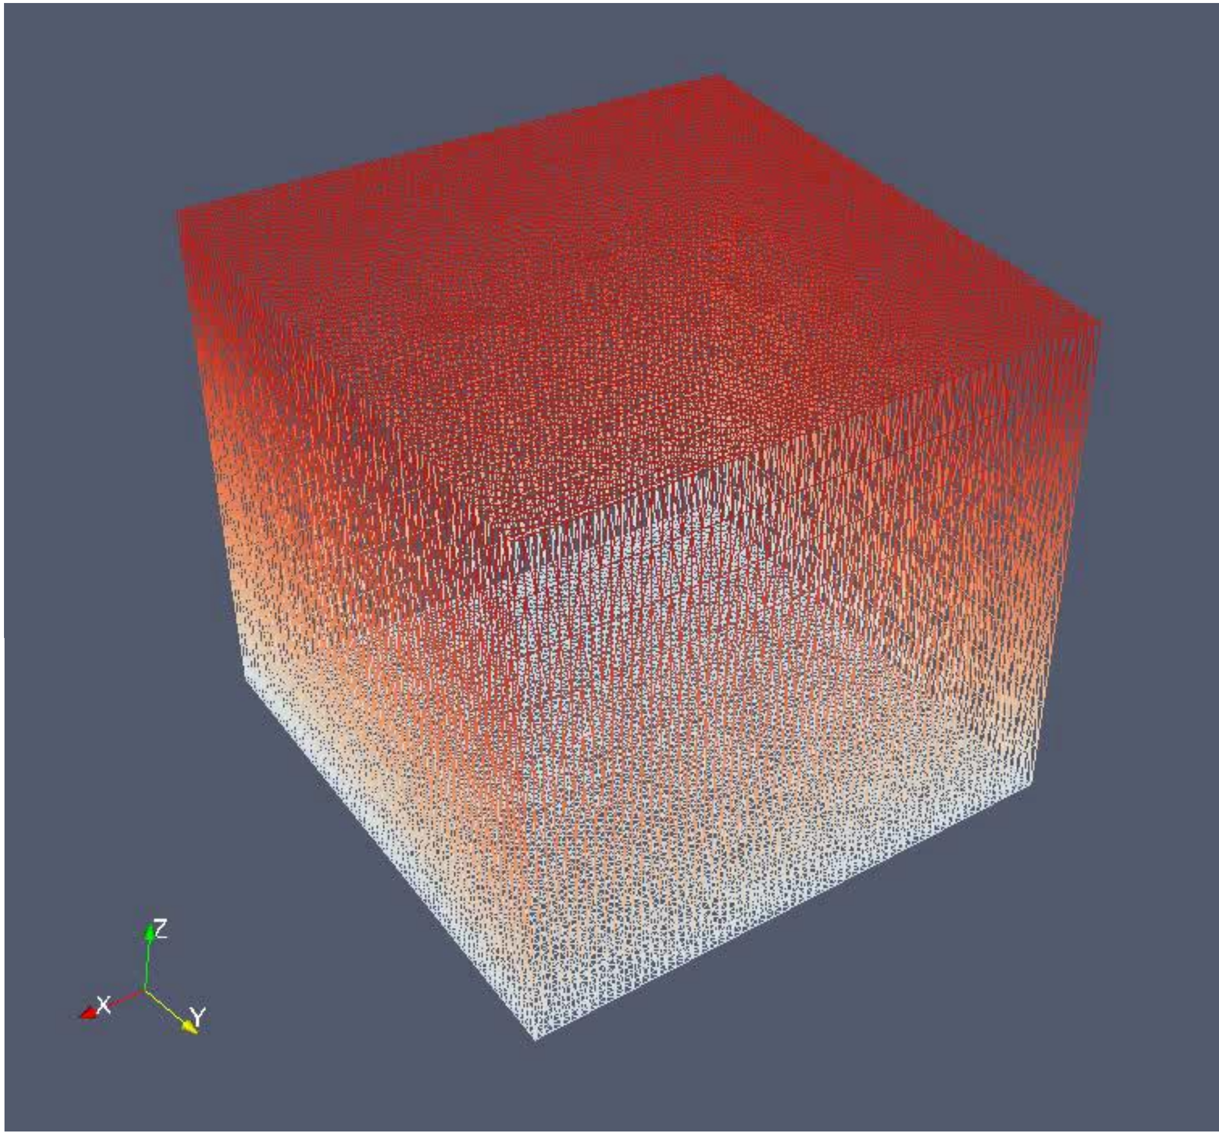

GFD example: The lock exchange problem, checking head speeds

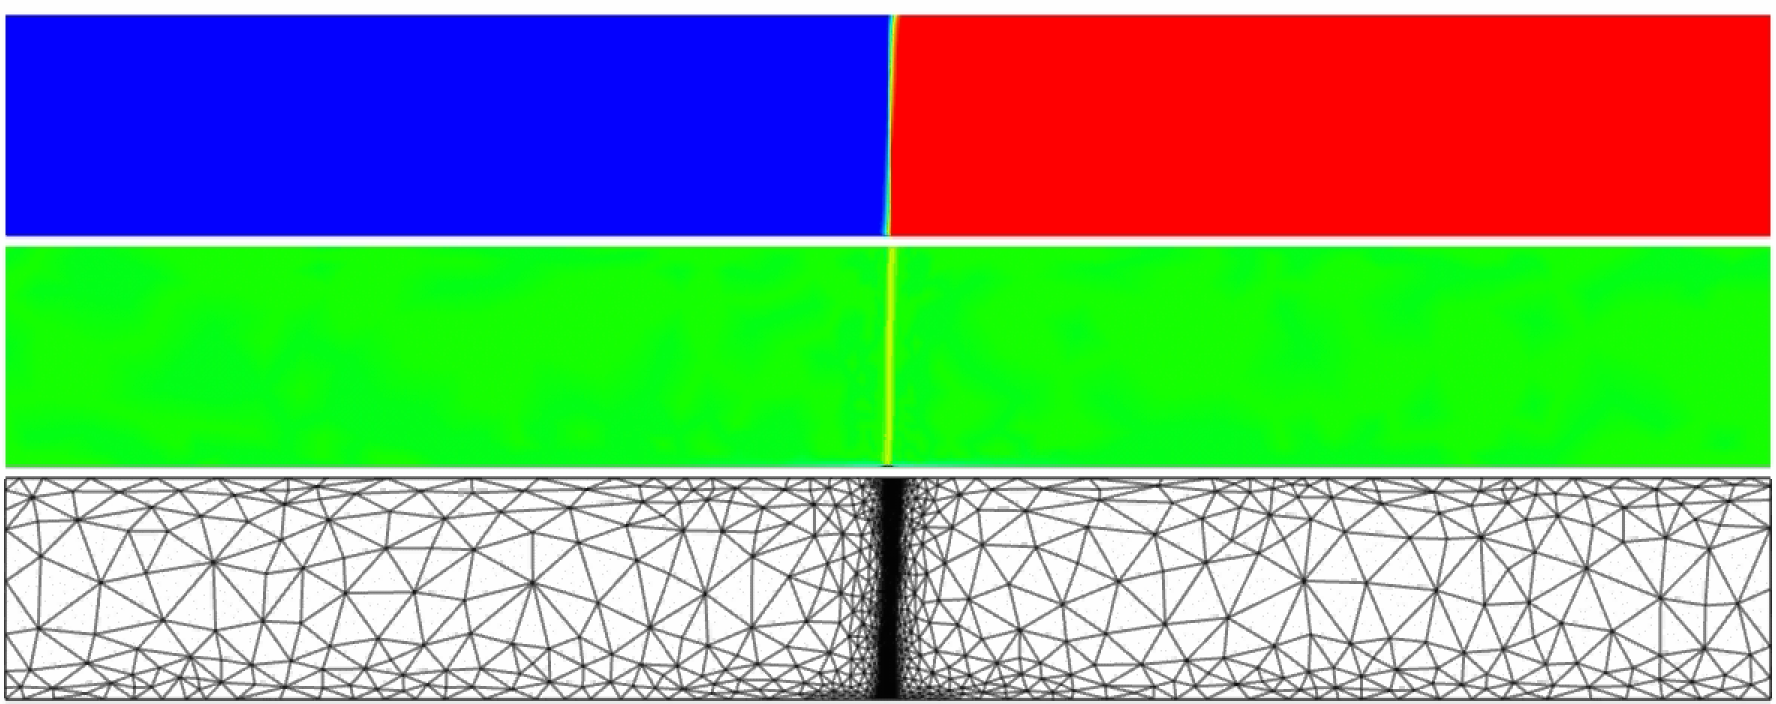

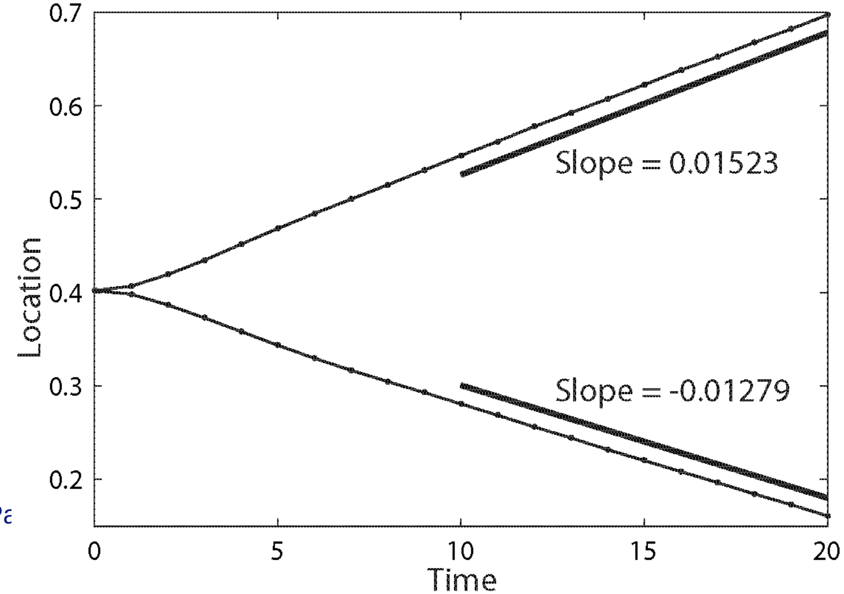

Compare speed of head at free and no slip boundaries against DNS results (Hartel et al, 2000): 0.01509 & -0.01284.

(Hiester et al., 2009)

Convergence with mesh resolution, and mesh roll up in a KH billow

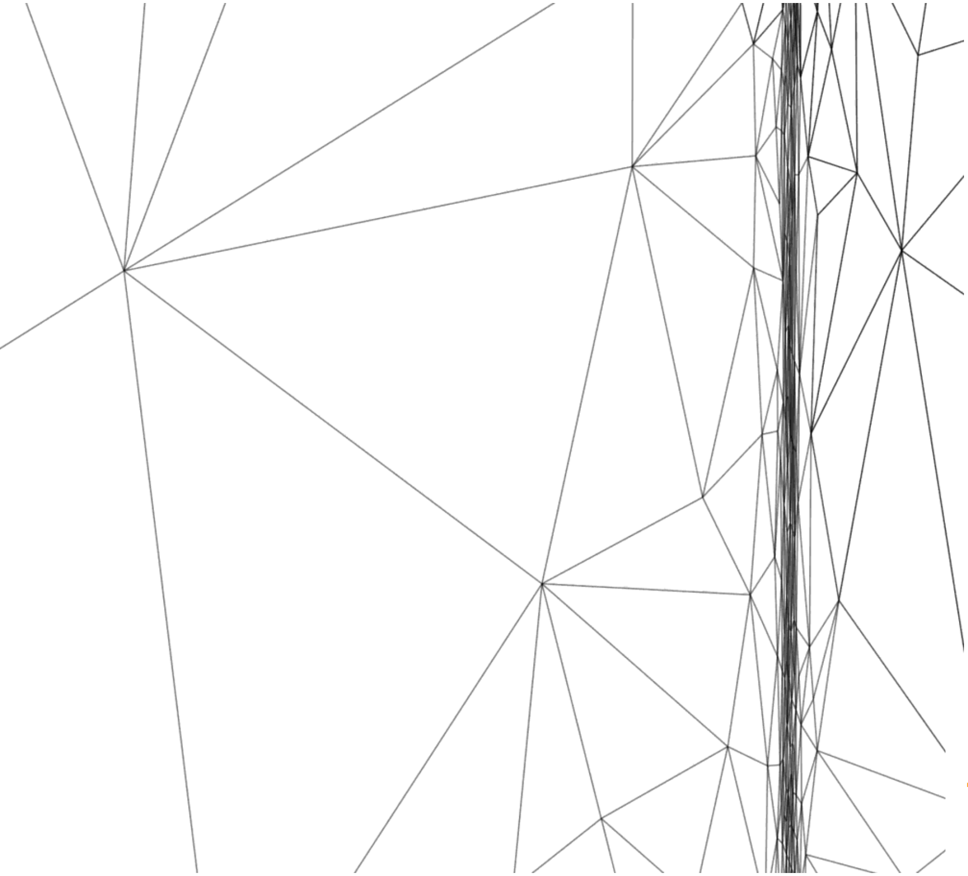

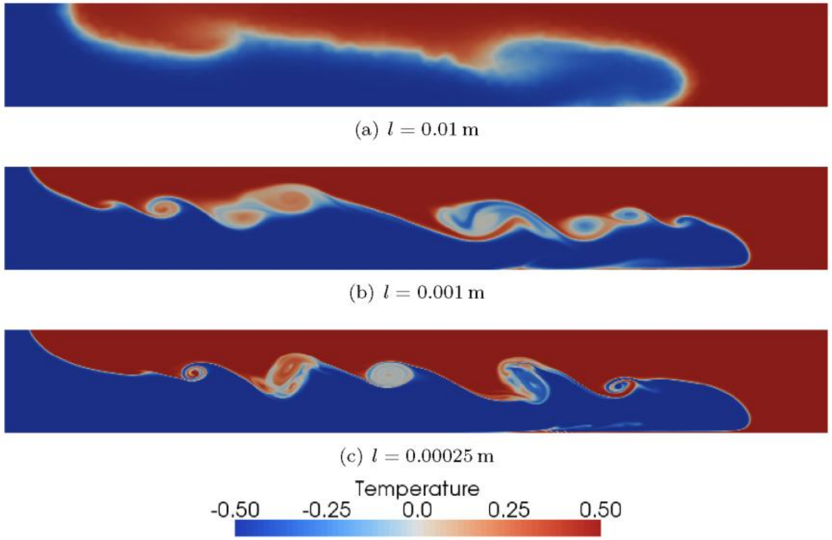

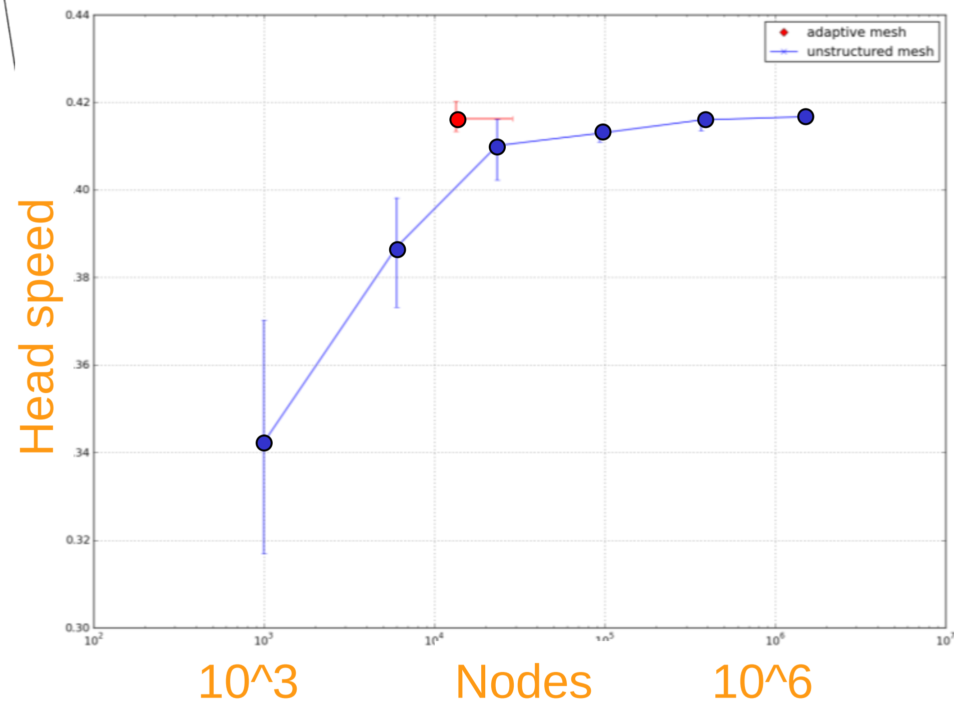

Should help substantially with spurious numerical mixing – currently being investigated

Some conclusions on mesh adaptivity options/parameter choices from a suite of tests:

- Very important to get correct resolution in both no- and free-slip boundaries
- Period of adapt important, so that we don't advect the interface into a region of coarse resolution
- Advecting the metric forward in time to guess where resolution will be needed can help
- Use of Galerkin projection based interpolation not observed to improve head speeds
- But may for mixing rates which are currently being examined
- Adaptivity here typically <10% of run time and gives savings of 1-2 order of magnitude in number of nodes for given accuracy

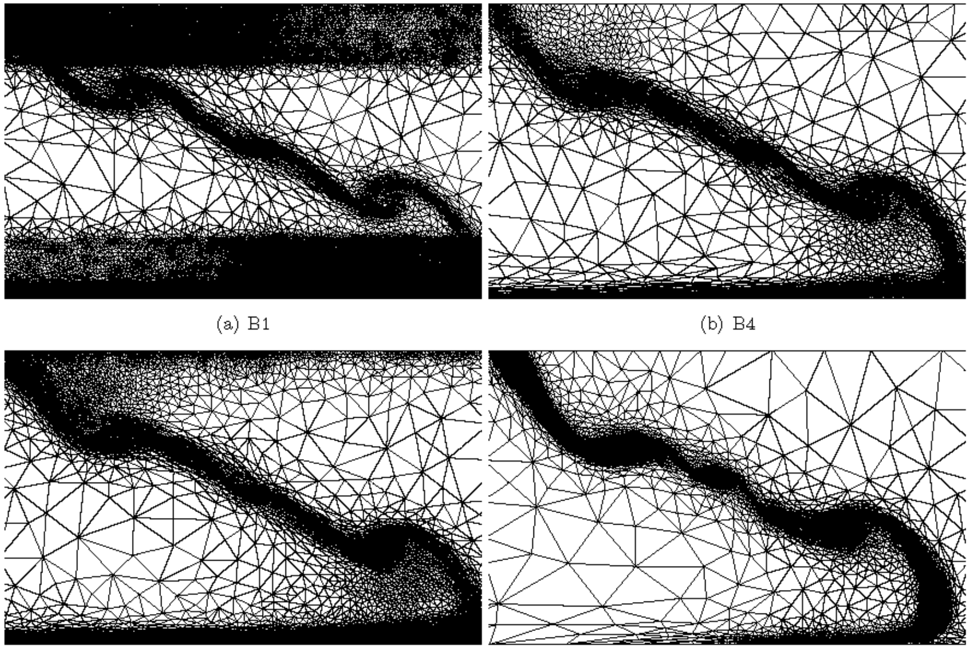

(a) B1 (b) B4

(c) B5 (d) B13

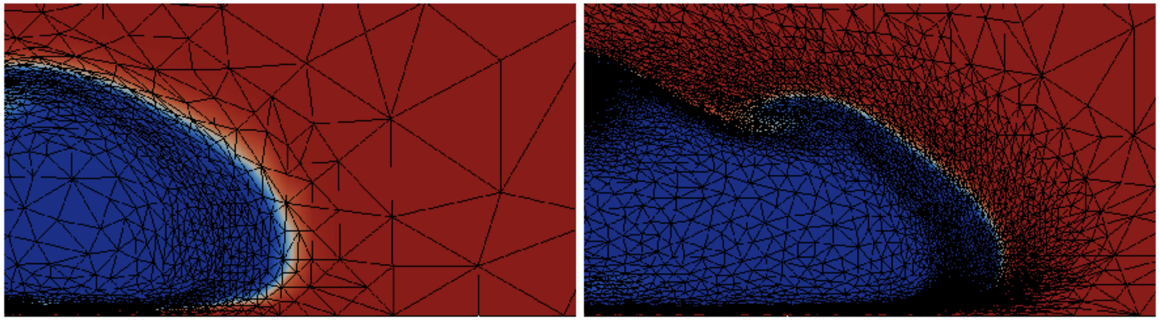

(a) no metric-advection, A4 (b) metric-advection, A5

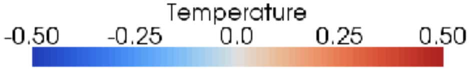

Investigation of head speed, entrainment and mixing in buoyancy driven gravity currents (Hiester et al.)

Head speed in adaptive results shows weak dependence on slope, as seen in experiments

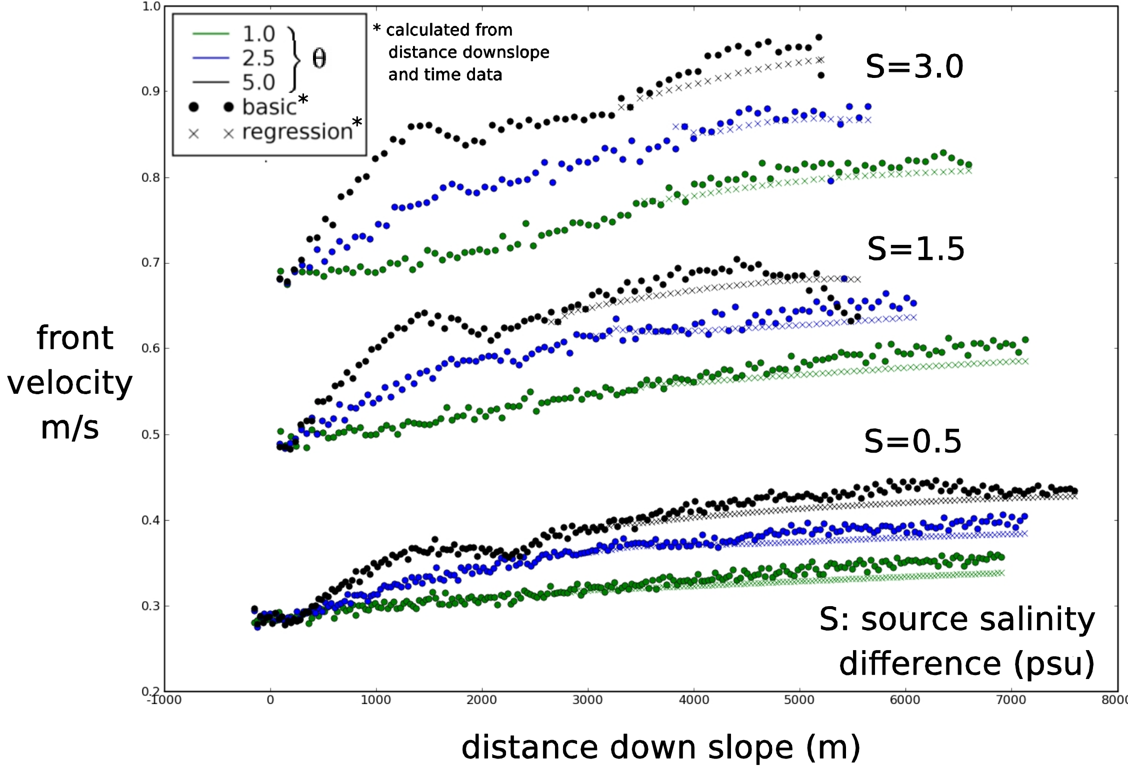

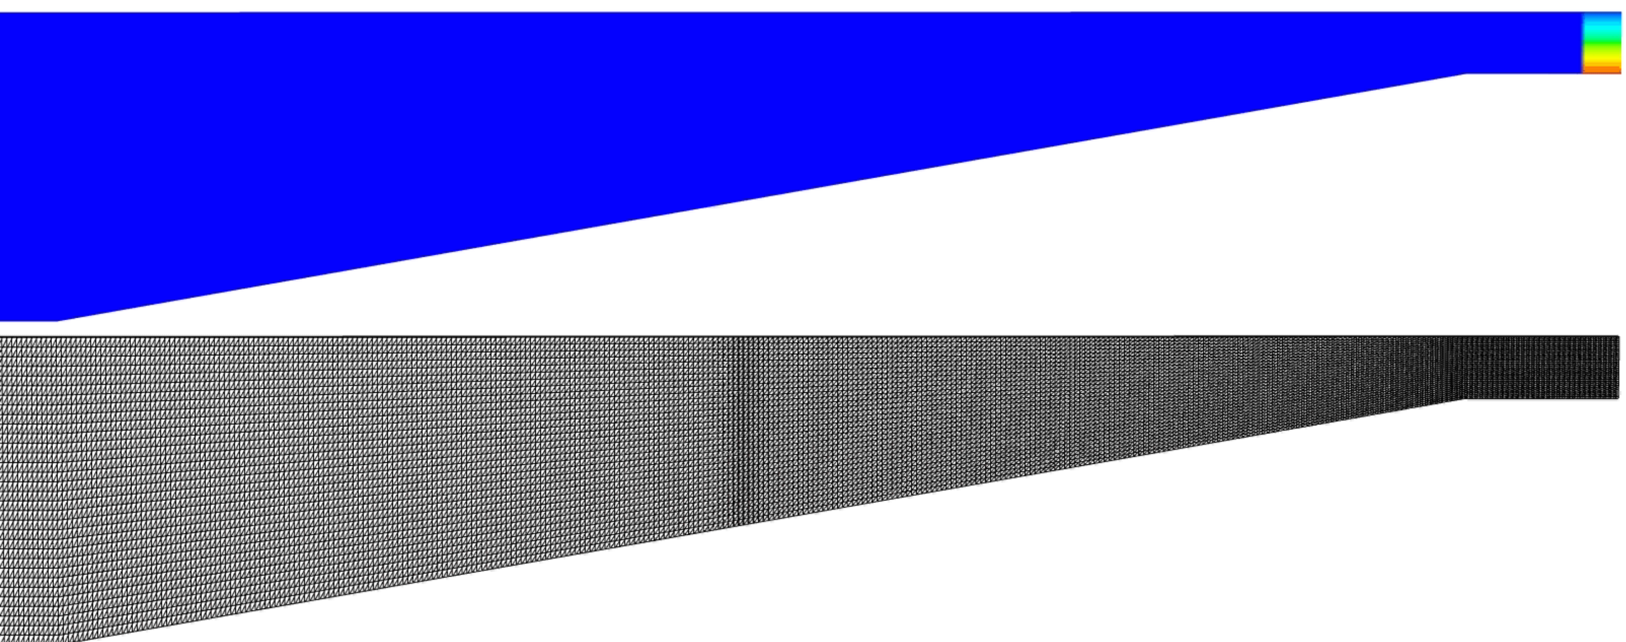

Setup following Ozgokmen et al.

