

Multigrid and Mixed-Precision: Fast Solvers for Weather and Climate Models.

Christopher Maynard

- Scientific Software Engineer
Met Office, Exeter, UK
- Associate Professor of Computer
Science.
University of Reading



Semi-implicit time-integration schemes, commonly used in Numerical Weather Prediction and Climate models, require a global matrix inversion of some kind. The linear solvers employed to do so must be fast and capable of running on highly parallel and complex supercomputers. Consequently there is a complex interplay between the algorithm and its implementation. In this presentation the use of mixed-precision arithmetic and a Geometric Multigrid Algorithm in the Met Office's Unified Model and LFRic Model are described and performance analysed.

It was the best of times, it was the worst of times ...

Apologies to Charles Dickens ...

Mixed-precision arithmetic in the ENDGame dynamical core of the Unified Model, a numerical weather prediction and climate model code
C.M. Maynard and D.N. Walters. Comp. Phys. Comm. V244 Nov 2019
69--75

Performance of multigrid solvers for the mixed finite element dynamical core, LFRic
C.M. Maynard, T. Melvin, E.H. Müller in Prep.

Numerical algorithms have a defined accuracy. How fast they converge to continuous differential equations

Computers use floating-point arithmetic
Variable accuracy *c.f.* to real numbers
Not associative
Accumulated round-off error
More precision → bigger data type

$\pi = 3.1400000001$ Precise but not accurate

$3 < \pi < 4$ Accurate but not precise (John Gustafson)

Most scientific applications, especially weather and climate use 64-bit arithmetic

Is this necessary? 32-bit faster (memory/cache CPU, GPU etc)

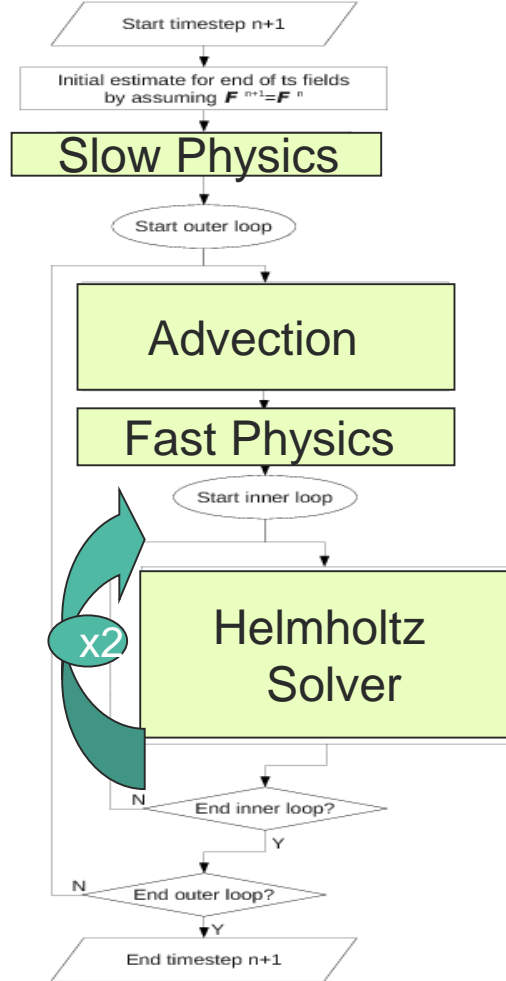



Lon-Lat grid \rightarrow polar singularity
Near poles grid points very close together
Explicit time-stepping scheme unfeasibly
short time-step for NWP
Semi-implicit schemes treat fast acoustic-
gravity modes implicitly
In combination with semi-Lagrangian
advection, SI allows stable integration
around pole
Long, but computationally expensive time-
steps
Global matrix inversion

UM timestep (ENDGame)

- 1x Slow physics
- 2x Advection
- 2x Fast Physics
- 4x Helmholtz solver
- 5x Dynamics residuals

SI ... expensive, but can take long timesteps



Equation takes the form

$$A \cdot \mathbf{x} = \mathbf{b}$$

Where A is a large, sparse matrix
 \mathbf{b} contains forcing terms

N1280 Lon-Lat mesh

~10Km resolution at mid-latitudes

$$(2 \times 1280) \frac{\times (3 \times 1280)}{2} \times 70 \approx 350M$$

For Semi-implicit time-stepping scheme, solver is part of a larger, non-linear system solution procedure

Accuracy of the solve is dictated **stability** of time-stepping scheme

FD $\sim \nabla p \sim 2^{\text{nd}}$ order \rightarrow limit to effect of accuracy of solve on pressure

Once solver error is sufficiently small, discretisation errors dominate

Post-conditioned BiCGStab

Inputs: \mathbf{x} , \mathbf{r} , ϵ_{tol} , δ_{min}

$\delta = \max(\|\mathbf{r}\|, \delta_{min});$

$\mathbf{p} = \mathbf{v} = \mathbf{0};$

$\mathbf{r} = \mathbf{r} - \mathbf{A}\mathbf{x}; \mathbf{r}_0 = \mathbf{r};$

$\alpha = \omega = n = 1;$

for $k = 1, 2, \dots$ **do**

$\rho = (\mathbf{r}, \mathbf{r}_0); \beta = \alpha \rho / (n\omega);$

$\mathbf{t} = \mathbf{r} - \beta \omega \mathbf{v}; \mathbf{s} = \mathbf{C}\mathbf{t};$

$\mathbf{p} = \mathbf{s} + \beta \mathbf{p}; \mathbf{v} = \mathbf{A}\mathbf{p};$

$n = (\mathbf{v}, \mathbf{r}_0); \alpha = \rho / n;$

if $\omega < 10^{-12}$ **then**

 | Convergence problem ω too small.

end

$\mathbf{s} = \mathbf{r} - \alpha \mathbf{v}; \tilde{\mathbf{s}} = \mathbf{C}\mathbf{s}; \mathbf{t} = \mathbf{A}\tilde{\mathbf{s}};$

$\omega = (\mathbf{t}, \mathbf{s}) / \|\mathbf{t}\|^2;$

$\mathbf{x} = \mathbf{x} + \alpha \mathbf{p} + \omega \tilde{\mathbf{s}}; \mathbf{r} = \mathbf{s} - \omega \mathbf{t};$

$n = \rho; \epsilon = \|\mathbf{r}\| / \delta;$

If $\epsilon < \epsilon_{tol}$ **exit;**

Halting criterion: norm of residual vector

$$\|\mathbf{r}\| = \|\mathbf{A} \cdot \mathbf{x}_i - \mathbf{b}\|$$

Stop when $\epsilon = \frac{\|\mathbf{r}\|}{\delta} < \epsilon_{tol}$

If $\epsilon_{tol} \gg \epsilon_{32}$

Where 32-bit Unit-of-least-precision (ULP) is $\sim 1.0 \times 10^{-7}$

Then 32-bit arithmetic is sufficient.

64-bit arithmetic won't improve *accuracy* of solution

$$\epsilon_{tol} = \{10^{-3}, 10^{-4}\}$$



Examine effect of precision
on convergence


$$\|\mathbf{r}\| = \|\mathbf{A} \cdot \mathbf{x}_i - \mathbf{b}\|$$

c.f. 32- 64- and 128-bit
arithmetic

32-bit takes more iterations
for residual fall

Iteration gap

Still converges




$$\mathcal{K}_p = \text{span}\{p, Ap, \dots\}$$

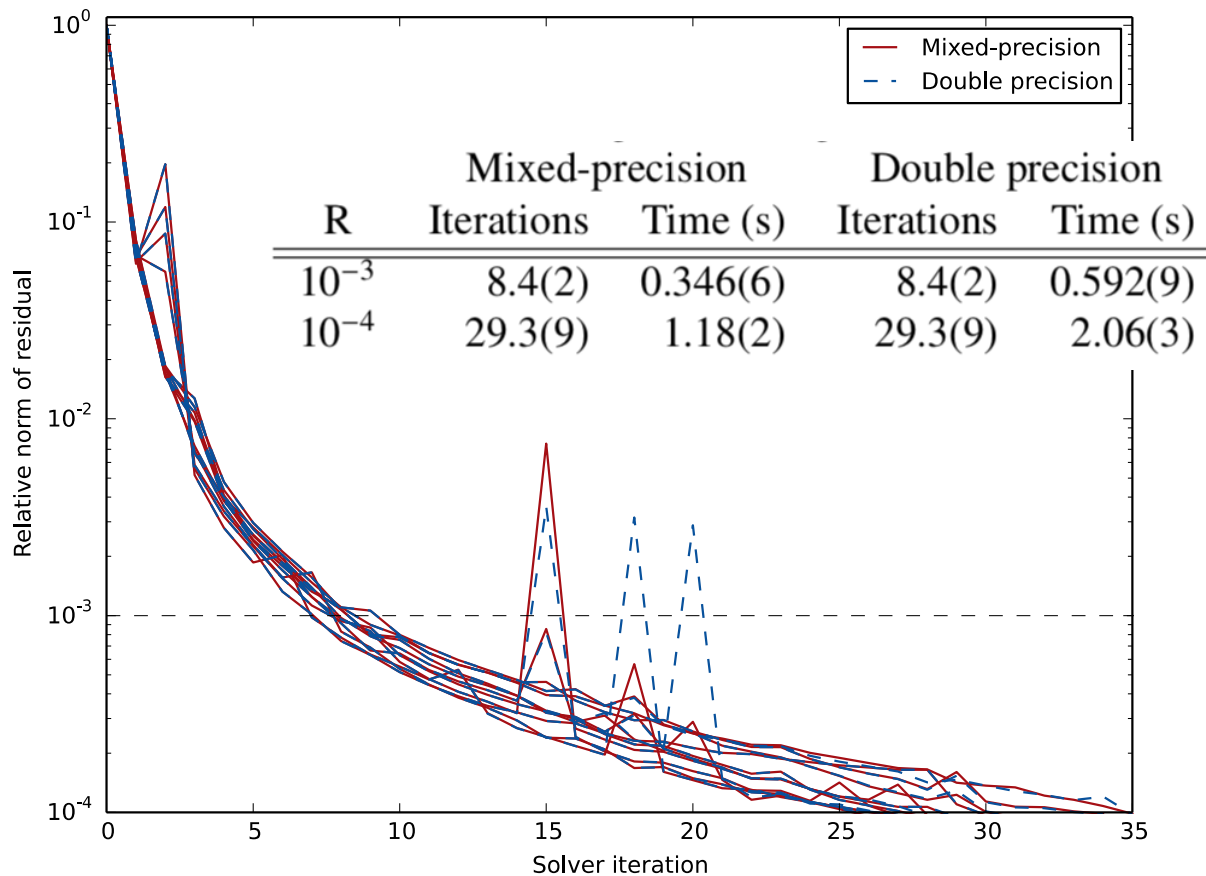
$$\mathcal{O}(\mathcal{K}_p) = p \cdot Ap$$

