Met Office

University of Reading

Scalable Linear Solvers for Next Generation Weather and Climate

Christopher Maynard

- Met Office, Exeter, UK and University of Reading

Thomas Melvin

- Met Office, Exeter, UK

Eike Müller

- University of Bath

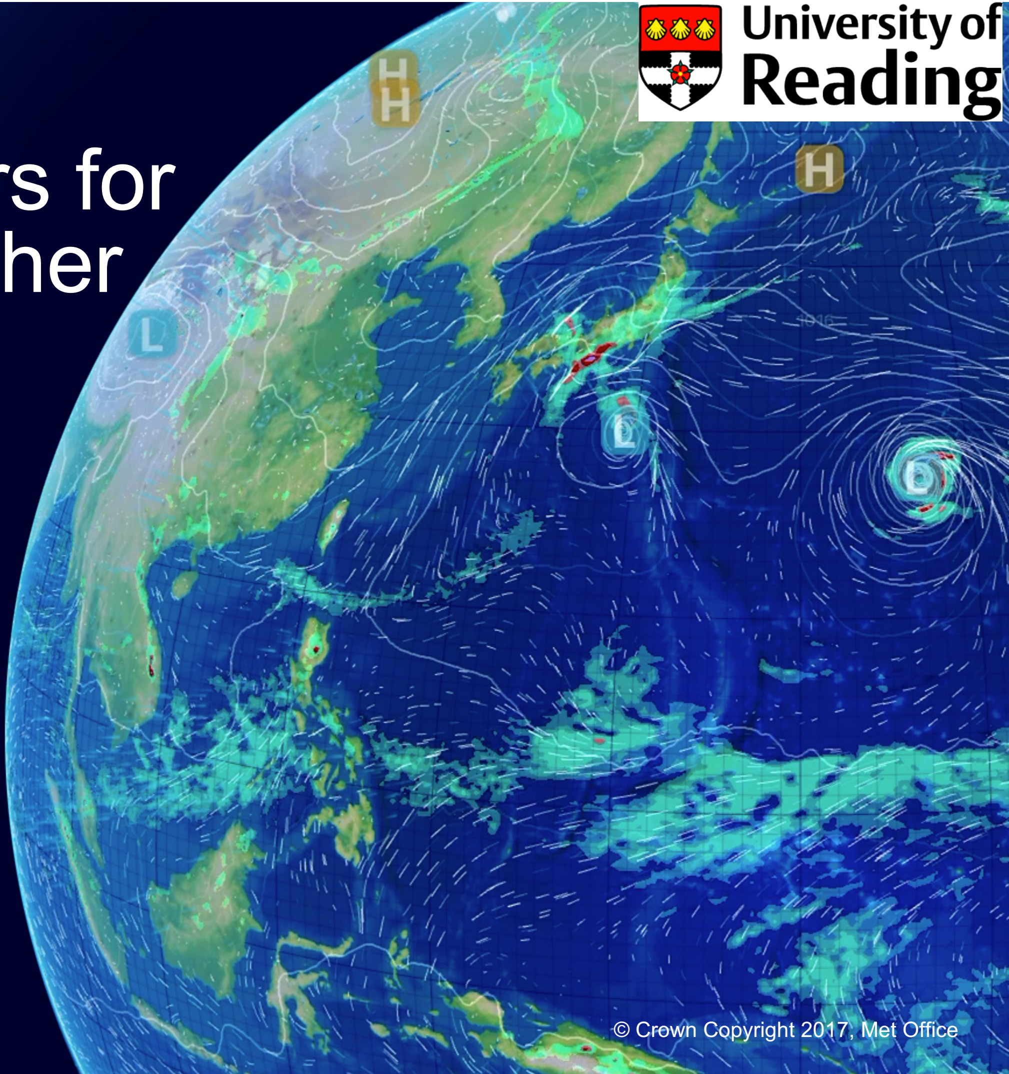

© Crown Copyright 2017, Met Office

Contents

Gungho, LFRic and PSyclone
Time-stepping and the solver
Performance Analysis

Acknowledments
Nerc: Gungho partner Universities
STFC: PSyclone development
Met Office: LFRic development team
Dynamics Research
Subgrid Physics developers

1. *PSyclone and its uses in LFRic* I. Kavcic **MS02** Wed 1300-1330
2. *On using a DSL Approach Performance Portability of the LFRic Weather and Climate Model* CMM **MS08** Wed 1630-1700
3. *Building a Performance Portable Software System for the Met Office's Weather and Climate Model, LFRic* **CSM07** (Poster session)

© Crown Copyright 2019, Met Office

Some names

工合 Gungho: Mixed finite element dynamical core

LFRic: Model infrastructure for next generation modelling

PSyClone: Parallel Systems code generation used in LFRic and Gungho

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

© Crown Copyright 2019, Met Office

The Unified Model

Uses Lon-Lat grid
Scientifically very good
Good computational performance

Very High Resolution scaling
6.5 Km resolution

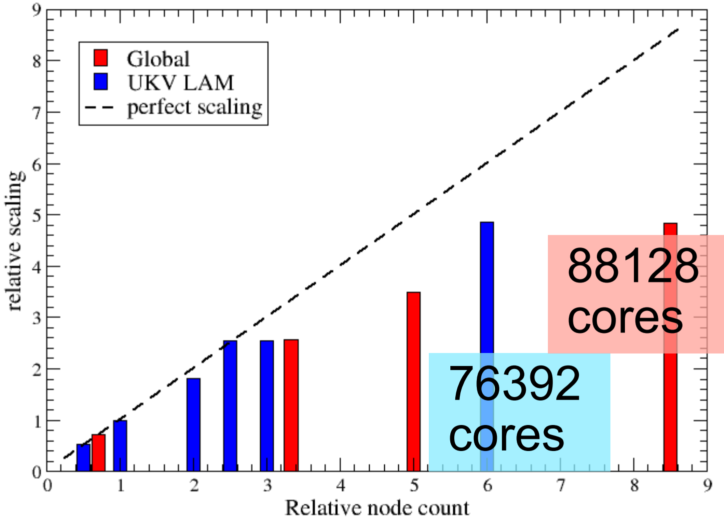

The finger of blame …
Lon-lat grid is preventing scaling
10km resolution (mid-latitudes) → 10m at poles

© Crown Copyright 2019, Met Office

What we are trying to do

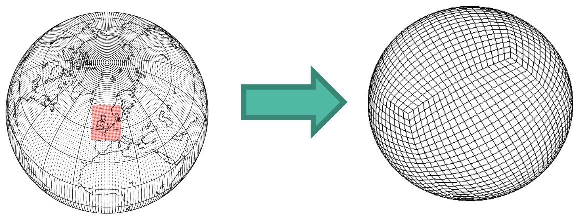

Move to quasi-uniform mesh to remove polar singularity

Maintain ‘good’ aspects of current model

- No computational modes
- Accurate dispersion
- Semi-Implicit timestepping
- Reuse subgrid parametrizations

