ExCALIBUR 10

EXCALIBUR FROM THE MET OFFICE PERSPECTIVE

Nigel Wood, Met Office

On behalf of many

And with thanks to our unsung project managers:

Christine Gifford, George Ware & Rachel Kidd

October 2024

Met Office

Mighty oaks from little acorns grow

September 2018, Stephen Belcher CSc Met Office (proposal for Strategic Priorities Fund Wave 2)

Harnessing the power of exascale supercomputers: Supercomputer architectures are undergoing revolutionary change as we move into the era of exascale computing. And current simulation codes, including codes for weather forecasting and climate projection, engineering simulation, planetary physics, solid Earth, drug design, will no longer work on the new machines. There is a pressing need to develop new software techniques, new numerical algorithms, train a new generation of computer scientists, and re-engineer simulation codes. The UK has international leadership in simulation, and the time is right to launch a multi-disciplinary programmne to place the UK at the forefront of this revolution. (Note that this theme refers to exascale computing for *simulation* and rather than *artificial intelligence* as the software and hardware requirements are very different)

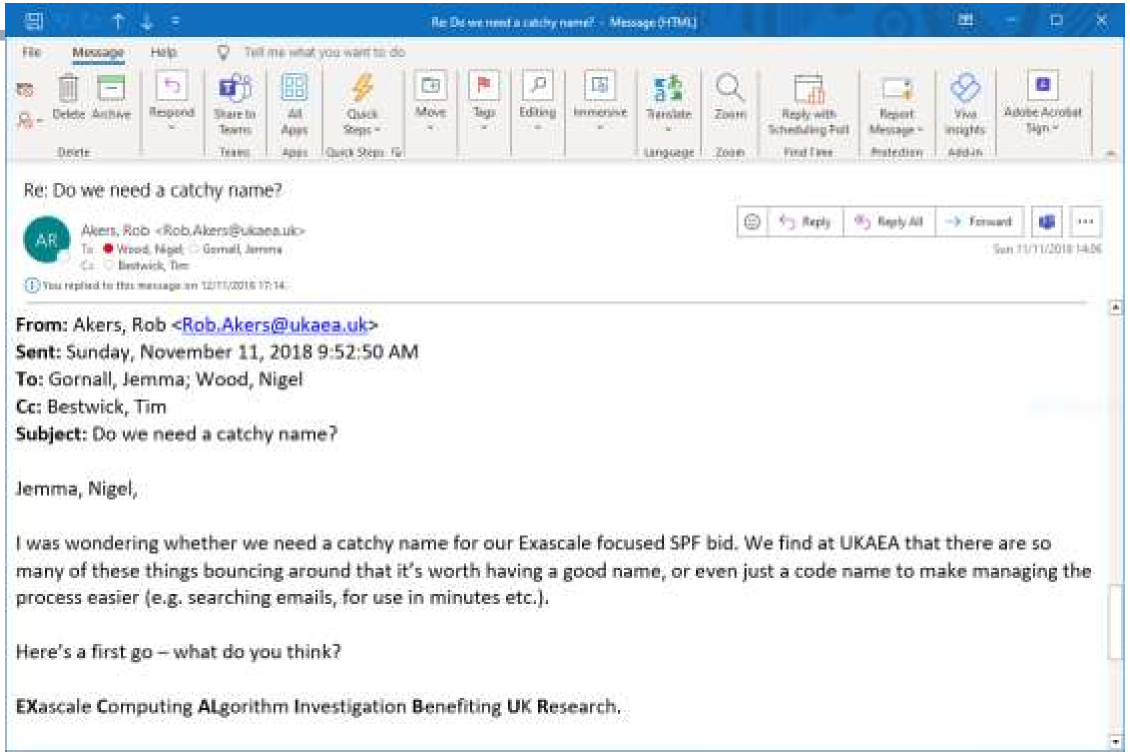

Re: Do we need a catchy name?

From: Akers, Rob <Rob.Akers@ukaea.uk>
Sent: Sunday, November 11, 2018 9:52:50 AM
To: Gornall, Jemma; Wood, Nigel
Cc: Bestwick, Tim
Subject: Do we need a catchy name?

Jemma, Nigel,

I was wondering whether we need a catchy name for our Exascale focused SPF bid. We find at UKAEA that there are so many of these things bouncing around that it's worth having a good name, or even just a code name to make managing the process easier (e.g. searching emails, for use in minutes etc.).