$$\Delta O_P = \frac{\text{abs}(O_P - O_{128})}{\text{abs}(O_P + O_{128})}$$


Zero → orthogonal
 One → no orthooinality
 32-bit (S) diverges earlier
 than 64-bit (D)



Solver implemented as *mixed* precision
 Pressure field was kept as 64-bit
 32-bit increments
 Ease of interfacing to model
 11 N1280 operational cfgs
 First time-step, first solve
 96 nodes Cray XC40
 12 MPI ranks/3 OMP threads



Tighten tol to 10^{-5}
Slow convergence –
hundreds of iters
BiCGstab does not
guarantee monotonic
convergence
Jumps in value of
residual
BiCGStab is breaking
down
Mixed-precision fares
worse – sometimes fails





Occasional problems at 10^{-4}

Slow convergence (hundreds of iters) – or even failures (divide by near zero)

Scalars → zero symptomatic of algorithm failing

In Mixed-precision global sums reverted to 64-bit arithmetic

Negligible cost (global sum is latency bound – sum is for single scalar)

Prevents failure, but slow convergence remains

In operations fixed iteration count limit imposed with full restart of solver

Ill conditioned problem arises from issues with “noise” in horizontal wind fields near poles

Original cfgs run with 10^{-3} tol, but problems in other parts of model

Tighter solver convergence helps but has its own problems

Solutions? i) Polar cap (transport across the poles)

ii) Multigrid (see later)

Efficiency (speed), accuracy and stability are all important considerations

Reduced precision can provide significant performance benefits (almost 2x for 32-bit versus 64-bit)

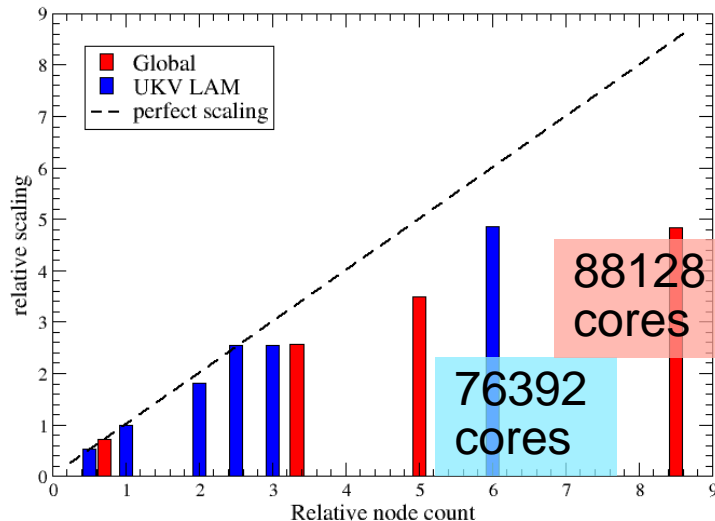
UM operations at Met Office runs in mixed-precision

Care is needed as complex interplay between round-off and other numerical errors

Especially where Numerical algorithms experience other problems

Uses Lon-Lat grid
Scientifically very good
Good computational performance

Very High Resolution scaling
6.5 Km resolution



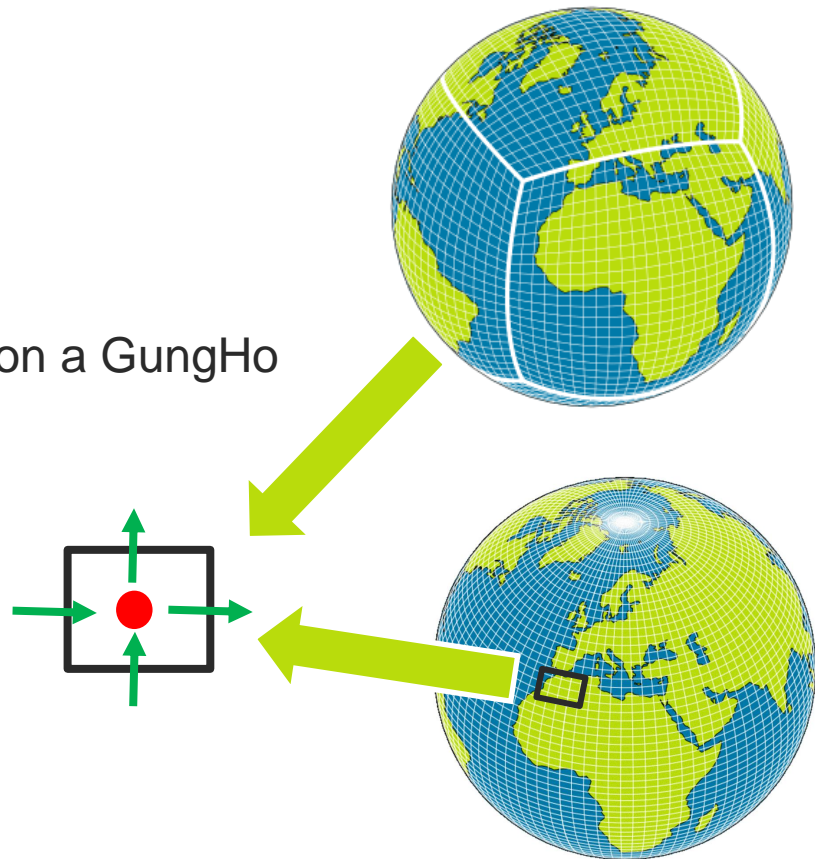
The finger of blame ...

Lon-lat grid is preventing scaling

10km resolution (mid-latitudes) → 10m at poles



GungHo Issues

- How to maintain **accuracy** of current model on a GungHo grid?
- Principal points about current grid are:
 - **Orthogonal, Quadrilateral, C-grid**
- Mixed Finite elements
 - Same layout as current C-grid
 - Not orthogonality condition
 - Gives consistent discretisation



Mixed Finite Element method gives


- Compatibility: $\nabla \times \nabla \varphi = 0, \nabla \cdot \nabla \times \mathbf{v} = 0$
- Accurate balance and adjustment properties
- No orthogonality constraints on the mesh
- Flexibility of choice mesh (quads, triangles) and accuracy (polynomial order)




Mixed Finite Element Method

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

\mathbb{W}_0	Pointwise scalars	
\mathbb{W}_1	Circulation Vectors	Vorticity
\mathbb{W}_2	Flux Vectors	Velocity
\mathbb{W}_3	Volume integrated Scalars	Pressure, Density
\mathbb{W}_θ	Pointwise scalars	Potential Temperature





Inspired by iterative-semi-implicit semi-Lagrangian scheme used in UM

Scalar transport uses high-order, upwind, explicit Eulerian FV scheme

Wave dynamics (and momentum transport) use iterative-semi-implicit, lowest order mixed finite element method (equivalent to C-grid/Charney-Phillips staggering)

$$\delta_t \mathbf{u} = -\overline{(2\Omega + \nabla \times \mathbf{u}) \times \mathbf{u} + \nabla (K + \Phi) + c_p \theta \nabla \Pi}^\alpha$$

$$\delta_t \rho = -\nabla \cdot \left[\mathcal{F} \left(\rho^n, \bar{\mathbf{u}}^{1/2} \right) \right]$$

$$\delta_t \theta = -\mathcal{A} \left(\theta^n, \bar{\mathbf{u}}^{1/2} \right)$$

$$\bar{F}^\alpha \equiv \alpha F^{n+1} + (1 - \alpha) F^n$$

Quasi-Newton Method: $\mathcal{L}(\mathbf{x}^*) \mathbf{x}' = -\mathcal{R}(\mathbf{x}^{(k)})$.