Lab work by Britter et al.

Salinity driven gravity current with linear background temperature stratification

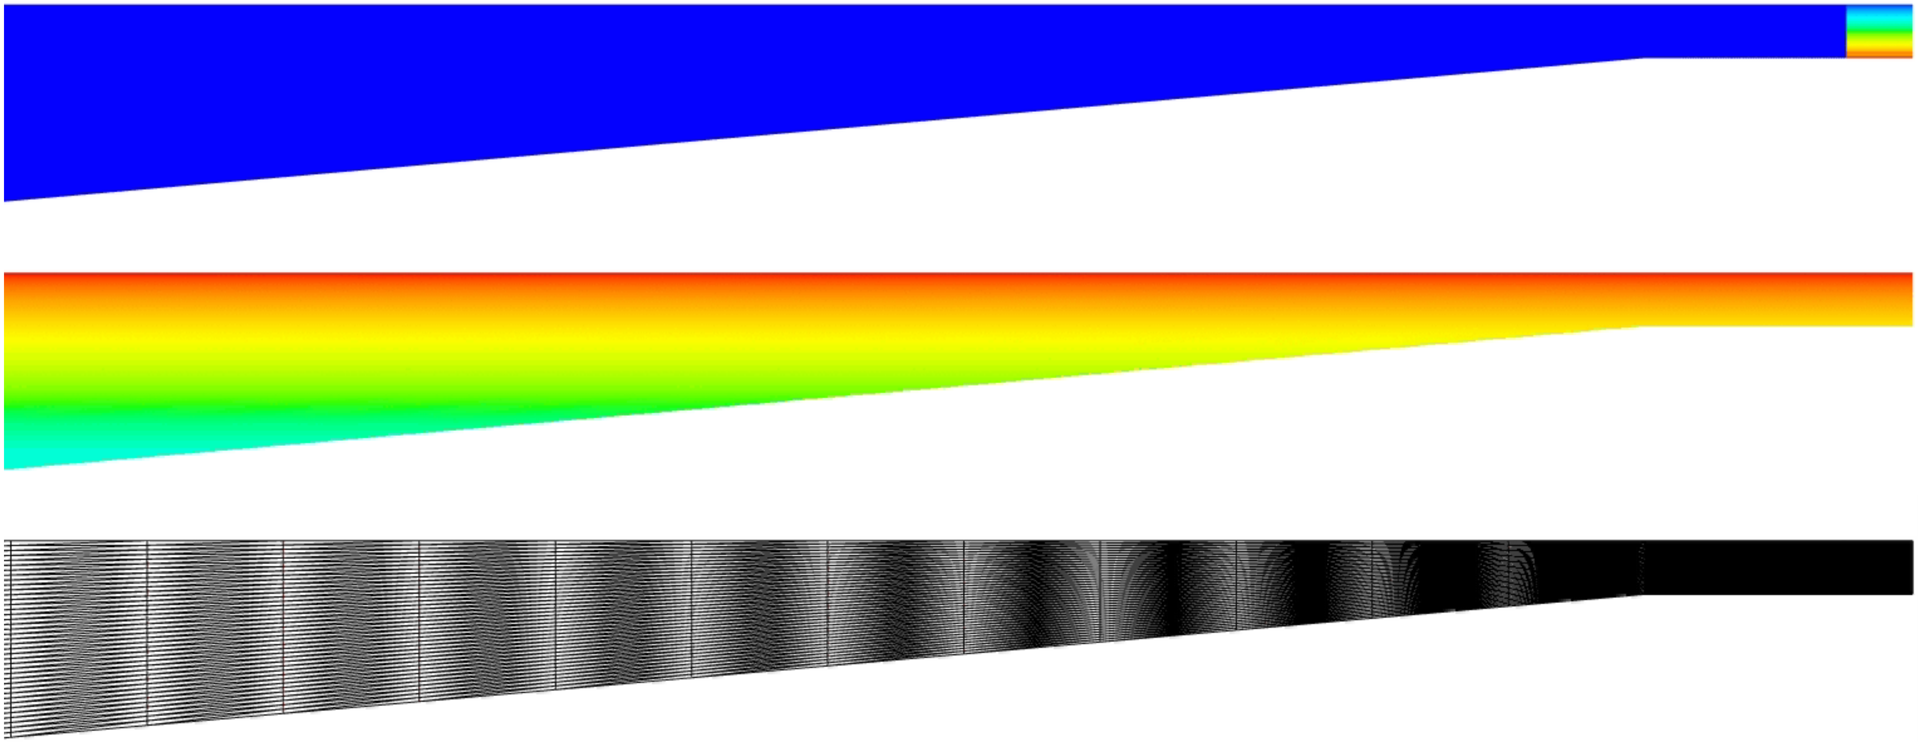

Sediment transport with 4 grain size classes

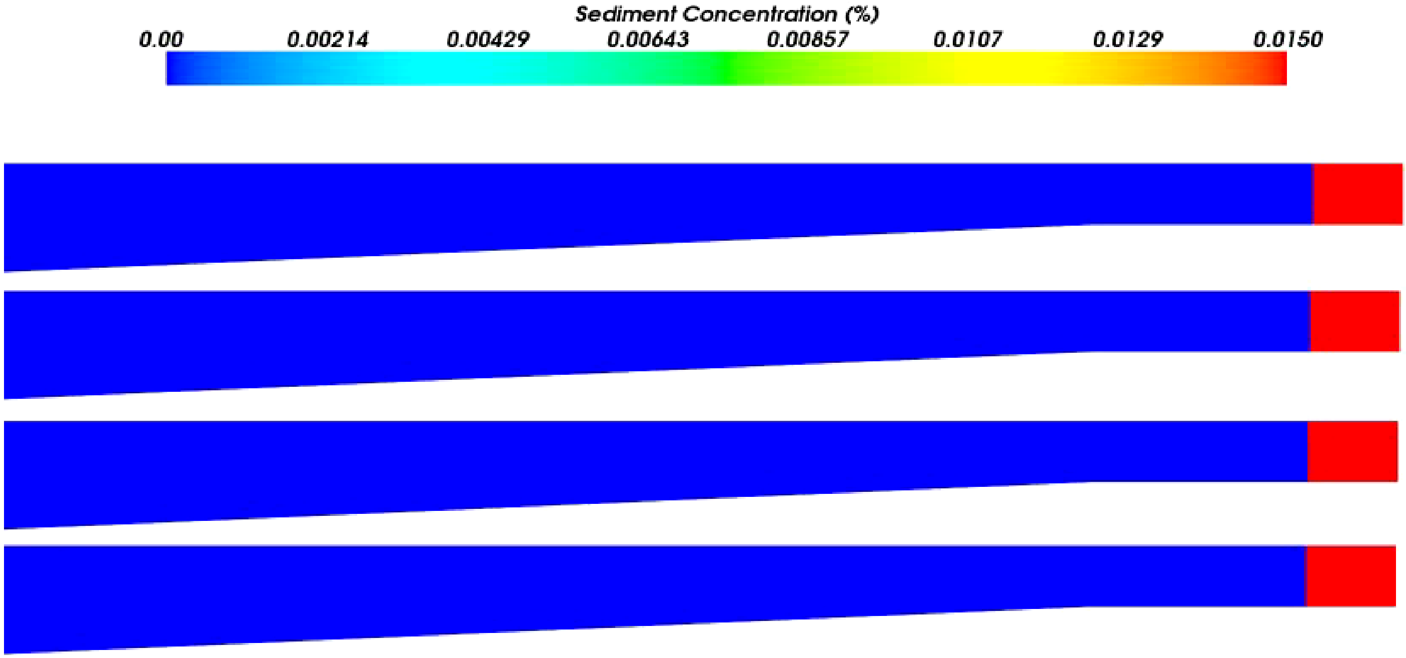

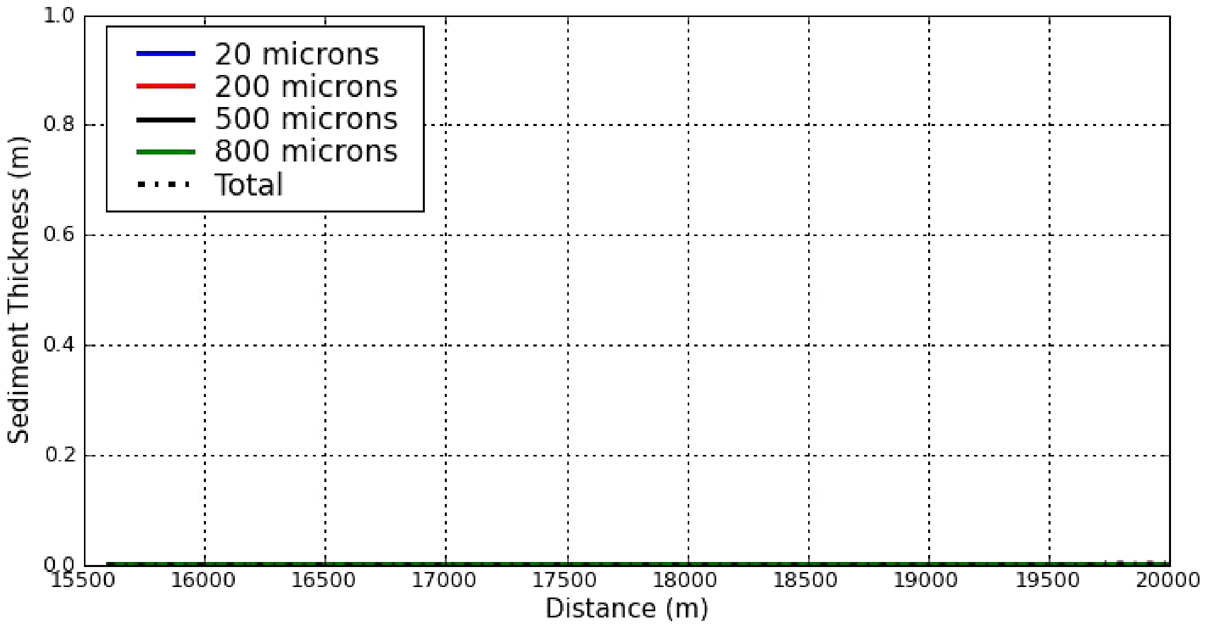

Sediment transport with temperature induced buoyancy reversal

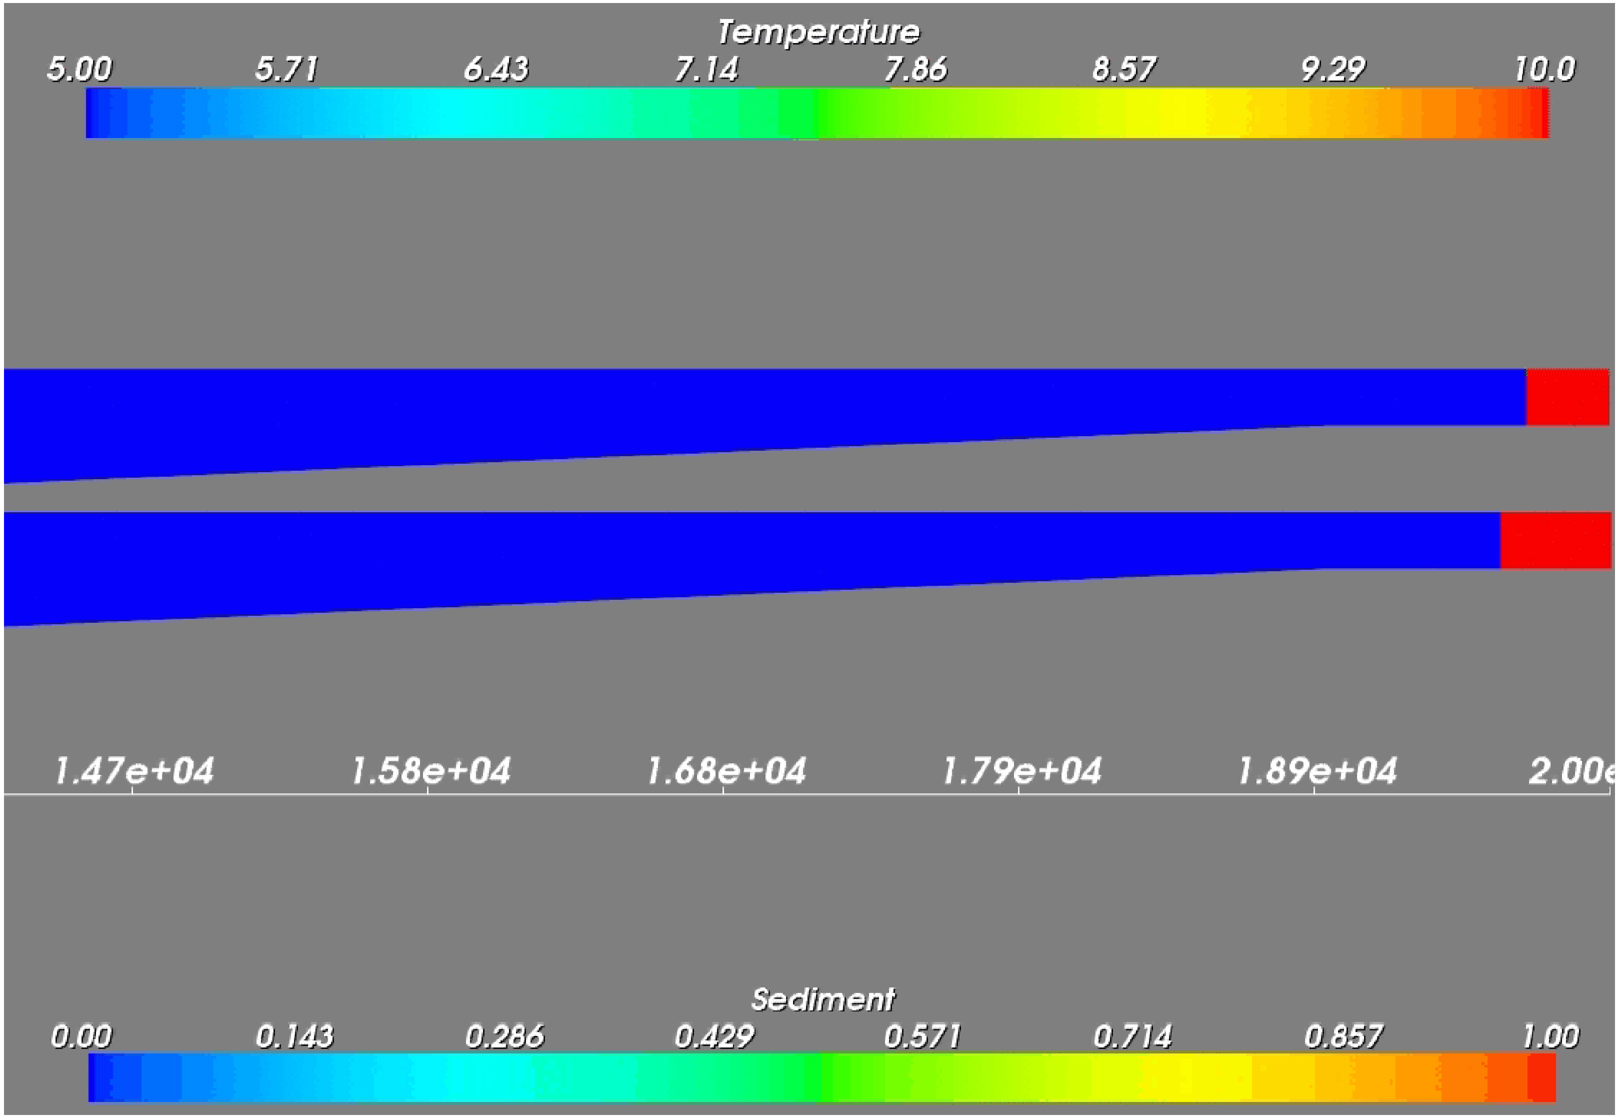

Many important boundary layers occur in geophysical applications – e.g. westward intensified boundary currents

Stommel (1948) derived a streamfunction equation following a number of assumptions about the flow dynamics:

$$r\nabla^2\psi + \beta\frac{\partial\psi}{\partial x} = -\frac{\partial\tau}{\partial y}$$

for which we have an analytical solution for comparison, cf. Hecht et al., 2000.

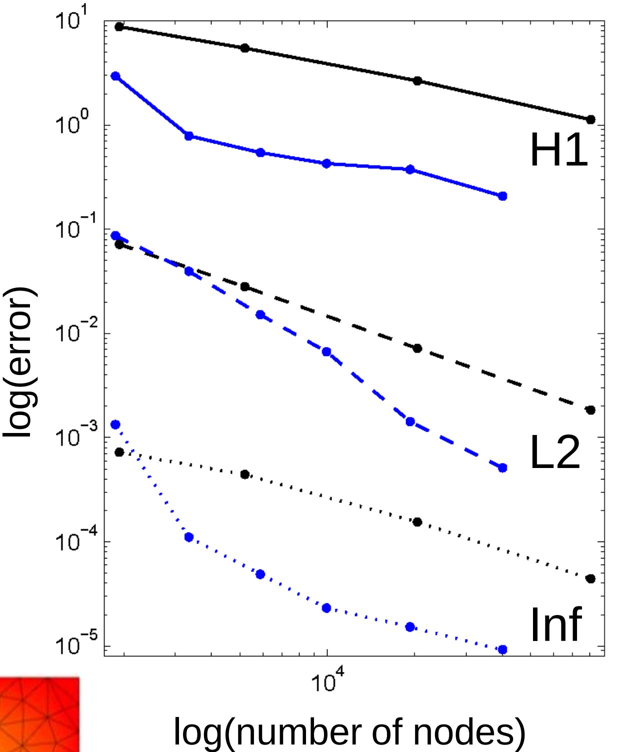

Errors in 3 norms for: uniform resolution (black) and adaptive resolution (blue)

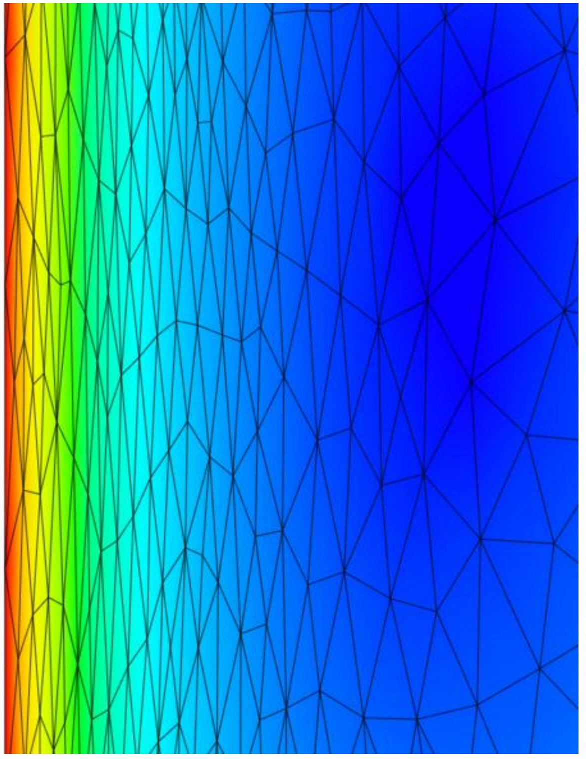

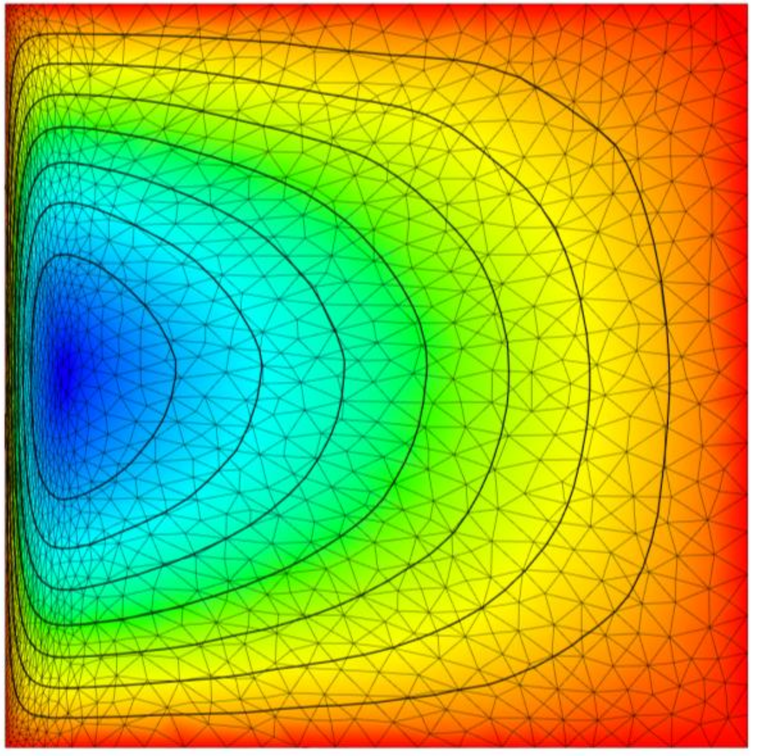

Intense western boundary current resolved with anisotropic adaptive resolution

Effect of varying the maximum allowed aspect ratio in the adapted mesh

Approximately a two orders of magnitude improvement in the error/cost relationship between uniform and anisotropic adaptive refinement

