Met Office Hadley Centre

Developing the Next Generation Dynamical Cores

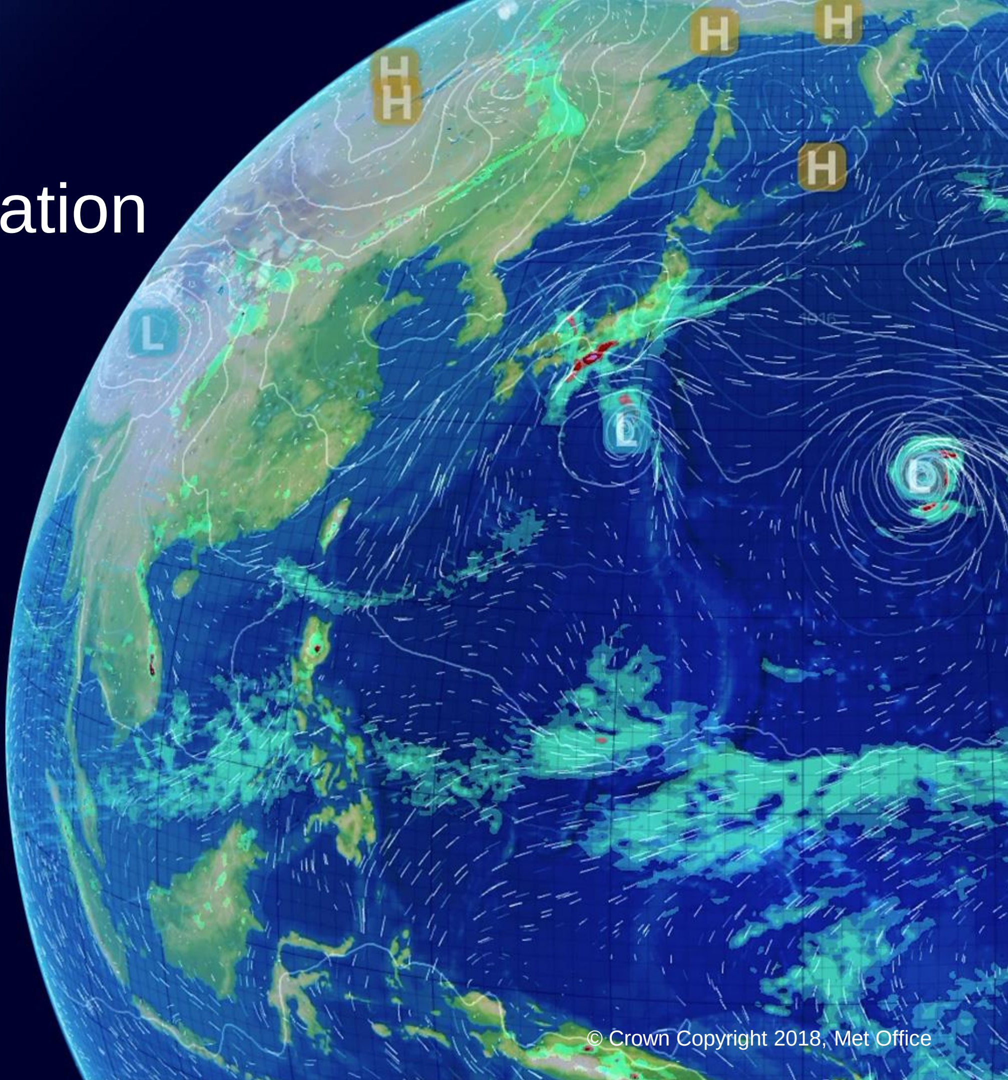

Ben Shipway + GungHo + LFRic + PSyclone + GHASP teams

ICAP/InDust meeting.

8 June 2018

Met Office
Hadley Centre

Developing the Next Generation ~~Dynamical Cores~~ Atmospheric models

Ben Shipway + GungHo + LFRic + PSyclone + GHASP teams

ICAP/InDust meeting.

8 June 2018

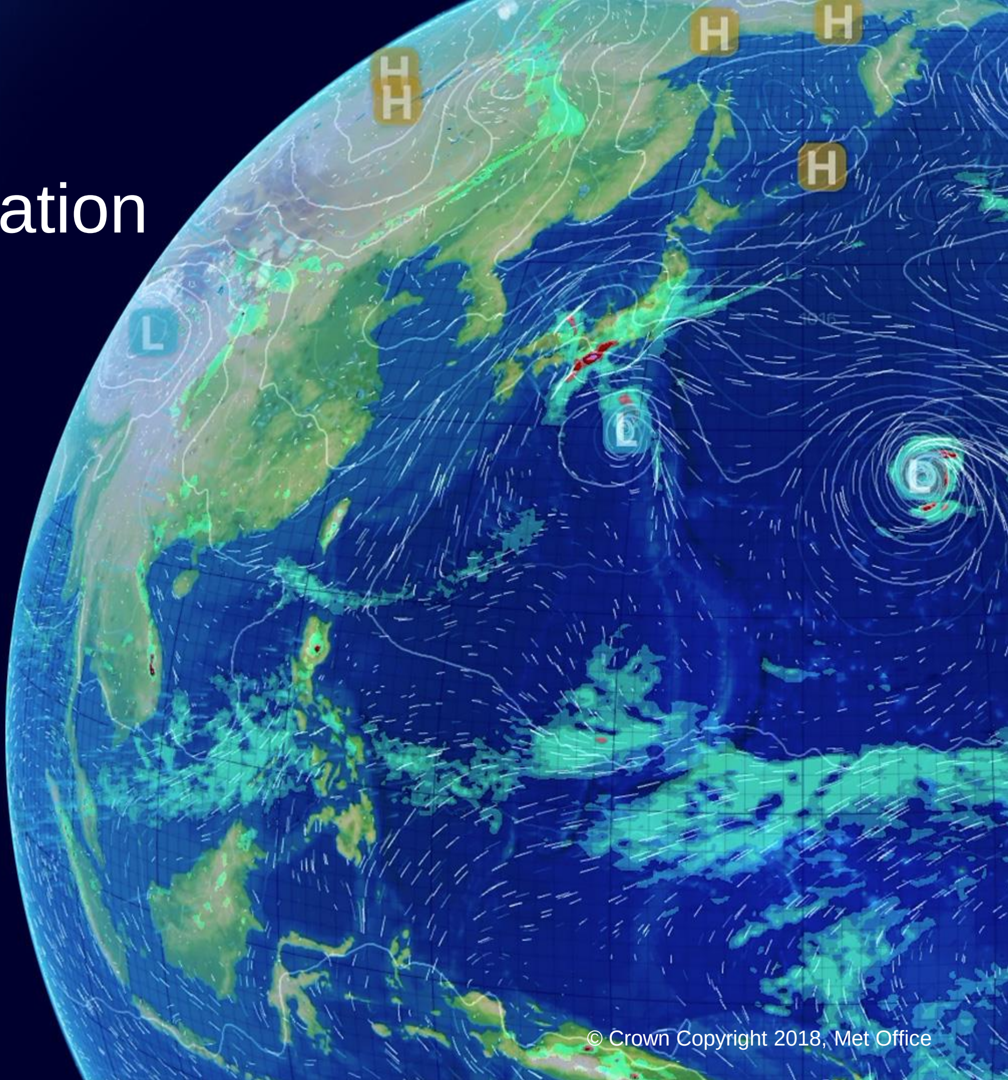

Overview

Scenes

Prologue

Act I: Why are we doing this?

- Scene 1. A new dynamical core?
- Scene 2. Exascale looms

Act II: How are we doing this?

- Scene 1. GungHo dynamical core
- Scene 2. LFRic infrastructure
- Scene 3. Psyclone

Act III: A new NWP & Climate model

- Scene 1. GHASP

The Key Protagonists

GungHo, A project and a new dynamical core

LFRic, A new model infrastructure and a new modelling system?

PSyClone, Code generation software for computational science

GHASP, The GungHo Atmospheric Science Project

Exascale, computing with ExaFlops and a catch-all term for next-generation HPC architectures

Finite Element Method (FEM), a numerical discretization akin to Finite Difference or Finite Volume

ENDGame, The current dynamical core

Met Office Hadley Centre

Prologue

- 2009: New Chief Scientist from academia worried about how the Unified Model scales

- Initiates a joint Met Office/UK academia project (GungHo) to research the redesign of the dynamical core

Story begins at the ENDGame

- ENDGame operational 2014
- Greatly improved scalability…
- …but not enough for Exascale

Andy Malcolm, Paul Selwood

- At 25km resolution, grid spacing near poles = 75m

- At 10km reduces to 12m!