Linearized around reference state (previous time-step state) $\mathbf{x}^* \equiv \mathbf{x}^n$

Solve for increments on latest state: $\mathbf{x}' \equiv \mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}$

Semi-Implicit system contains terms needed for acoustic and buoyancy terms

$$\mathcal{L}(\mathbf{x}_{\text{phys}}^*) \mathbf{x}'_{\text{phys}} = \begin{cases} \mathbf{u}' - \mu \left(\frac{\mathbf{n}_b \cdot \mathbf{u}'}{\mathbf{n}_b \cdot \mathbf{z}_b} \right) \mathbf{z}_b \\ \quad + \tau_u \Delta t c_p (\theta' \nabla \Pi^* + \theta^* \nabla \Pi'), \\ \rho' + \tau_\rho \Delta t \nabla \cdot (\rho^* \mathbf{u}'), \\ \theta' + \tau_\theta \Delta t \mathbf{u}' \cdot \nabla \theta^*, \\ \frac{1 - \kappa}{\kappa} \frac{\Pi'}{\Pi^*} - \frac{\rho'}{\rho^*} - \frac{\theta'}{\theta^*}, \end{cases}$$

Solver Outer system with Iterative (GCR) solver

$$\begin{pmatrix}
 M_2^{\mu, C} & & -P_{2\theta}^{\Pi*} & -G^{\theta*} \\
 D^{\rho*} & M_3 & & \\
 P_{\theta 2}^{\theta*} & & M_{\theta} & \\
 & -M_3^{\rho*} & -P_{3\theta}^{\theta*} & M_3^{\Pi*}
 \end{pmatrix}
 \begin{pmatrix}
 \tilde{u}' \\
 \tilde{\rho}' \\
 \tilde{\theta}' \\
 \tilde{\Pi}'
 \end{pmatrix}
 =
 \begin{pmatrix}
 -\mathcal{R}_u \\
 -\mathcal{R}_{\rho} \\
 -\mathcal{R}_{\theta} \\
 -\mathcal{R}_{\Pi}
 \end{pmatrix}$$

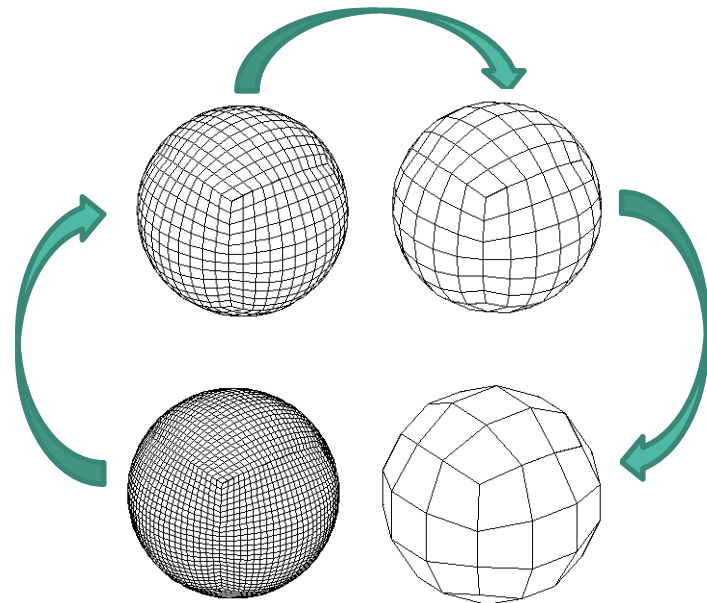
- Contains all couplings
- Preconditioned by approximate Schur complement for the pressure increment
- Velocity and potential temperature mass matrices are lumped

- Helmholtz system $H\Pi' = R$ solved using a single Geometric-Multi-Grid V-cycle with block-Jacobi smoother

$$H = M_3^{\Pi^*} + \left(P_{3\theta}^* \dot{M}_\theta^{-1} P_{\theta 2}^{\theta^*, z} + M_3^{\rho^*} M_3^{-1} D^{\rho^*} \right) \left(\dot{M}_2^{\mu, C} \right)^{-1} G^{\theta^*}.$$

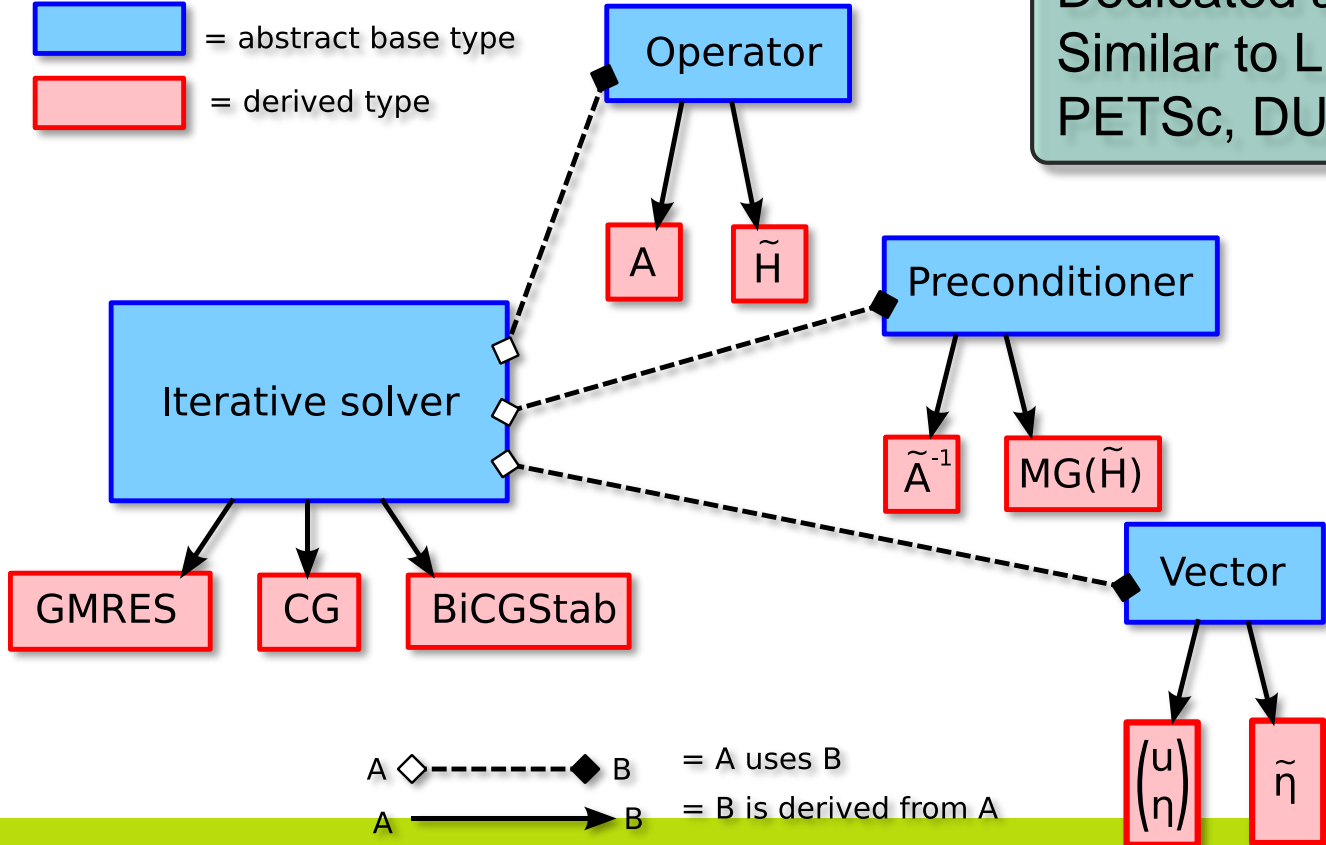
- Block-Jacobi smoother with small number (2) of iterations on each level
- Exact (tridiagonal) vertical solve: \hat{H}_z^{-1}

$$\tilde{\Pi}' \leftarrow \tilde{\Pi}' + \omega \hat{H}_z^{-1} \left(\mathcal{B} - H\tilde{\Pi}' \right)$$

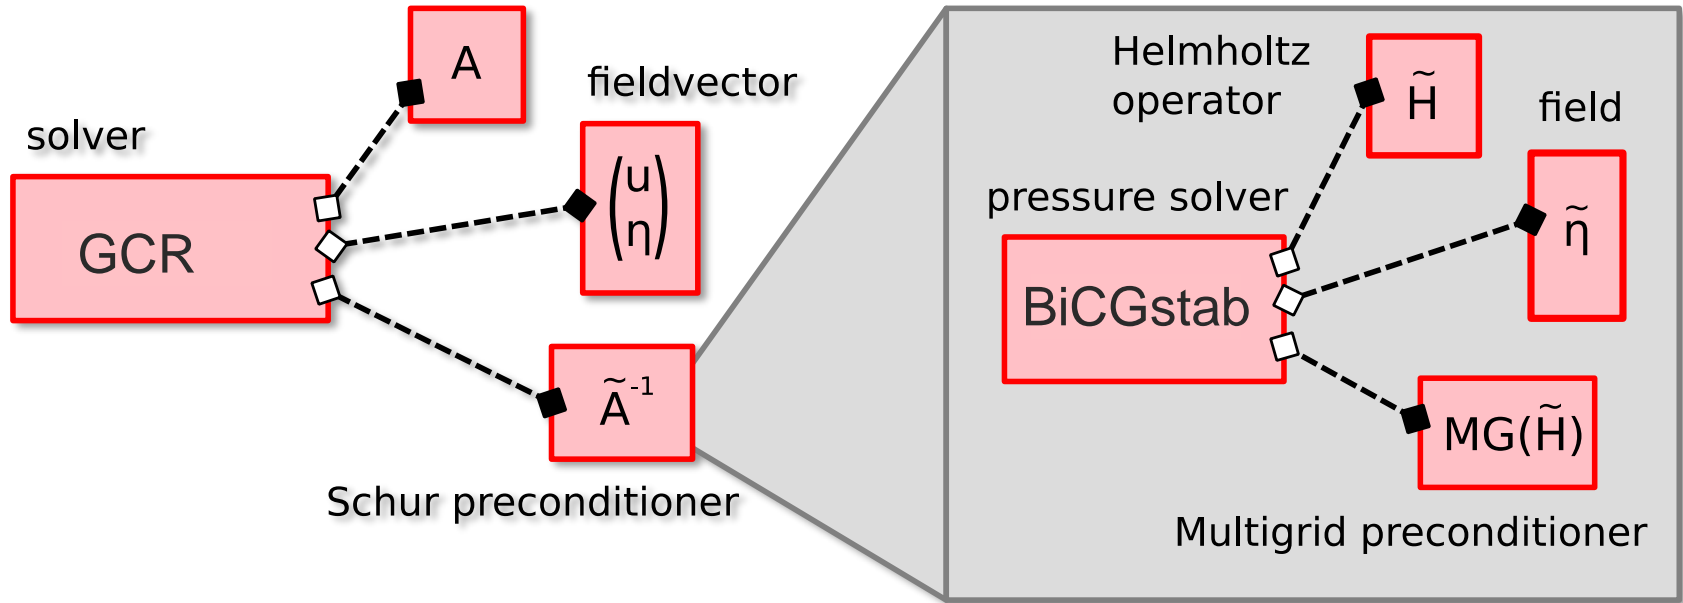


Dedicated abstraction in F2K3 OO
 Similar to Lin. Alg Libs e.g.
 PETSc, DUNE-ISTL, Trillinos