Improve inherent conservation

Improve scalability

© Crown Copyright 2019, Met Office

Mixed Finite Elements

Mixed Finite Element method gives

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

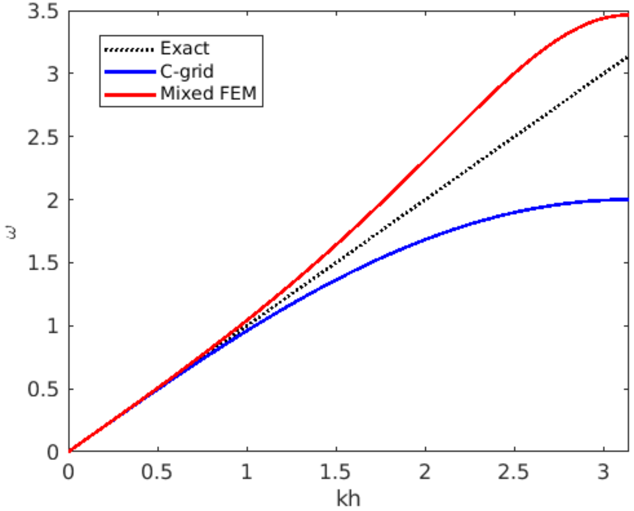

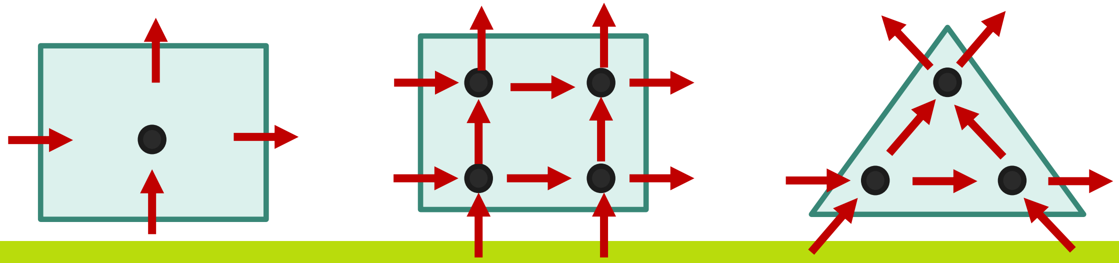

© Crown Copyright 2019, Met Office

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
$\mathbb{W}_\theta$	Pointwise scalars	Potential Temperature

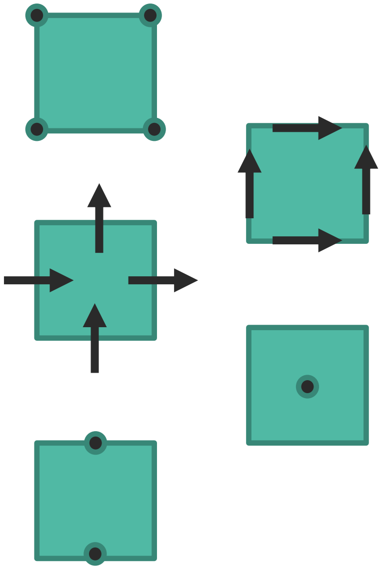

© Crown Copyright 2019, Met Office

Gungho Discretisation

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)

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

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

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

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

© Crown Copyright 2019, Met Office

Time-stepping

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

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

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

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

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

© Crown Copyright 2019, Met Office

Time-stepping II

Solver Outer system with Iterative (GCR) solver

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

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

© Crown Copyright 2019, Met Office

Multigrid

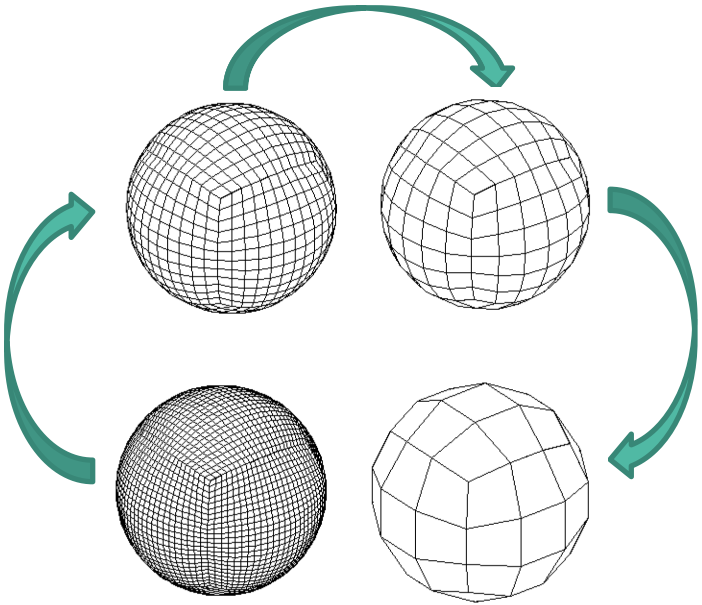

- Helmholtz system $H\Pi' = R$ solved using a single Geometric-Multi-Grid V-cycle with block-Jacobi smoother