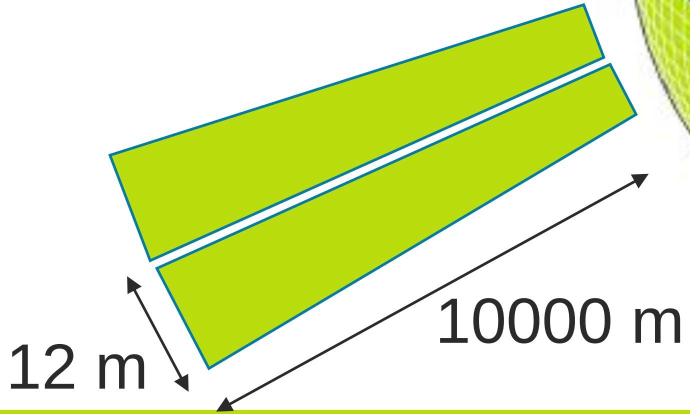

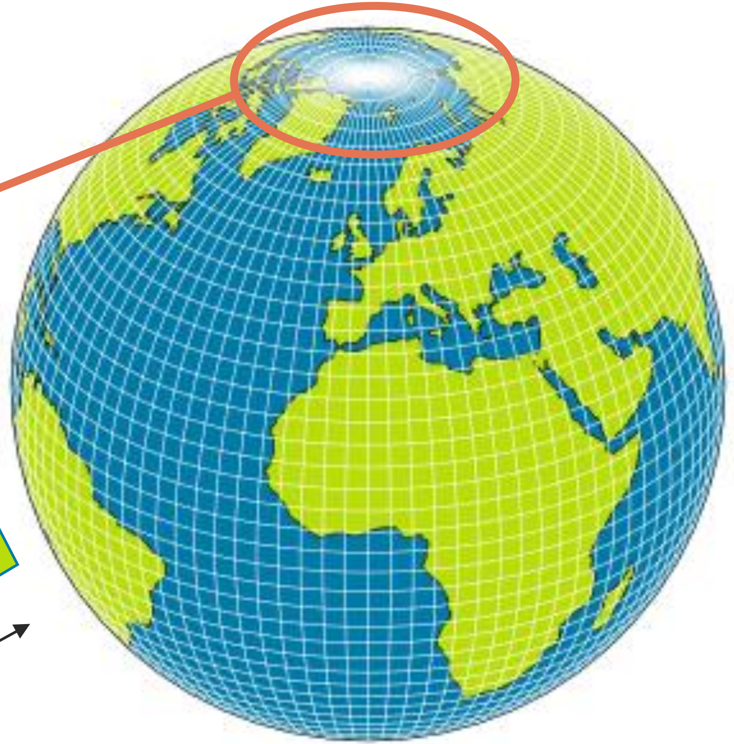

Exascale looms

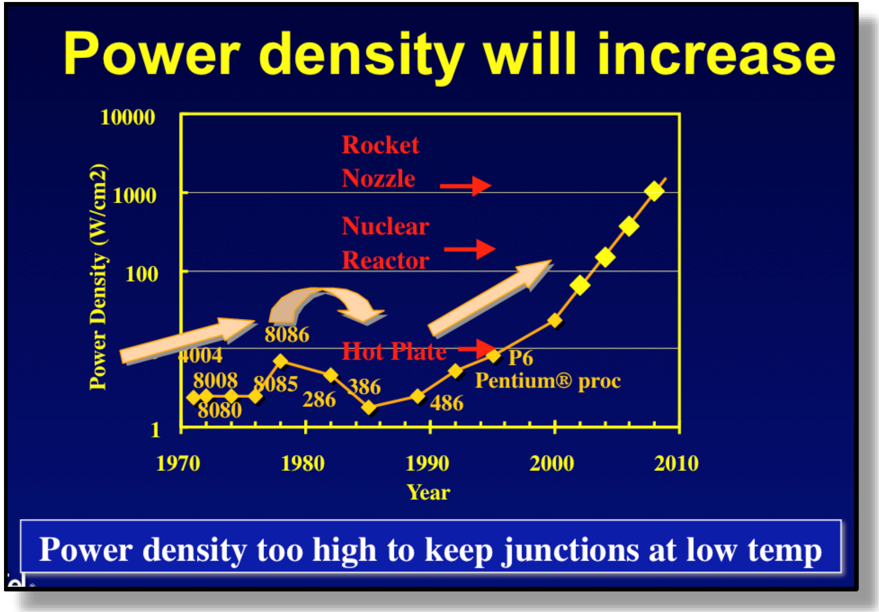

Intel, 2000

Met Office
Hadley Centre

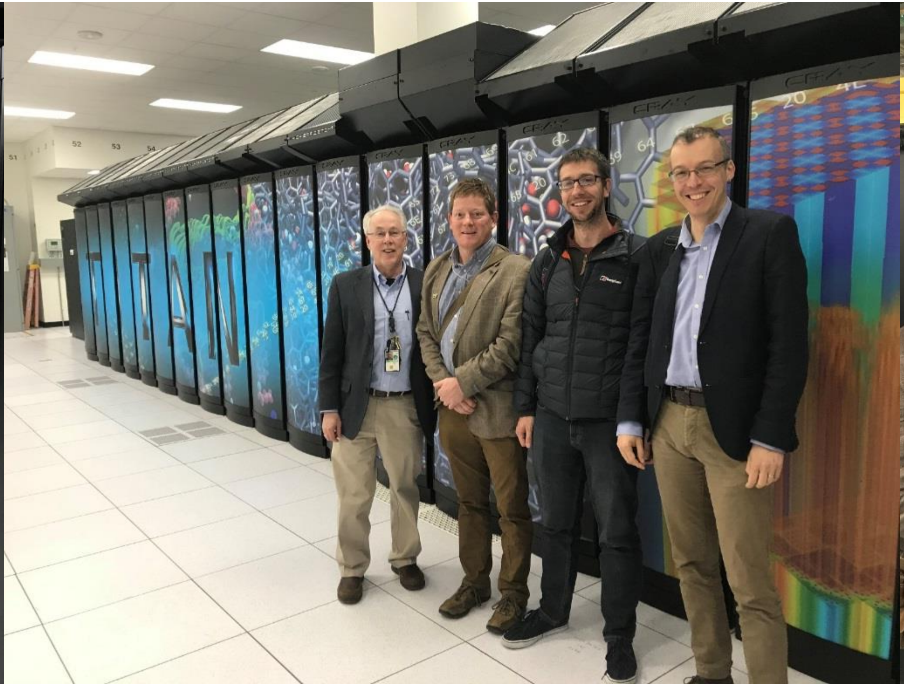

How are we going to do this?