= abstract base type
 = derived type

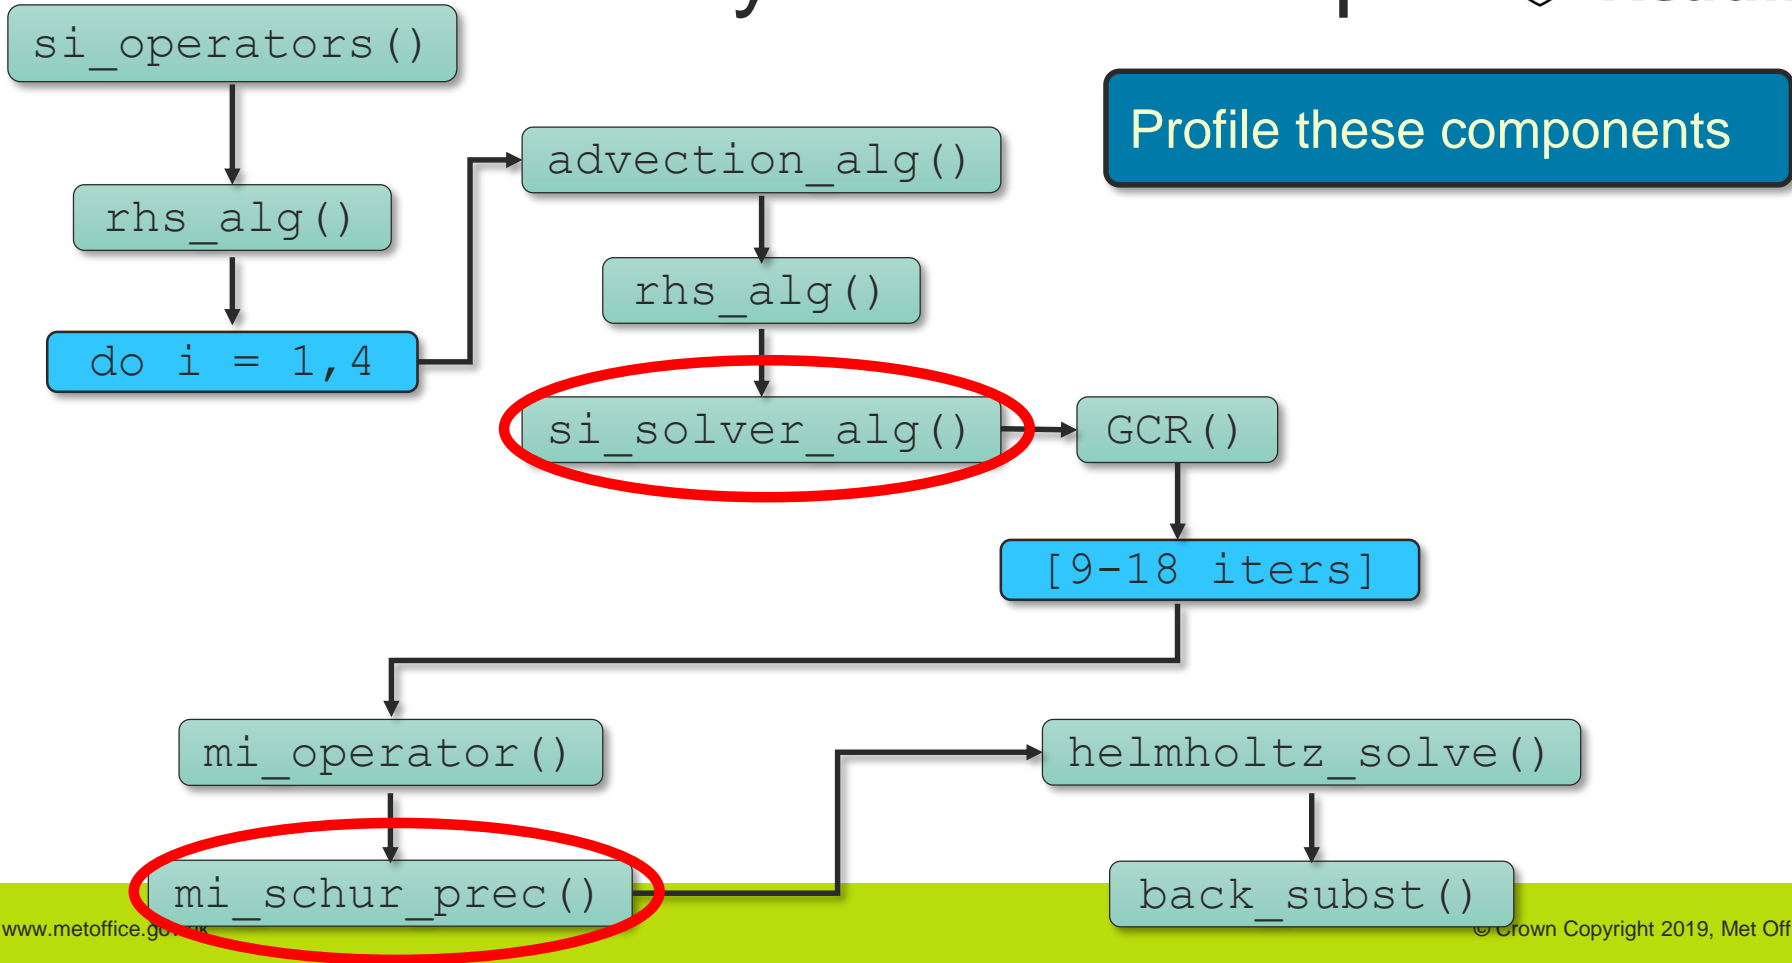


$A \diamond \text{---} \blacklozenge B$ = A uses B
 $A \text{---} \blacktriangleright B$ = B is derived from A



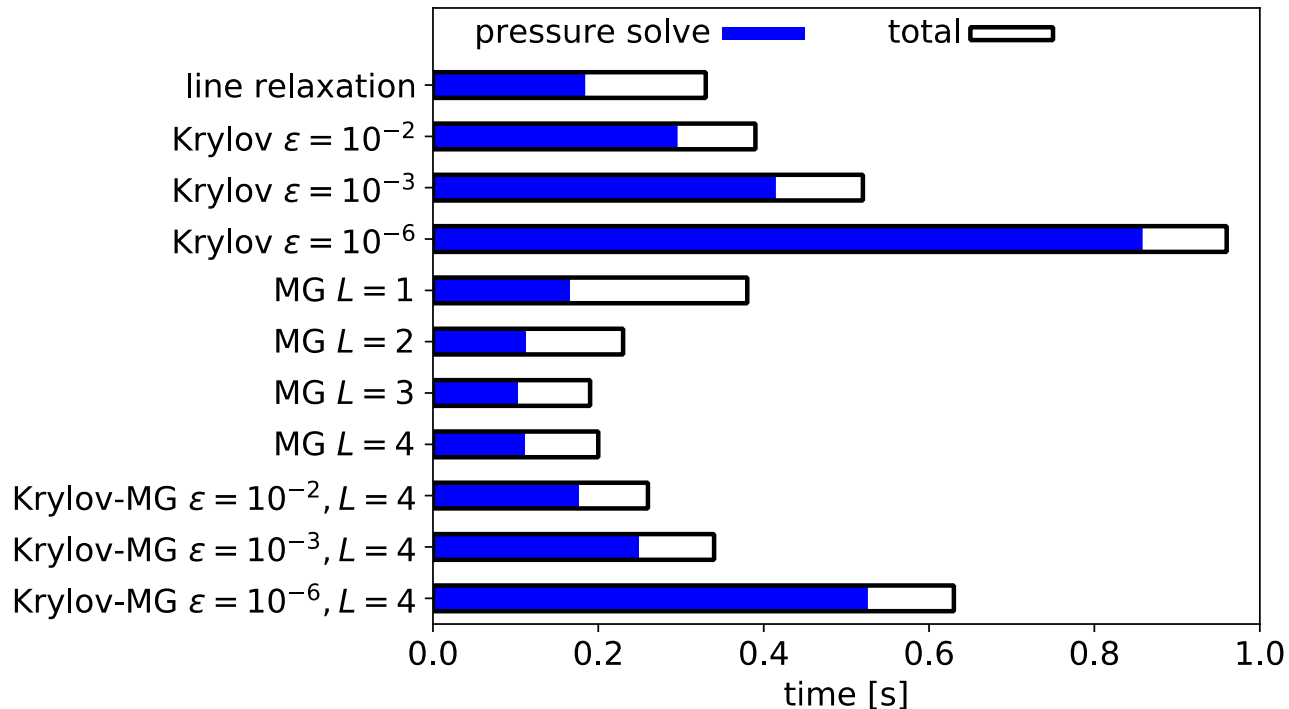
Allows for easy implementation of sophisticated nested solver
Multigrid preconditioner - reduce work for iterative solver
- faster and less global sums (better scaling)

Anatomy of a time-step



C192 cubed sphere
 with 30 L (~50Km)
 Baroclinic wave test
 Met Office Cray
 XC40 64 nodes
 (2304 cores) Mixed
 mode 6 MPI/6 OMP
 threads

c.f. $\|r\| = \|Ax - b\|$ Of
 Krylov 10^{-2}
 Before and after MG
 3-level V-cycle



SI \rightarrow long time-step as possible
Stability is limited by *vertical* stability.