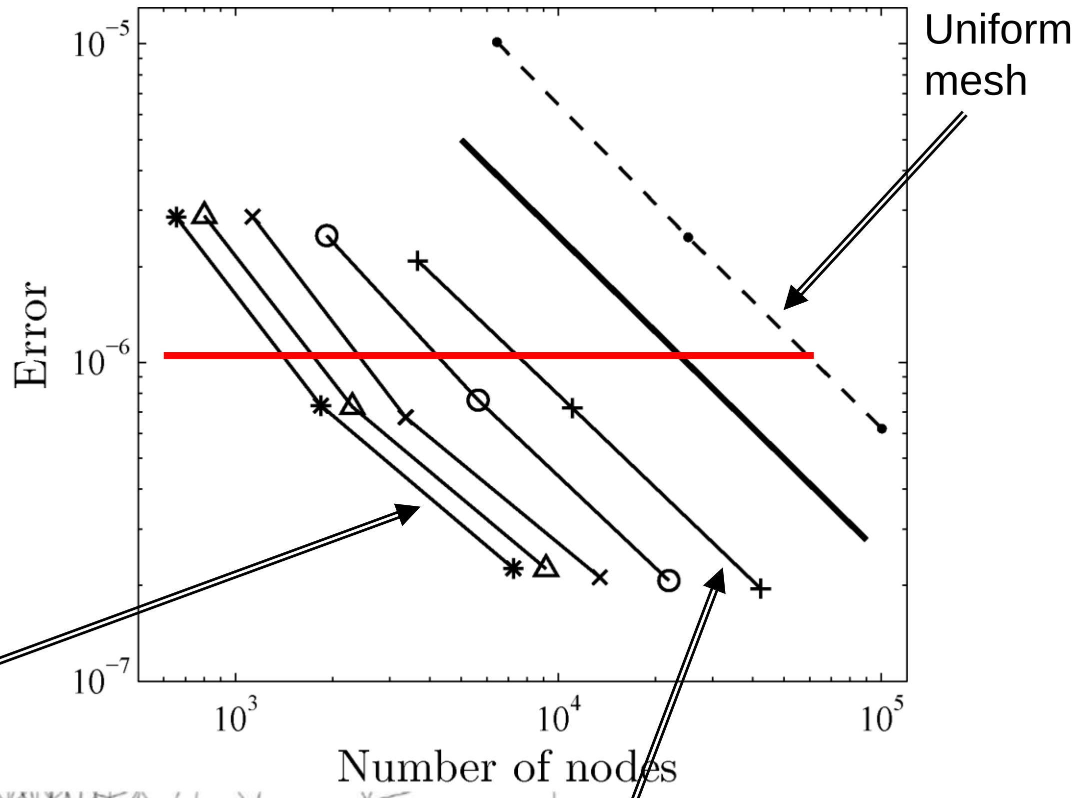

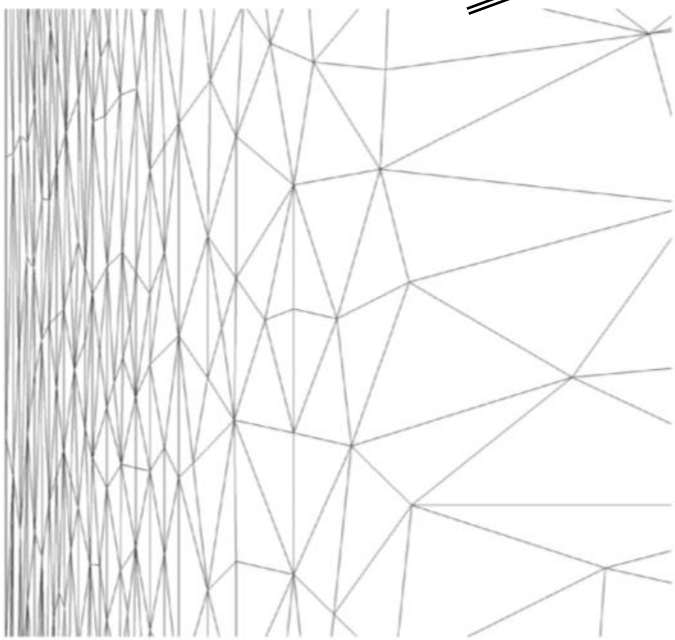

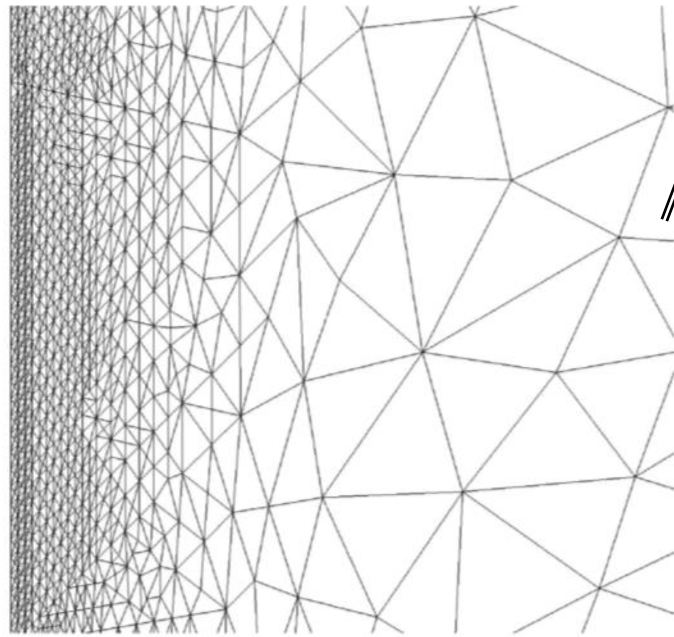

Advecting a tracer field through the boundary layer (cf. Hecht et al, Hanert et al)

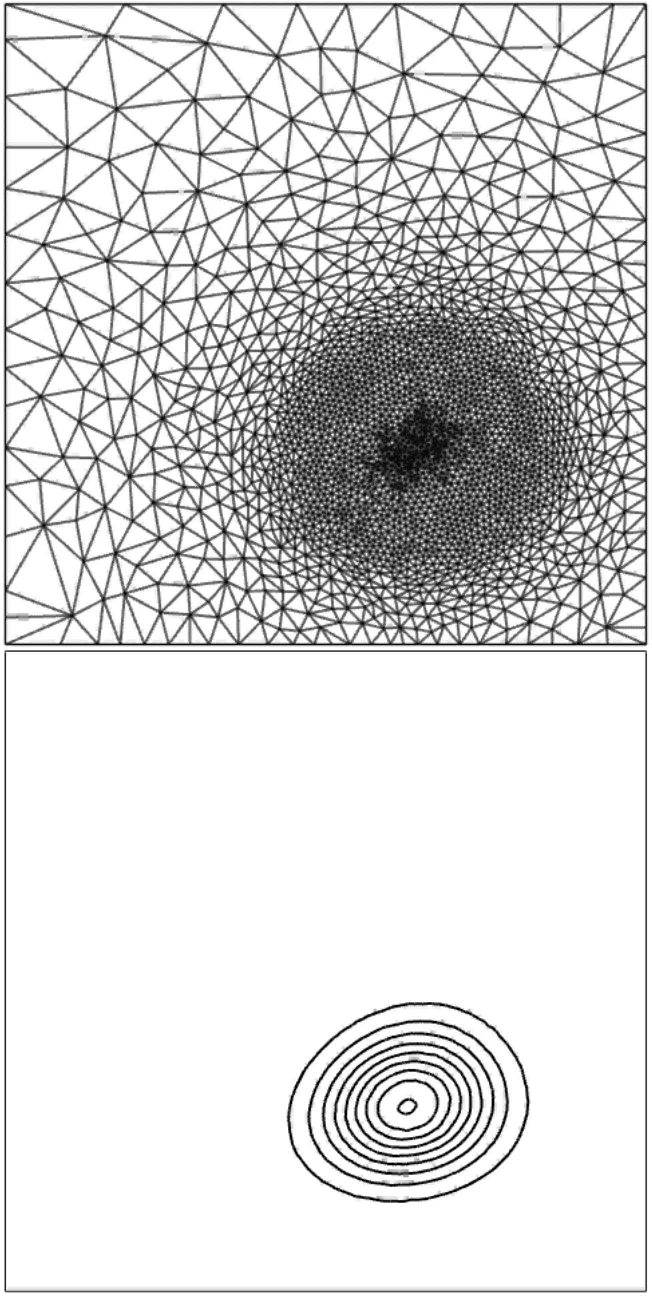

Initial Gaussian tracer field advected with the velocity field from the Stommel gyre – stretching through the boundary layer makes this a tough problem for the advection method.

Number of nodes against time shown right with 5 different error weights.

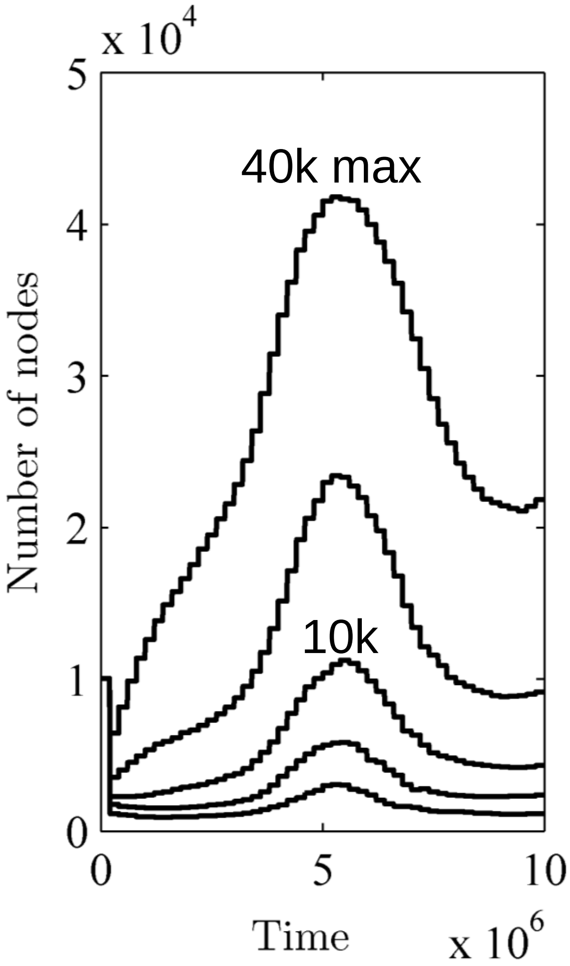

Comparisons between uniform fixed and adaptive simulations

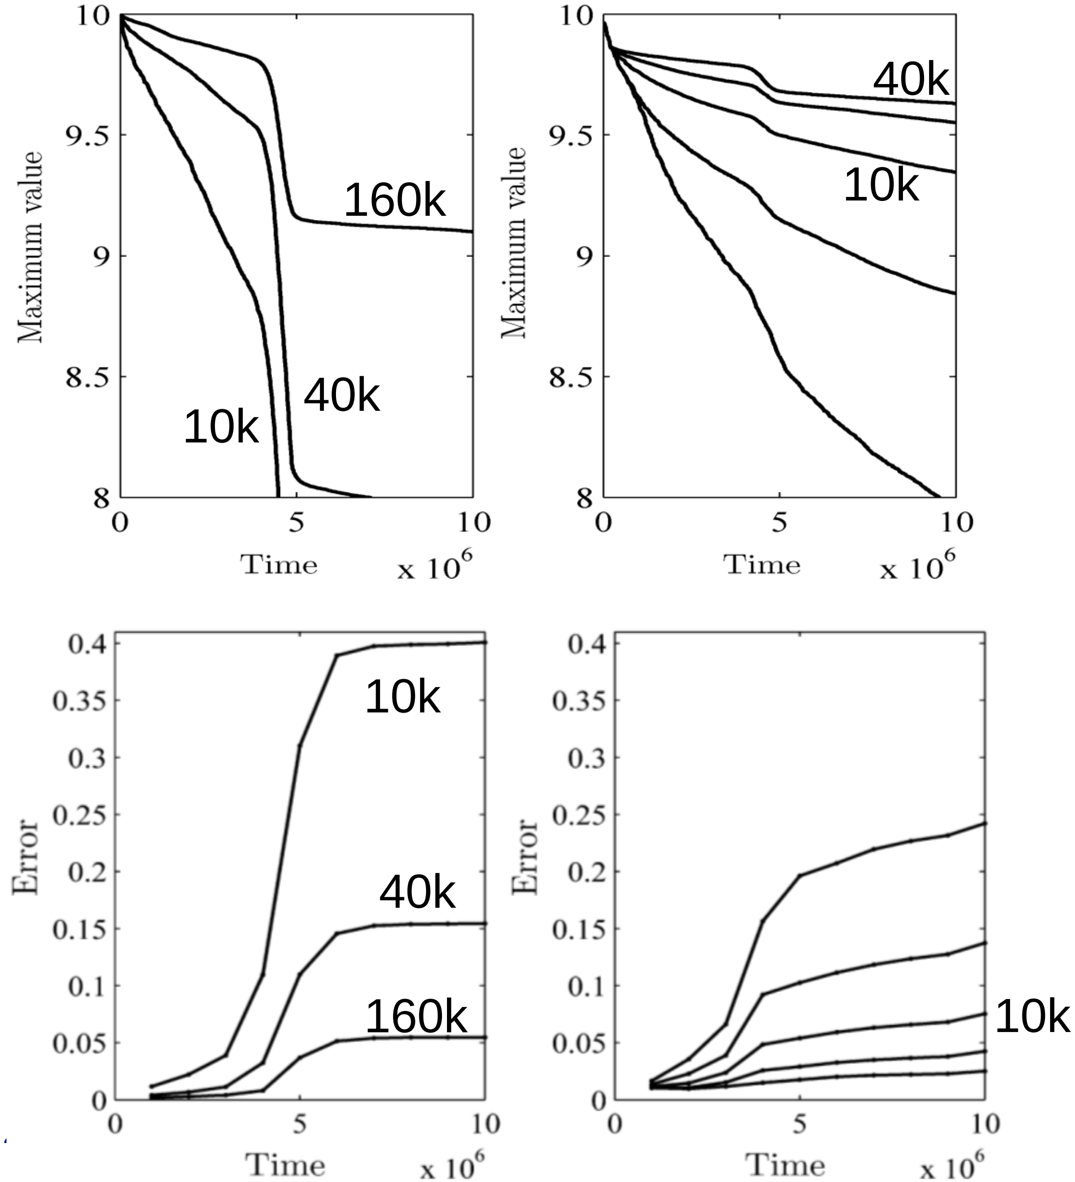

Maximum tracer concentration against time – should be constant at 10.

L2 norm of error compared to an exact solution obtained by integrating back particle paths, against time.

Observe an order of magnitude saving in nodes

Aside: advection of a top hat using DG

We expect DG to become our method of choice for tracers in many applications

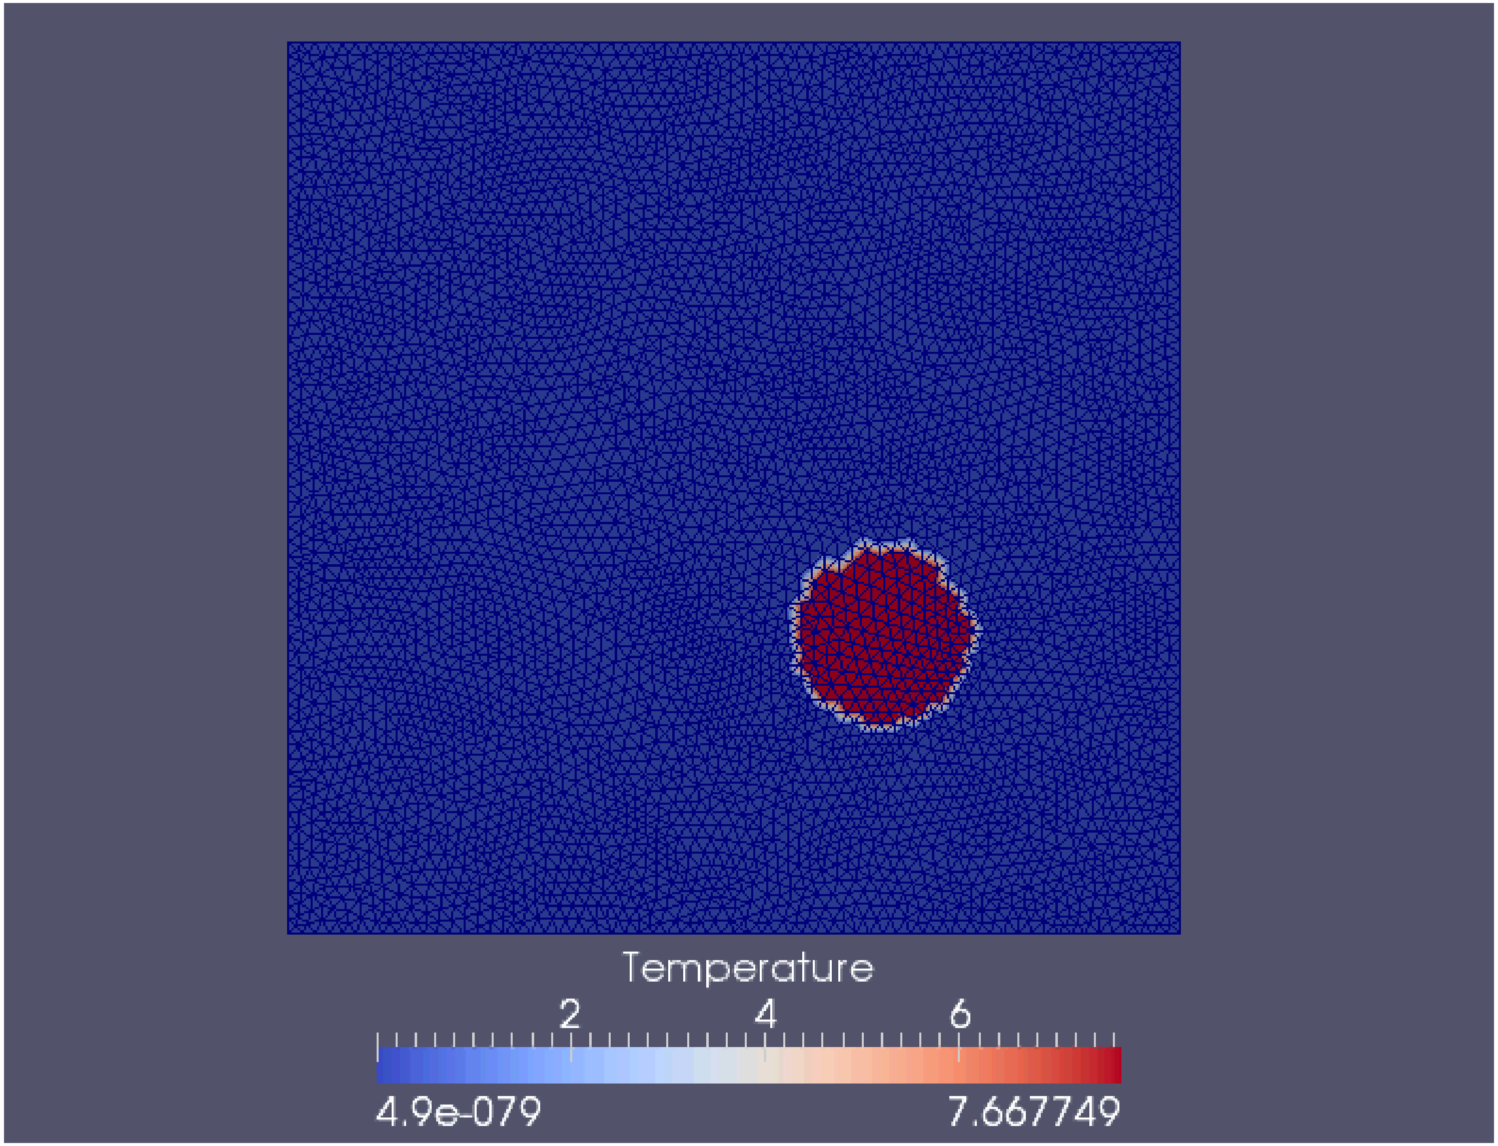

Wind driven barotropic gyre, western boundary current and eddies resolved with an adapting anisotropic mesh

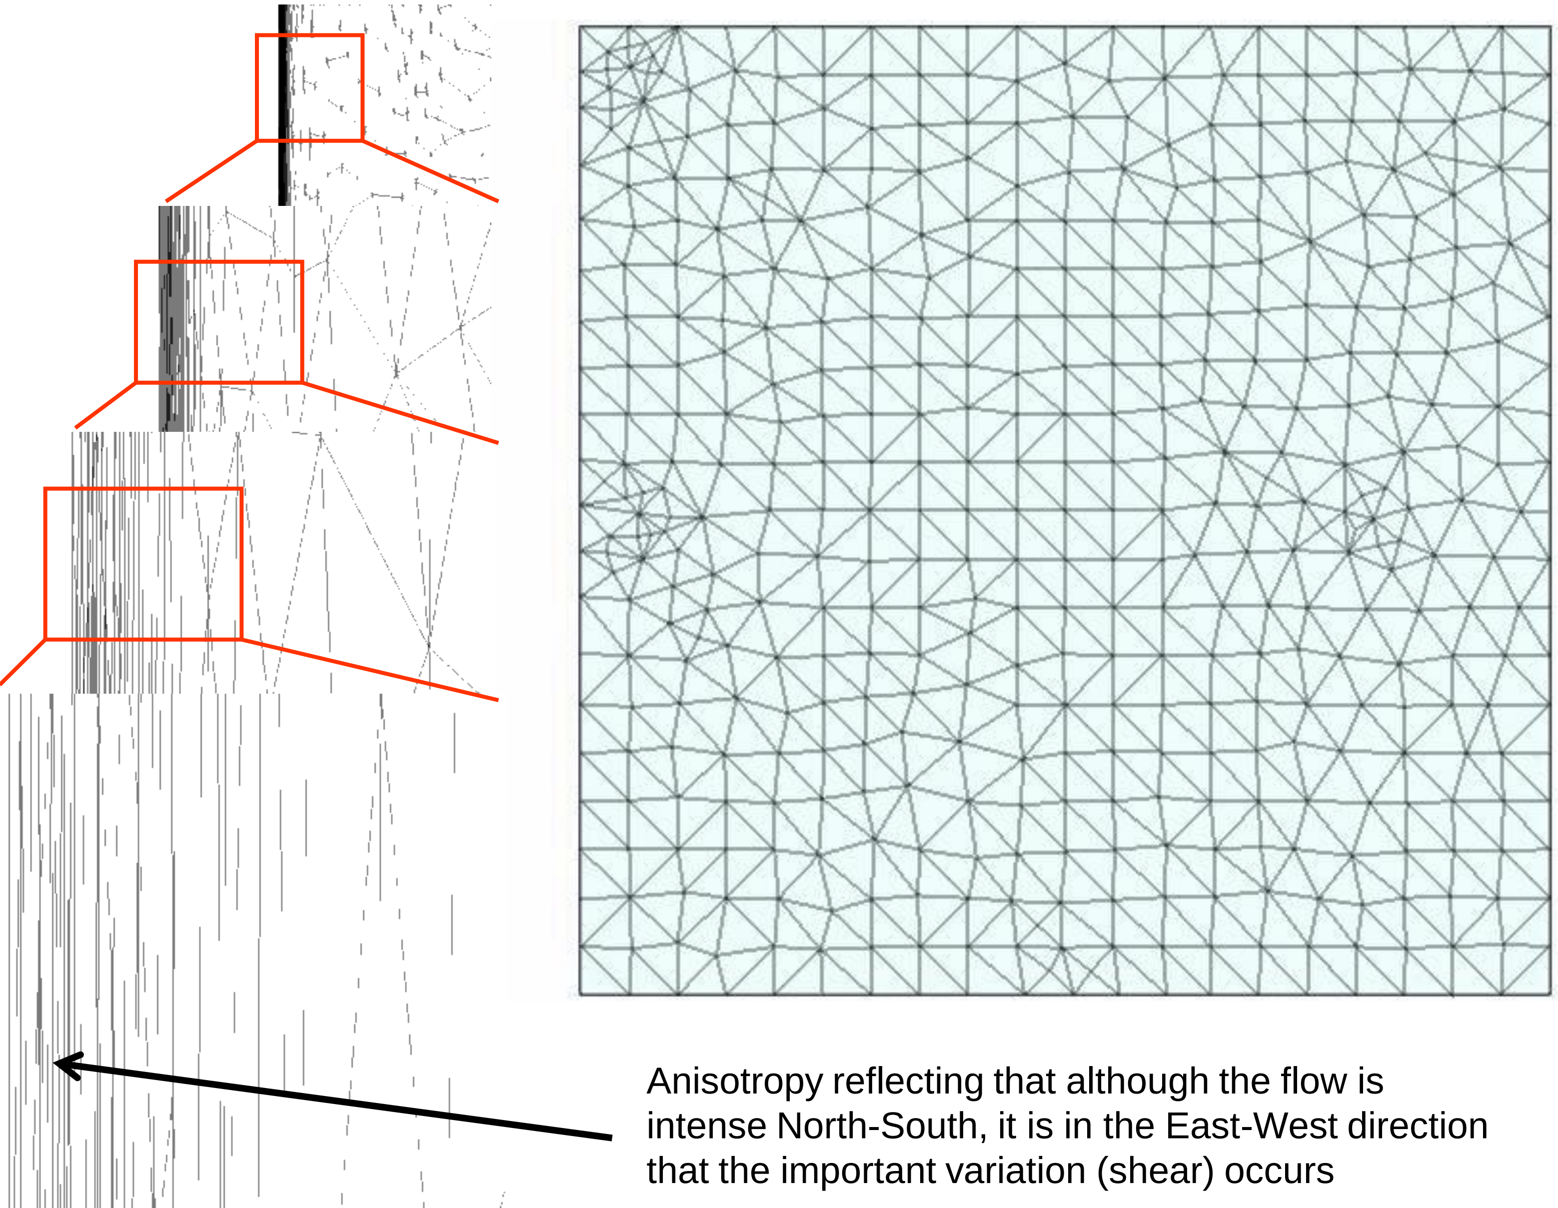

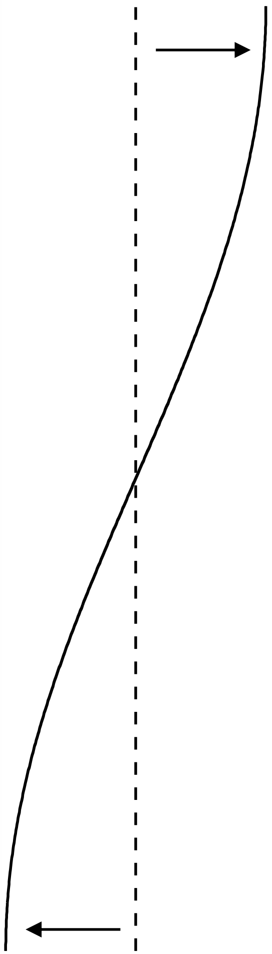

Anisotropy reflecting that although the flow is intense North-South, it is in the East-West direction that the important variation (shear) occurs