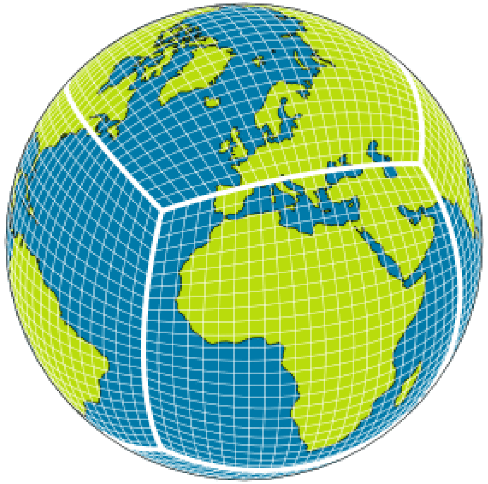

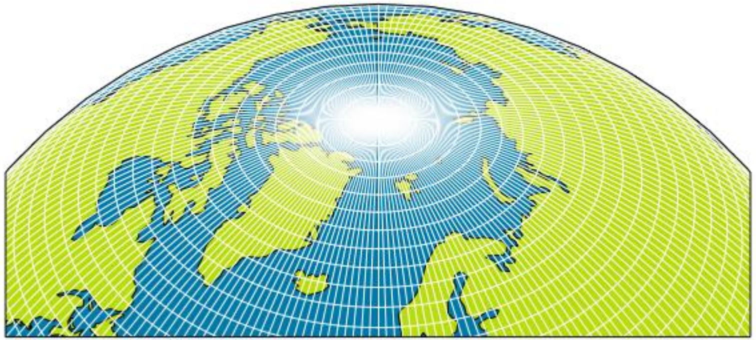

Met Office Hadley Centre

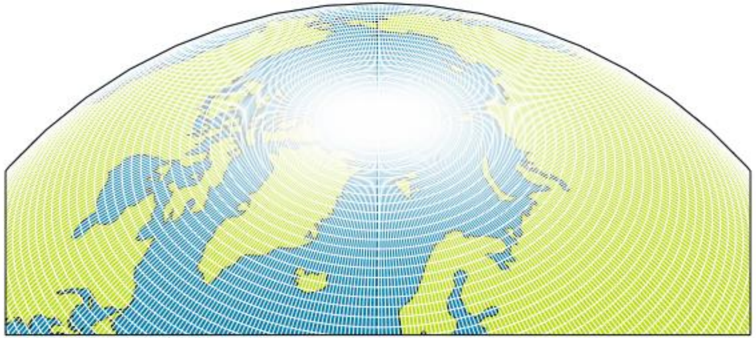

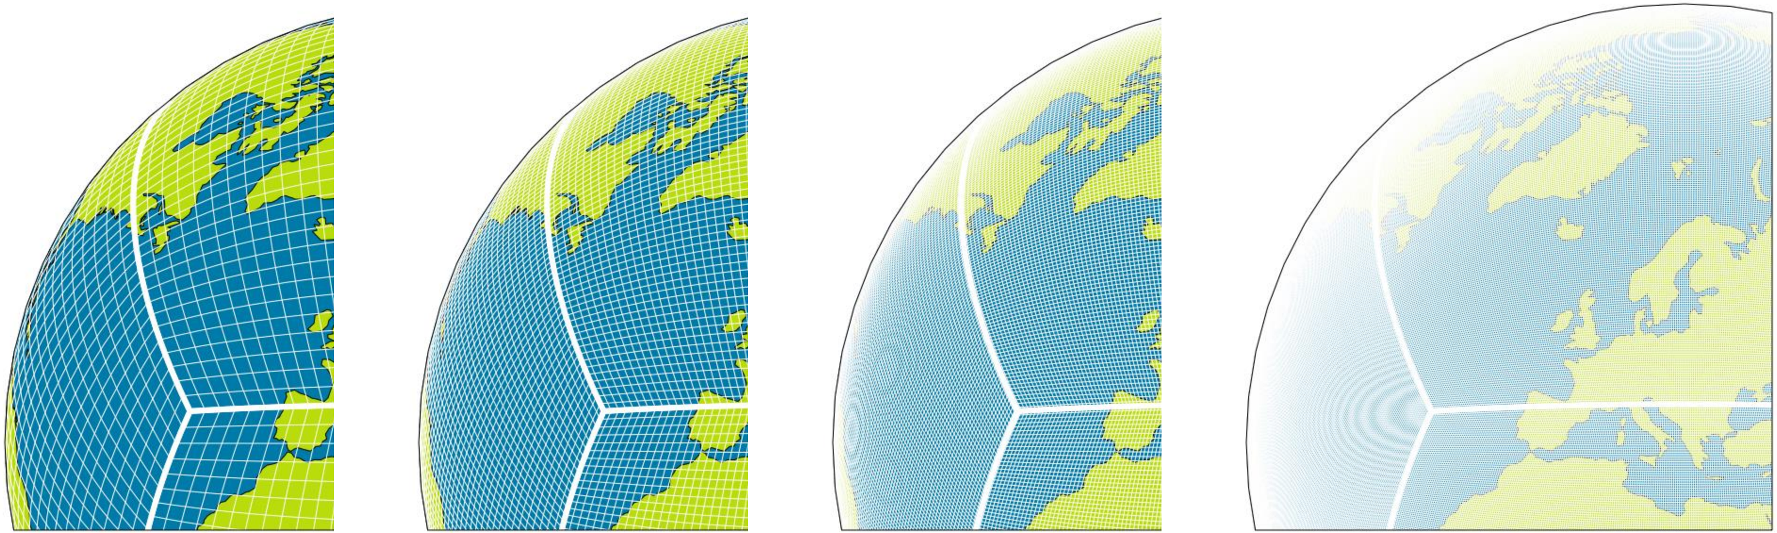

Mixed FEM

- How to maintain accuracy of current model on a GungHo grid? Staniforth & Thuburn (2012)
- Principal points about current grid are:
  - Orthogonal, Quadrilateral, C-grid
- These allow good numerical aspects:
  - Lack of spurious modes
  - Mimetic properties
  - Good dispersion properties

Mixed FEM

Choose function spaces $\mathbb{W}_i$ $(i = 0, .., 3)$ such that the following mapping holds

$$\mathbb{W}_0 \xrightarrow{\nabla} \mathbb{W}_1 \xrightarrow{\nabla\times} \mathbb{W}_2 \xrightarrow{\nabla.} \mathbb{W}_3$$

- For basis function $P \in \mathbb{W}_0$ then $\nabla P \in \mathbb{W}_1$ and can choose $\mathbb{W}_0$ to enforce compatibility: e.g $\nabla \times \nabla P \equiv 0$.

Function spaces correspond to geometric objects:

$\mathbb{W}_0$ Pointwise scalar functions

$\mathbb{W}_1$ Circulation vectors (e.g. vorticity)

$\mathbb{W}_2$ Flux vectors (e.g. mass flux)

$\mathbb{W}_3$ Volume integrated scalars (e.g. density)

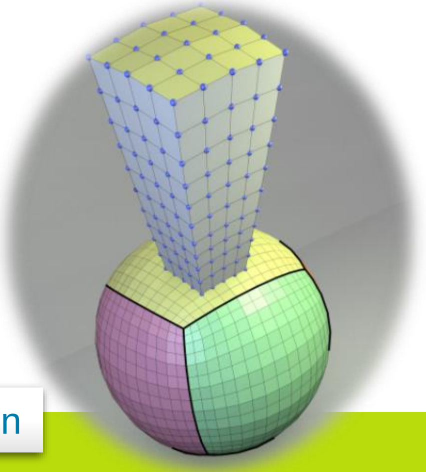

Tom Melvin

Mixed FEM

$$\begin{array}{ccccccc} \theta & & \xi & & \mathbf{u} & & \rho \\ \mathbb{W}_0 & \xrightarrow{\nabla} & \mathbb{W}_1 & \xrightarrow{\nabla\times} & \mathbb{W}_2 & \xrightarrow{\nabla\cdot} & \mathbb{W}_3 \end{array}$$

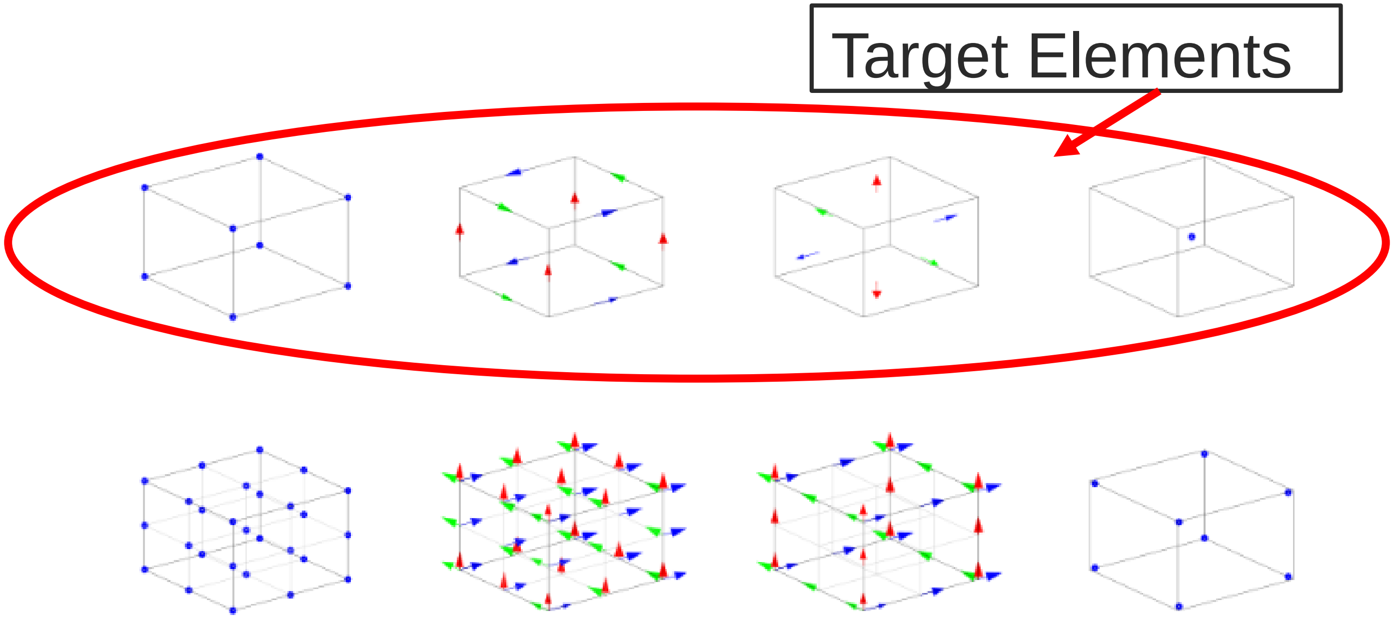

Tom Melvin

GungHo vs ENDGame

	ENDGame	GungHo
Grid	Lat-Long	Cubed-Sphere
Prognostic Variables	ρ, θ, Π, u, v, w	ρ, θ, Π, **u**
Prognostics Equations	Advective form	Flux/Advective/Vector Invariant form
Spatial Discretisation	2^{nd} Order FD	Mixed FEM
Temporal Discretisation	Iterative Semi-Implicit	Iterative Incremental Semi-Implicit
Advection	Semi-Lagrangian	COSMIC (Dimensionally split, Eulerian)

GungHo Transport

- Flux-based Semi-Lagrangian
- Key area where there is a deliberate departure from ENDGame