C192 \sim 50Km, $\Delta t = 1200$

$$\text{CFL}_H = c_s \frac{\Delta t}{\Delta x} \quad c_s = 340 \text{m/s}$$

$\text{CFL}_H \sim 8$

C1152 \sim 9Km and $\Delta t = 205\text{s} \rightarrow \text{CFL} \sim 8$

Baroclinic wave test (Again 30L)

Kr 10^{-2} cf 3-level MG

Up to 1536 *nodes*

Lower is better

MG is at least 2x faster
than Kr

Both show good scaling


X-axis is logarithmic

96 : 1536 ~ 16x

55296 cores

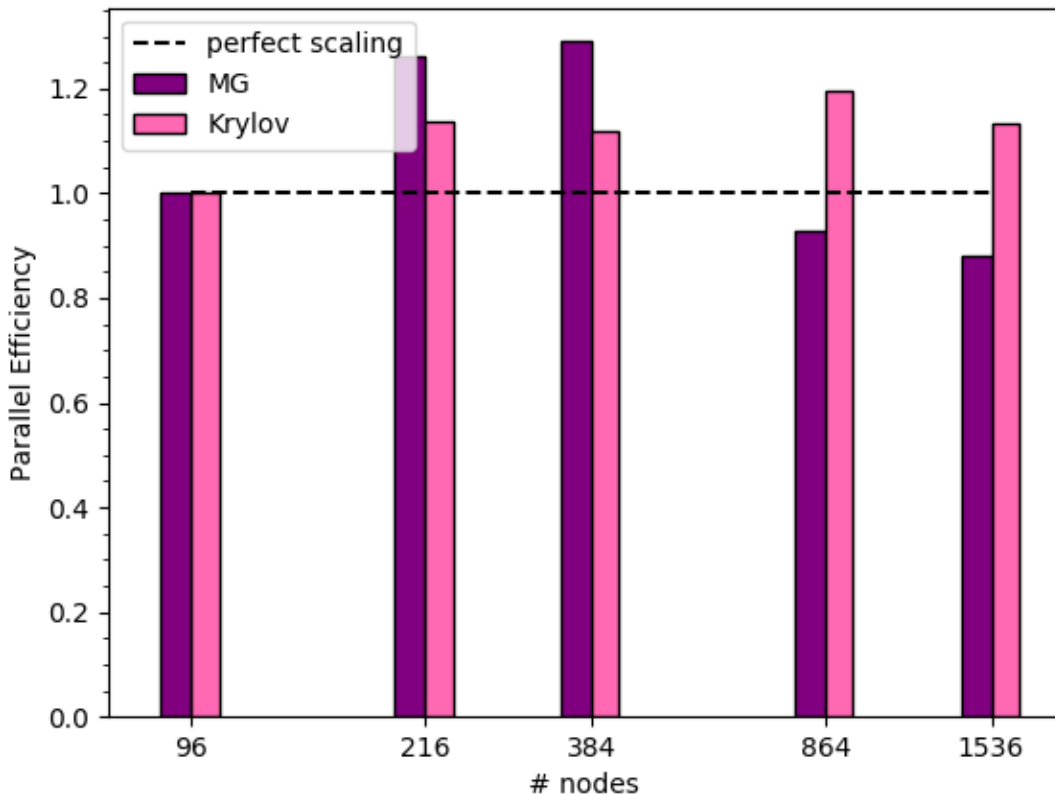
$LV = \{48, 32, 24, 16, 12\}^2$

Strong scaling on C1152 for 100 time-steps




Higher is better
Scaled from 96 nodes
Both show good scaling
KR is better because 96
node is slow!

Parallel Efficiency of Schur-precon



Lower is better
Data produced by
CrayPAT
96 node MG runs out of
memory
Less comms for MG
Large variation due to
Aries adaptive routing

Strong scaling of Halo-Exchange






Lower is better

Both algs have global sums in outer solve, plus limited diagnostic
Kr still has GS for inner solver

10^{-2} → only a few iterations.


Very large variation due to Aries adaptive routing

Strong scaling of Global-Sum




Lower is better
MG is much more
efficient
Much less work
Used Schur-precon
scaling to estimate MG
96 node cost
No comms, hence good
scaling

Strong scaling of Matrix-vector




Lower is better
X-axis is linear
Data are reversed
Shows cost of computation
Scales linearly with problem size
Smallest problem size not much work *c.f.* with comms
Fischer *et al*, suggests strong scaling limit is around LV~10000 (my interpretation)
doi 10.2514/6.2015-3049

Matrix-vector versus problem size





ENDGame WC



UM ENDGame N1280

Multigrid for Helmholtz solve.

Faster than Krylov (BiCGStab)

Big effect on solver

Polar singularity (Advection) spoils scaling

Lower precision helps both equally
MG can help avoid convergence issues

LFRic solver, currently 64-bit

Mixed-precision solver planned


Coarse grids can be done in lower precision, especially comms (8-bit)

Complex interplay between accuracy, efficiency, algorithm and implementation


64-bit arithmetic is expensive. Lower precision can, with care be used without compromising accuracy – depending on algorithm and implementation

Choice of algorithm, such as Multigrid to avoid global sums or Redundant computation to reduce communication are in some some being deployed to exploit architectural features - scaling

Accelerator architectures will require specific algorithmic choices



C96 2 day Aquaplanet
Surface moisture,
Mid-level cloud
SW heating



Gungho: Mixed finite element dynamical core



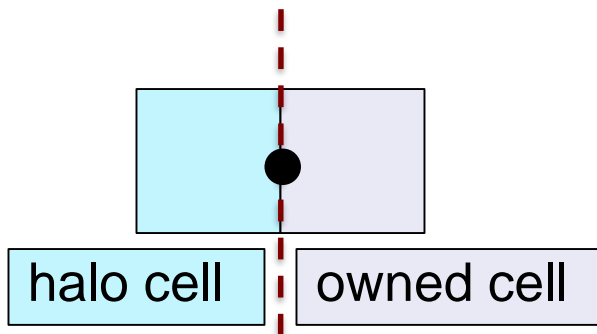
LFRic: Model infrastructure for next generation modelling




PSyclone: Parallel Systems code generation used in LFRic and Gungho

UM | Unified Model

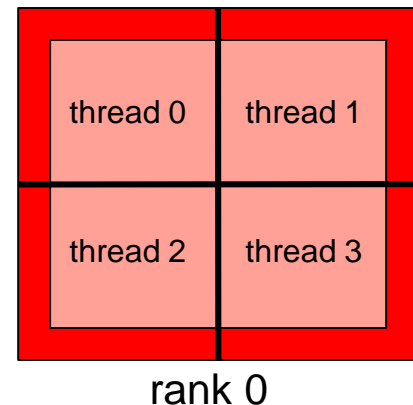
UM: Current modelling environment (UM parametrisations are being reused in LFRic

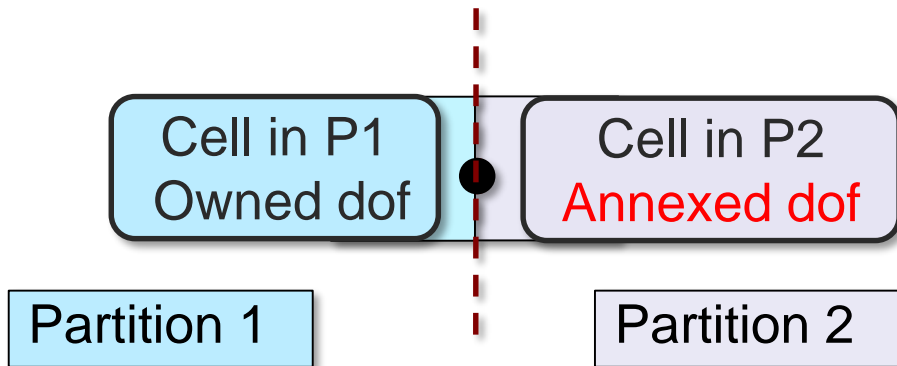


Dof living on shared (partitioned) entity (edge).
 Receive contribution from owned and halo cell.
 Redundant compute contribution in halo to shared dof.
 Less communication