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

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

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

© Crown Copyright 2019, Met Office

The solver

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

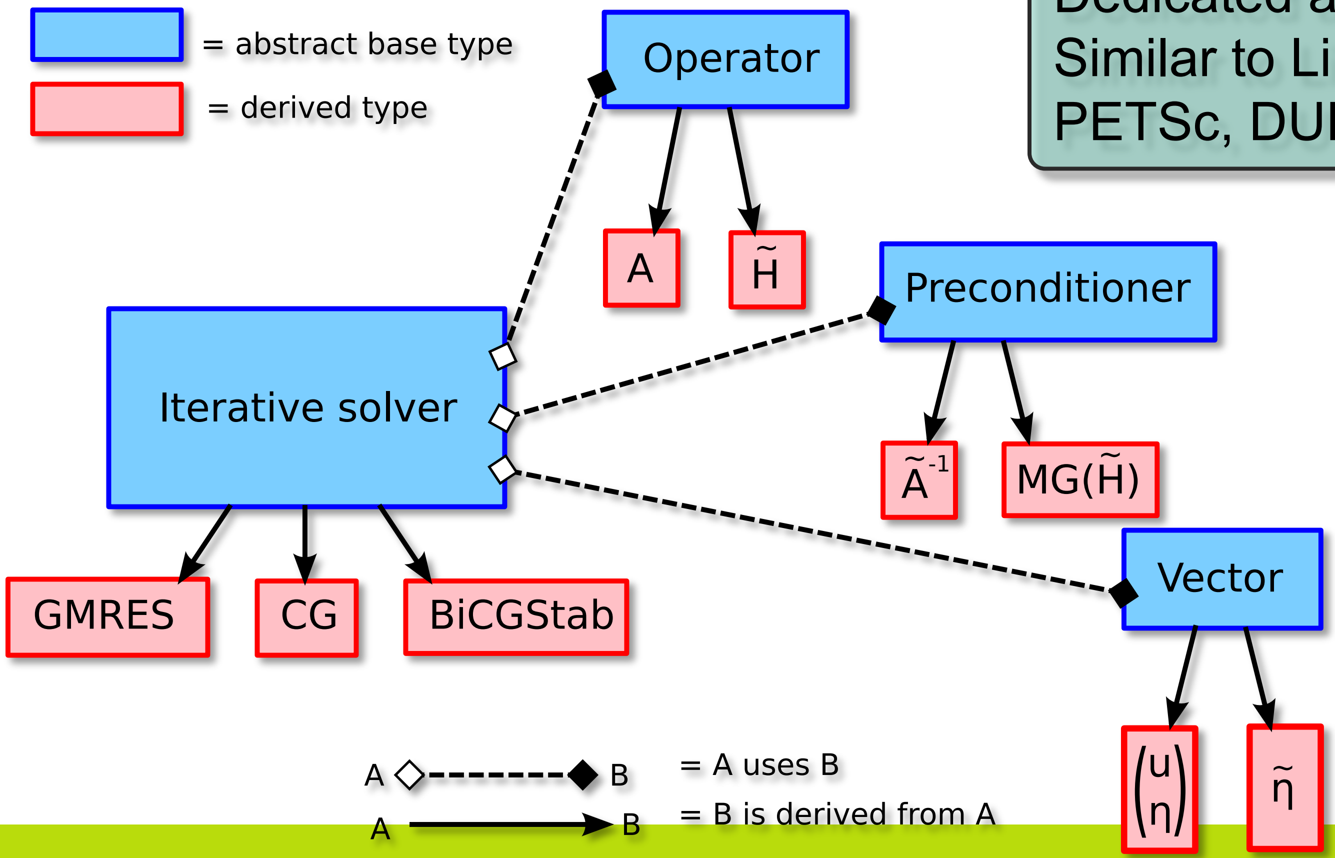

© Crown Copyright 2017, Met Office

The solver II

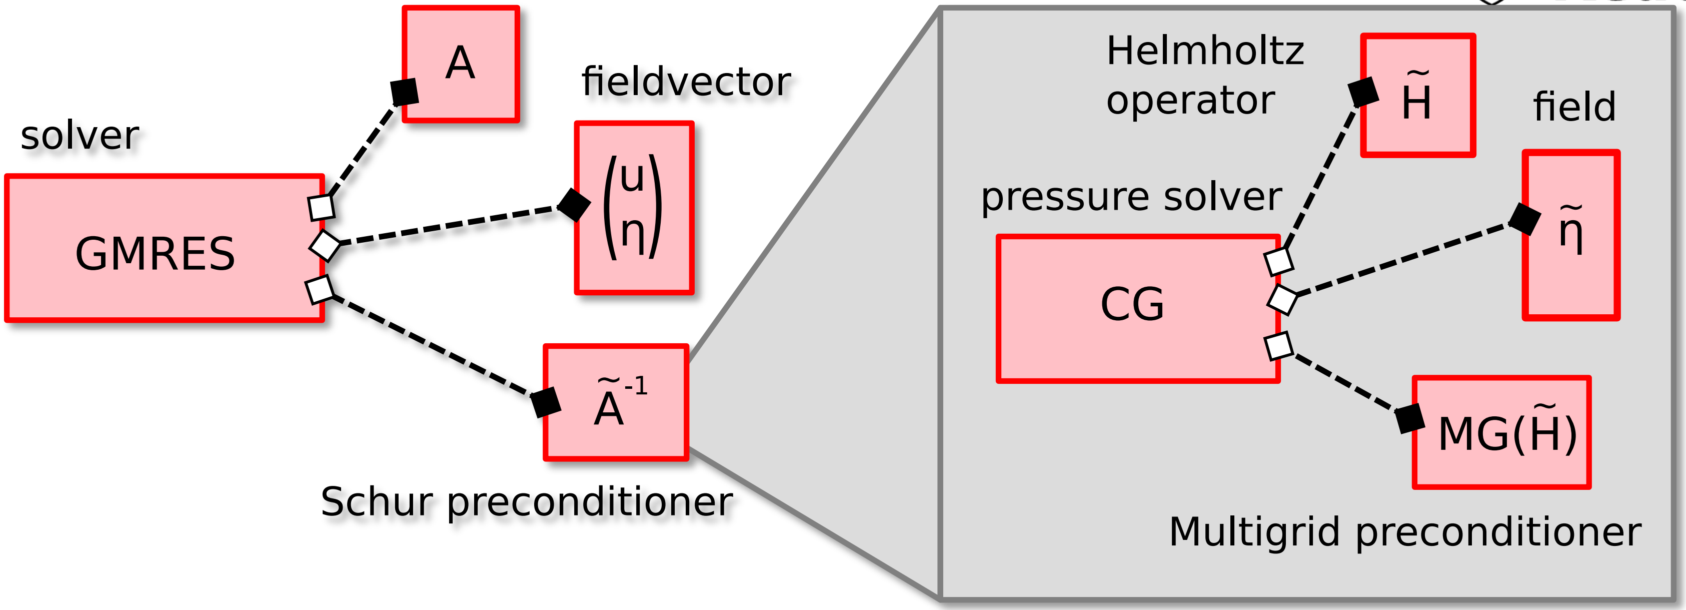

Allows for easy implementation of sophisticated nested solver
Multigrid preconditioner - reduce work for iterative solver
- faster and less global sums (better scaling)

© Crown Copyright 2017, Met Office

Initial Results

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. $||r|| = ||\mathbf{A}x - b||$ Of Krylov 10^{-2}
Before and after MG 3-level V-cycle