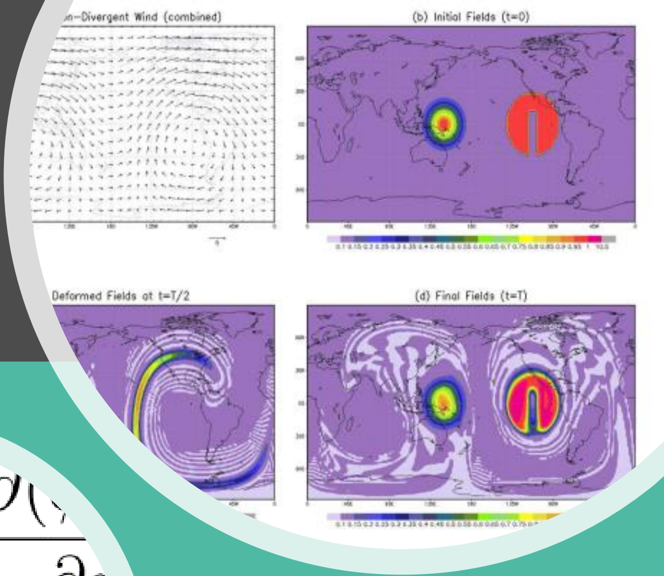

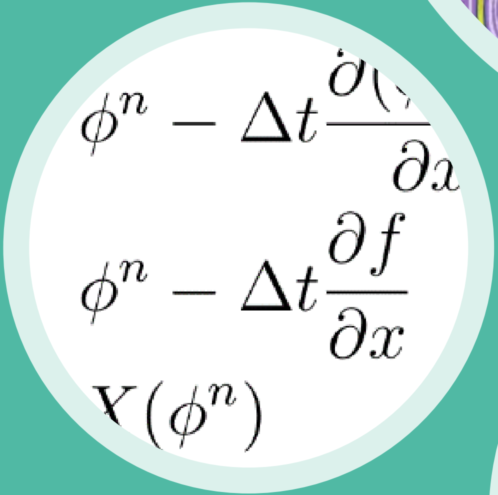

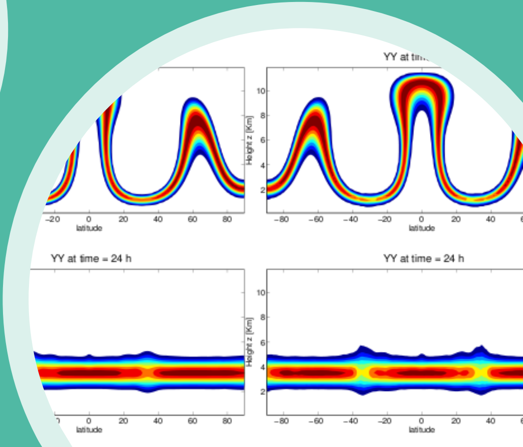

Importance of conservation?

- High resolution LAM simulations
- Strong updrafts at the grid scale
- 'Pointwise' semi-Lagrangian transport cannot 'see' convergence properly
- Significant impact on the precipitation

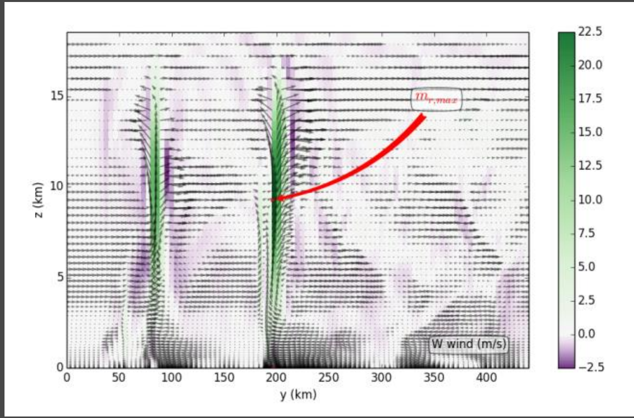

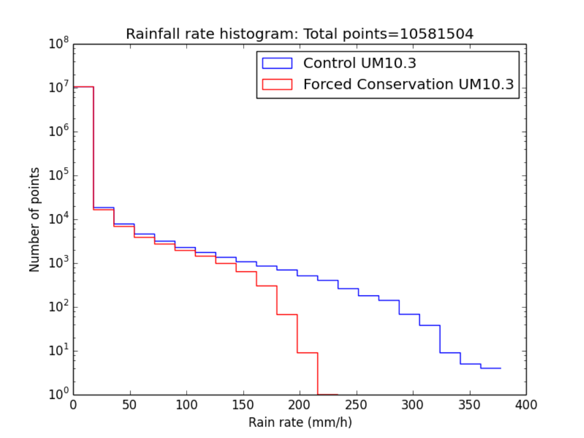

Age of Air

- One of several tests designed to assess impact of transport formulation and features such as conservation

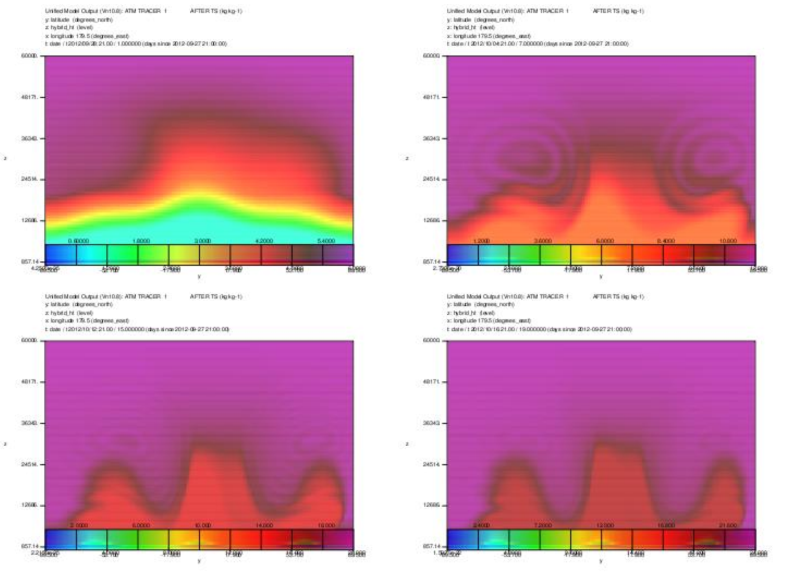

Figure 7: An idealised climate with age of air at different times; flow-field=idelaised climate, arge of air field at t=0, and source=1

Mohamed Zerroukat

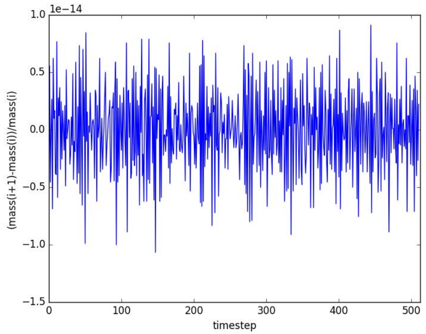
1.0 1e−14
0.5
0.0
−0.5
−1.0
−1.5
(mass(i+1)-mass(i))/mass(i)
0
100
200
300
400
500
timestep

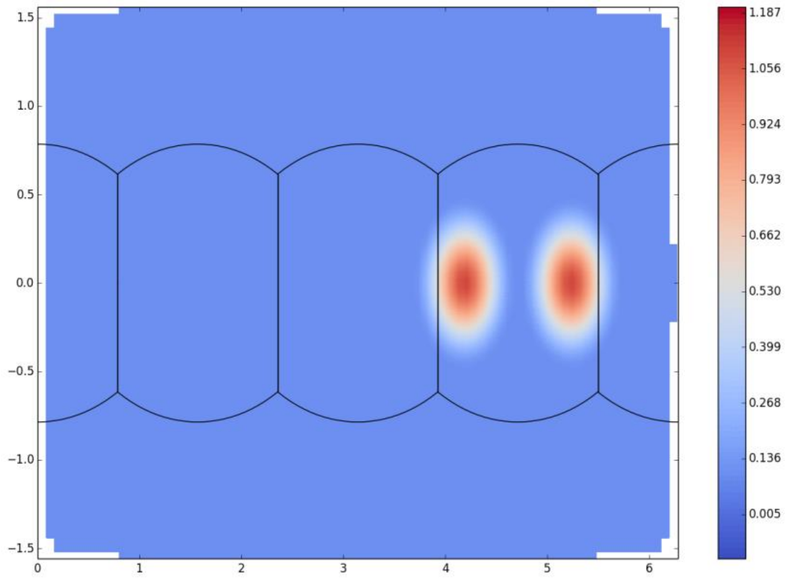
1.5
1.0
0.5
0.0
−0.5
−1.0
−1.5
0
1
2
3
4
5
6
1.187
1.056
0.924
0.793
0.662
0.530
0.399
0.268
0.136
0.005