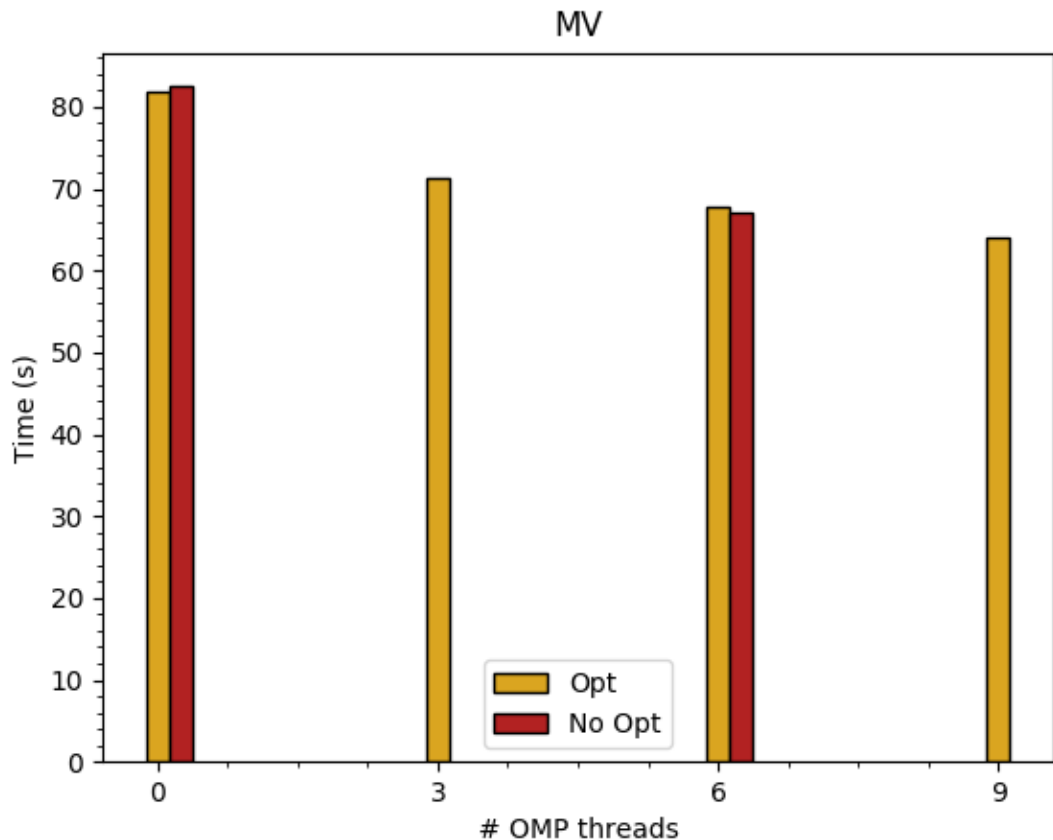


MPI only, 4 MPI ranks all have halos
 Hybrid, 1 MPI task has a halo, 4 OpenMP threads share halo
 boundary-to-area scaling
 → Less work for OpenMP threads

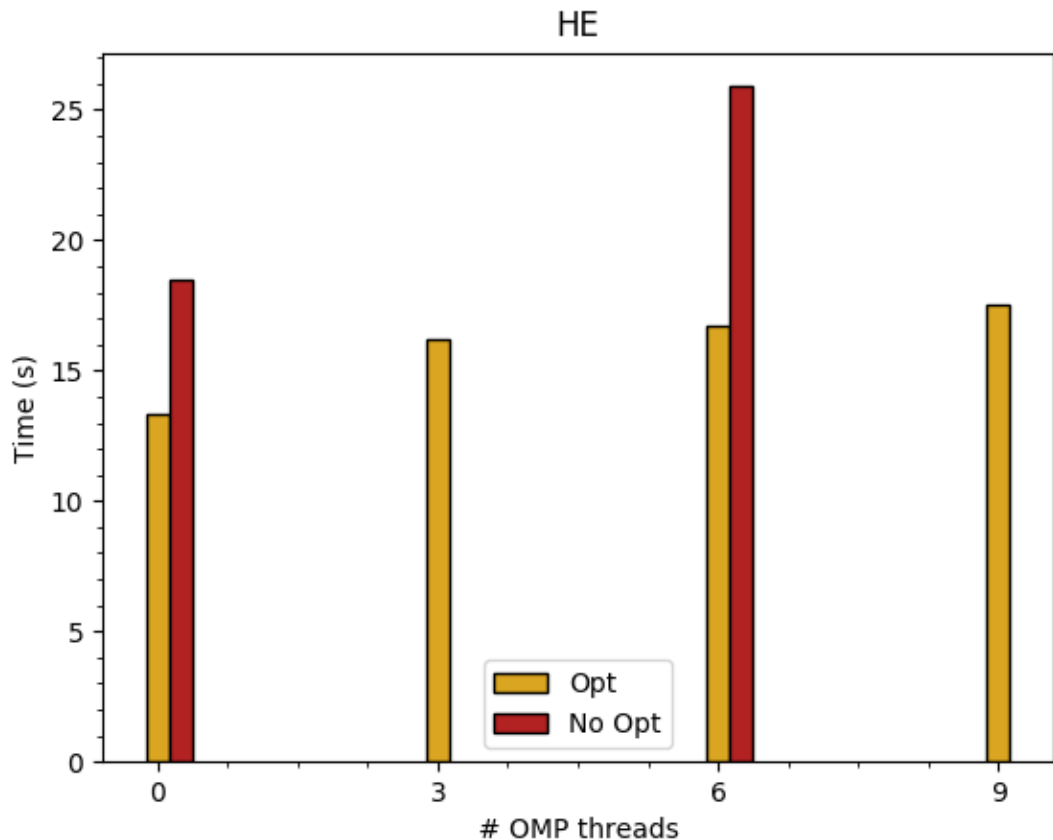




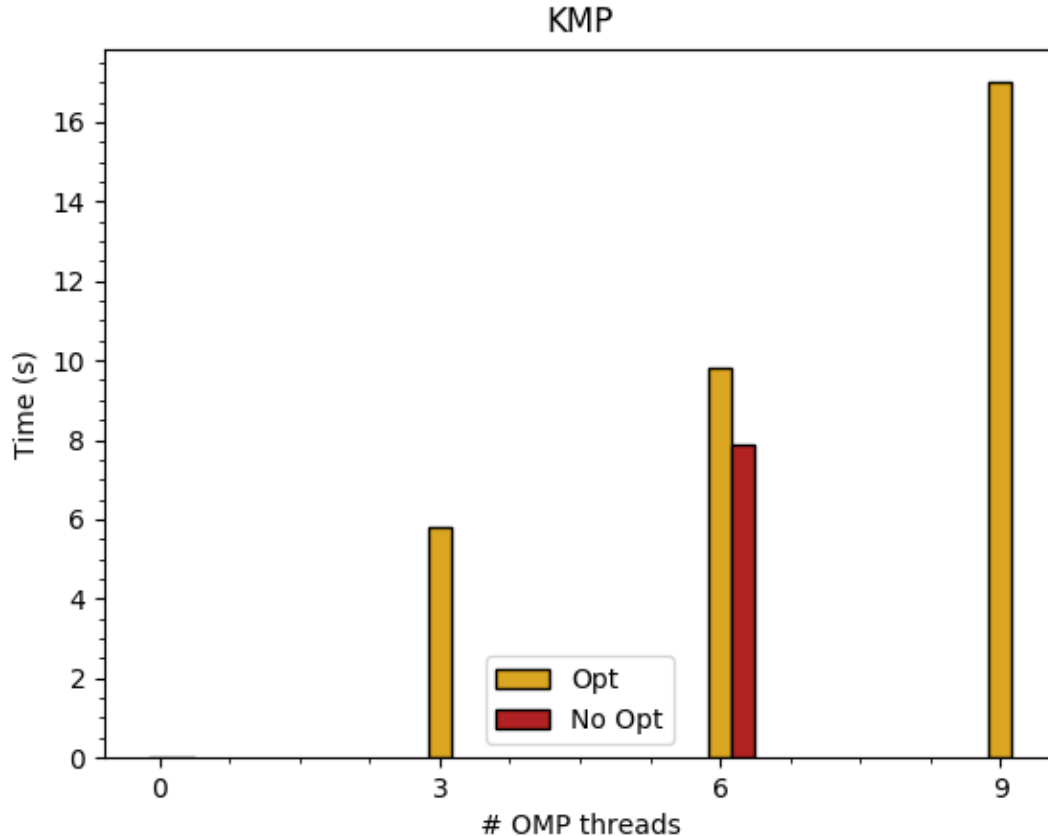
Point-wise computations (e.g. set field to a scalar) loop over dofs
Looping to owned dofs → halo exchange required for P2
Looping to annexed dofs is now transformation in Pycclone
Small increase in redundant computation
Large reduction in number of halo exchanges required



C288 running on 96 nodes
"0" thread is 36 MPI ranks
per node
3/12, 6/6 and 9/4 (Dual
socket
Profile by CrayPAT
Pure computation
OMP is faster as it has less
work



More threads \rightarrow fewer MPI ranks send/receive *bigger* messages
Import to tune
Rendezvous/Eager protocol limit (larger)



Intel 17 compiler
This seems very large to me
Can't compiler F2K3 OO
objects with Cray or PGI
Single kernel results suggest
Cray is better
`OMP_WAIT_POLICY=active`