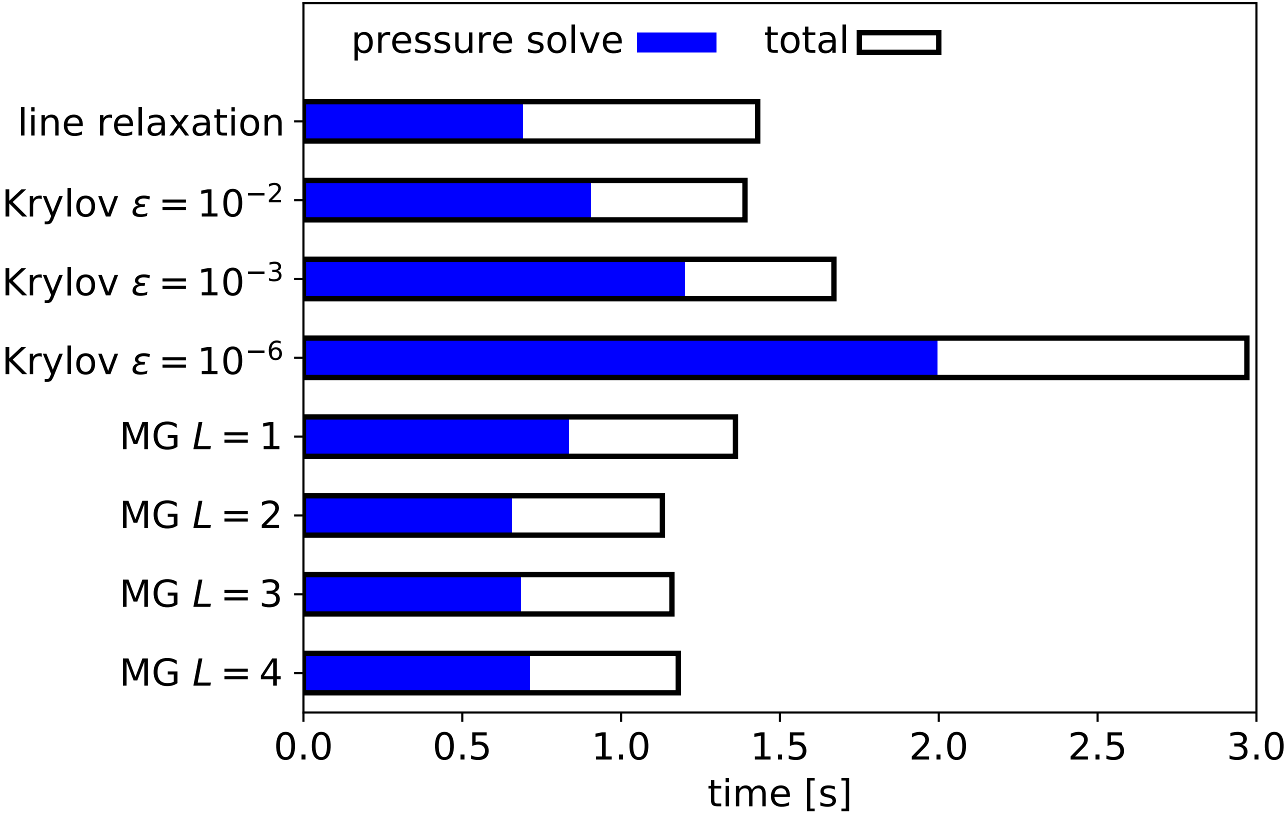

© Crown Copyright 2019, Met Office

Anatomy of a time-step

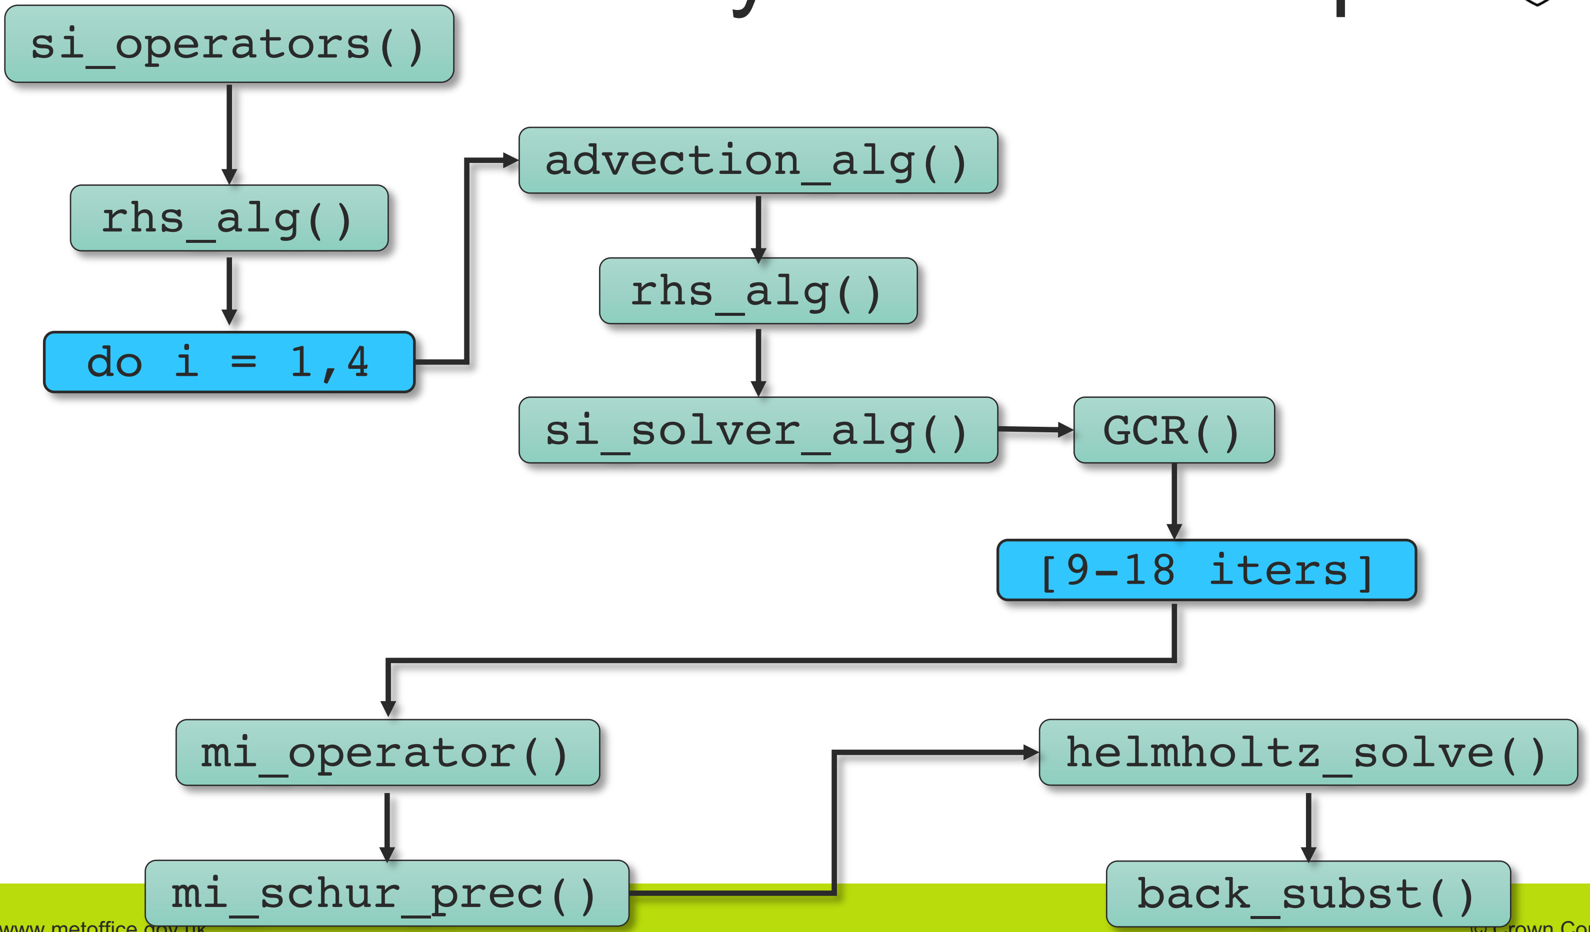

© Crown Copyright 2019, Met Office

Strong Scaling

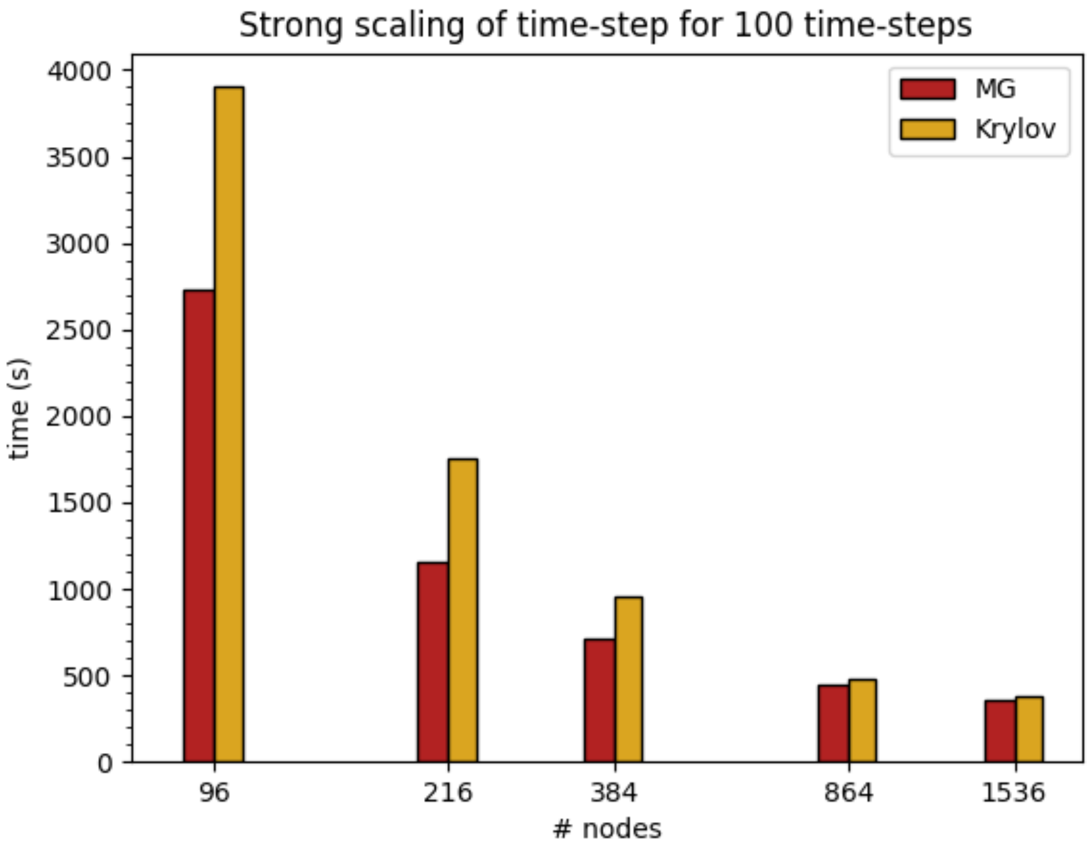

C1152 cubed sphere ~ 9 Km
Mixed mode 6 MPI/6 OMP

96 nodes LV = 48x48 x30L ~ 75 K dofs
1536 nodes LV = 12x12 x30L ~4300 dofs

Multigrid is faster, but …
Krylov subspace solver scales better but always slower …