Wind driven 3D baroclinic gyre

Preliminary work comparing with MITgcm & NEMO

Problem

- Domain: 1000x1000x2 km
- Double gyre wind forcing
- Linear temperature stratification and linear EoS
- Beta plane
- Viscosity/diffusivity: 100, 0.001 m^2s^{-1}
- 3 year simulation time

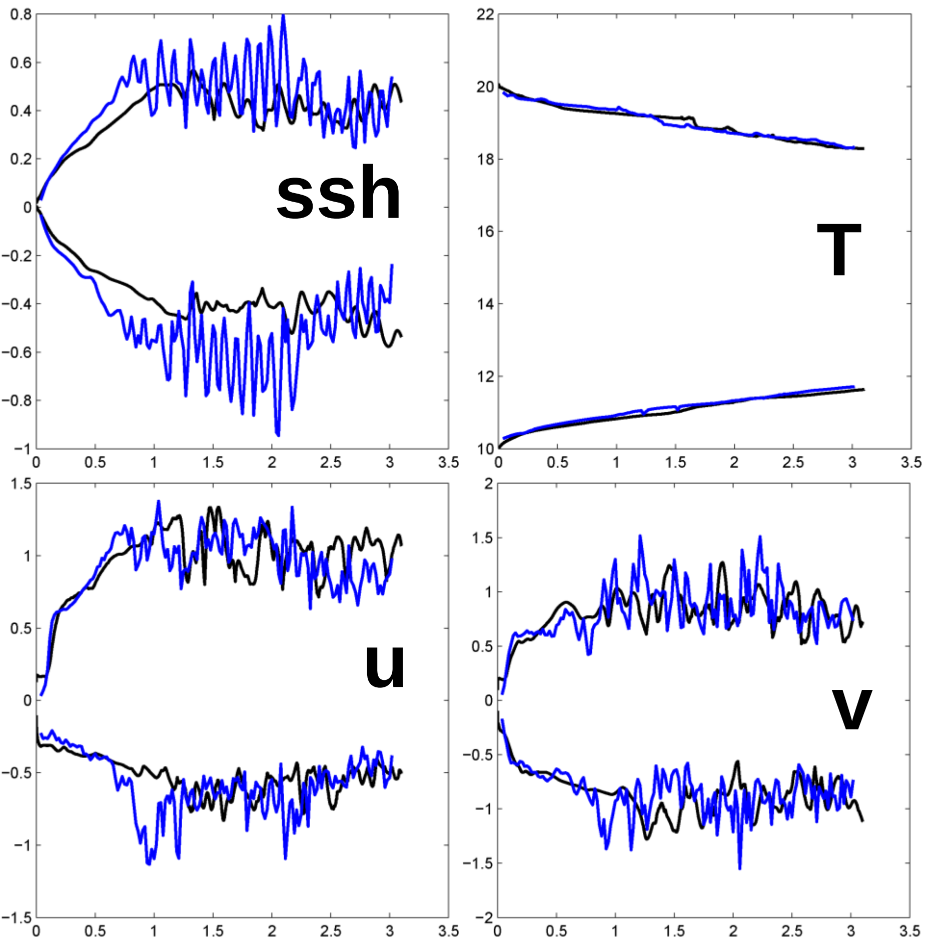

MITgcm (blue)

- Resolution: 100x100x20
- Timestep: 600s
- Set-up as per tutorial problem
- Hydrostatic mode at present

ICOM (black)

- Coordinate mesh: 100x100x20
- Timestep: 21,600s
- Non-hydrostatic
- P1dg-P2 element for momentum-pressure
- Bassi-Rebay for DG viscosity
- P1 SUPG for temperature

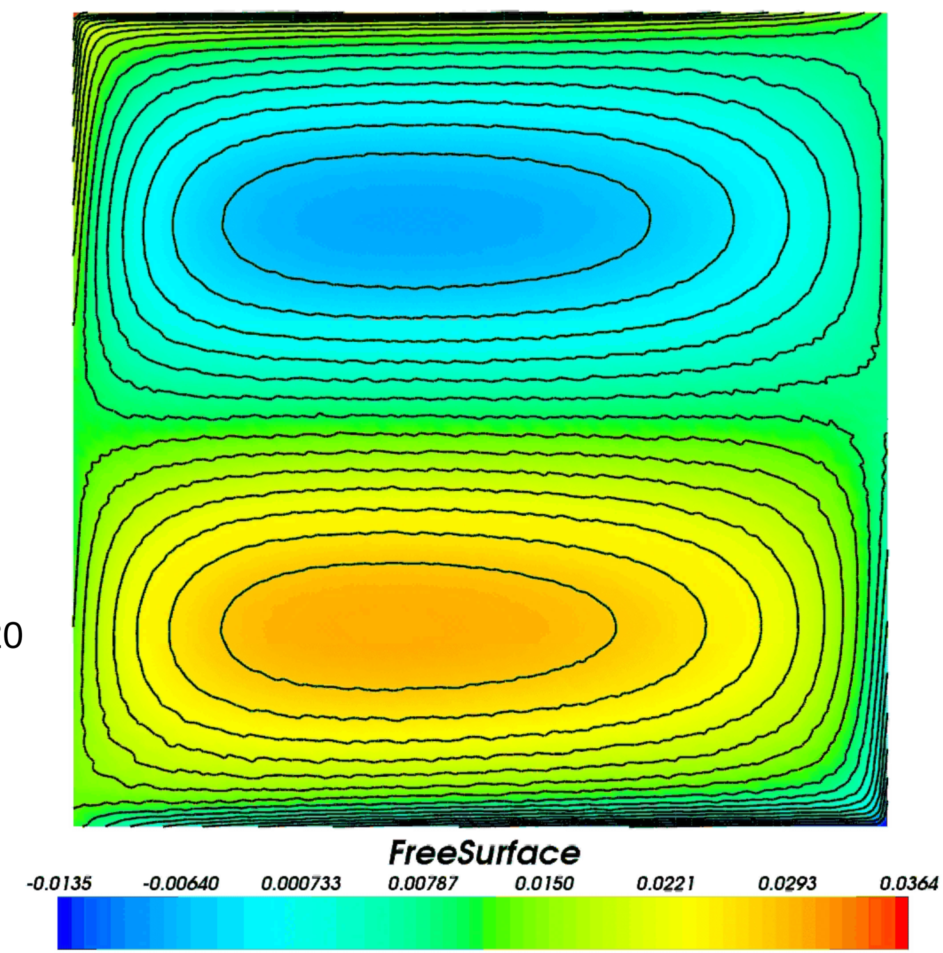

Wind driven baroclinic gyre –
higher horizontal viscosity

Problem

- As before, apart from:
- Viscosity/diffusivity: 2000, 0.00001 m^2s^{-1}
- In hope that we get a steady solution, can check convergence and stress diapycnal mixing rates

- We're using an element choice which uses a discontinuous representation of velocity and higher order pressure
- We would hope therefore that we will be more accurate for the same coordinate mesh, also more expensive of course!

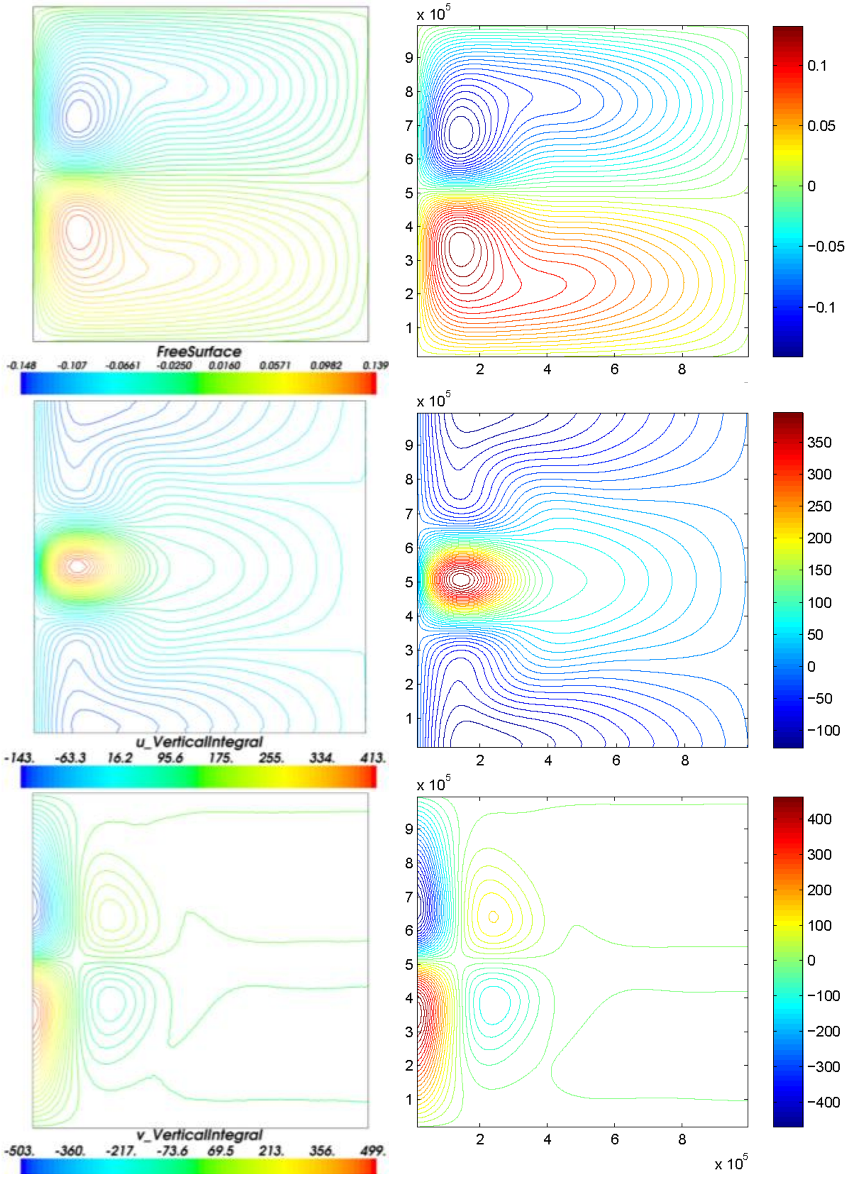

Some initial statistics at 3 yrs

	Min h	Max h	Min uint	Max uint	Min vint	Max vint
MIT 50x40	-0.1477	0.1397	-143.6	406.1	-487.4	476.9
MIT 100x80	-0.1482	0.1393	-140.6	409.4	-495.3	487.2
ICOM 50x40	-0.1487	0.1395	-142.7	413.4	-503.0	498.9

Use bifurcations in the dynamics to understand effective resolution and numerical diffusion

Problem

- As before but now
- Domain: 2000x2000x2 km
- Viscosity/diffusivity: 500, 0.00001 $m^2 s^{-1}$

- Choose resolution here so that there is equivalence in the number of velocity d.o.f.s
- MITgcm = 200x200x80, dt = 600
- ICOM = 70x70x28, dt = 18,000

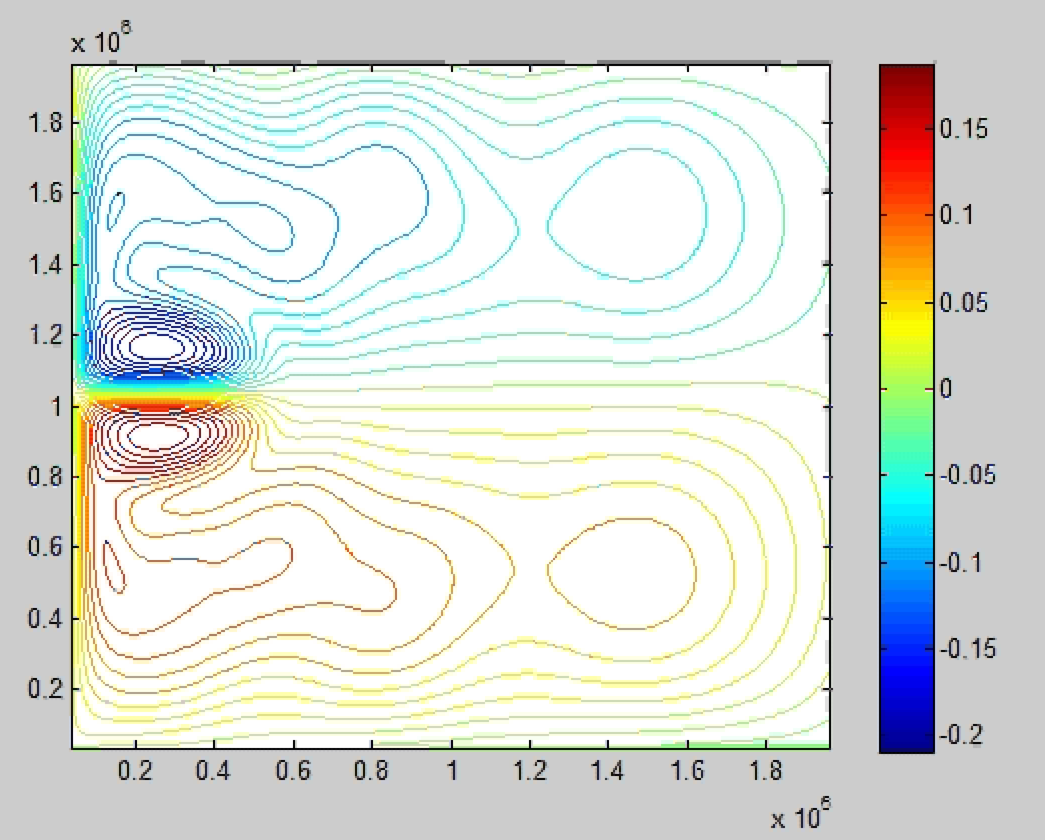

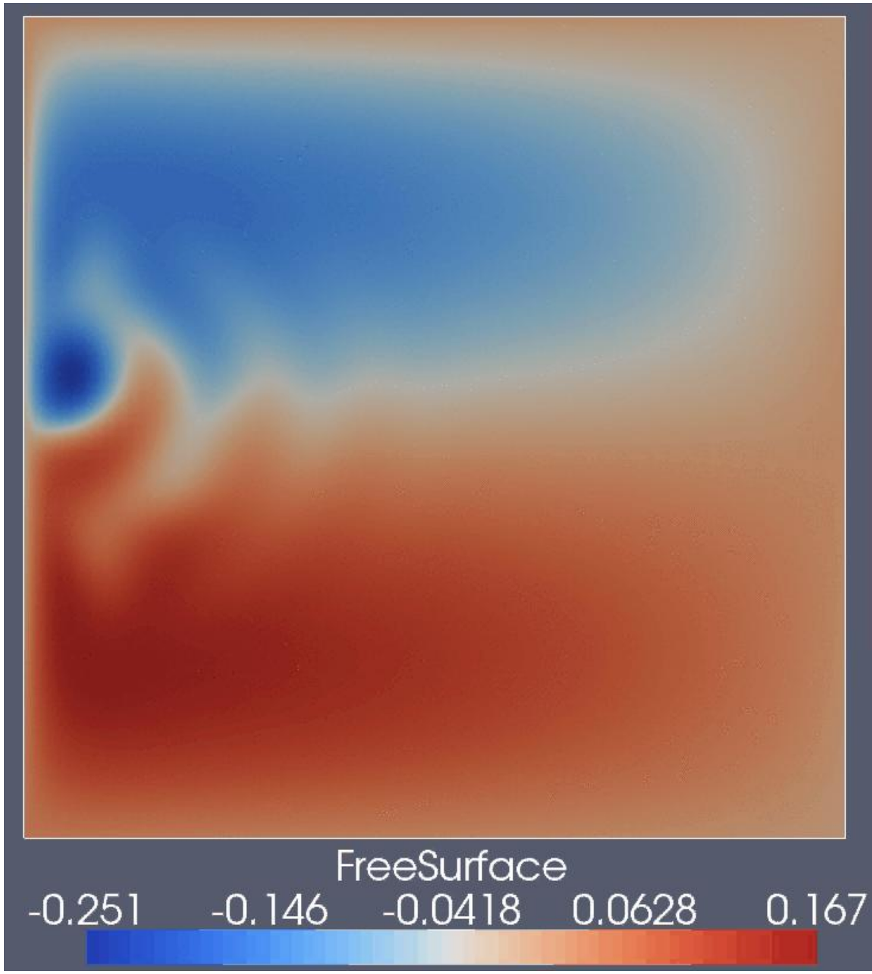

With this configuration we observe a periodic break-up of the recirculations in MITgcm but not in ICOM

What happens when we vary resolution?

In MITgcm we revert to a steady solution if we coarsen mesh

In ICOM we start to look closer to the MIT results when we double horizontal resolution

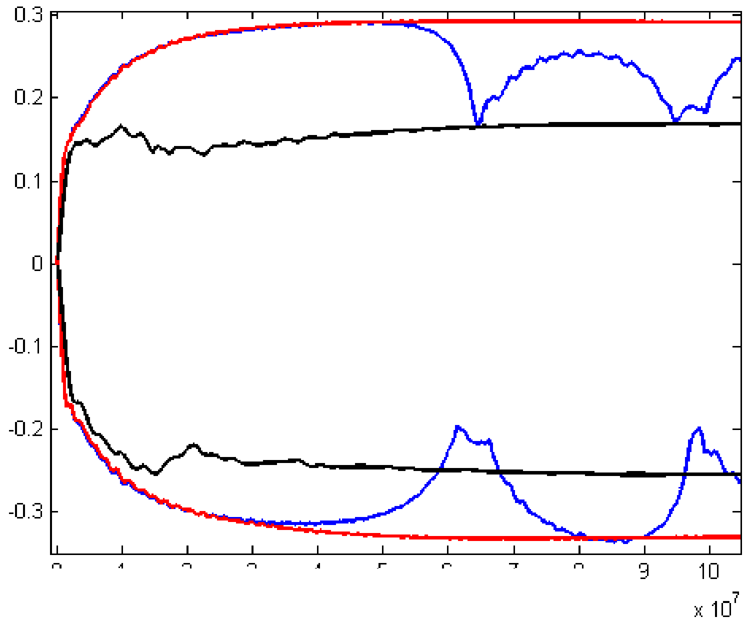

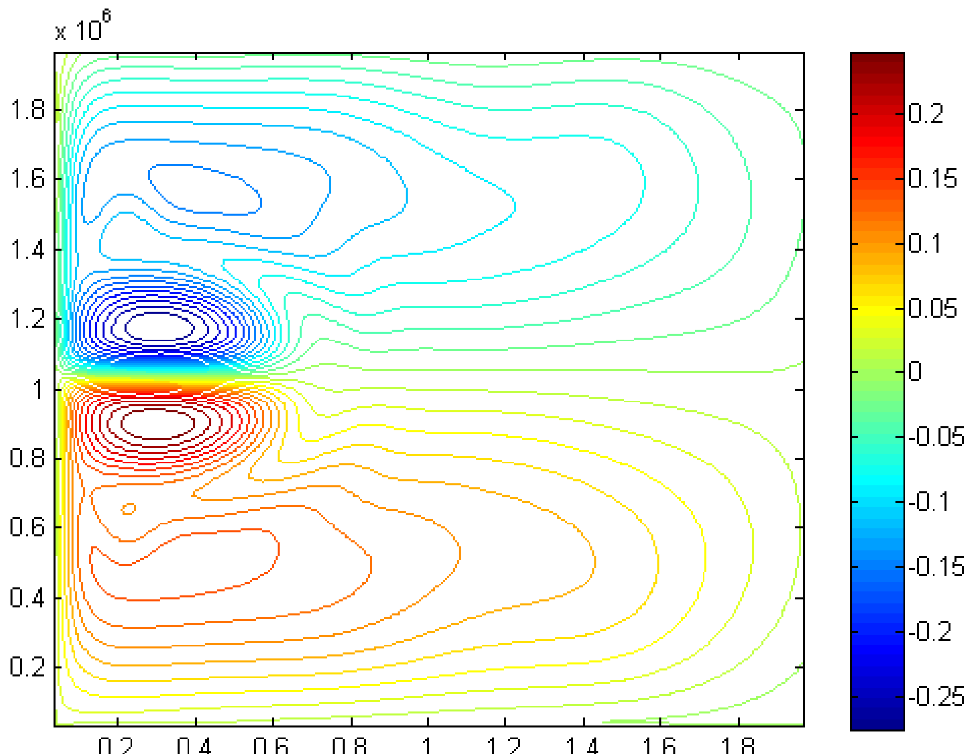

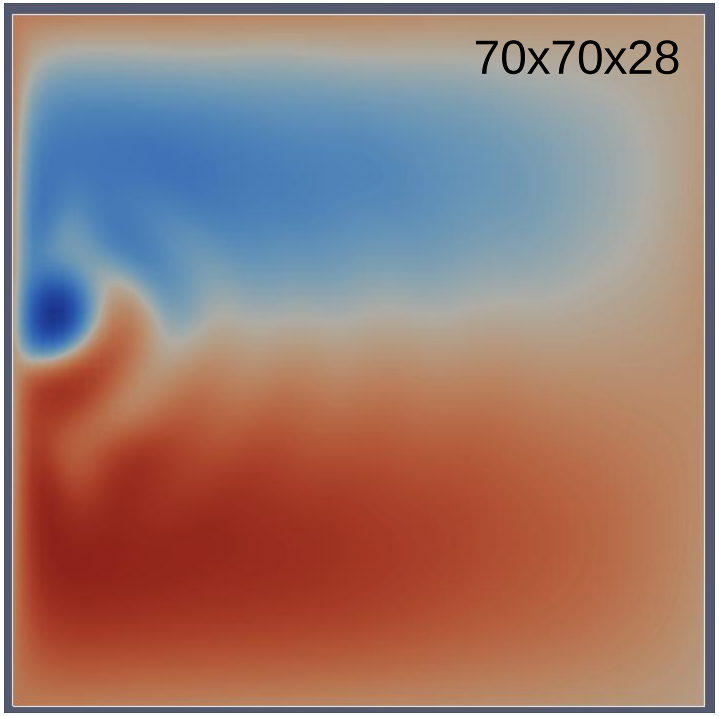

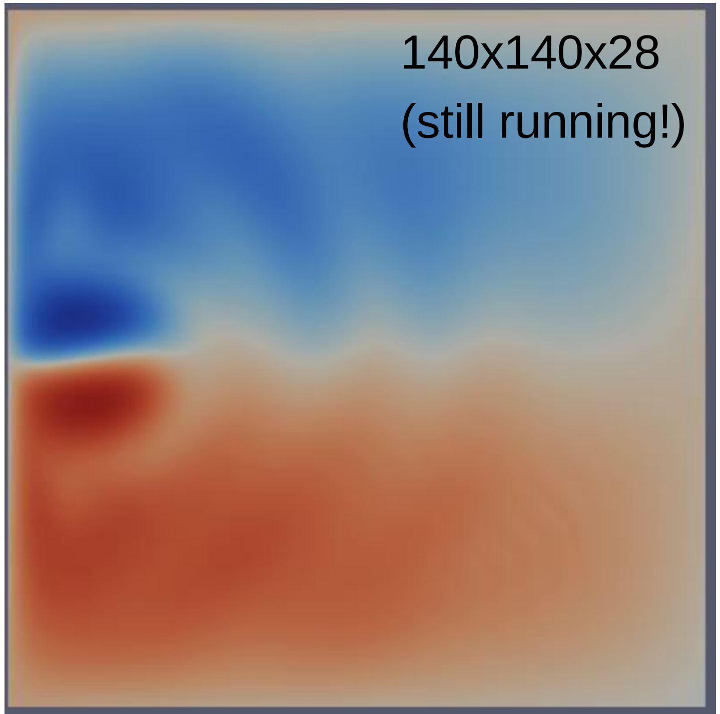

Scaling performance on a Cray XT4 (remember that here we are using different grids, such that the #velocity d.o.f. is consistent, and making use of ICOM's ability to use larger time steps, also note that user-incompetence with MITgcm and NEMO may be a factor!)