Fortran – high level language
 Abstraction of the numerical mathematics
 Implementation and architecture is hidden
 Code – text which conforms to the semantics and syntax of the language definition
 Compiler transforms code into

```

real(kind=r_def), dimension(nqp_h), intent(in) :: wqp_h
real(kind=r_def), dimension(nqp_v), intent(in) :: wqp_v

!Internal variables
integer :: df, df2, k, ik
integer :: qp1, qp2

real(kind=r_def), dimension(nqp_h), intent(in) :: chi1_e, chi2_e, chi3_e
real(kind=r_def) :: integrand
real(kind=r_def), dimension(nqp_h), intent(in) :: dj
real(kind=r_def), dimension(nqp_h, nqp_v) :: jac

!loop over layers: Start from 1 as in this loop k is not an offset
do k = 1, nlayers
  ik = k + (cell-1)*nlayers

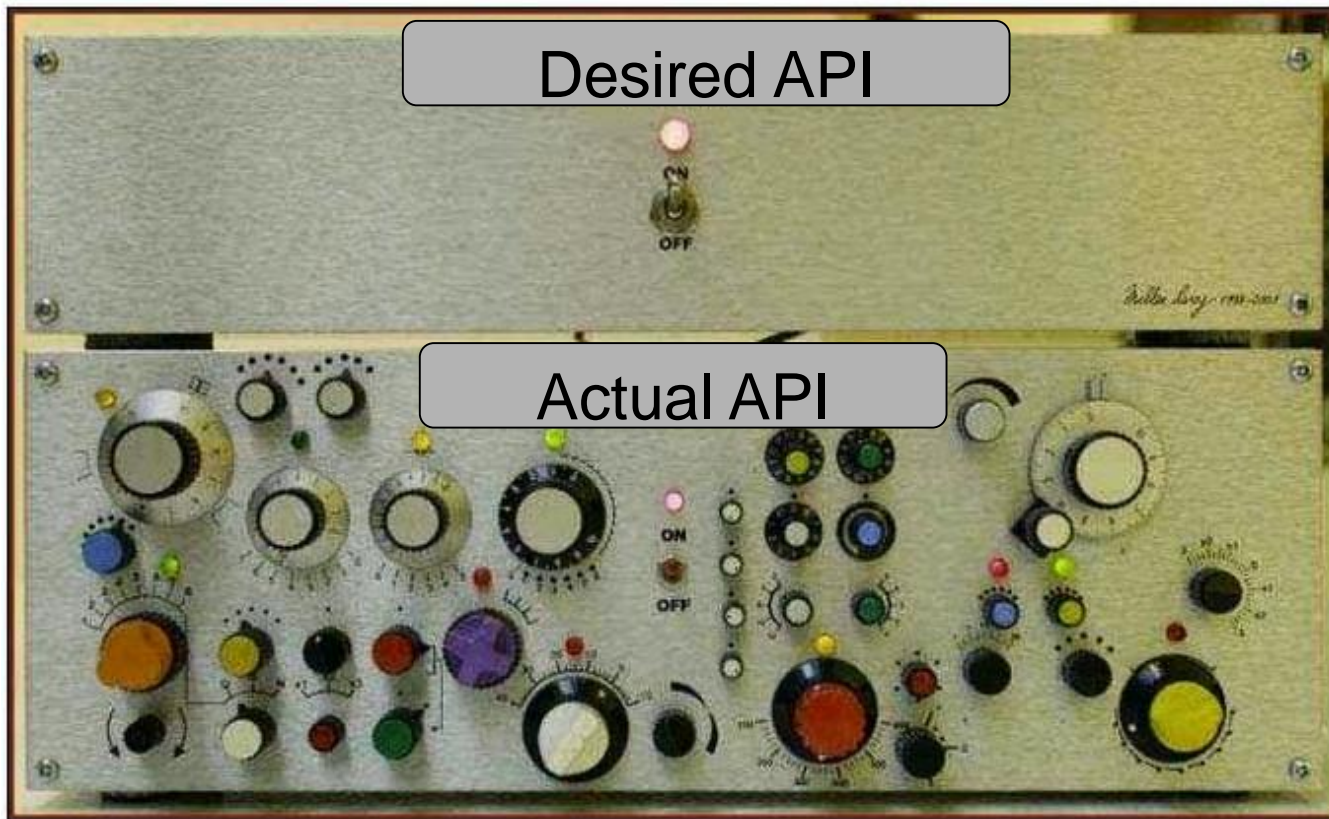
  ! indirect the chi coord field here
  do df = 1, ndf_chi
    chi1_e(df) = chi1(map_chi(df) + k - 1)
    chi2_e(df) = chi2(map_chi(df) + k - 1)
    chi3_e(df) = chi3(map_chi(df) + k - 1)
  end do

  call coordinate_jacobian(ndf_chi, nqp_h, nqp_v, chi1_e, chi2_e, chi3_e, &
    diff_basis_chi_jac_di)

```


Separation of concerns

Abstraction is *broken* by parallel/performance/memory features exposed
 Hacked back together with
 MPI, OMP, Open ACC, OpenCL, CUDA, PGAS, SIMD, compiler directives
 Libraries, languages (exts), directives and compiler (specific) directives



Scientific programming
Find numerical solution (and estimate of the uncertainty) to a (set of) mathematical equations which describe the action of a physical system

Parallel programming and optimisation are the methods by which large problems can be solved faster than real-time.



Alg layer – high level expression of operations on global fields

Kernel layer – low level Explicit operation on a single column of data

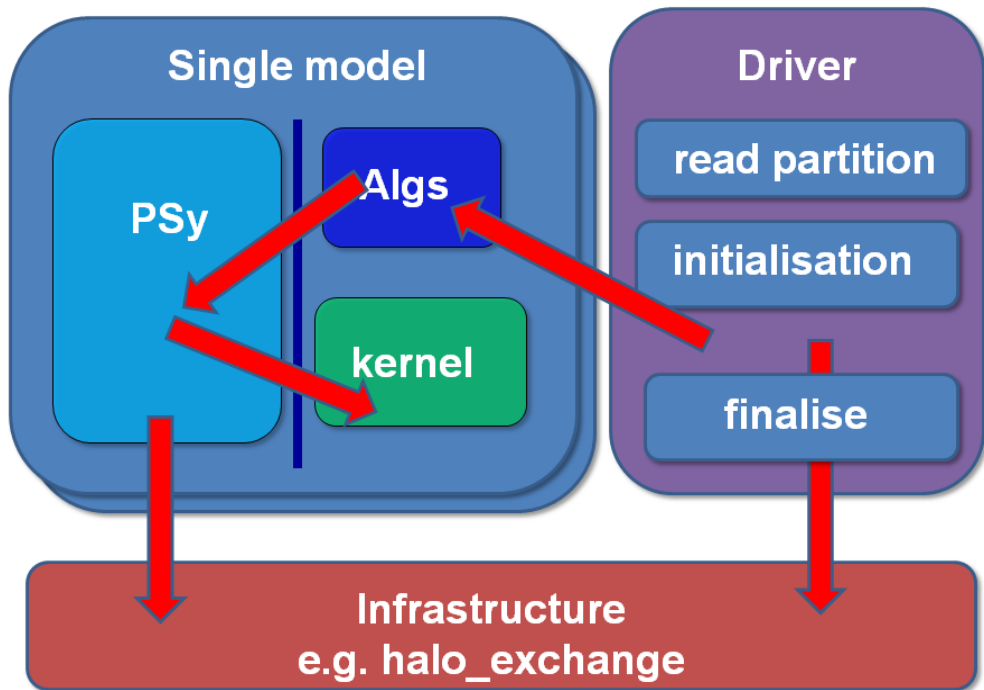
Code has to follow set of rules (PSyKAI API is DSL)

Parallelisation System

Horizontal looping and parallel code.

Can generate parallel code according to rules

- PSyKAI



```
call invoke(                                     &
    held_suarez_kernel_type(                   &
        rhs_heldsuarez(igh_u),                &
        rhs_heldsuarez(igh_t),                &
        state_n(igh_u),                       &
        state_n(igh_t),                       &
        state_n(igh_d),                       &
        chi, qr) ,                             &

    enforce_bc_kernel_type(                   &
        rhs_heldsuarez(igh_u) )               &
)
```

invoke() *Do this in parallel*
kernels single column operations
fields data parallel global fields

Multiple kernels in single invoke → scope of ordering/parallel communication, etc

Embed metadata
as (compilable)
Fortran, but it
doesn't get
executed
Data Access
descriptors
Explicitly
describe kernel
arguments
Richer
information than
Fortran itself

```
!> The type declaration for the kernel. Contains the metadata needed by the Psy layer
type, public, extends(kernel_type) :: exner_gradient_kernel_type
  private
  type(arg_type) :: meta_args(3) = (/                                &
    arg_type(GH_FIELD,    GH_INC,  W2),                            &
    arg_type(GH_FIELD,    GH_READ, W3),                            &
    arg_type(GH_FIELD,    GH_READ, ANY_SPACE_9)                    &
  /)
  type(func_type) :: meta_funcs(3) = (/                              &
    func_type(W2, GH_BASIS, GH_DIFF_BASIS),                        &
    func_type(W3, GH_BASIS),                                       &
    func_type(ANY_SPACE_9, GH_BASIS, GH_DIFF_BASIS)                &
  /)
  integer :: iterates_over = CELLS
  integer :: gh_shape = GH_QUADRATURE_XYoZ
  ! gh_shape replaces evaluator_shape
  integer :: evaluator_shape = QUADRATURE_XYoZ
contains
  procedure, nopass :: exner_gradient_code
end type
```


Python code generator
Parser, transformations, generation
Controls parallel code (MPI/OpenMP and OpenACC)
Potentially other programming models
e.g. OpenCL for FPGA