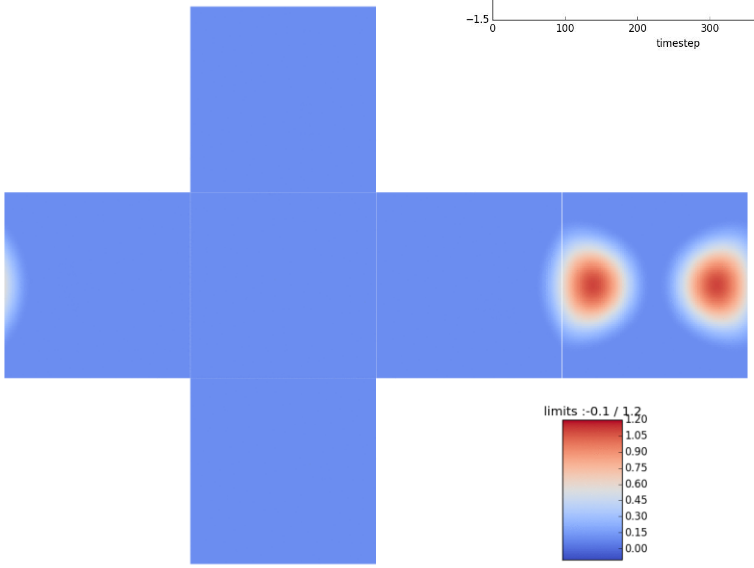
limits :-0.1 / 1.2
1.20
1.05
0.90
0.75
0.60
0.45
0.30
0.15
0.00

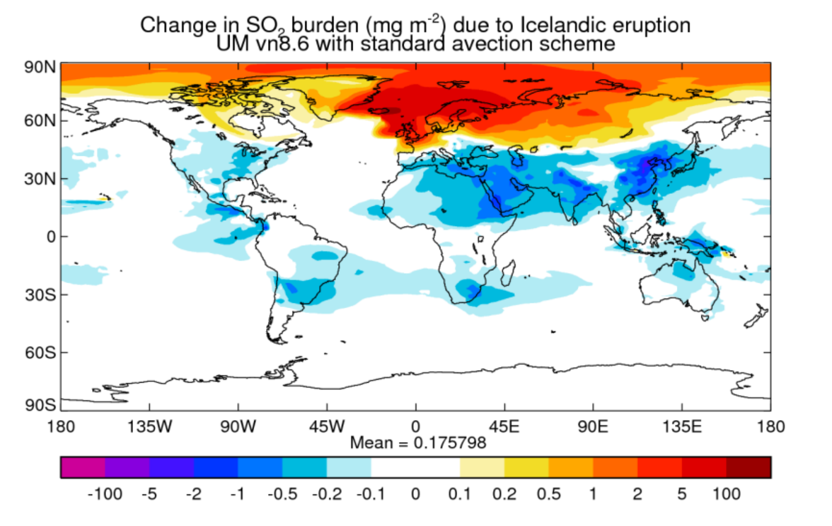

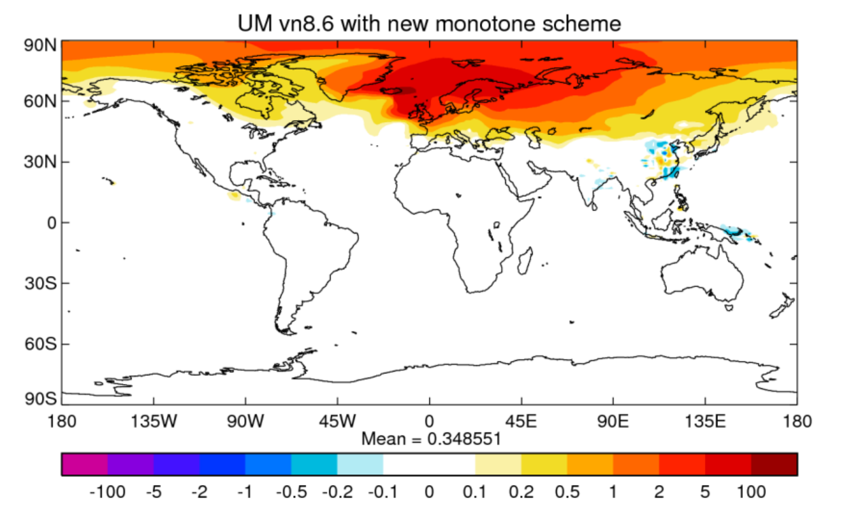

Monotonicity

Met Office
Hadley Centre

LFRic Infrastructure

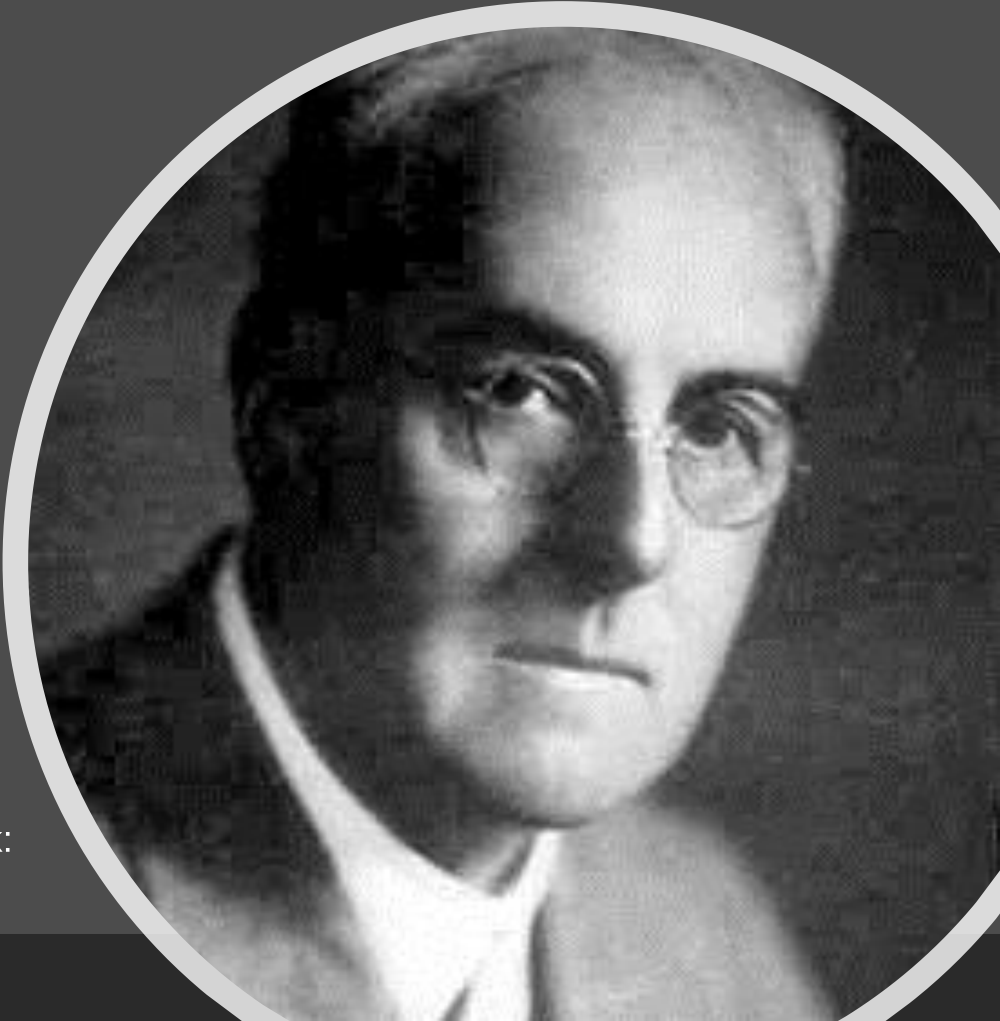

In 1922, L.F. Richardson published the book:
Weather Prediction by Numerical Process

Separation of Concerns

"PSyKAL"
model

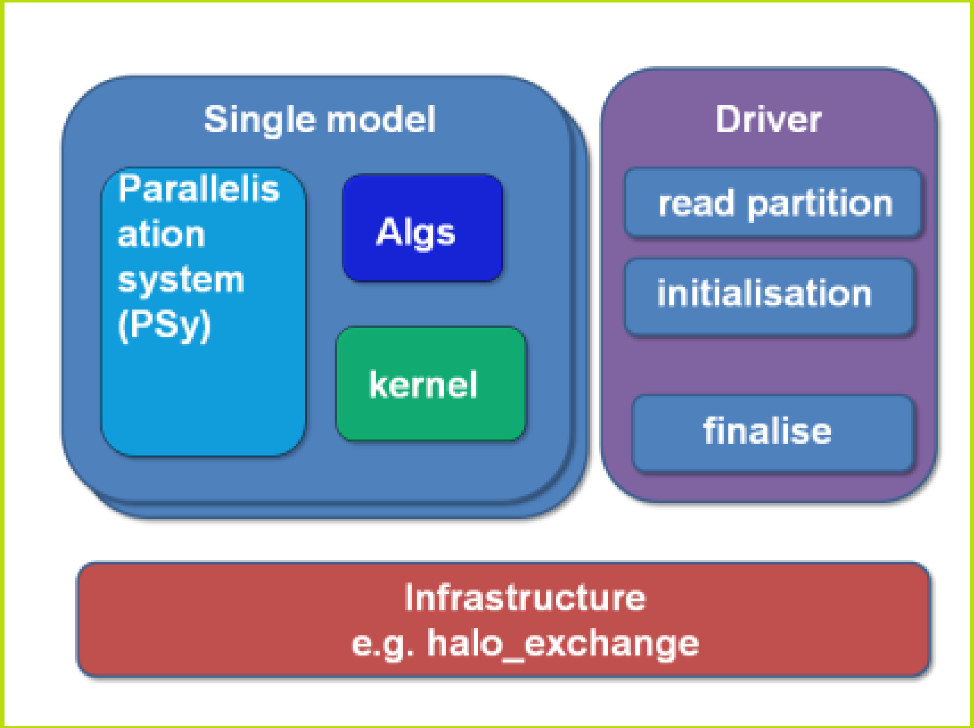

Met Office Hadley Centre

An example code

y = a exp(b)