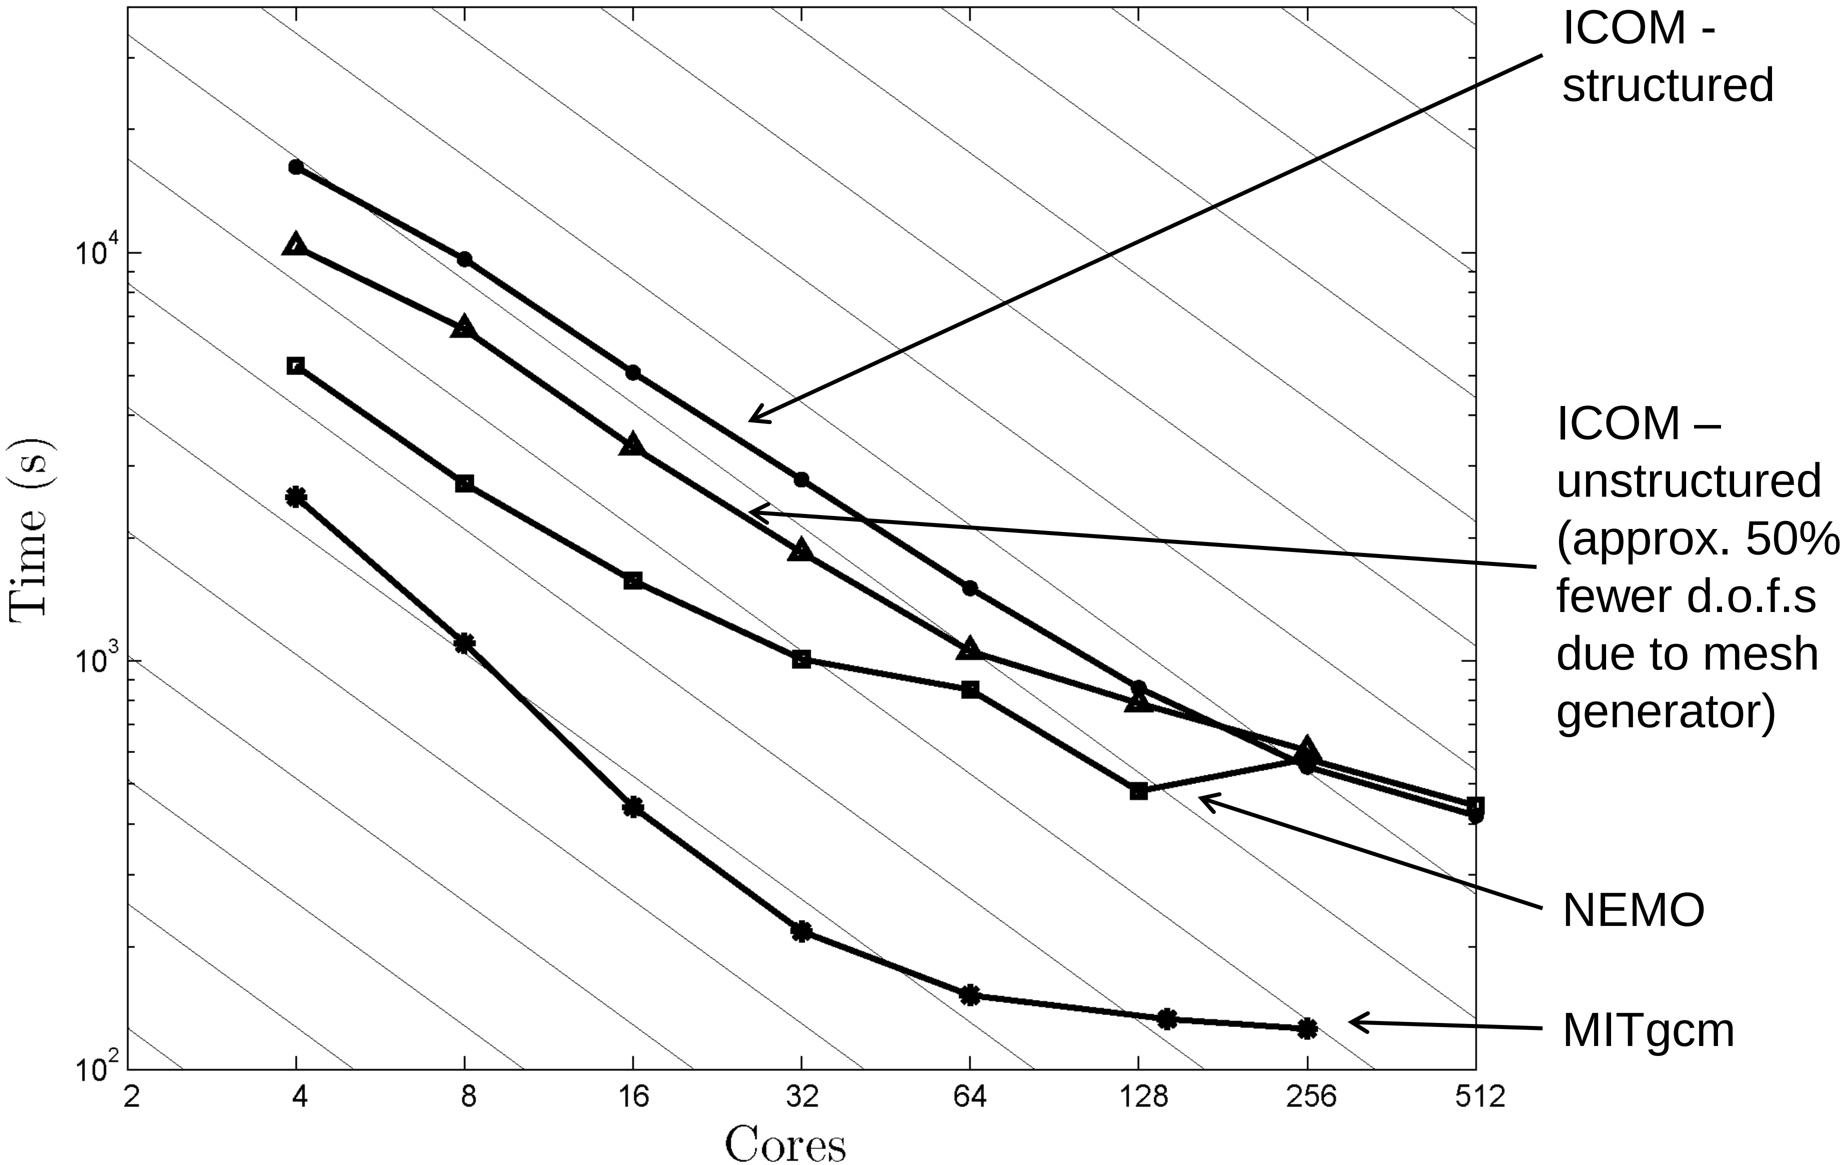

Next steps in this model inter-comparison

- Further use of bifurcations in solution dynamics to understand effects of resolution in these models
- Additional solution metrics, e.g. mixing
- Currently using a poor scalar advection method in ICOM (we know that CG+SUPG is not good, but CV performs poorly with stratification) -> DG
- Scaling and time to solution metrics on different hardware and using different decomposition methods, identify and fix bottlenecks
- Impacts of fixed but variable resolution meshes and then dynamic adaptivity
- More complex domains and real-world forcing/initialisation

A 2+1-dimensional approach to adaptivity which allows for columns to be preserved if necessary

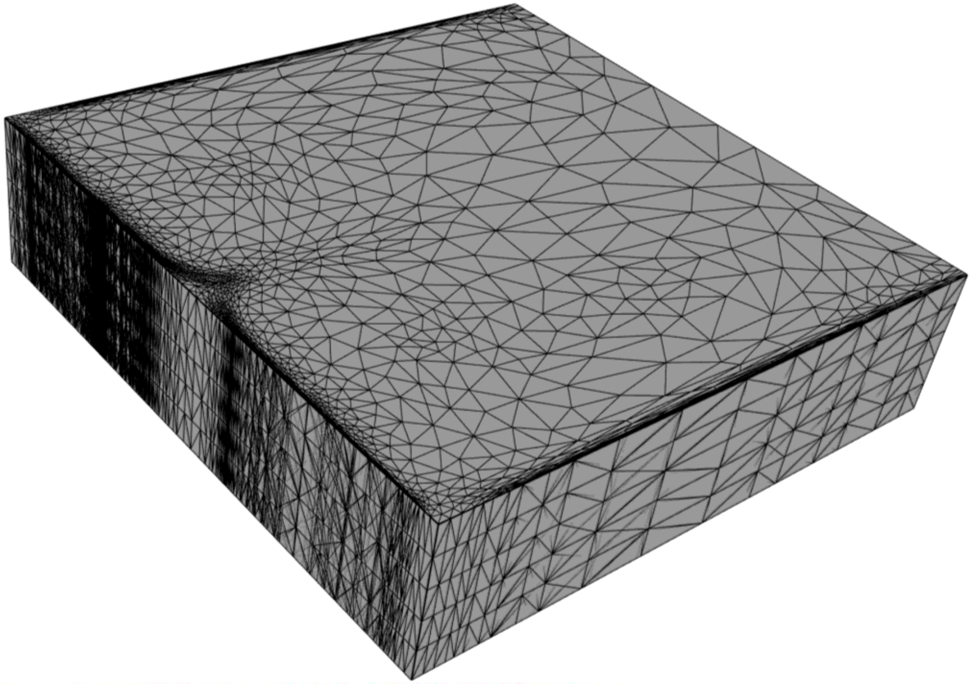

Of course our ultimate aim is that this would be an emergent property of a suitably intelligent error measure

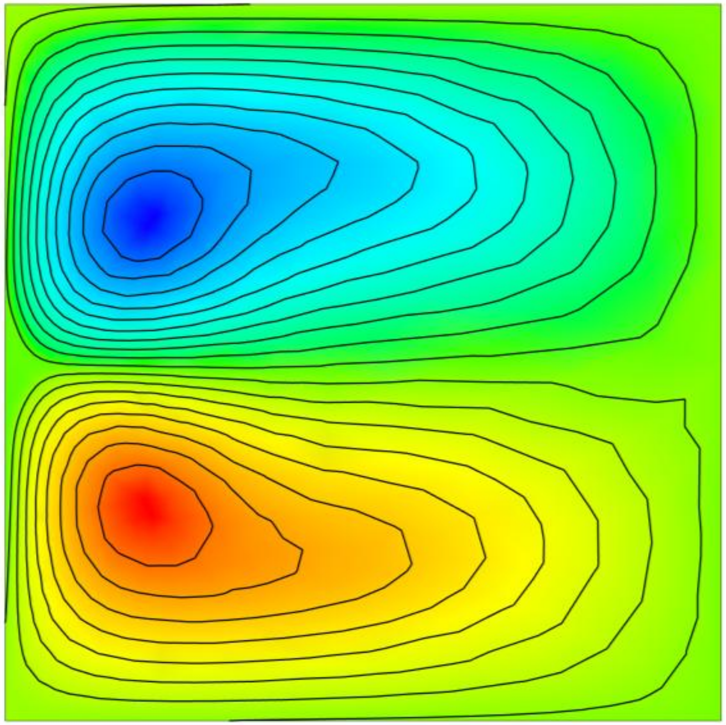

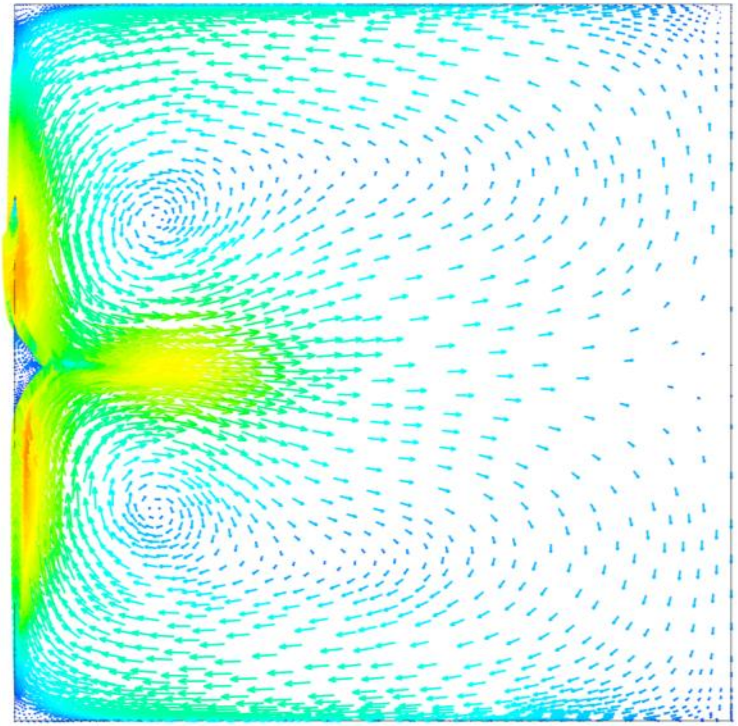

An idealised adaptive barotropic problem and a fixed mesh baroclinic problem

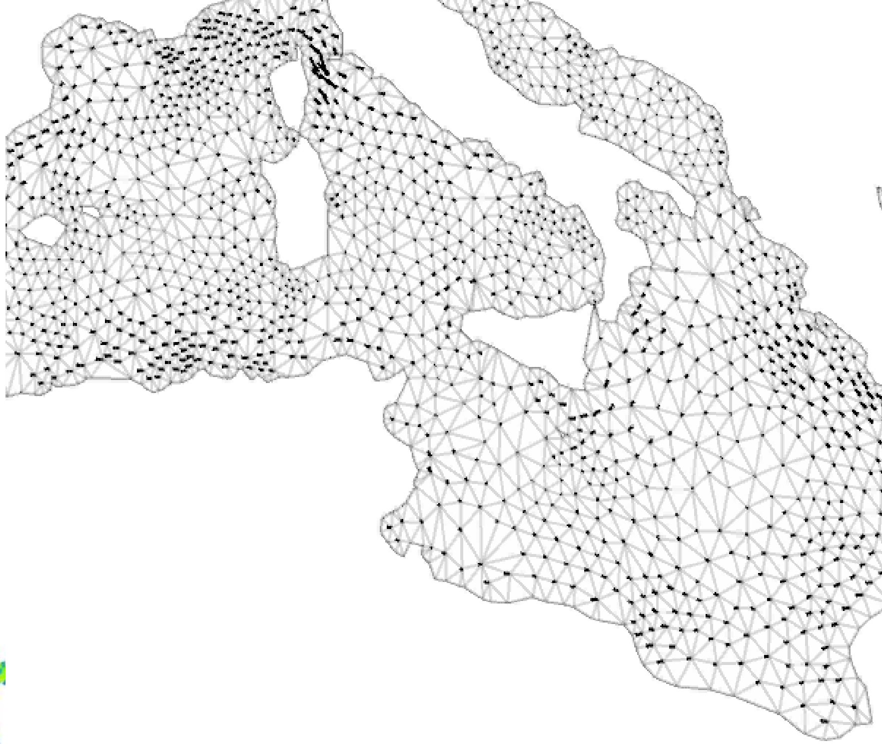

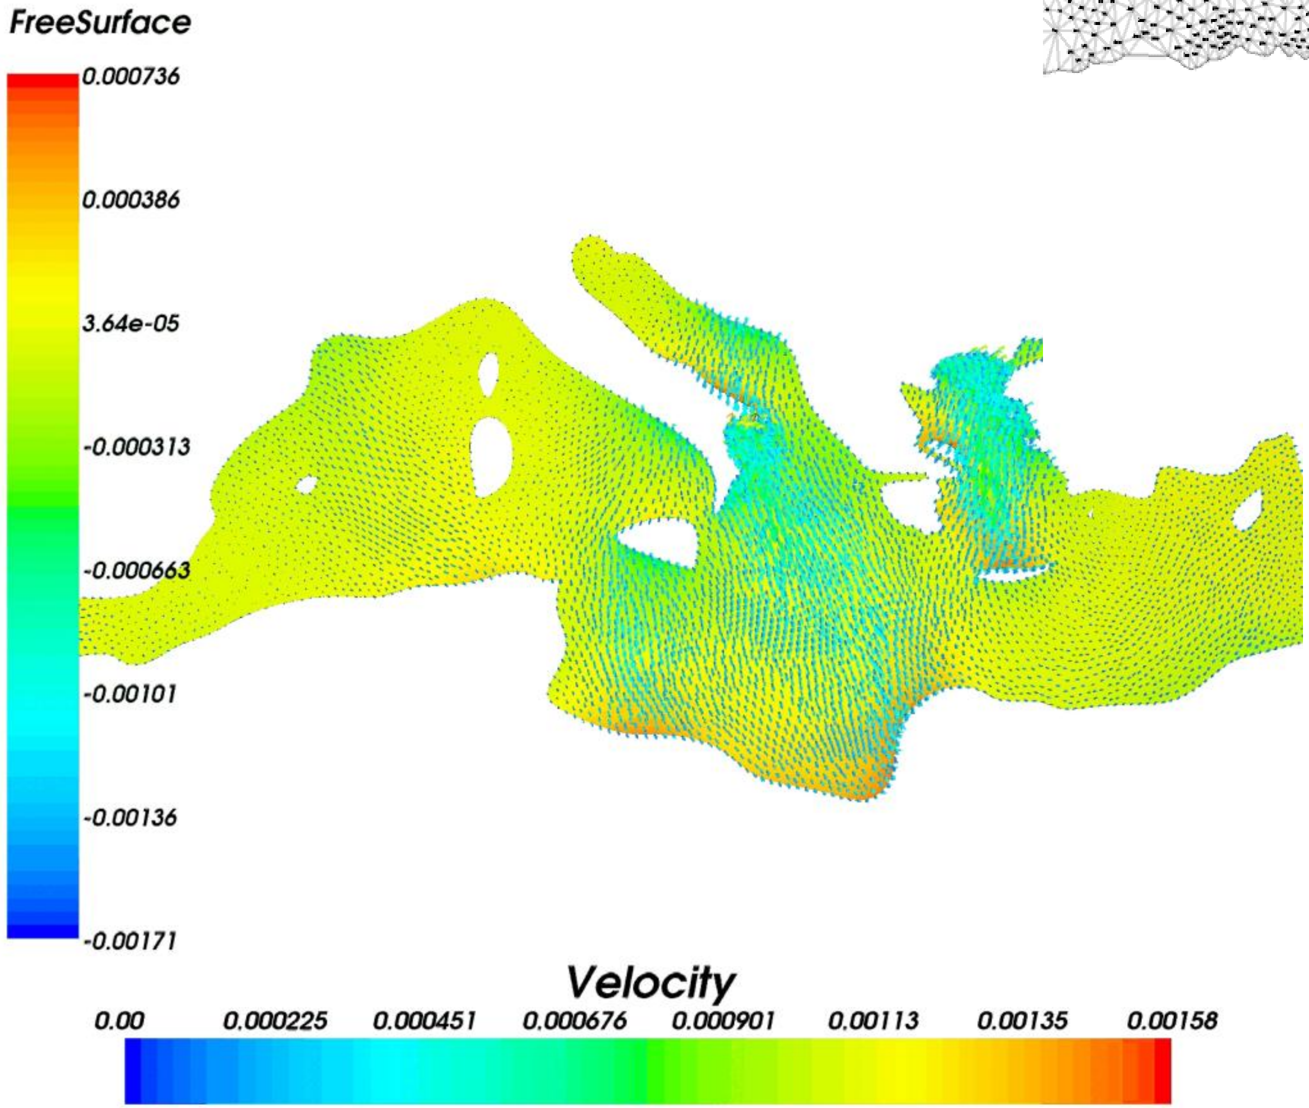

- Computed on the sphere in Cartesian space using super-parametric elements to account for curvature
- ECMWF ERA40 forcing with
- Large&Yeager bulk formulae
- Currently flat bottomed and linear stratification
- Parallelisation being debugged now

More real world comparisons: Evolution of the upper mixed layer with an adapting mesh over a column