© Crown Copyright 2019, Met Office

Met Office

Anatomy of a time-step

University of Reading

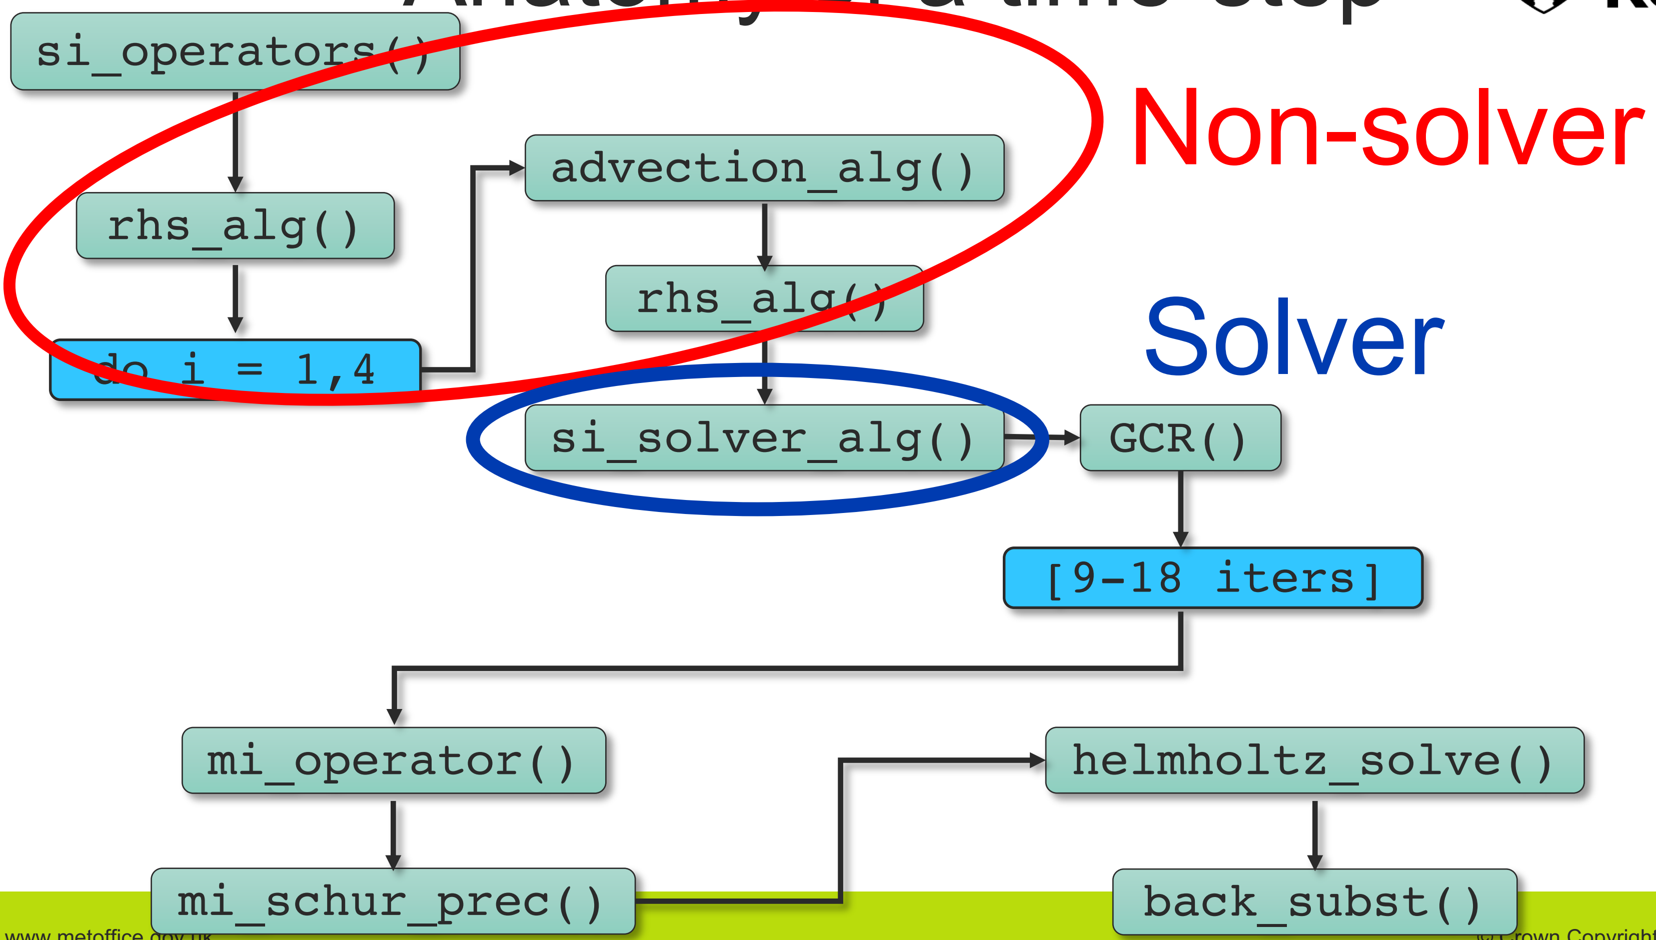

© Crown Copyright 2019, Met Office

Parallel efficiency

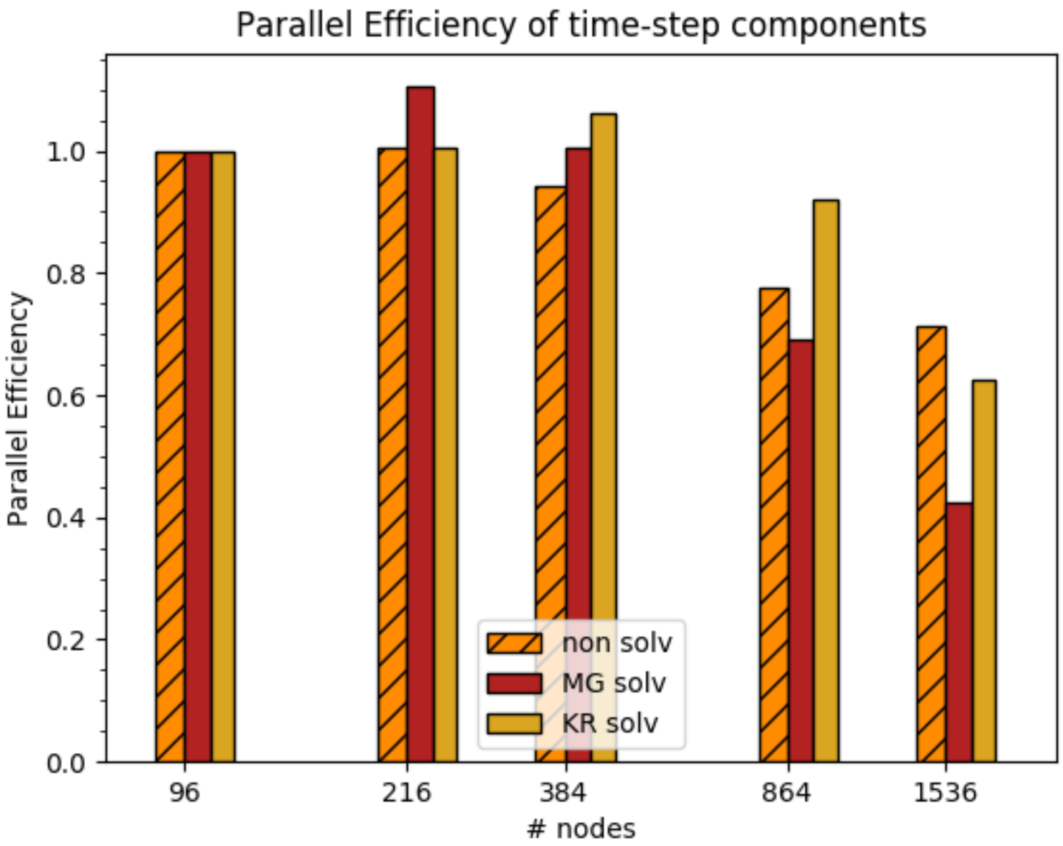

C.F. 96 node run
Non-solve is same for both, scales OK, not great

MG and KR solver start to drop over

MG scaling looks worse because starts better, always faster