Here's a first go – what do you think?

EXascale **C**omputing **AL**gorithm **I**nvestigation **B**enefiting **UK** **R**esearch.

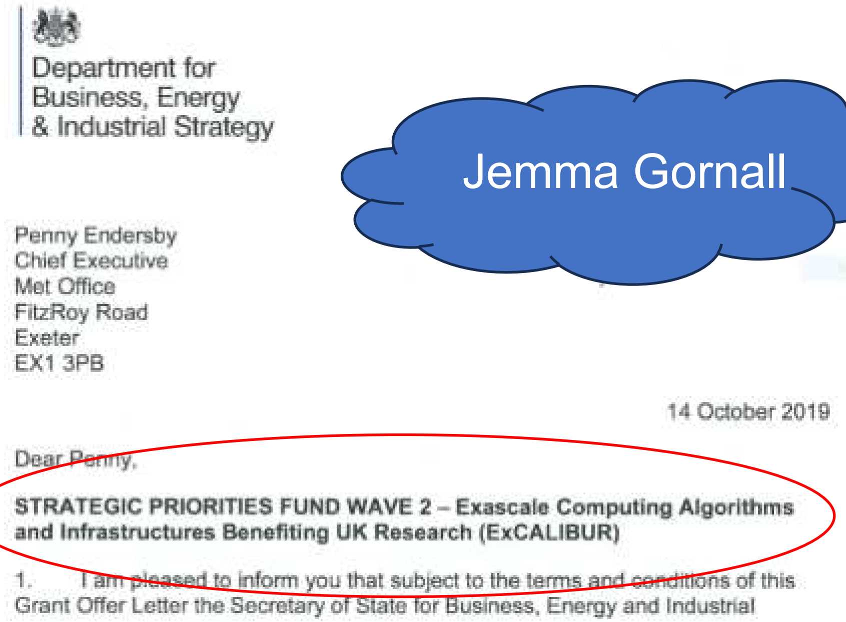

Department for Business, Energy & Industrial Strategy

Penny Endersby
Chief Executive
Met Office
FitzRoy Road
Exeter
EX1 3PB

14 October 2019

Dear Penny,

STRATEGIC PRIORITIES FUND WAVE 2 – Exascale Computing Algorithms and Infrastructures Benefiting UK Research (ExCALIBUR)

1. I am pleased to inform you that subject to the terms and conditions of this Grant Offer Letter the Secretary of State for Business, Energy and Industrial

Jemma Gornall

Nearly £0.5M per word ☺

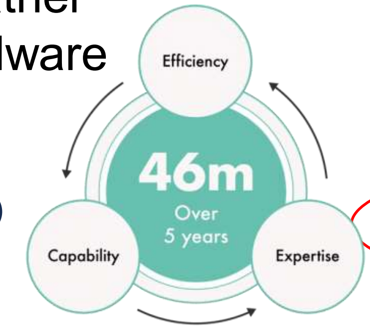

UK's Weather & Climate prediction system

- Operational forecasts
  - ➢Global (resolution approx. 10km)
  - ➢Regional (resolution approx. 1.5km)