Serial code

```
…
do i = 1,N
 do k = 1,nlayers
   y(k,i) = a(k,i) * exp( b(k,i) )
 end do
end do
…
```

Distributed memory (MPI)

```
…
call halo_exchange(a) # is this needed?
call halo_exchange(b) # is this needed?

do i = 1, i_halo_extent
 do k = 1,nlayers
   y(k,i) = a(k,i) * exp( b(k,i) )
 end do
end do
…
```

An example code

$\underline{y} = \underline{a}\ \exp(\ \underline{b})$

Serial code

```
…
  do i = 1,N
   do k = 1,nlayers
     y(k,i) = a(k,i) * exp( b(k,i) )
   end do
  end do
…
```

Shared memory (OpenMP)

```
…
  do ithread = 1, nthreads
  !$omp parallel default(shared), private(i)
  !$omp do schedule(static)
   do i = istart(ithread), iend(ithread)
    do k = 1,nlayers
     y(k,i) = a(k,i) * exp( b(k,i) )
    end do
   end do
  !$omp end do
  !$omp end parallel
  end do
…
```

An example code

$\underline{y} = \underline{a}\ \exp(\ \underline{b})$

LFRic algorithm code

```
Use my_kernel_mod, only: my_kernel
…
   call invoke( my_kernel( y, a, b) )
…
```

Generated PSY layer code

```
Use my_kernel_mod, only: my_kernel
… DIRECTIVES + HALO EXCHANGES …
… UNPACK DATA AND INDIRECTION MAPS …
  do cell = 1, ncells
     call invoke_my_kernel( nlayers,           &
            y(:,cell), a(:,cell), b(:,cell), &
            map(cell) )
```

LFRic kernel code

```
module my_kernel_mod

… METADATA …

subroutine my_kernel(nlayers, y, a, b
                     map)
…
do k = 1, nlayers
   y( map + k )  =  a( map + k ) * exp( b( map + k ) )
end do
…
```

Met Office
Hadley Centre

PSyClone

Recap: Separation of Concerns

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

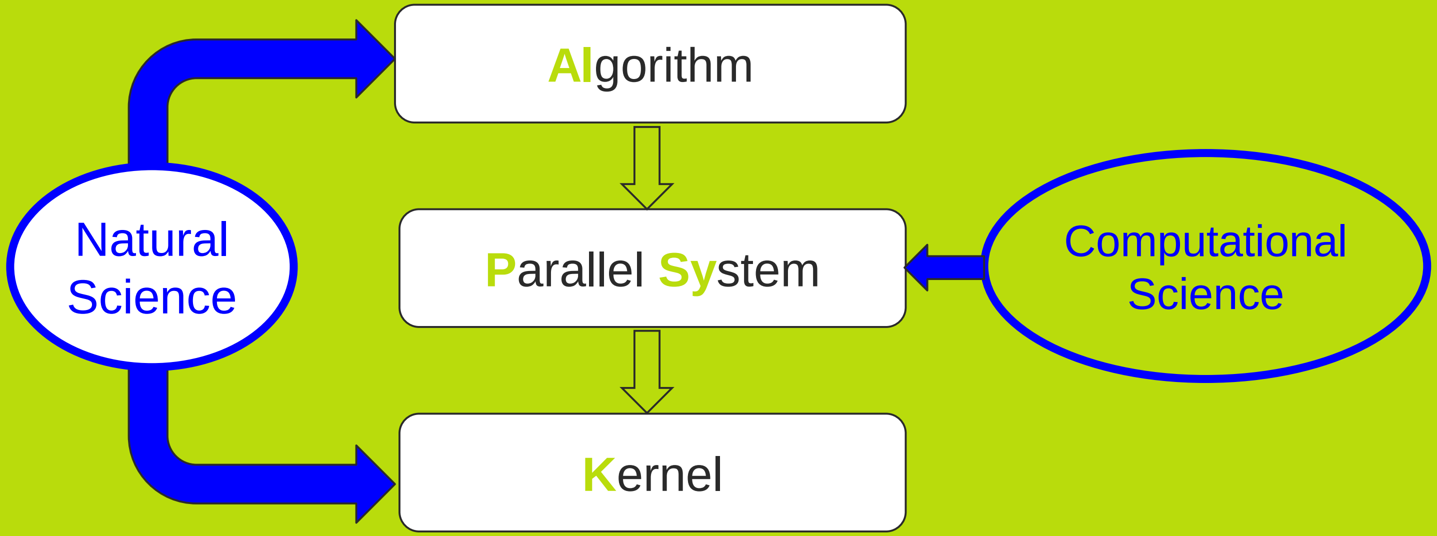

Rupert Ford and Andy Porter, Hartree Centre

Generating the PSy Layer

- A domain-specific compiler for embedded DSL(s)
  - Finite Difference/Volume, Finite Element
  - Currently Fortran -> Fortran
  - Supports distributed- and shared-memory parallelism
  - Should reduce programmer errors (both correctness and optimisation)

- A tool for use by HPC experts
  - Hard to beat a human
  - Optimisations encoded as a ‘recipe’ rather than baked into the scientific source code
  - Different recipes for different architectures

Rupert Ford and Andy Porter, Hartree Centre

LFRic + GungHo + PSyclone

It works!

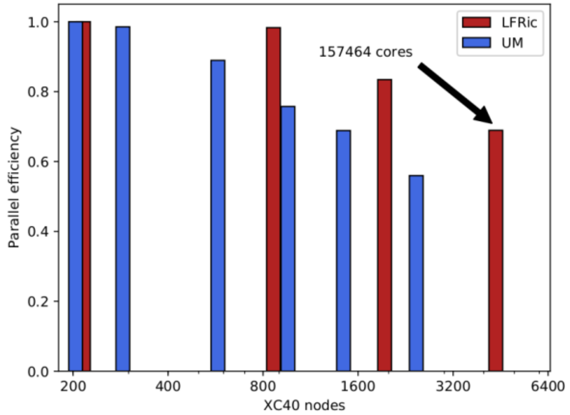

Figure 1: Scaling of the parallel efficiency of an N2048 configuration of the Unified Model time-step compared to *Dromedary* release of LFRic (baroclinic wave) for a C1944 configuration time-step.

Chris Maynard

Met Office
Hadley Centre

The GungHo Atmospheric Science Project

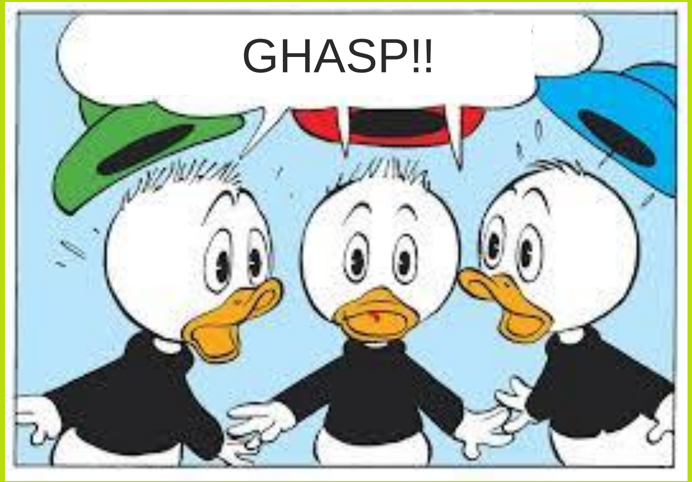

Met Office
Hadley Centre

GHASP

Basic **coupling infrastructure** enables lowest order finite element fields to be used as input to the physics schemes.

Minimal updates to UM code. Preprocessor directive available where necessary:
#ifdef LFRIC

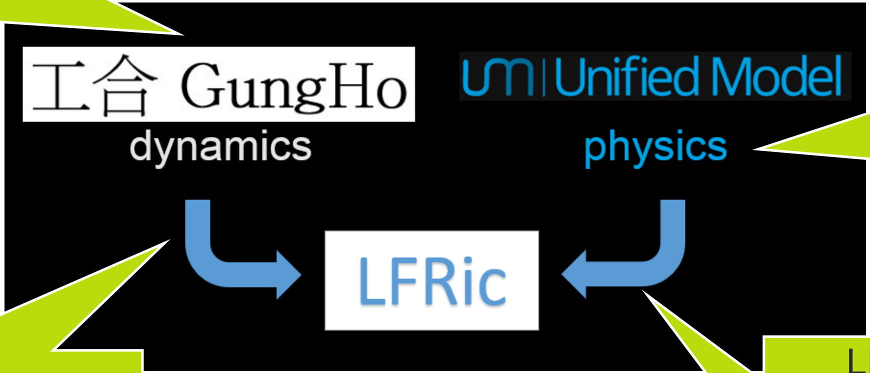

UM **boundary layer, radiation and microphysics** codes have been run with idealized forcing from GungHo dynamical core.