© Crown Copyright 2019, Met Office

Anatomy of a time-step

Mixed op

Mixed Schur precon

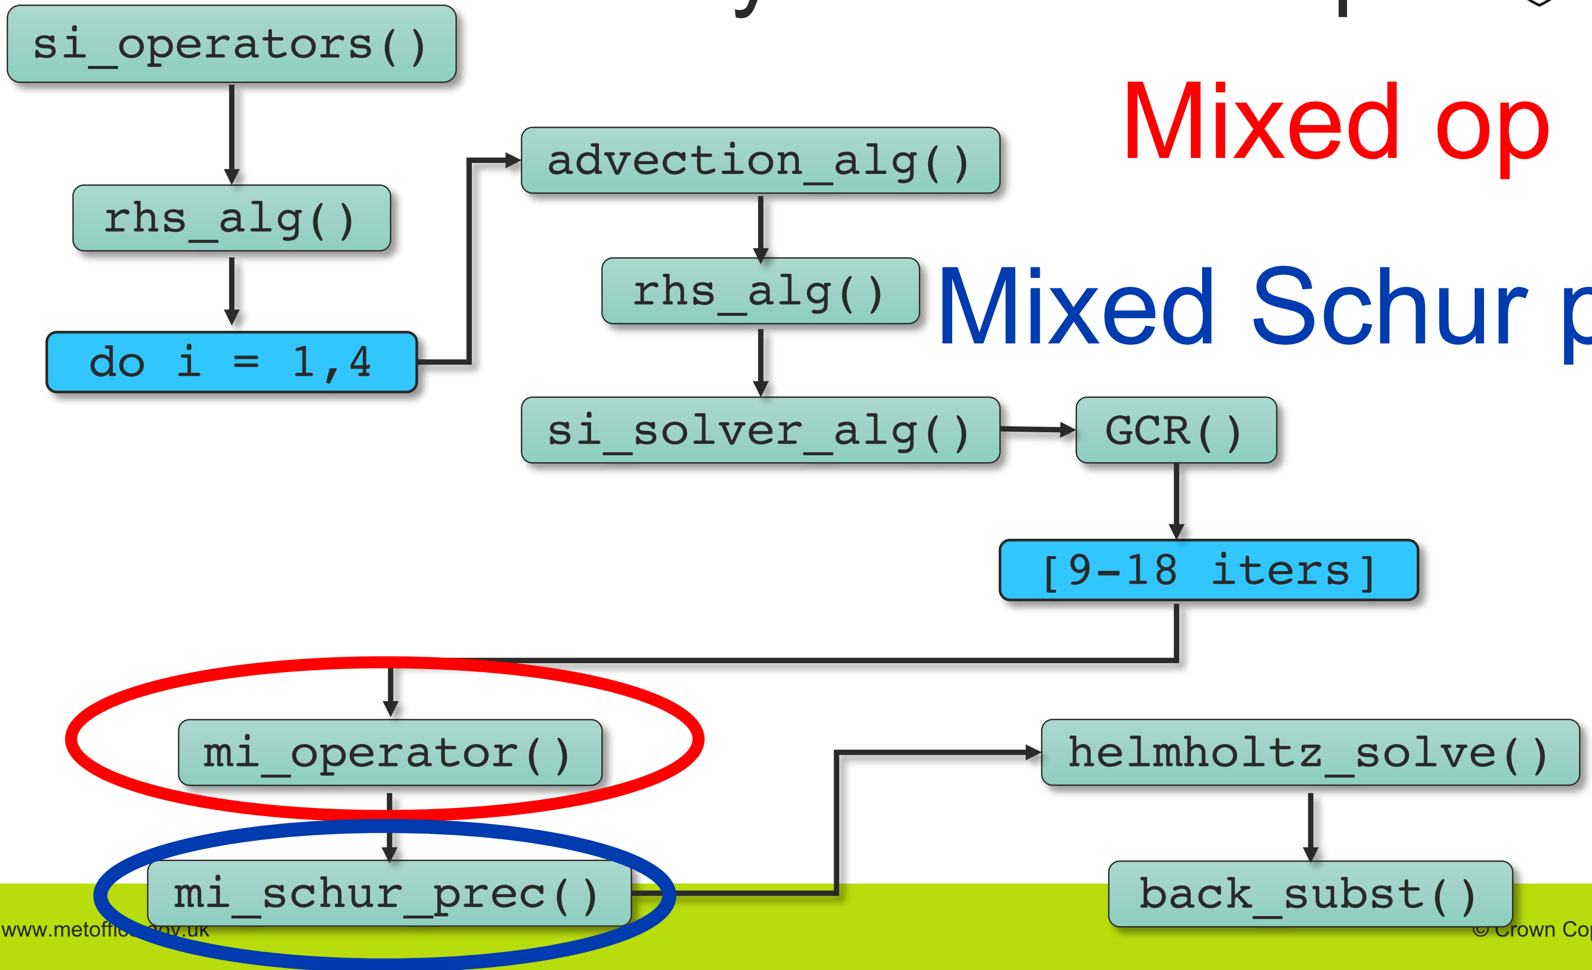

© Crown Copyright 2019, Met Office

Semi-implicit solver

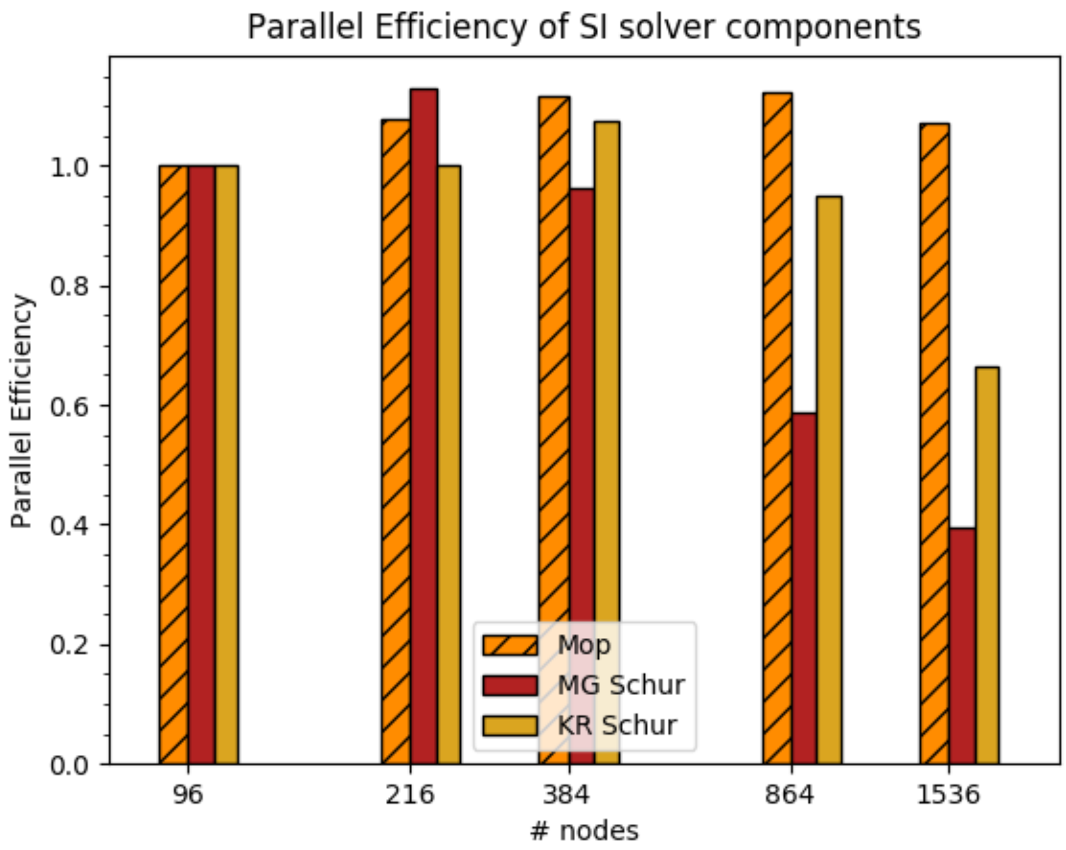

Mixed operator is same for both shows perfect scaling

MG and KR Mixed Schur precon, MG looks to be scaling much worse.