10 km

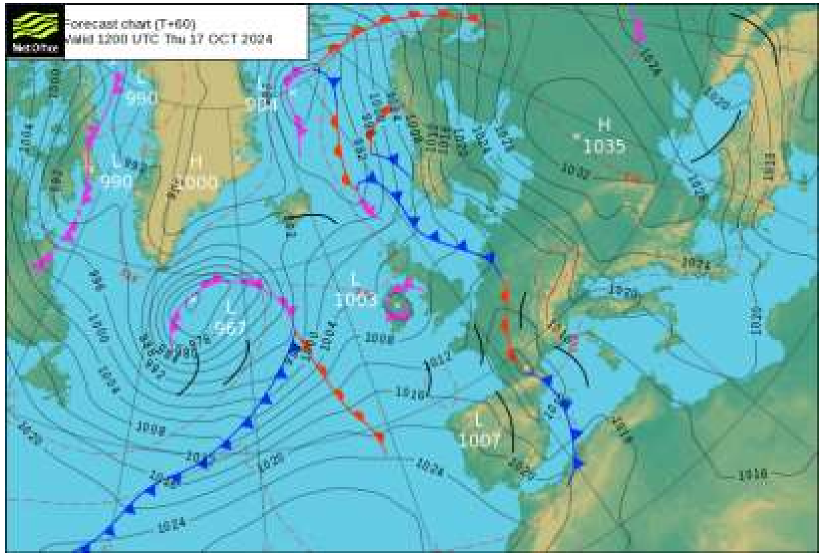

300 m

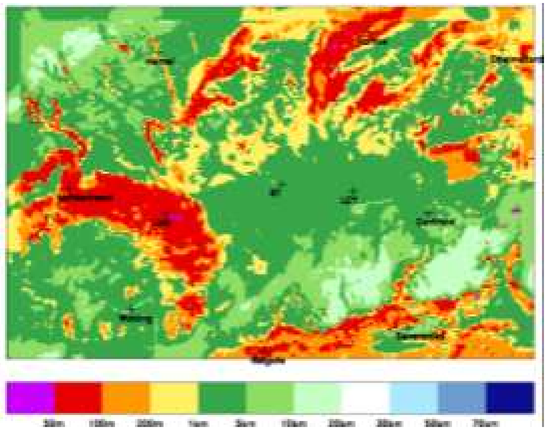

- Global and regional climate predictions
  - ➢Global resolution around 120km
  - ➢Regional around 4-1.5km
  - ➢Run for 10-100-… years