LFRic build system modified to extract, preprocess and build UM **source code directly from the UM repositories.**

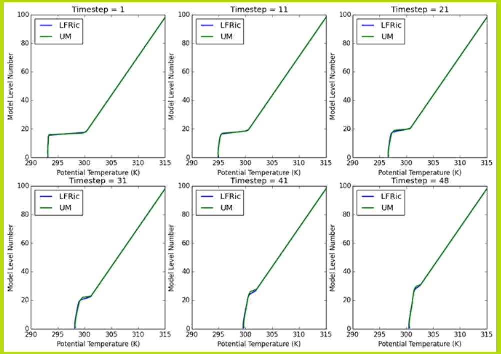

Timestep = 1
Timestep = 11
Timestep = 21
Timestep = 31
Timestep = 41
Timestep = 48
LFRic
UM
Model Level Number
Potential Temperature (K)

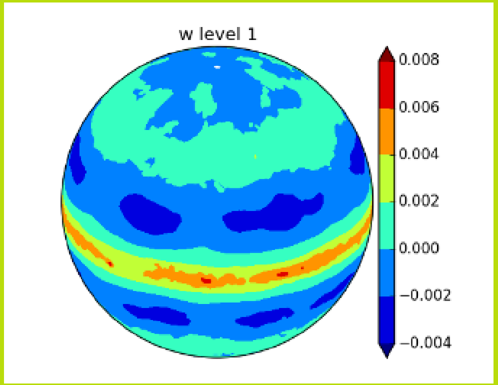

w level 1
0.008
0.006
0.004
0.002
0.000
−0.002
−0.004

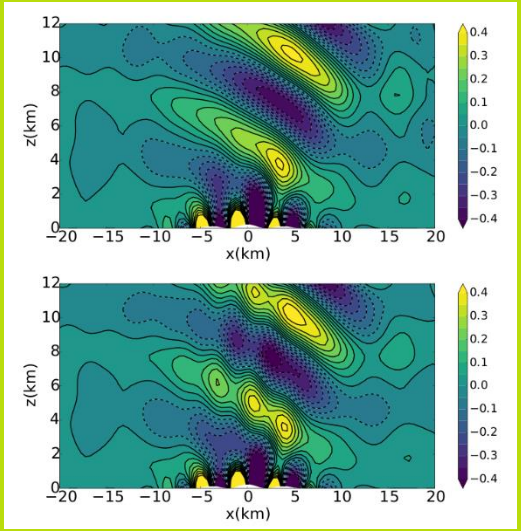

z(km)
x(km)
0.4
0.3
0.2
0.1
0.0
−0.1
−0.2
−0.3
−0.4

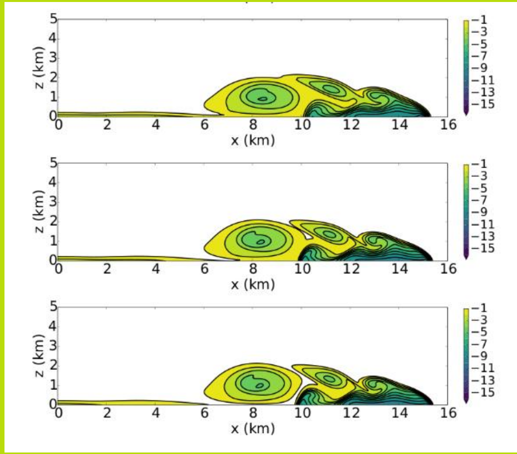

z (km)
x (km)
−1
−3
−5
−7
−9
−11
−13
−15

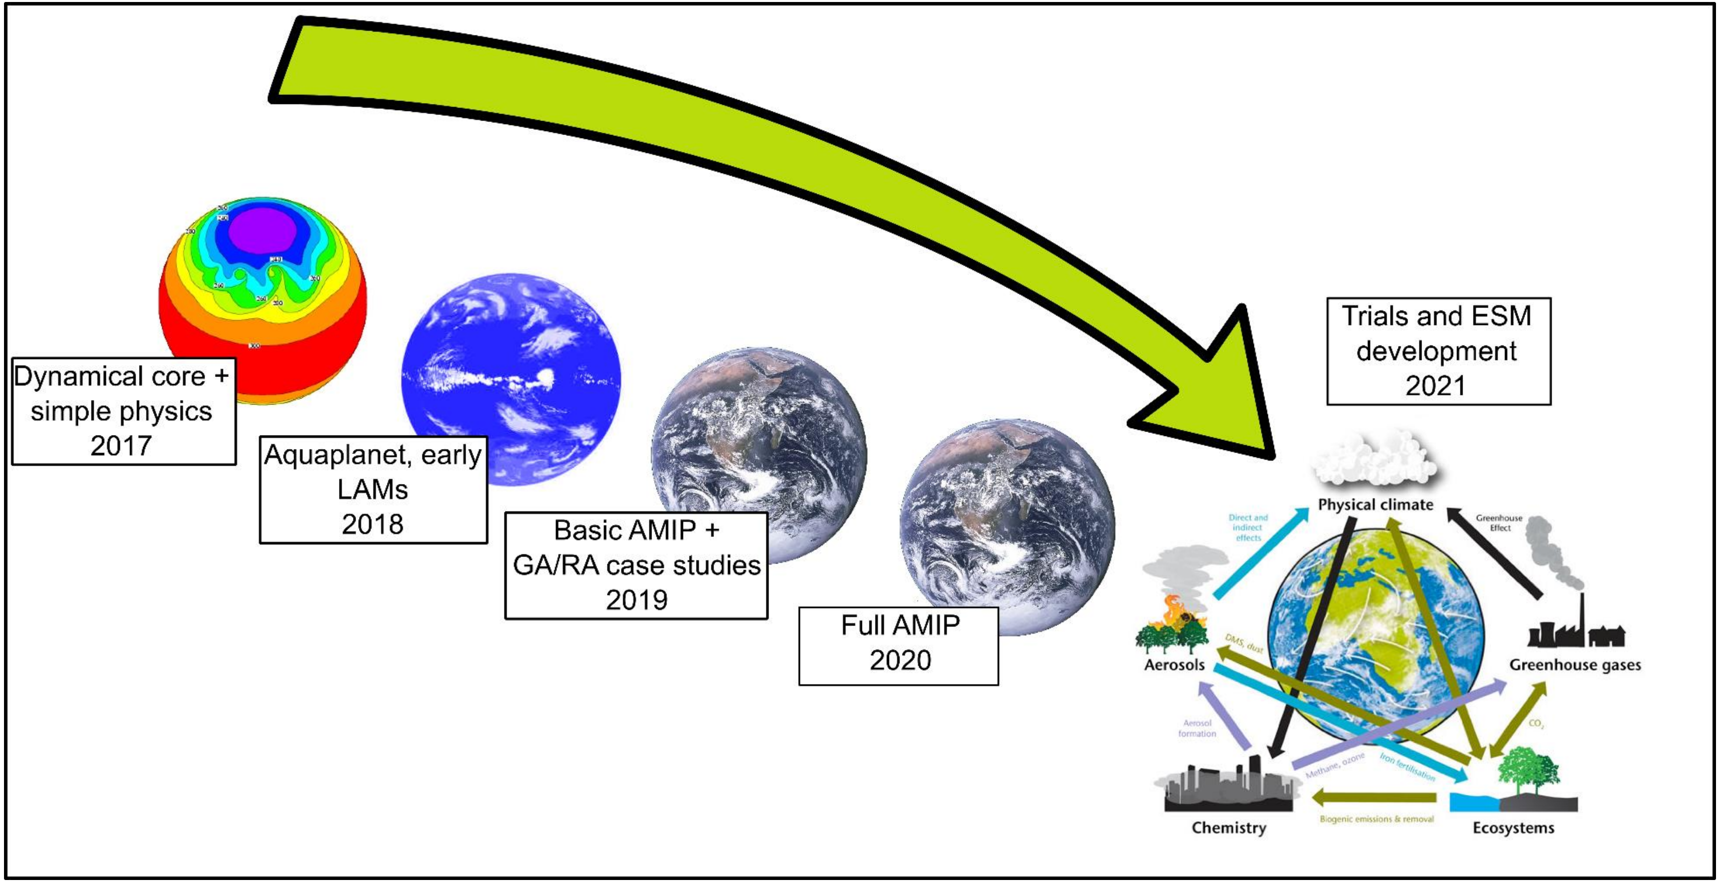
Dynamical core +
simple physics
2017
Aquaplanet, early
LAMs
2018
Basic AMIP +
GA/RA case studies
2019
Full AMIP
2020
Trials and ESM
development
2021
Physical climate
Direct and indirect effects
Greenhouse Effect
Aerosols
Greenhouse gases
DMS, dust
Aerosol formation
CO2
Methane, ozone
Iron fertilisation
Chemistry
Biogenic emissions & removal
Ecosystems