MG scaling looks worse because starts better, always faster

© Crown Copyright 2019, Met Office

Met Office

Anatomy of a time-step

University of Reading

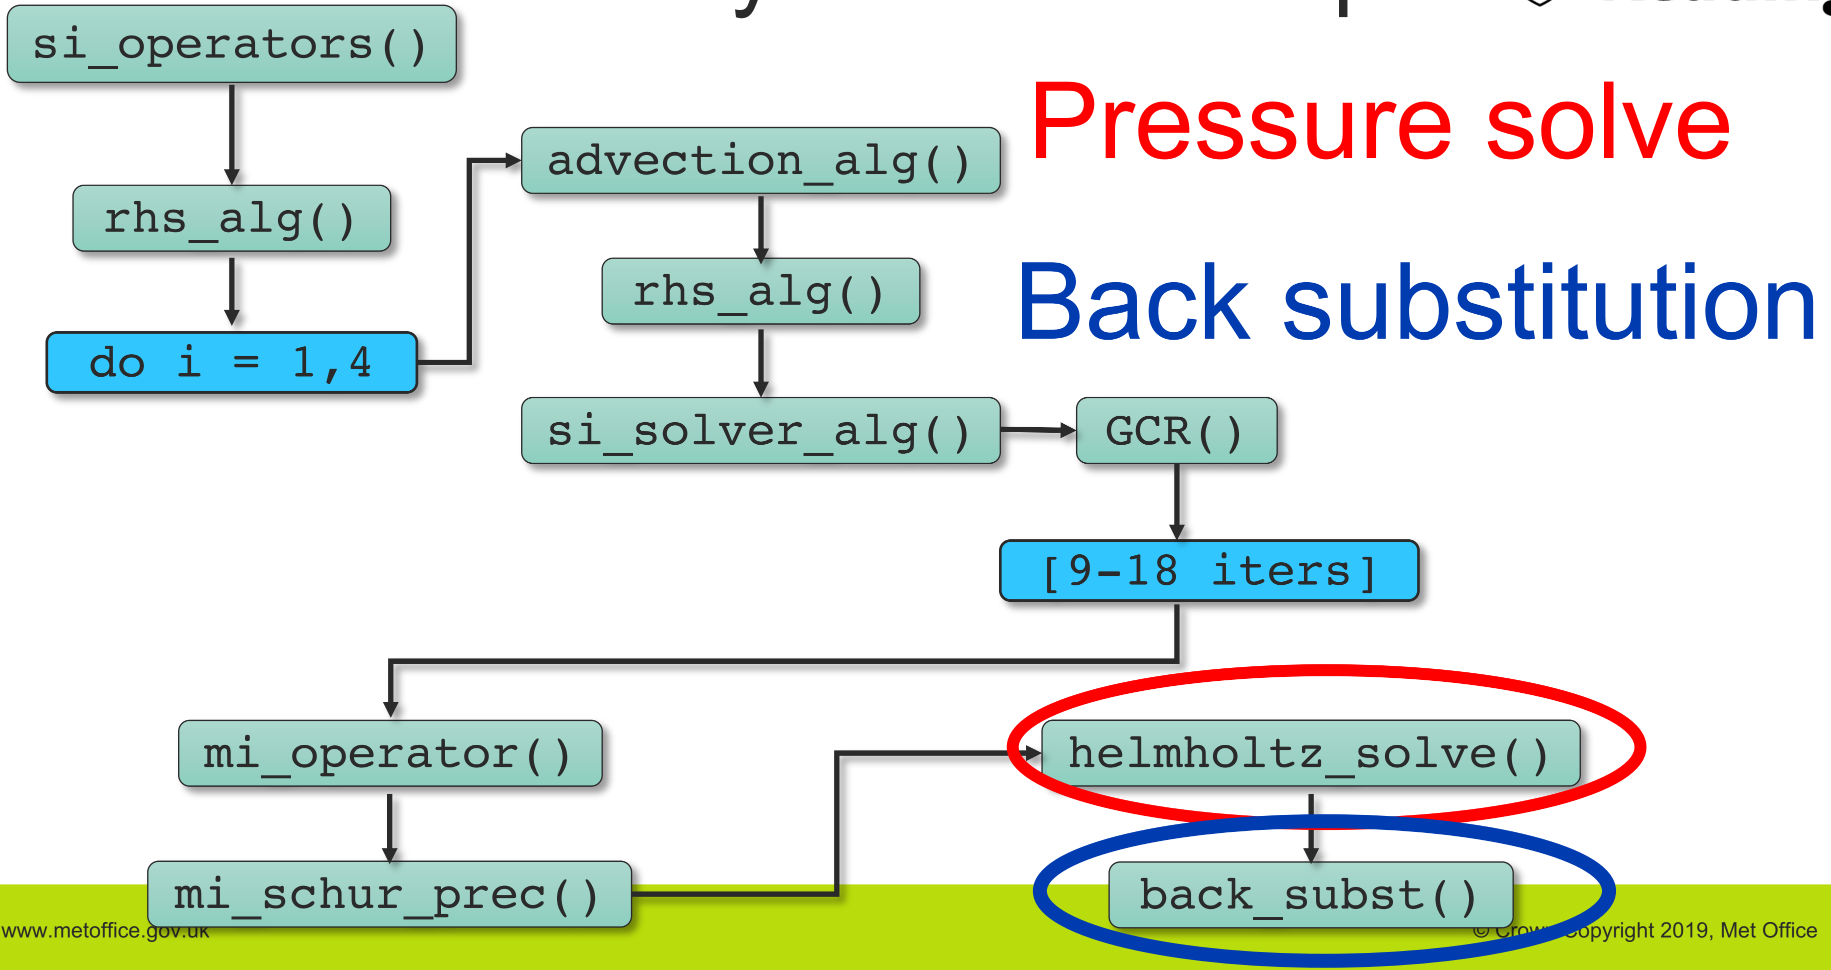

© Crown Copyright 2019, Met Office

Pressure solve

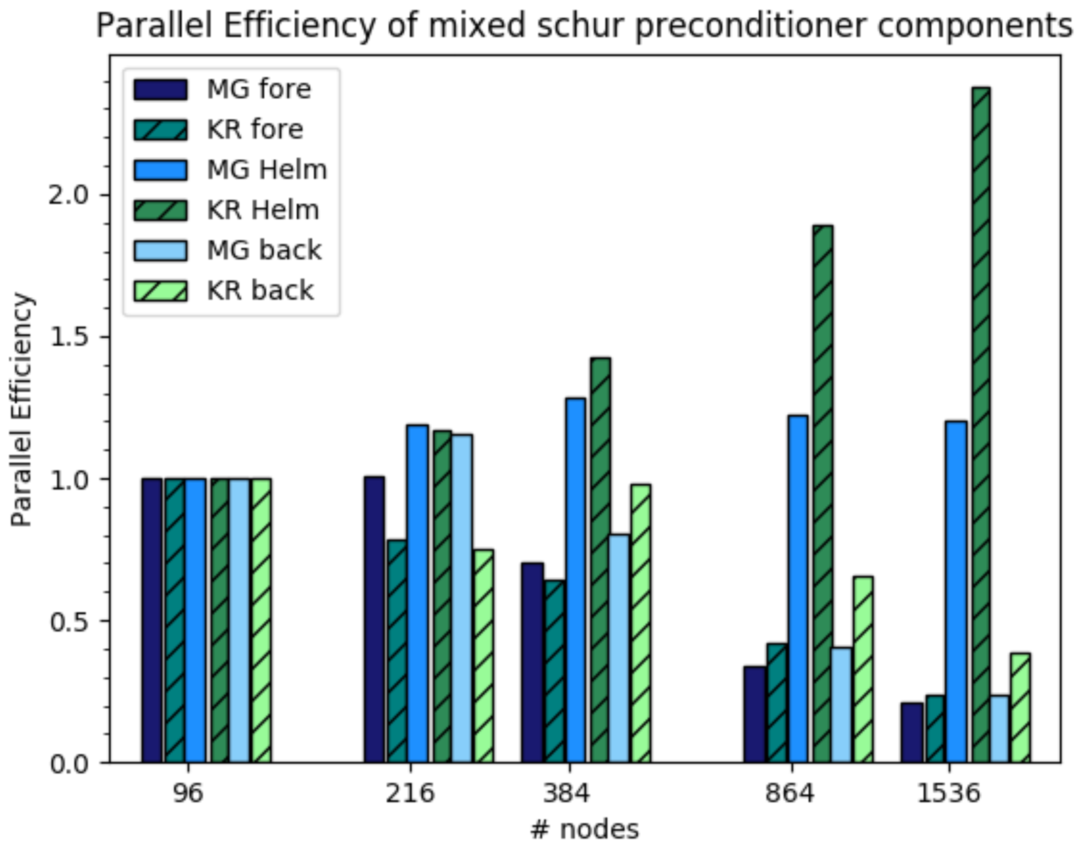

Forward and backward substitution for both KR and MG show poor scaling

KR helm (solver) shows super linear scaling (very slow on 96 nodes)

MG shows perfect scaling
→ Efficient everywhere. Always fastest

Run-time on 1536 nodes is tiny for solver – scaled away
What happed to global sums?