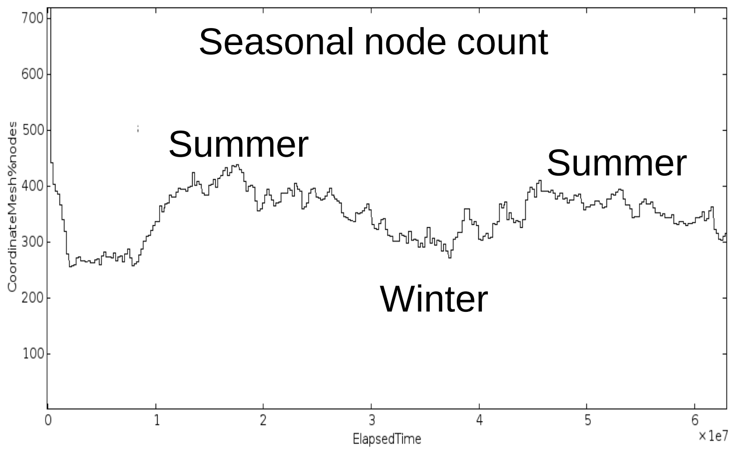

MLD at Ocean Station Papa

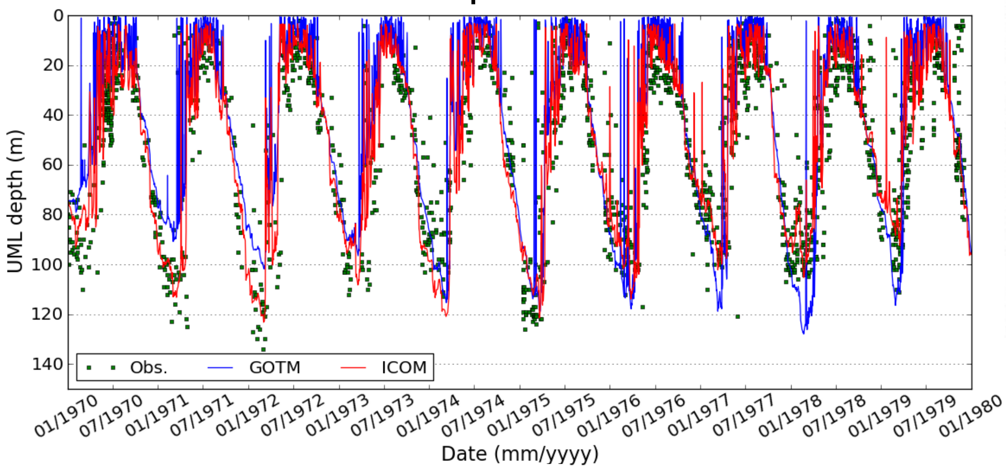

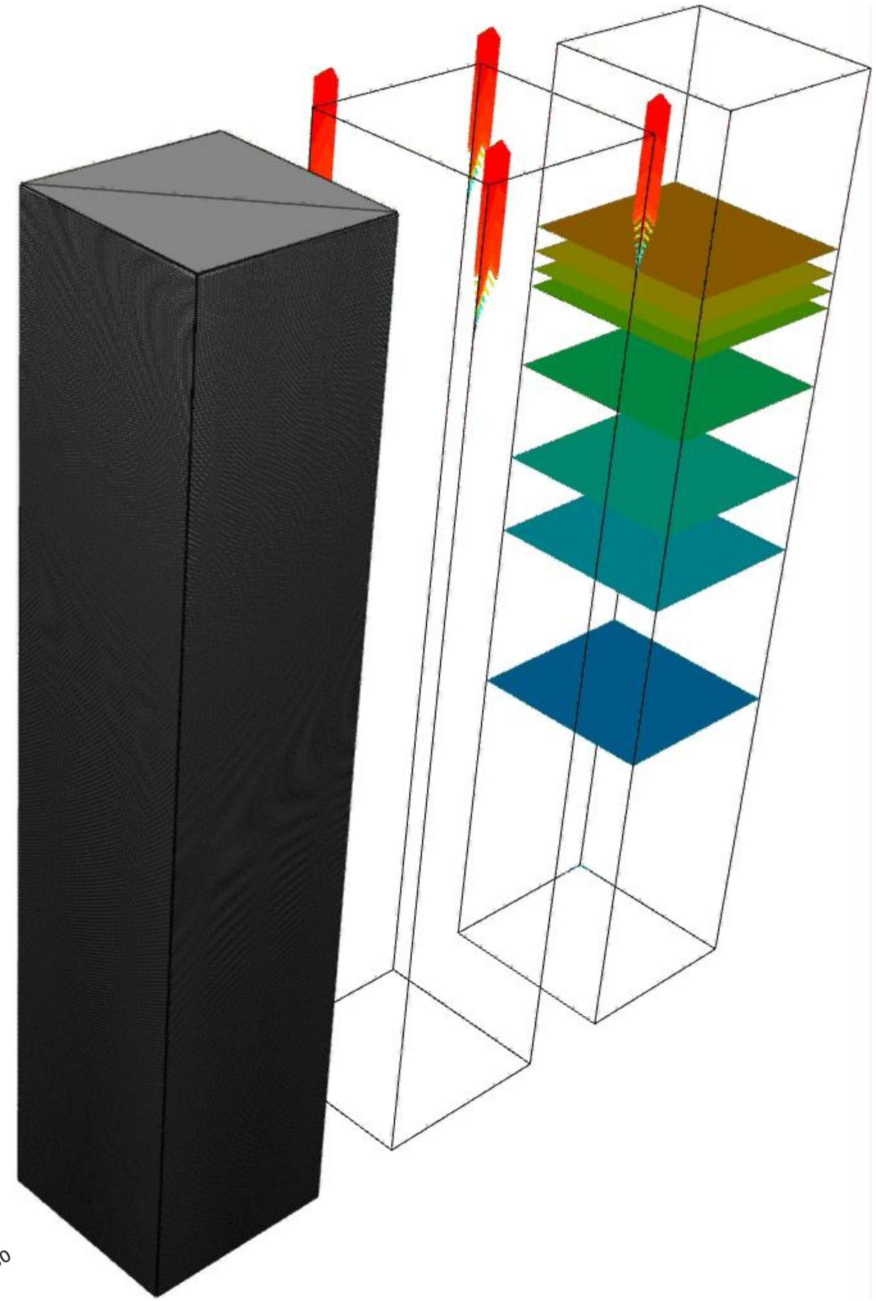

Hill et al., 2010

Obs

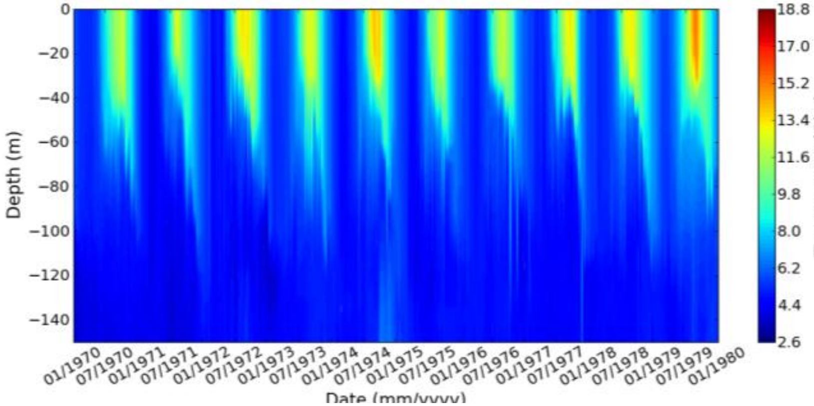

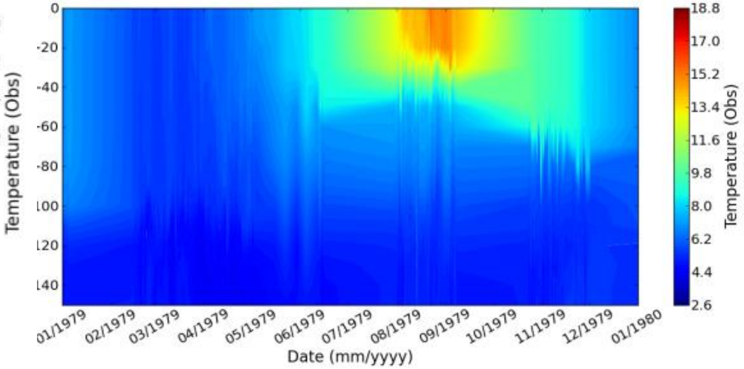

GOTM

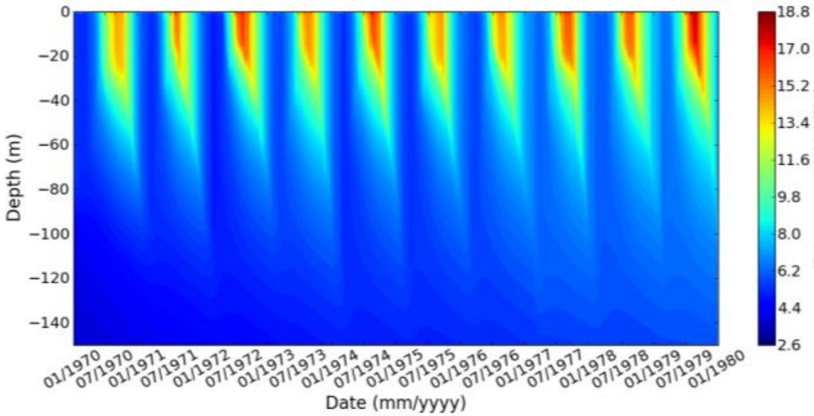

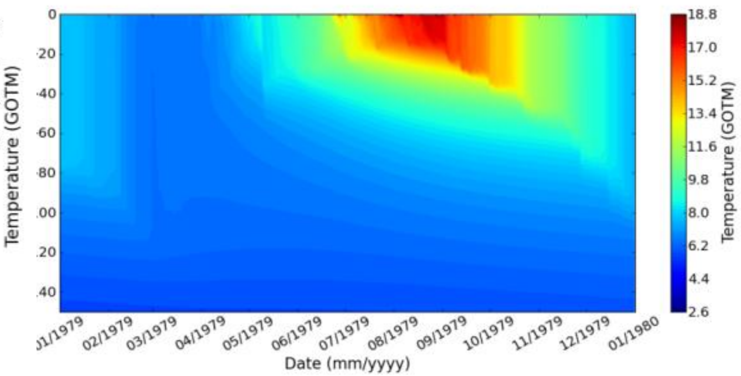

ICOM
(fixed)

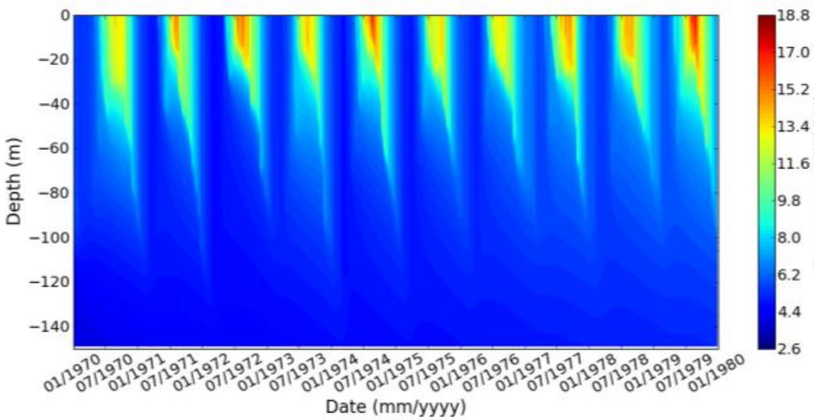

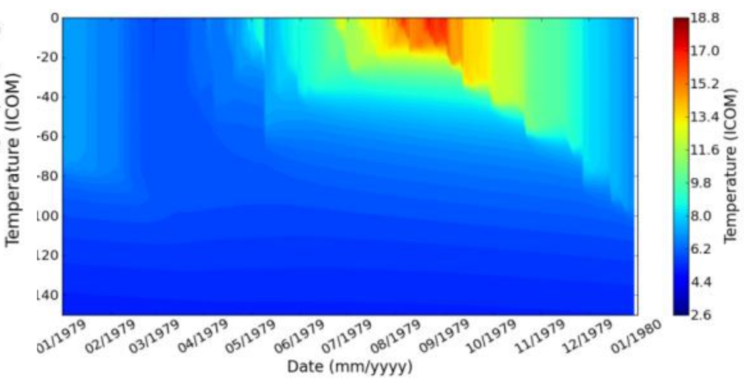

Comparing models, resolutions and the use of adaptivity

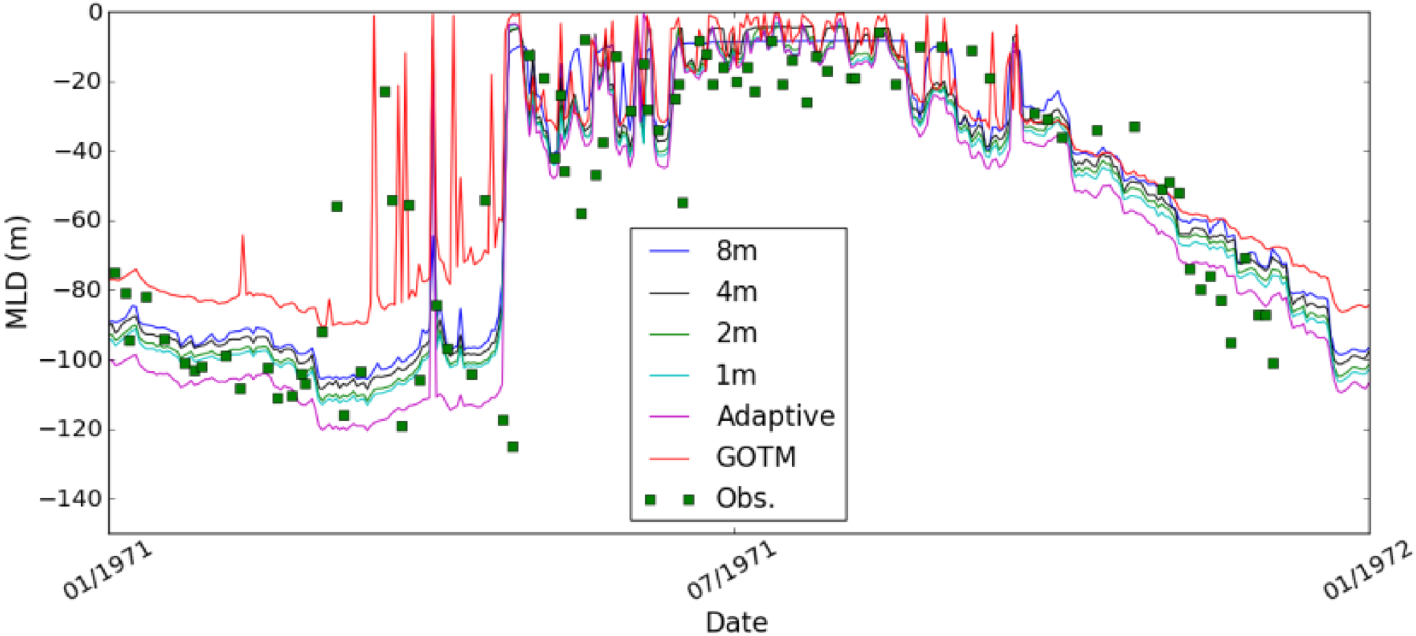

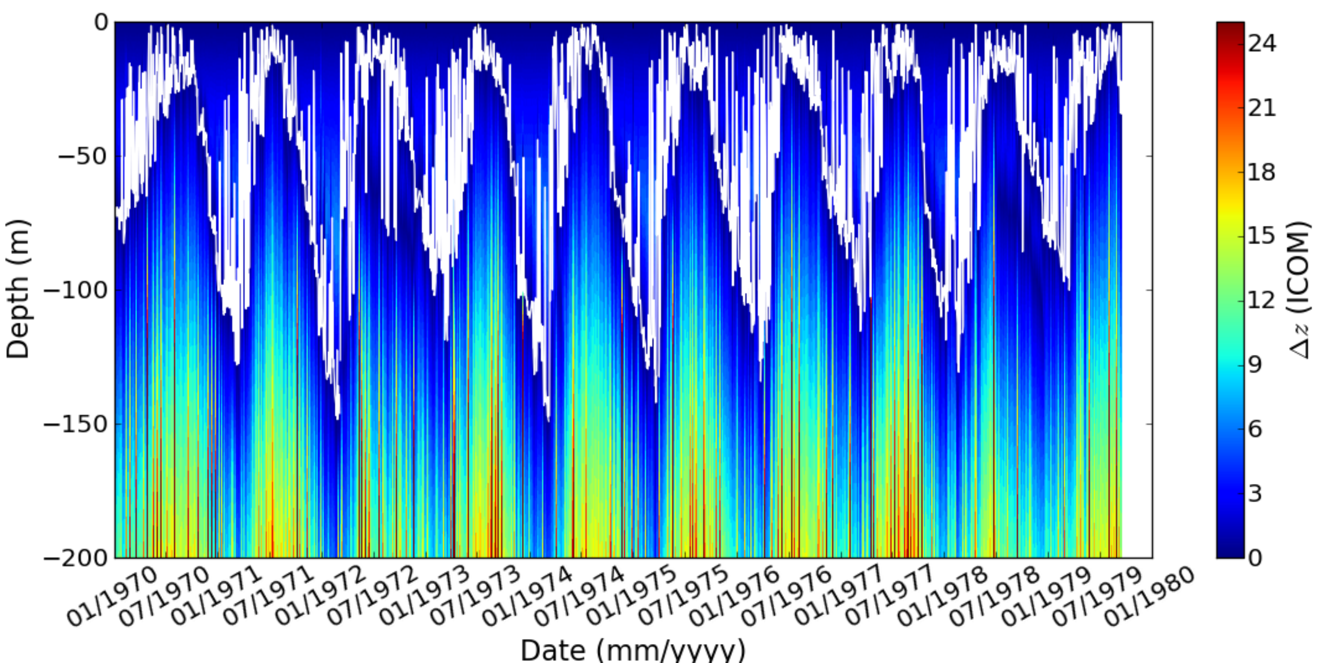

Biology module with four species (Ham, Popova et al., 2010)

Aim to use adaptivity to simulate sub-mesoscale processes and nutrient transports

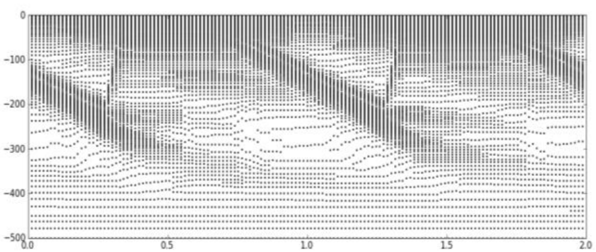

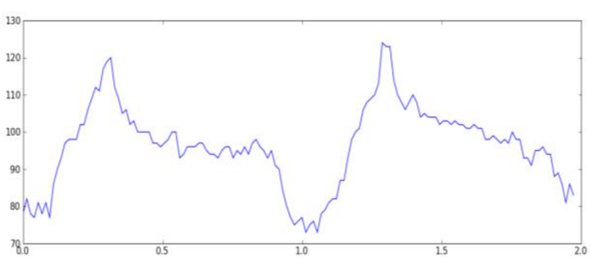

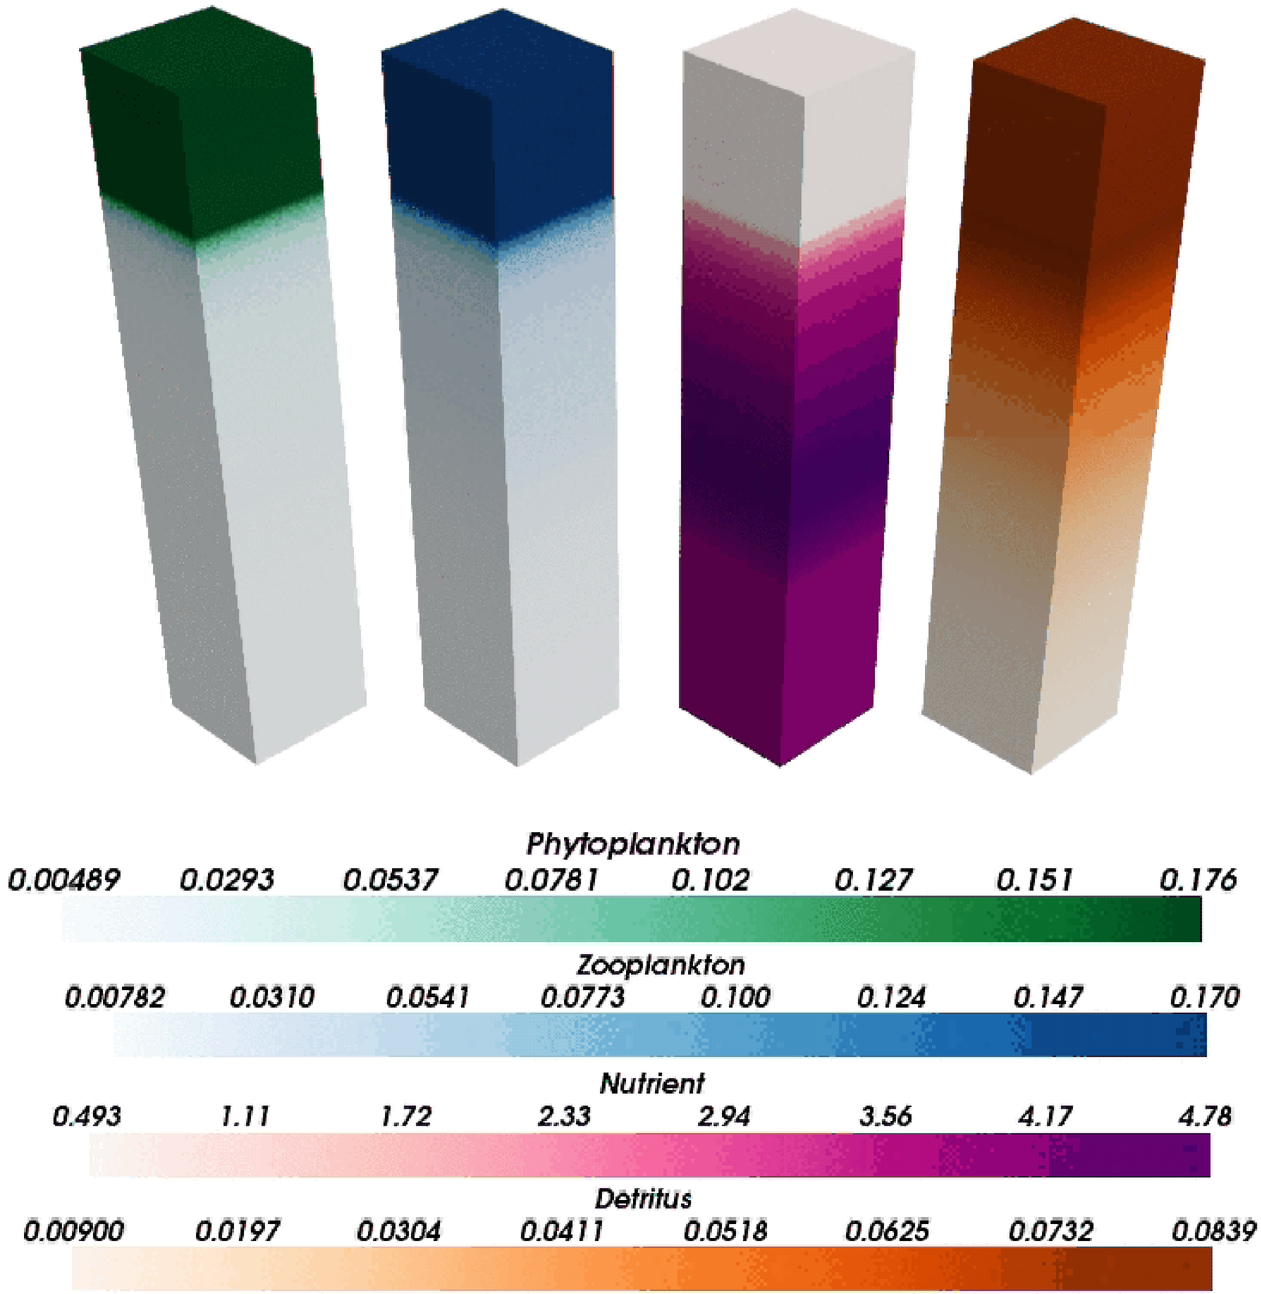

Simulating Geohazards:

Landslide generated wave around a conical island in 3D (Wilson et al., 2010)

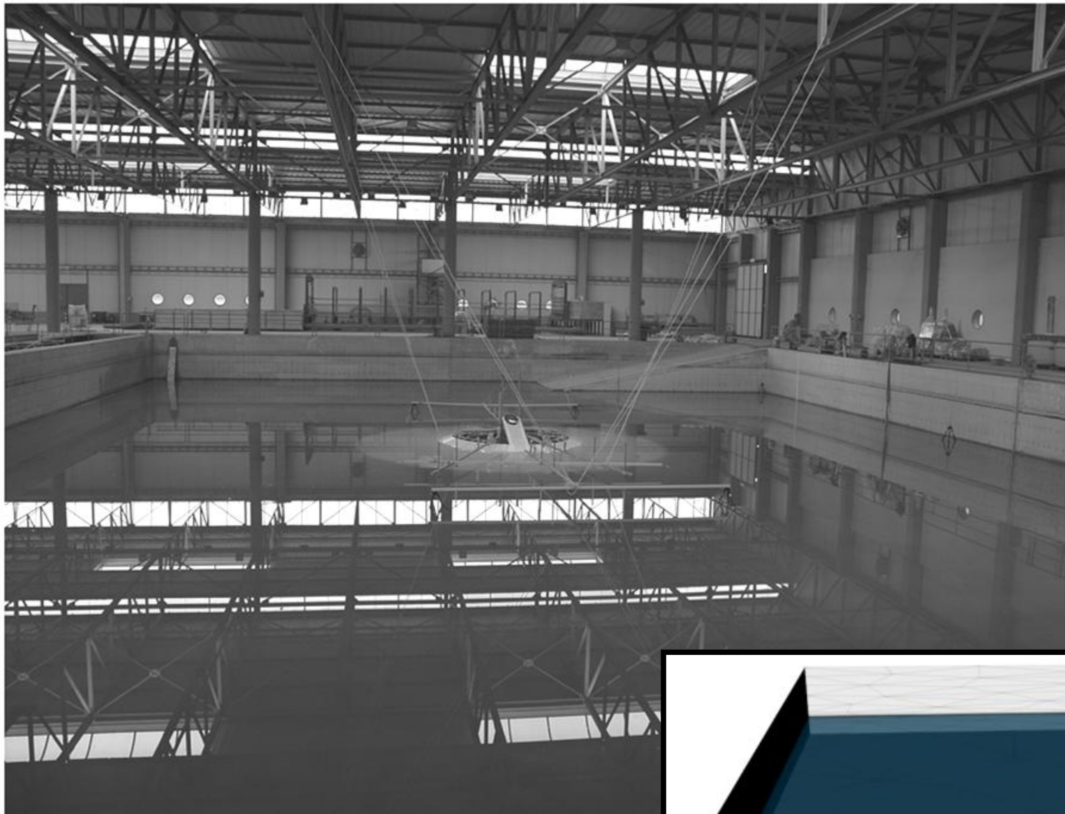

Landslide-generated tsunamis runup at the coast of a conical island: New physical model experiments, Di Risio, M, De Girolami, P., Bellotti, G., Panizzo, A., Aristodemo, F., Molfetta, M. G., Petrillo, A. F., Journal of Geophysical Research, v. 114, 2009