Extra slides

Met Office Hadley Centre

Split direction scheme –1D operators

- 1D advection operators exist in LFRic algorithm layer
- Stencils used to extract data from neighbouring cells

$$\frac{\partial \phi}{\partial t} + \frac{\partial(\phi u)}{\partial x} = 0$$

$$\longrightarrow$$

$$\begin{aligned} \phi^* &= \phi^n - \Delta t \frac{\partial(\phi^n u)}{\partial x} \\ &= \phi^n - \Delta t \frac{\partial f}{\partial x} \\ &= X(\phi^n) \end{aligned}$$

Met Office Hadley Centre

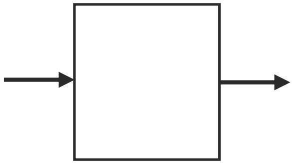

Flux calculation

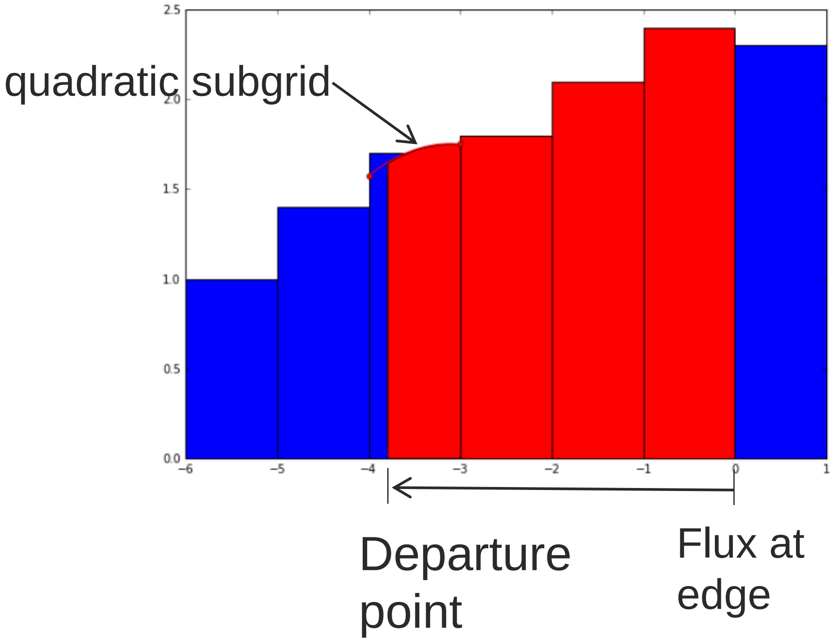

$$
\begin{aligned}
\text{flux} &= \frac{1}{\Delta t}\left[\int_{x_{dep}}^{x_{arr}} \phi dx\right] \\
&= \frac{1}{\Delta t}\left[\int_{-3.8}^{0} \phi dx\right] \\
&= \frac{1}{\Delta t}\left[\int_{-3.8}^{-3} \phi dx + \int_{-3}^{-2} \phi dx + \int_{-2}^{-1} \phi dx + \int_{-1}^{0} \phi dx\right] \\
&= \frac{1}{\Delta t}\left[\int_{-3.8}^{-3} \phi dx + M_3 + M_2 + M_1\right]
\end{aligned}
$$

- Subgrid approximation (quadratic, linear)
- Fluxes unique on each cell face
- Monotonicity limiters available

COSMIC

$$\frac{\partial \phi}{\partial t} + \frac{\partial (\phi u)}{\partial x} + \frac{\partial (\phi v)}{\partial y} = 0$$

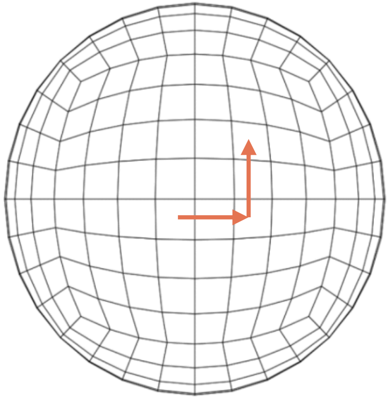

$$\phi^{n+1} = \tfrac{1}{2}\{Y(X(\phi^n)) + X(Y(\phi^n))\}$$

x step

$$\begin{aligned} \phi^* &= \phi^n - \Delta t \frac{\partial (\phi^n u)}{\partial x} \\ &= X(\phi^n) \end{aligned}$$

y step

$$\begin{aligned} \phi^{n+1} &= \phi^* - \Delta t \frac{\partial (\phi^* v)}{\partial y} \\ &= Y(\phi^*) \end{aligned}$$

Mass conservation

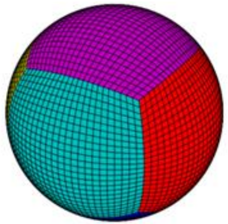

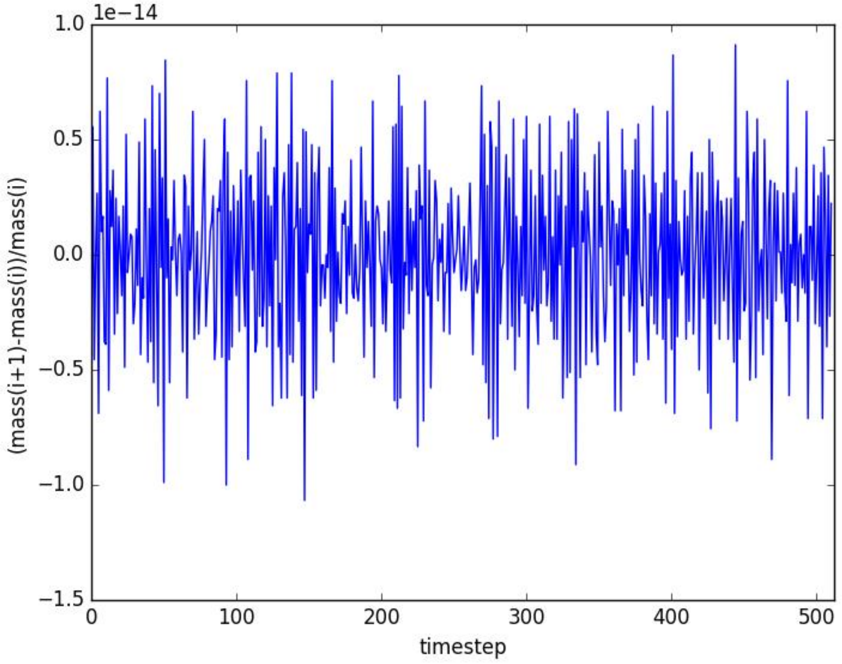

- Solid body rotation test

PSyClone Transformation Example

Taken from LFRic subroutine apply_helmholtz_lhs.
Multiple invoke's in an algorithm file
Multiple kernels in an invoke
Mixed builtin's and coded kernels
Slightly modified (replaced operator).

```
call invoke( setval_c(grad_p, 0.0_r_def),                               &
             scaled_matrix_vector_kernel_type(grad_p, p, div_star,      &
                                              hb_inv),                  &
             enforce_bc_kernel_type( grad_p ),                          &
             apply_variable_hx_kernel_type(                             &
                   Hp, grad_p, mt_lumped_inv, p,                        &
                   compound_div, p3theta, ptheta2, m3_exner_star,       &
                   tau_t, timestep_term) )
```

PSyClone Transformation Example

PSyclone’s internal representation.
A schedule.

```
Schedule[invoke='invoke_0' dm=True]
    Loop[type='dofs',field_space='any_space_1',it_space='dofs', upper_bound='ndofs']
        Call setval_c(grad_p,0.0_r_def)
    HaloExchange[field='grad_p', type='region', depth=1, check_dirty=False]
    HaloExchange[field='p', type='region', depth=1, check_dirty=True]
    HaloExchange[field='div_star', type='region', depth=1, check_dirty=True]
    HaloExchange[field='hb_inv', type='region', depth=1, check_dirty=True]
    Loop[type='',field_space='any_space_1',it_space='cells', upper_bound='cell_halo(1)']
        KernCall scaled_matrix_vector_code(grad_p,p,div_star,hb_inv) [module_inline=False]
    HaloExchange[field='grad_p', type='region', depth=1, check_dirty=False]
    Loop[type='',field_space='any_space_1',it_space='cells', upper_bound='cell_halo(1)']
        KernCall enforce_bc_code(grad_p) [module_inline=False]
    HaloExchange[field='mt_lumped_inv', type='region', depth=1, check_dirty=True]
    Loop[type='',field_space='w3',it_space='cells', upper_bound='ncells']
        KernCall apply_variable_hx_code(hp,grad_p,mt_lumped_inv,p,compound_div,p3theta,pthe
ta2,m3_exner_star,tau_t,timestep_term) [module_inline=False]
```