© Crown Copyright 2019, Met Office

How many global sums?

Met Office

University of Reading

- 100 ts x 4 sub ts = 400 calls to GCR outer solver
- 9-18 iters → 5800 calls to BiCGstab
- Solve to 10^{-2} ave 7 iters
- Total number of BiCGstab iters ~ 50K
- 5 Global sums per iter
- ~250K GS
- At 55K cores, GS latency < $10\mu s$ → 2.5 s
- Best case scenario, adaptive network slow that down

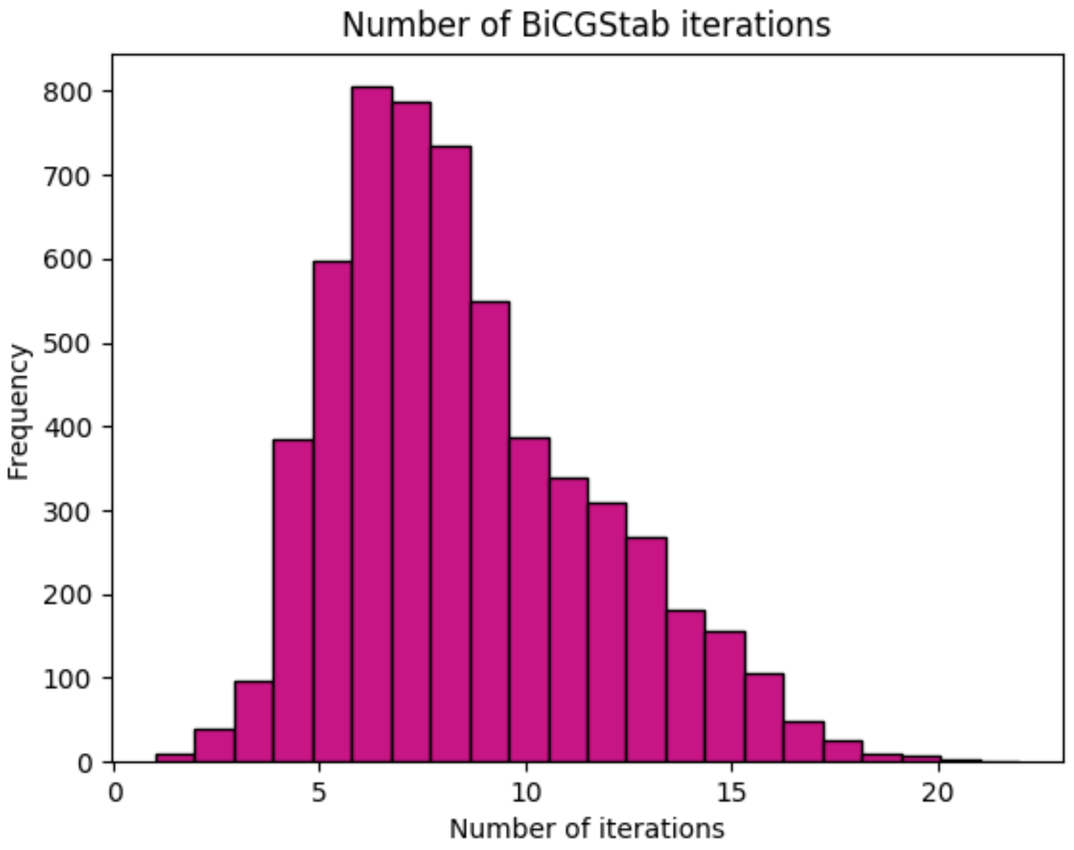

How many global sums

Solve 10^{-6} ave is 50-55 iters
Total BiCGStab iters ~ 300,000
→ 1.5 million Global sums
5x more than 10^{-2}
Here, multigrid has massive scaling advantage. 5-10x reduction in cost of GS
GCR still has GS

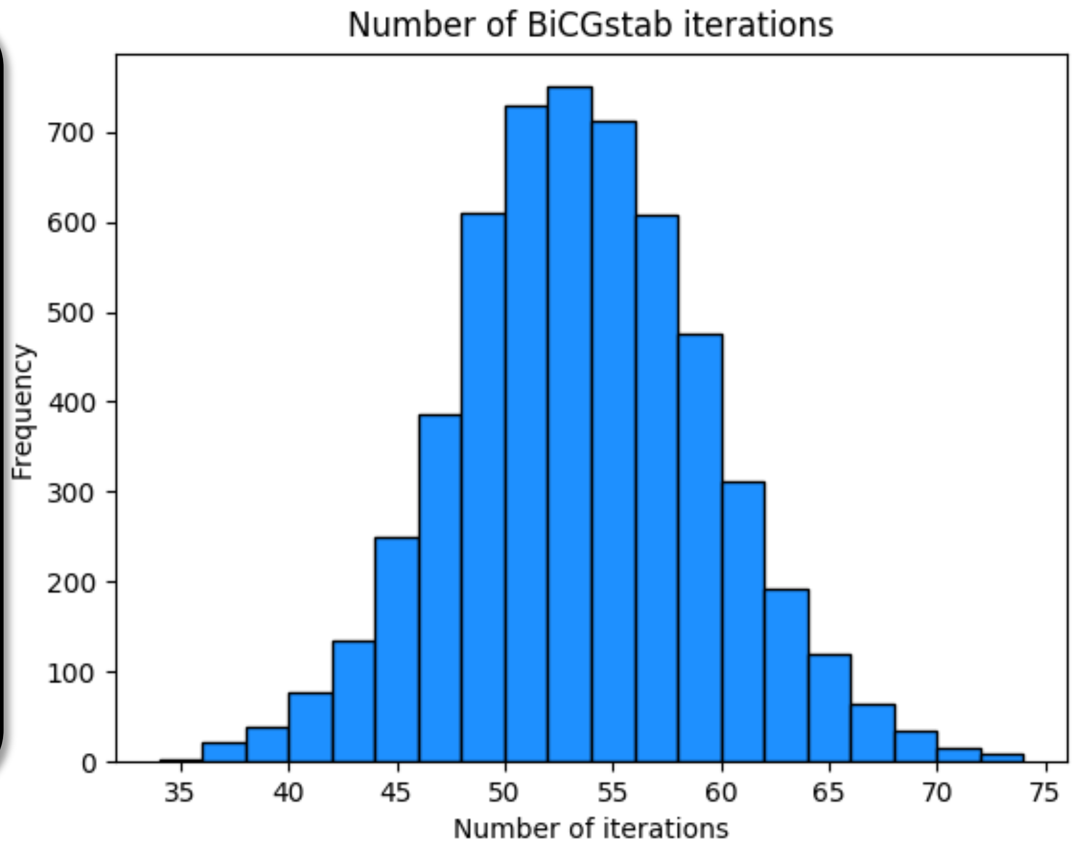

Conclusions

Implemented sophisticated solver framework
Allows for ease-of-use algorithmic changes
Including multigrid solver
Compared BiCGstab 10^{-2} to 3-level MG

MG always faster
BiCGStab super-scales – because dominated by computation on few nodes for large residual
Both solvers scale away, and become small in the profile → local comms then dominate

Is $||r|| = 10^{-2}$ realistic? With Physics & Orography?
Less ideal problem – multigrid has much bigger advantage as Kr solver will see many more Global Sums.

© Crown Copyright 2019, Met Office