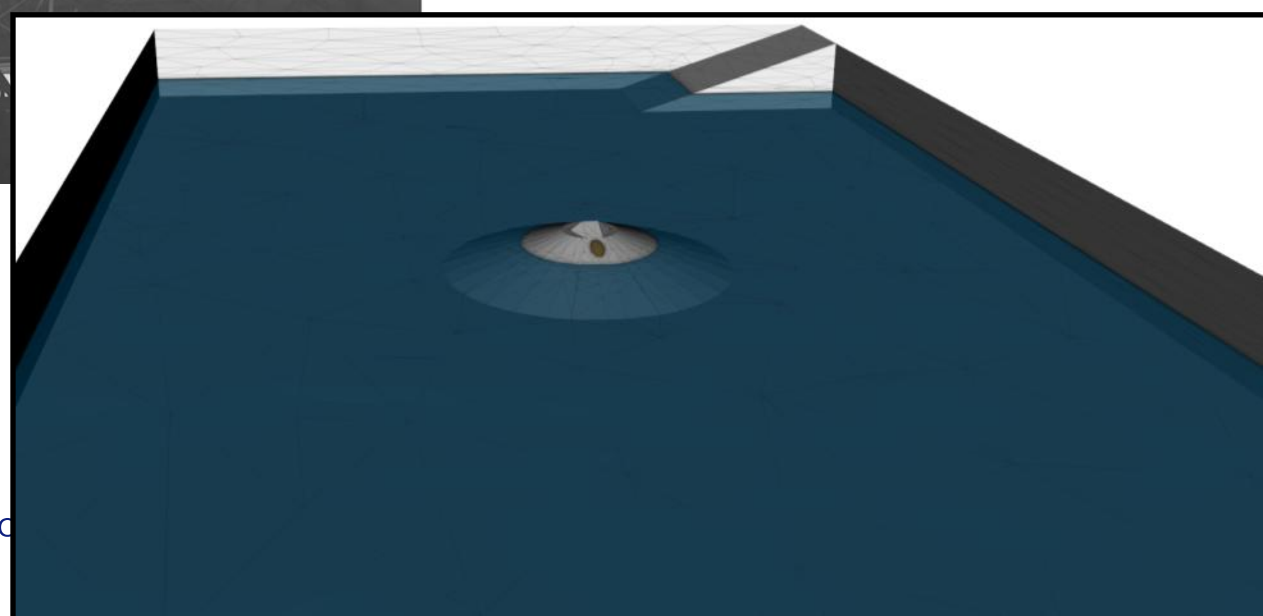

Using a multi-material approach to represent water, air and landslide

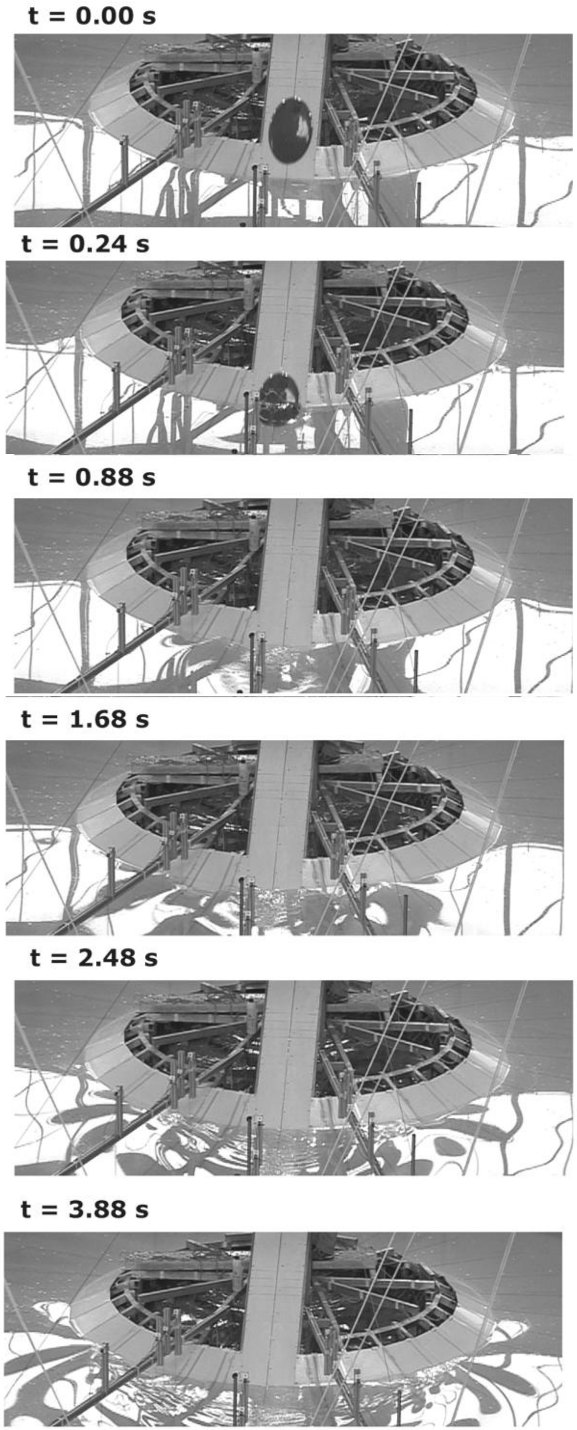

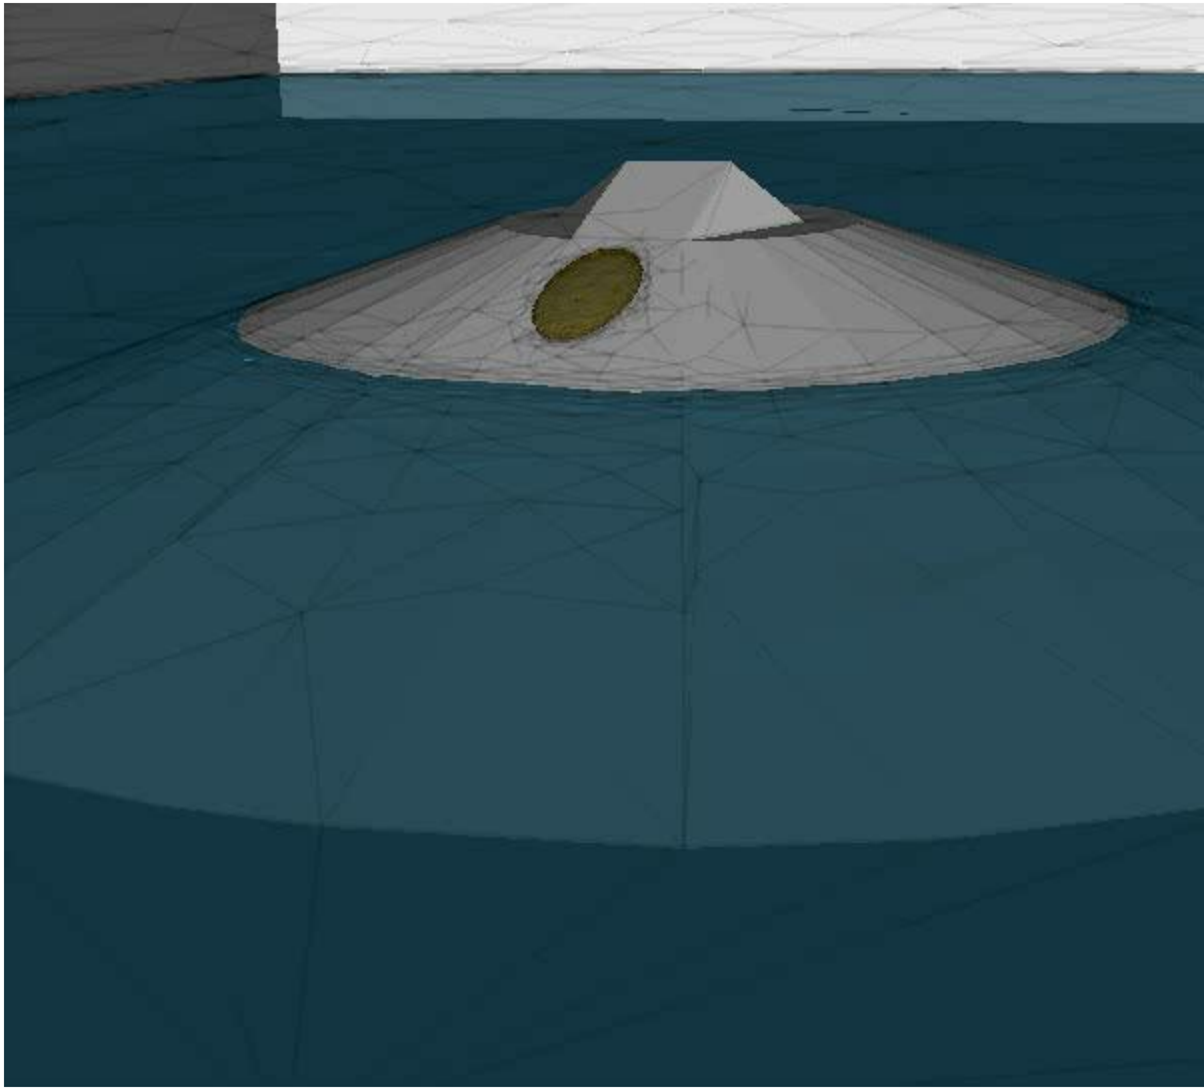

F

Lituya Bay landslide, 1958

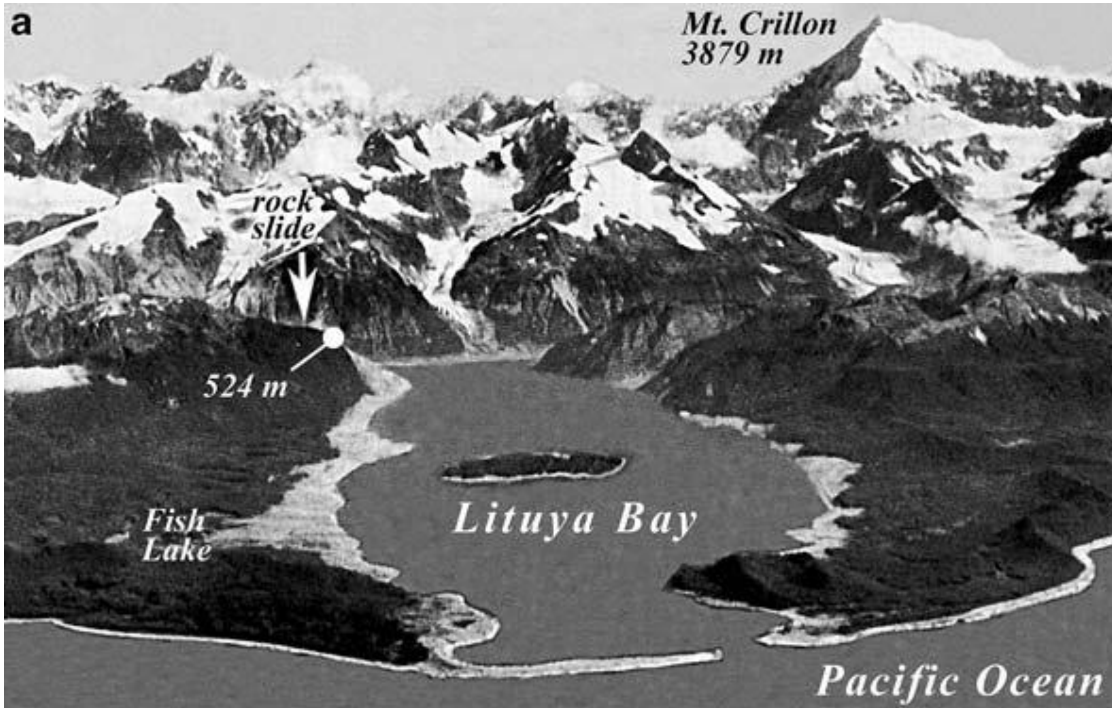

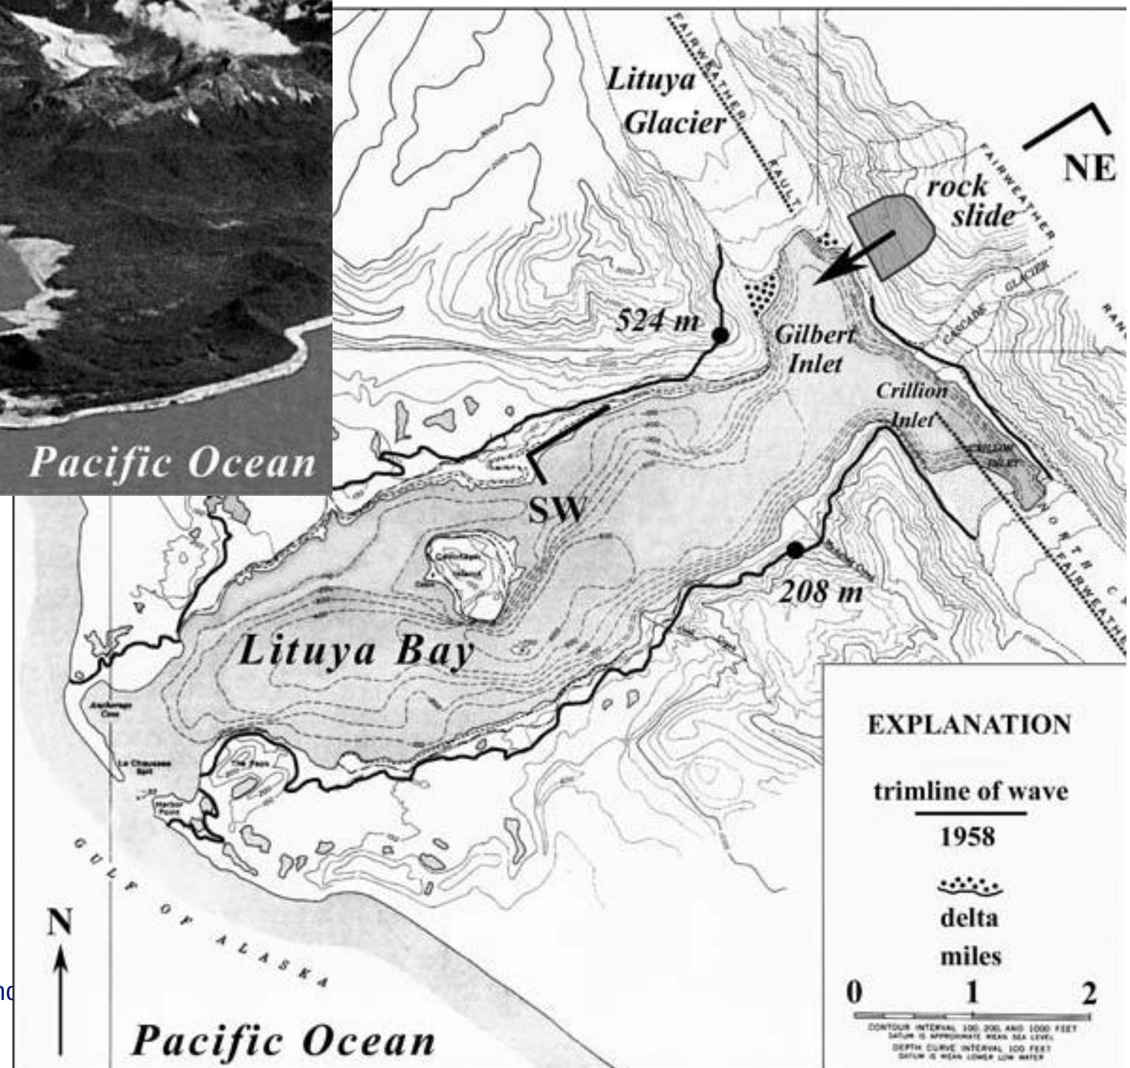

Fritz, H. M., Mohammed, F., Yoo, J. (2009) Lituya Bay Landslide Impact Generated Mega-Tsunami 50th Anniversary, Pure and Applied Geophysics, v. 166, pp. 153-175.

Lituya Bay analog validation problem in 2D

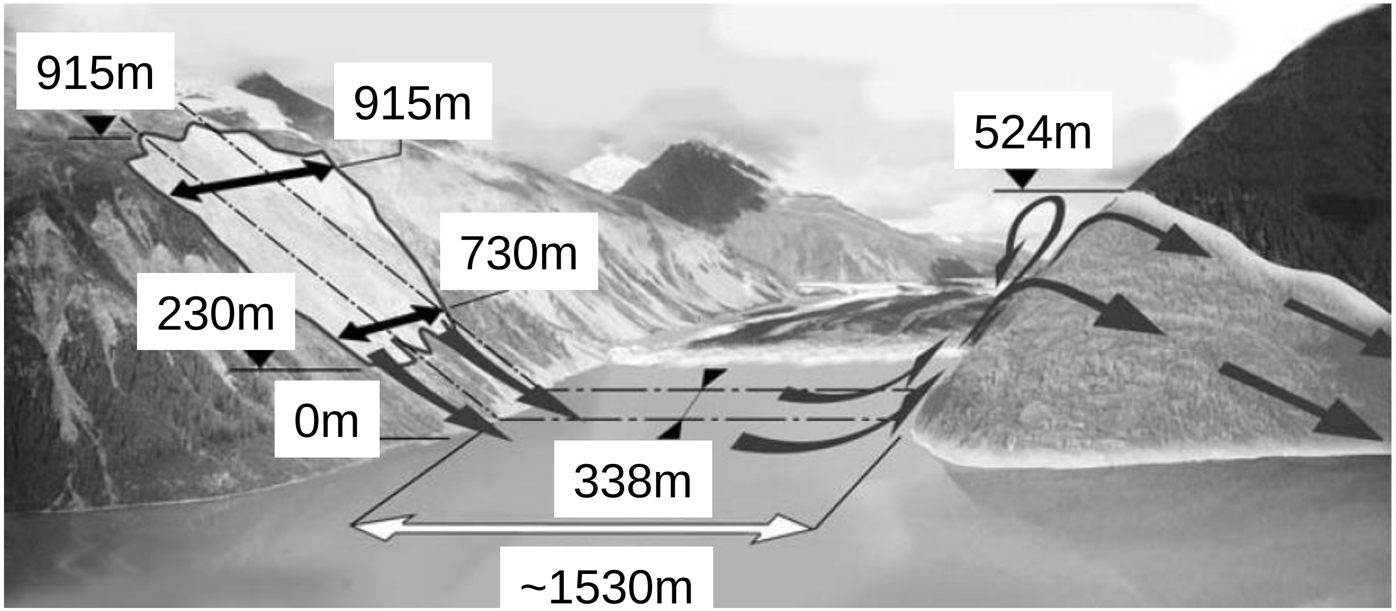

Fritz, H. M., Mohammed, F., Yoo, J. (2009) Lituya Bay Landslide Impact Generated Mega-Tsunami 50th Anniversary, Pure and Applied Geophysics, v. 166, pp. 153-175.

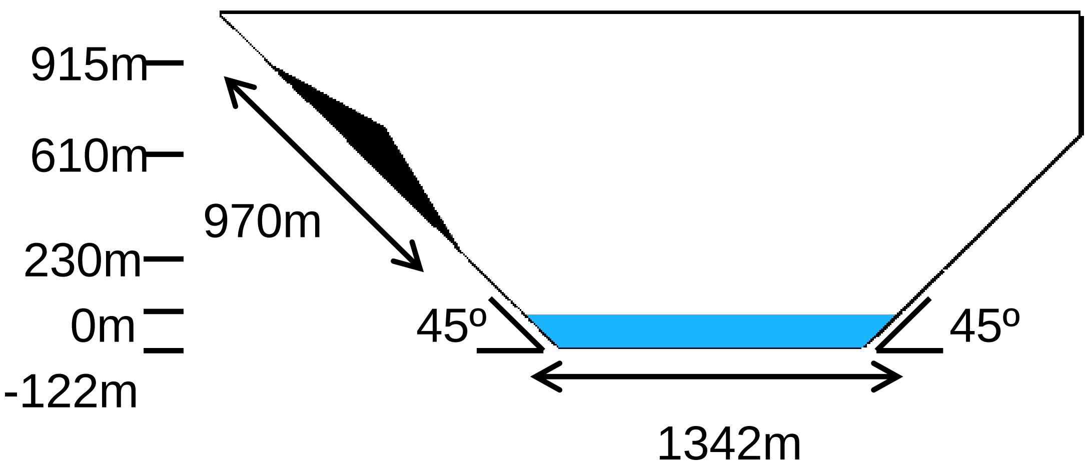

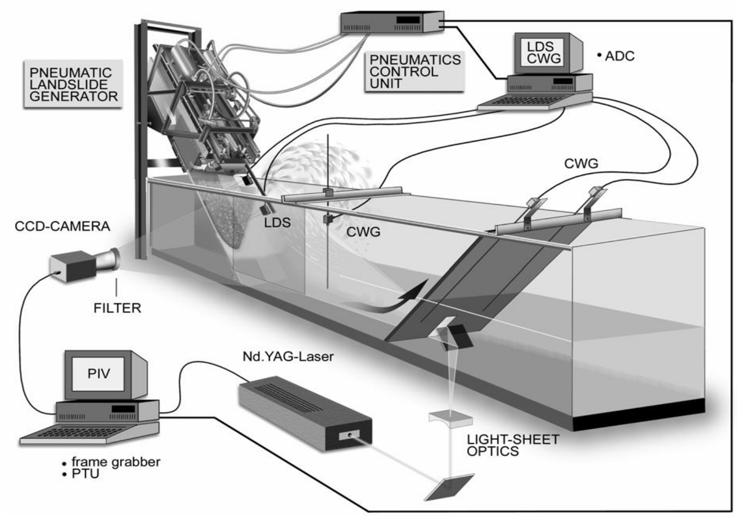

(Wilson et al., 2010)

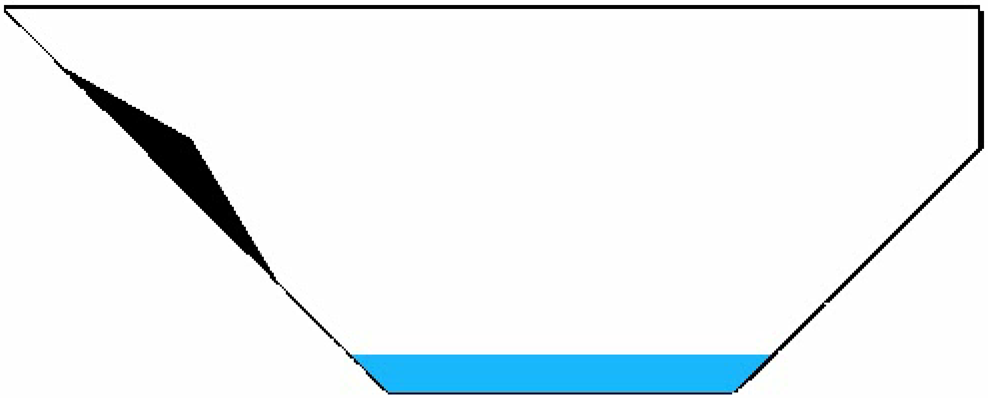

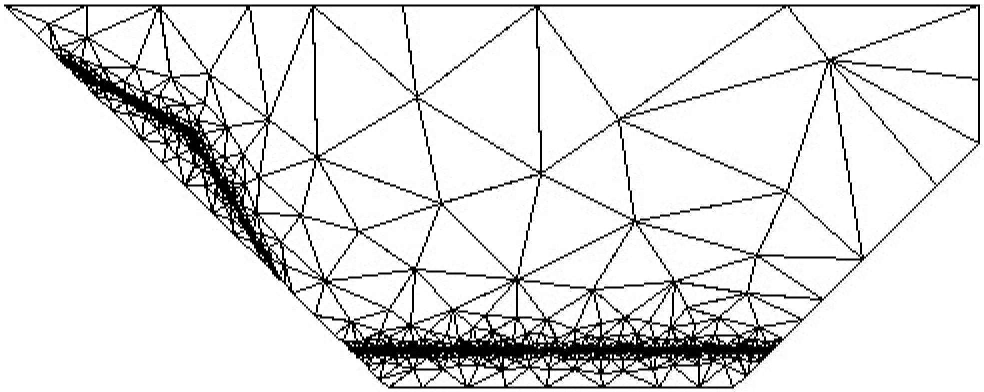

Comparing lab experiments (left) with numerics (right)