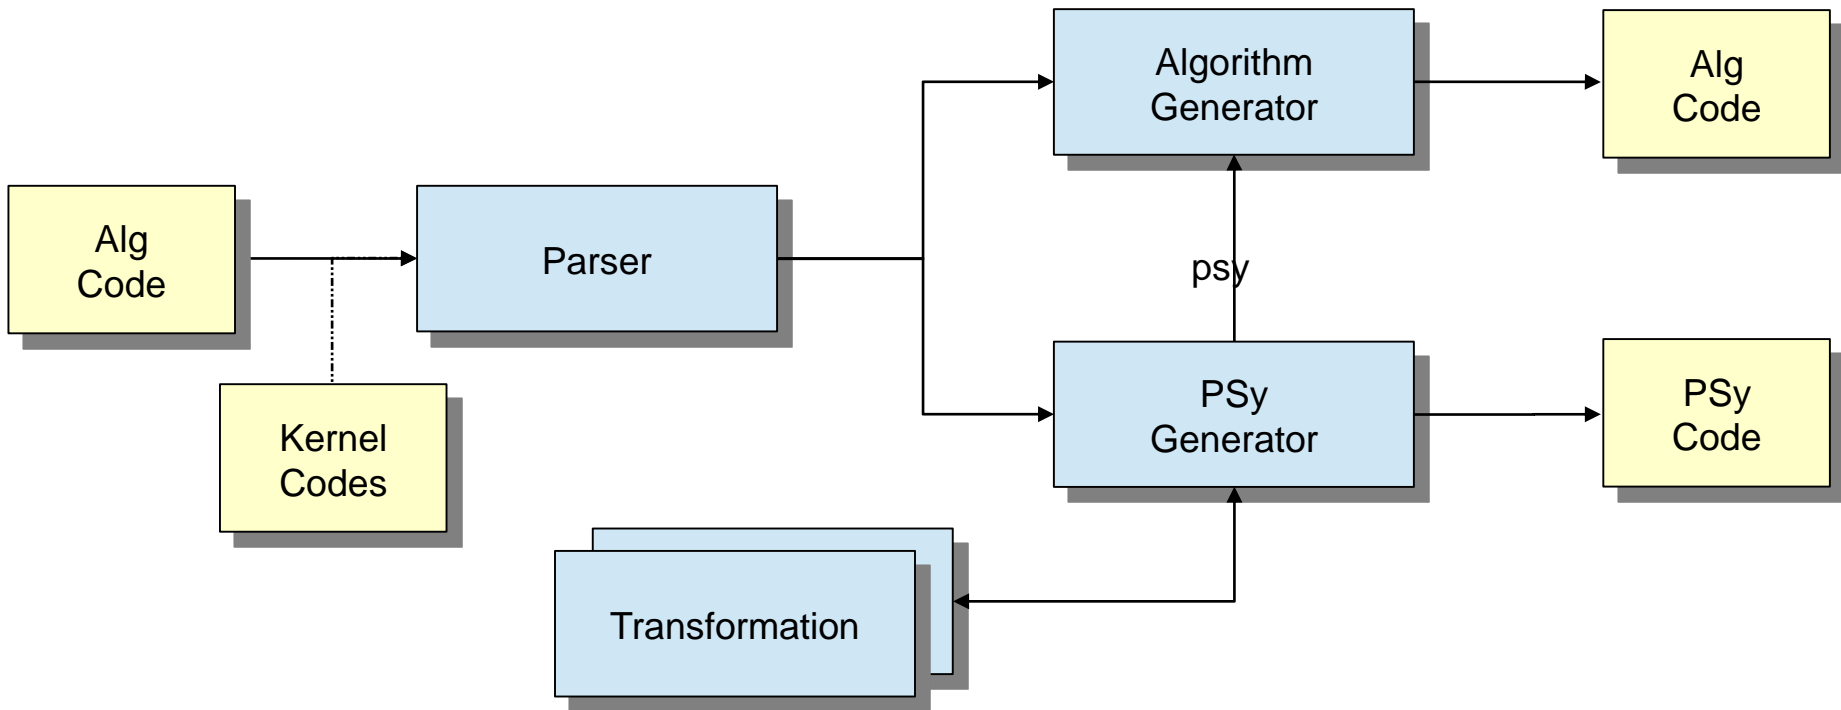
What is PSyclone (Brikipedia)?



Developed at STFC Hartree
R. Ford, A. Porter, S. Siso
J. Henrichs, BoM
I Kavcic, M Hambley, CMM (MO)
Works with PSyKAI API

"a former weathercaster turned crazy bad-guy with a craving for destruction"

<http://lego.wikia.com/wiki/PSyclone>



Update halos
YAXT→MPI

colouring from
infrastructure

OpenMP
workshare across
cells in colour

```
!
IF (chi_proxy(3)%is_dirty(depth=1)) THEN
  CALL chi_proxy(3)%halo_exchange(depth=1)
END IF
!
CALL rhs_heldsuarez_proxy%vspace%get_colours(ncolour, ncp_colour, cmap)
!
DO colour=1,ncolour
  !$omp parallel default(shared), private(cell)
  !$omp do schedule(static)
  DO cell=1,ncp_colour(colour)
    !
    call held_suarez_code(nlayers, rhs_heldsuarez_proxy%data, &
      ! ...
    )
  END DO
  !$omp end do
  !$omp end parallel
END DO
!
! Set halos dirty for fields modified in the above loop
!
CALL rhs_heldsuarez_proxy%set_dirty()
```

kernel call for single
column. Args are
arrays and scalars



Single kernel invoke

```
Transforming invoke 'invoke_26_rtheta_kernel_type' ...
Schedule[invoke='invoke_26_rtheta_kernel_type' dm=False]
  Loop[type='',field_space='w0',it_space='cells', upper_bound='ncells']
    KernCall rtheta_code(rtheta,theta,wind) [module_inline=False]
```

Apply distributed memory

```
Transforming invoke 'invoke_26_rtheta_kernel_type' ...
Schedule[invoke='invoke_26_rtheta_kernel_type' dm=True]
  HaloExchange[field='rtheta', type='region', depth=1, check_dirty=True]
  HaloExchange[field='theta', type='region', depth=1, check_dirty=True]
  HaloExchange[field='wind', type='region', depth=1, check_dirty=True]
  Loop[type='',field_space='w0',it_space='cells', upper_bound='cell_halo(1)']
    KernCall rtheta_code(rtheta,theta,wind) [module_inline=False]
```



Simple python script to
apply Open MP
transformation
Can apply on whole
model
Or as fine-grained as
single file

```
from psyclone.transformations import Dynamo0p3ColourTrans, \
    Dynamo0p3OMPLoopTrans, \
    OMPParallelTrans

def trans(psy):
    ctrans = Dynamo0p3ColourTrans()
    otrans = Dynamo0p3OMPLoopTrans()
    oregtrans = OMPParallelTrans()

    # Loop over all of the Invokes in the PSy object
    for invoke in psy.invokes.invoke_list:

        print "Transforming invoke '{0}' ...".format(invoke.name)
        schedule = invoke.schedule

        # Colour loops unless they are on W3 or over dofs
        for loop in schedule.loops():
            if loop.iteration_space == "cells" and loop.field_space != "w3":
                schedule, _ = ctrans.apply(loop)

        # Add OpenMP to loops unless they are over colours
        for loop in schedule.loops():
            if loop.loop_type != "colours":
                schedule, _ = oregtrans.apply(loop)
                schedule, _ = otrans.apply(loop, reprod=True)

        # take a look at what we've done
        schedule.view()

    return psy
```



```
Transforming invoke 'invoke_26_rtheta_kernel_type' ...|
Schedule[invoke='invoke_26_rtheta_kernel_type' dm=True]
  HaloExchange[field='rtheta', type='region', depth=1, check_dirty=True]
  HaloExchange[field='theta', type='region', depth=1, check_dirty=True]
  HaloExchange[field='wind', type='region', depth=1, check_dirty=True]
  Loop[type='colours', field_space='w0', it_space='cells', upper_bound='ncolours']
    Directive[OMP parallel]
      Directive[OMP do]
        Loop[type='colour', field_space='w0', it_space='cells', upper_bound='ncolour']
          KernCall rtheta_code(rtheta,theta,wind) [module_inline=False]
```

Update halos
YAXT→MPI

colouring from
infrastructure


OpenMP
workshare across
cells in colour

```
!
IF (chi_proxy(3)%is_dirty(depth=1)) THEN
  CALL chi_proxy(3)%halo_exchange(depth=1)
END IF
!
CALL rhs_heldsuarez_proxy%vspace%get_colours(ncolour, ncp_colour, cmap)
!
DO colour=1,ncolour
  !$omp parallel default(shared), private(cell)
  !$omp do schedule(static)
  DO cell=1,ncp_colour(colour)
    !
    call held_suarez_code(nlayers, rhs_heldsuarez_proxy%data, &
      ! ...
    )
  END DO
  !$omp end do
  !$omp end parallel
END DO
!
! Set halos dirty for fields modified in the above loop
!
CALL rhs_heldsuarez_proxy%set_dirty()
```

kernel call for single
column. Args are
arrays and scalars


W_0 space (vertices)

Mesh



Data array (1-d)

- data (map (1, 4) + 0) ←
- data (map (1, 3) + 1) ←
- data (map (2, 2) + 0) ←
- data (map (2, 1) + 1) ←



Dofmap 2-d array

1	2	6	7	11	12	16	17
6	7	21	22	26	27	11	12
...	...	ndof per cell	

n cell

PSy layer
Kernel layer

Visit same dof more than once: loop over cells, levels, dofs
 Mesh and dofmap form an ordered set
 Change mesh topology (element), geometry (cubed sphere)
 → Change to mesh generation and partition
 → No change to science code