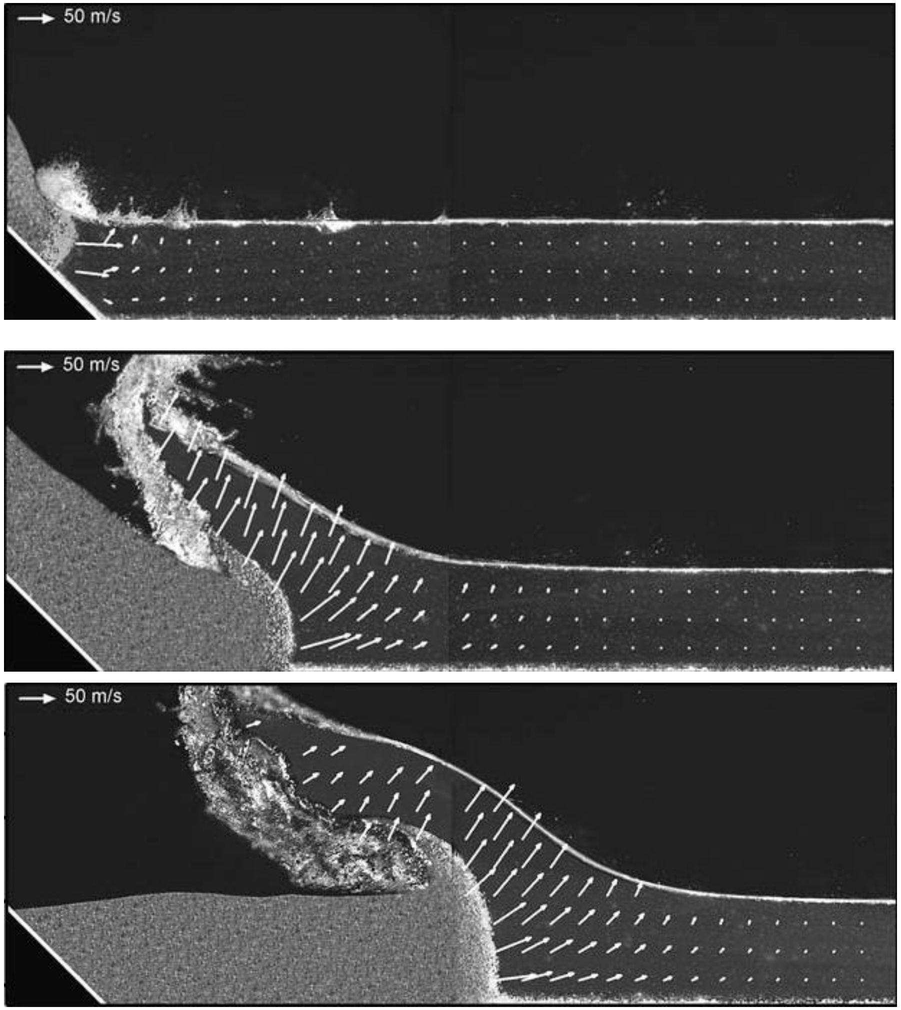

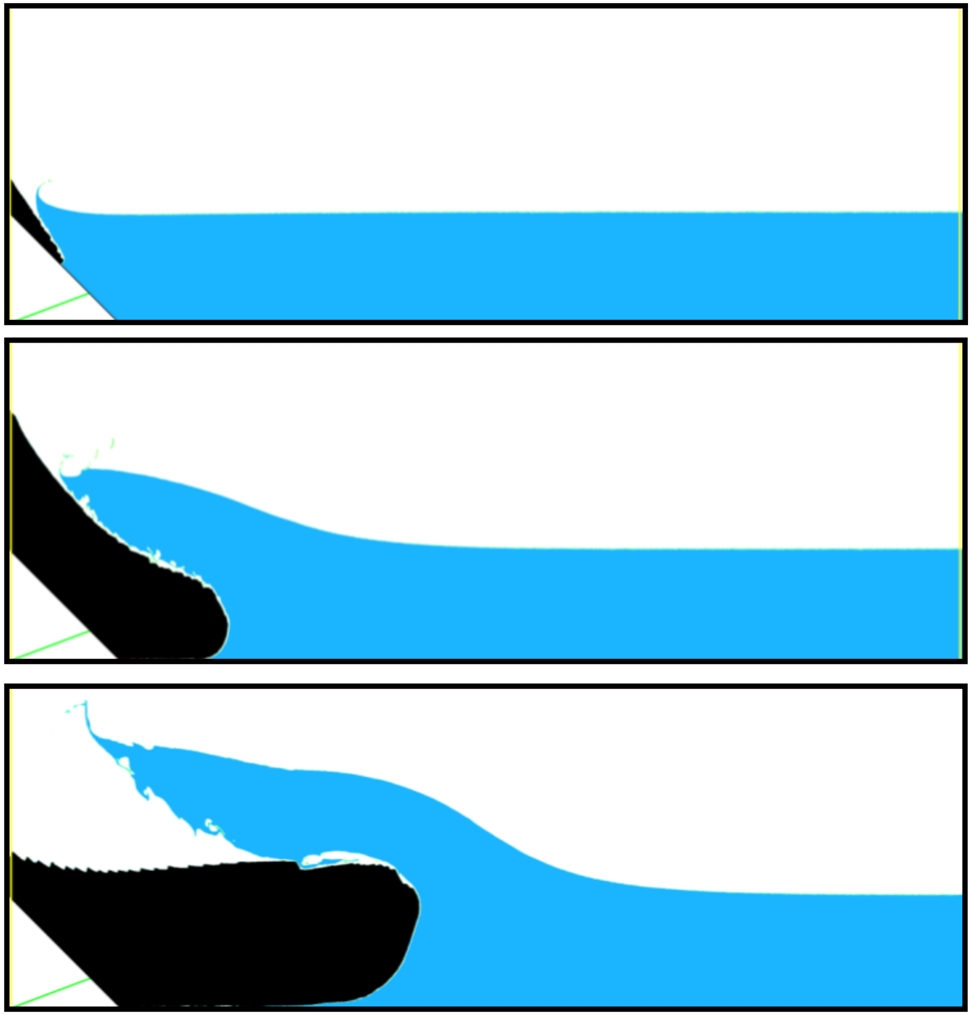

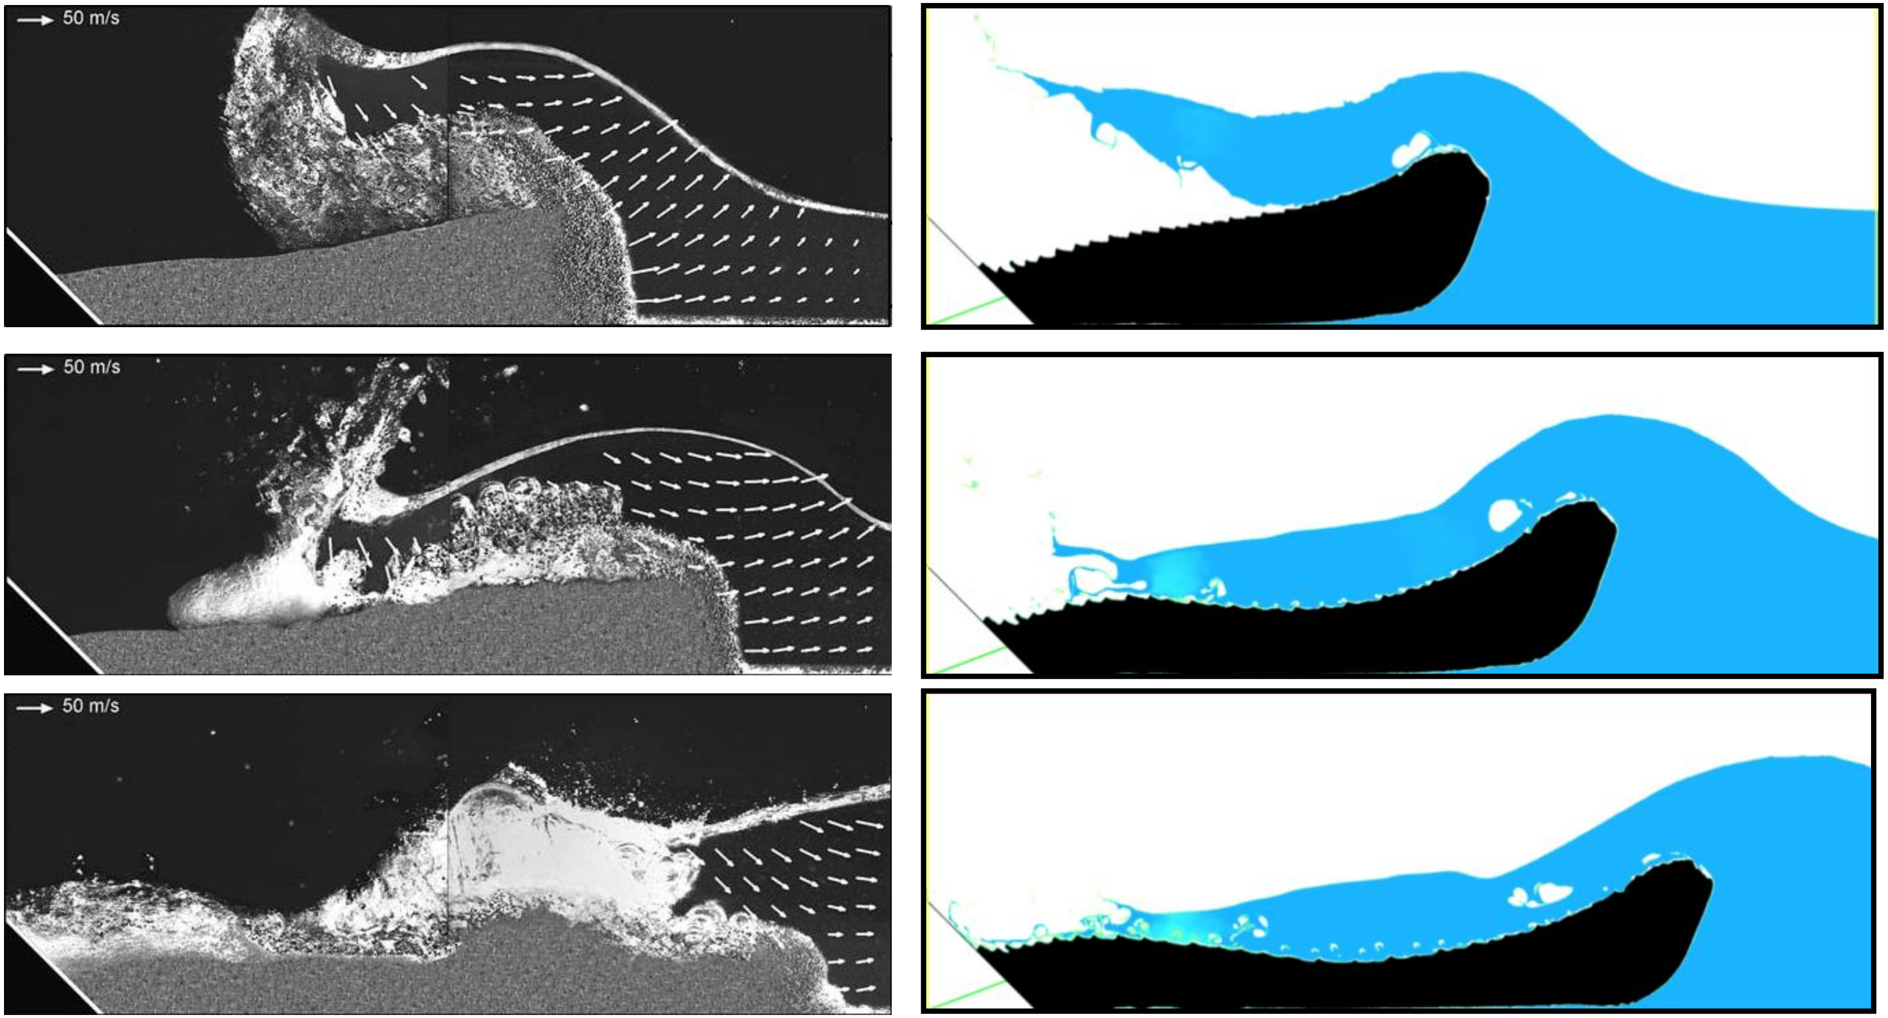
50 m/s
50 m/s
50 m/s

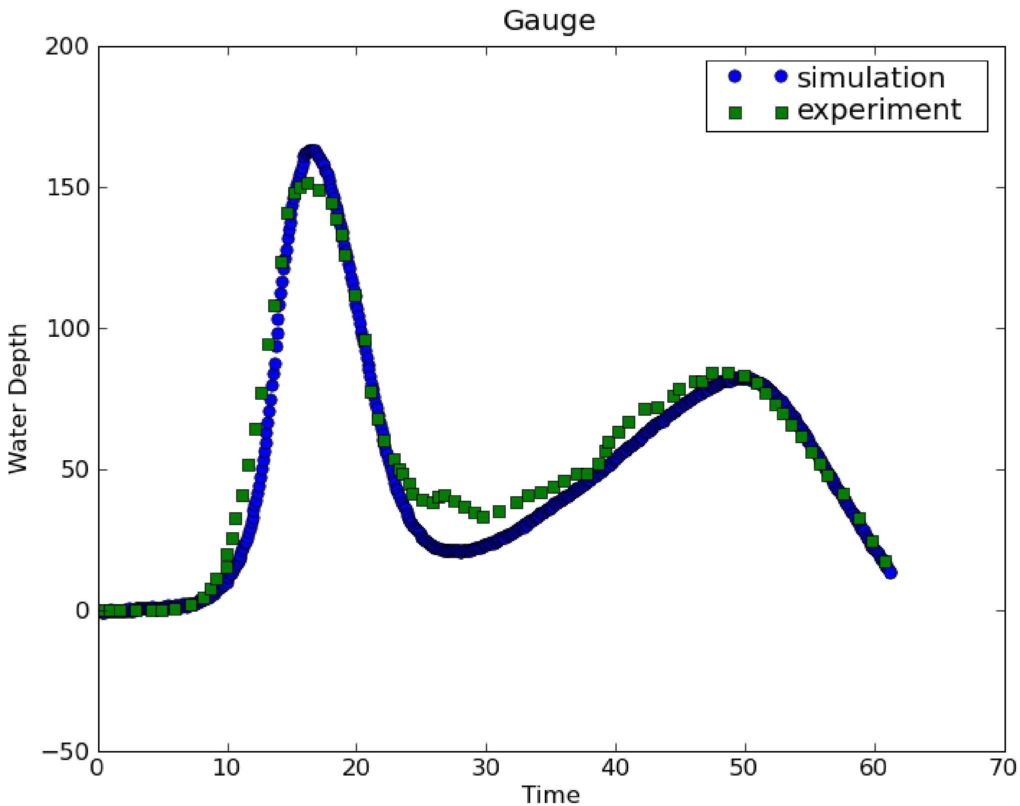
Gauge
simulation
experiment
Water Depth
Time
200
150
100
50
0
−50
0
10
20
30
40
50
60
70

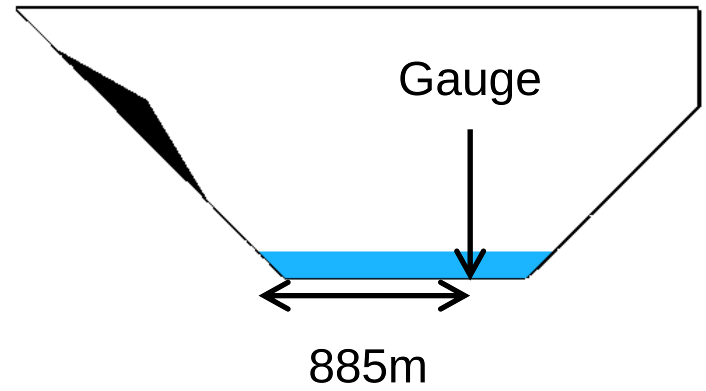
Gauge
885m

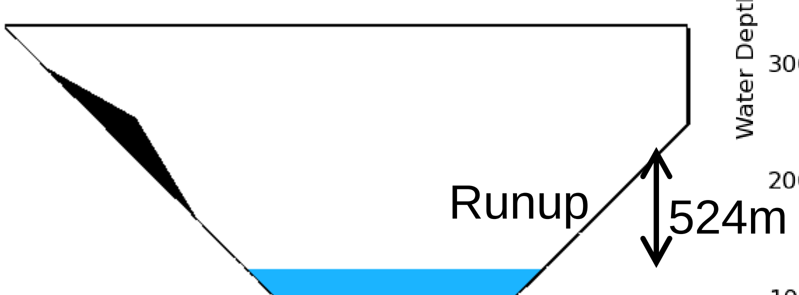
Runup
524m

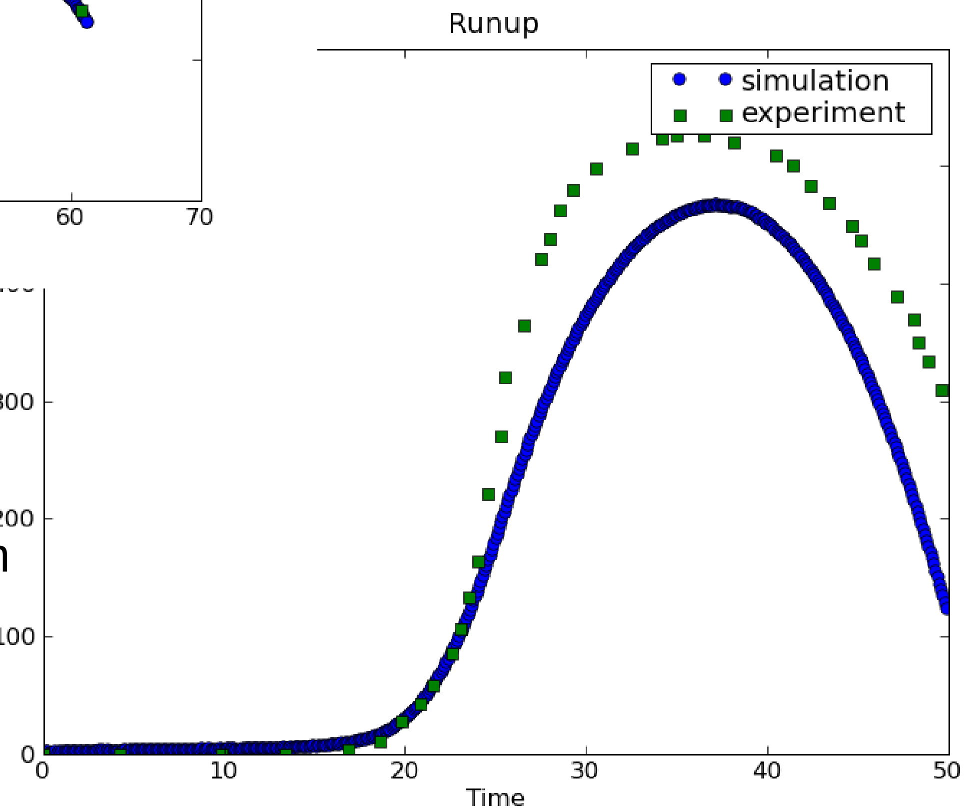
Runup
simulation
experiment
Water Depth
Time
300
200
100
0
0
10
20
30
40
50

Atmospheric pollution dispersal

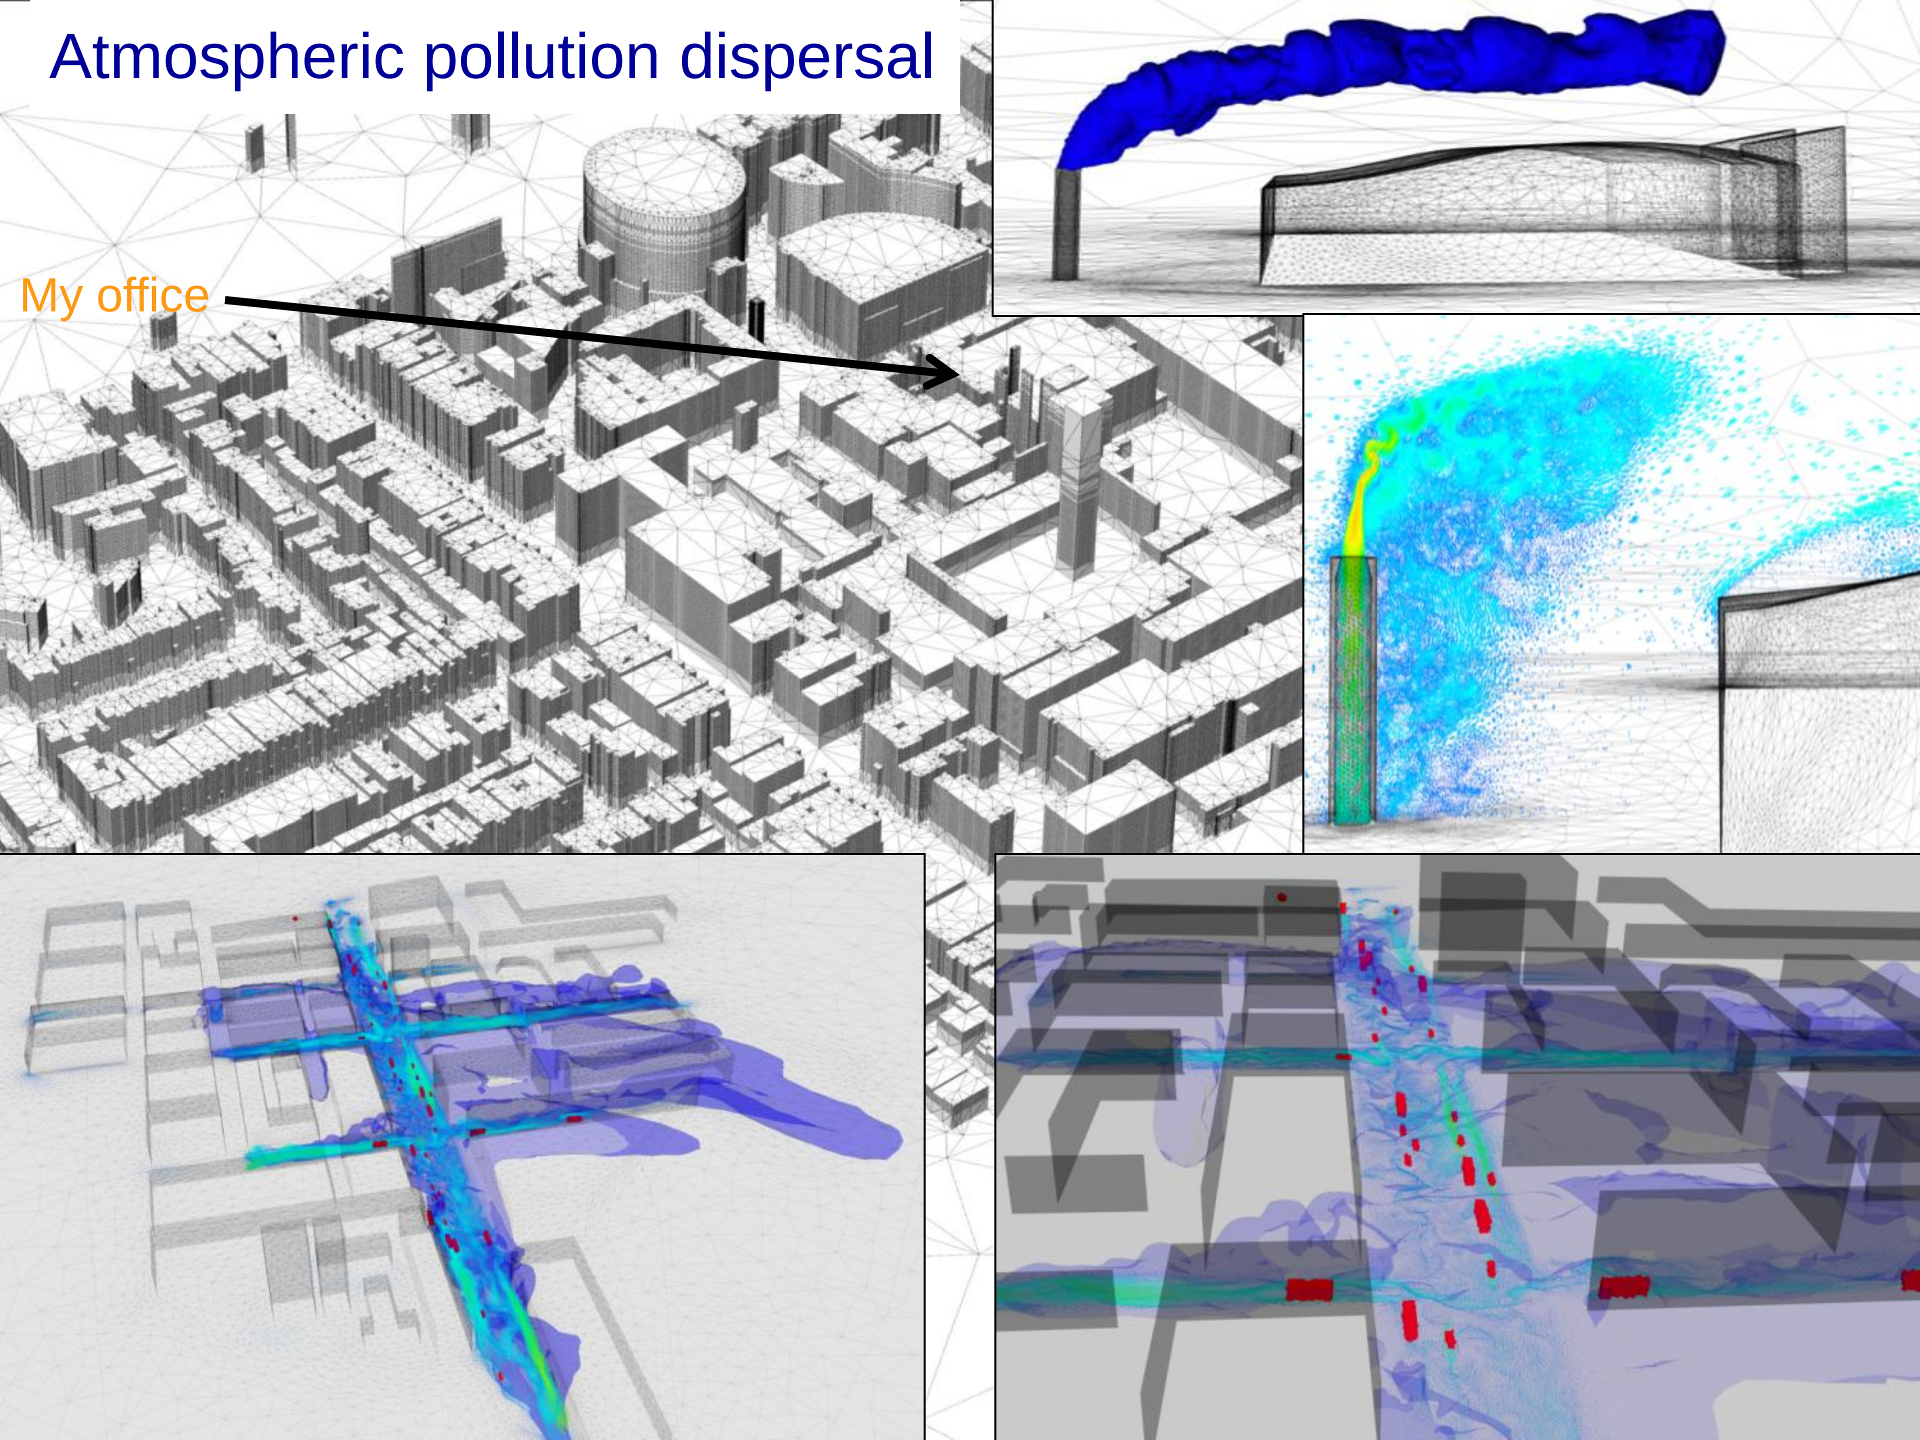

Final Comments

- Dynamic adaptivity (with anisotropy) seems to have compelling advantages for GFD, e.g. coupling estuary, shelf, deep ocean seamlessly; explicitly representing (more of the) processes crucial to the MOC
- Further validation on more realistic problems ongoing, including inter-comparisons with other models, rigorous accuracy/efficiency comparisons

Open Questions

- How (e.g. h, r, p, anisotropic), where & when to adapt?
- Stable, robust and accurate discretisation methods
- Resolving both advective and wave processes
- Diapycnal mixing rates
- How to link variable resolution with SGS models
- Coupling to other physics
- Time-stepping
- Which types of problems will unstructured meshes with/without adaptivity be an advantage/disadvantage, or the only realistically feasible approach?