300 km

- Seasonal predictions
  - ➢Resolution approx. 60km

Unified ⇒ Same solver, same parametrisations, same code base for all

“The quiet revolution”*: ≈ 1 day’s lead time per decade

5-day forecast in 2020 ~ as accurate as 4-day forecast in 2010

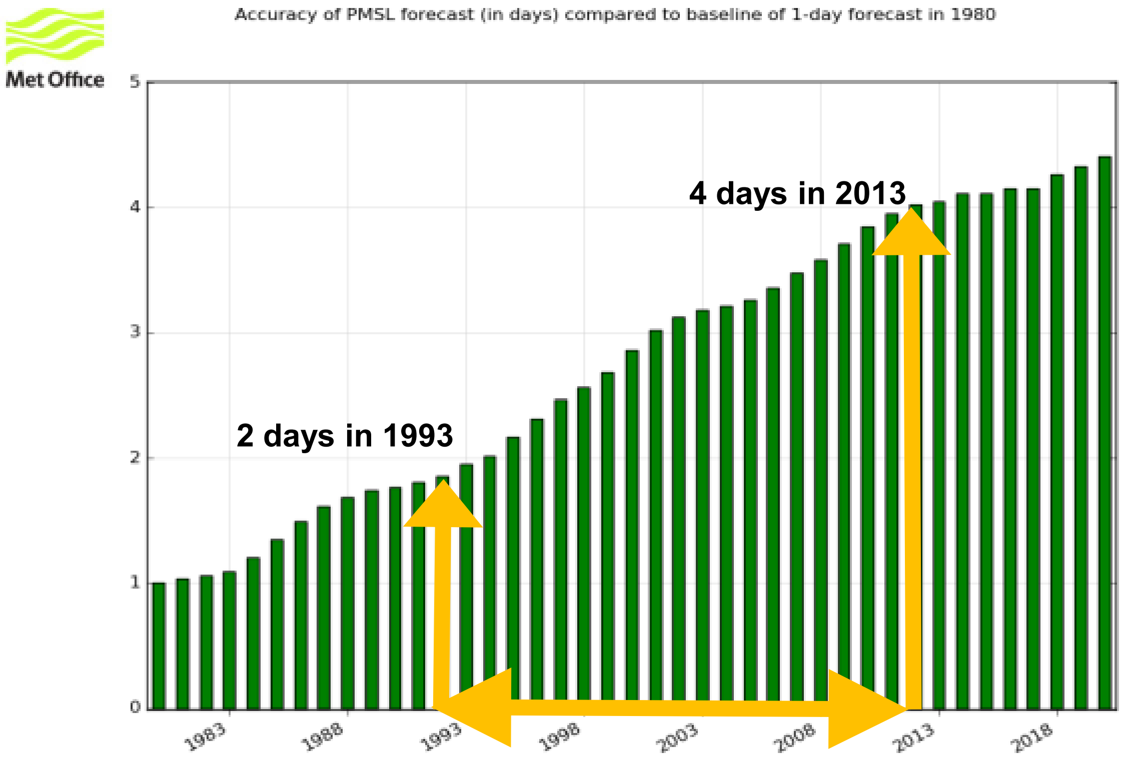

Sea level pressure

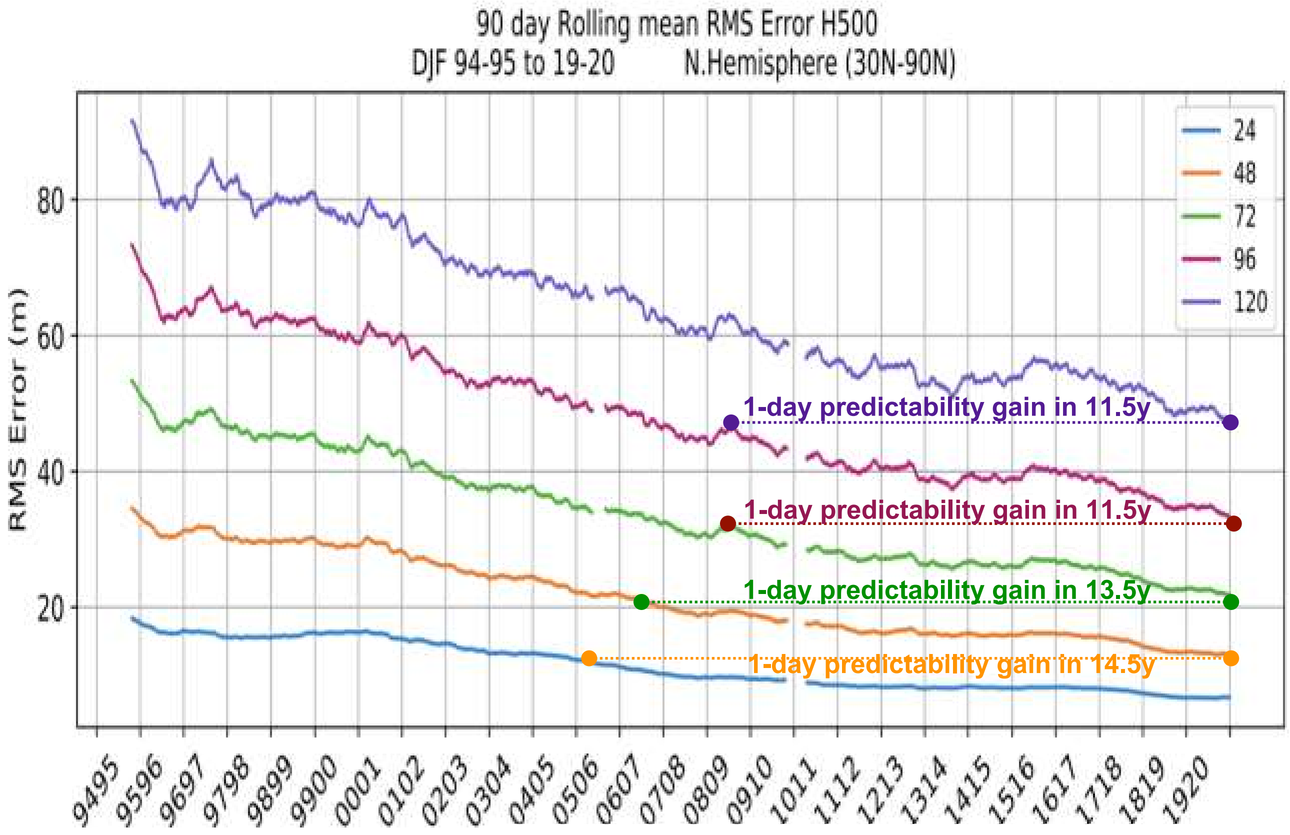

500 hPa height

“…impact of Numerical Weather Prediction among greatest of any area of science… comparable to simulation of human brain and evolution of early universe”*

*Bauer, Thorpe, Brunet (2015) Nature

ExCALIBUR 10

History of Met Office computers

Another view of KDF 9 computer, with line printer prominent in foreground. The Director-General is holding a specimen forecast output.

ExCALIBUR

Our challenge

Why Exascale?

- Currently we simulate the world's weather at 10 km intervals
- To complete the 7-day forecast in 1 hour needs a petascale machine (16 Pflops) (we use 19,000 cores)

- To get to 5 km means 2x2 more cells and a 2 times smaller interval in time
  - **O(10)** increase in compute power & data
- To get to 1 km means 10x10 more cells and a 10 times smaller interval in time
  - **O(1000)** increase in compute power & data

Weather & Climate prediction system

Schematic

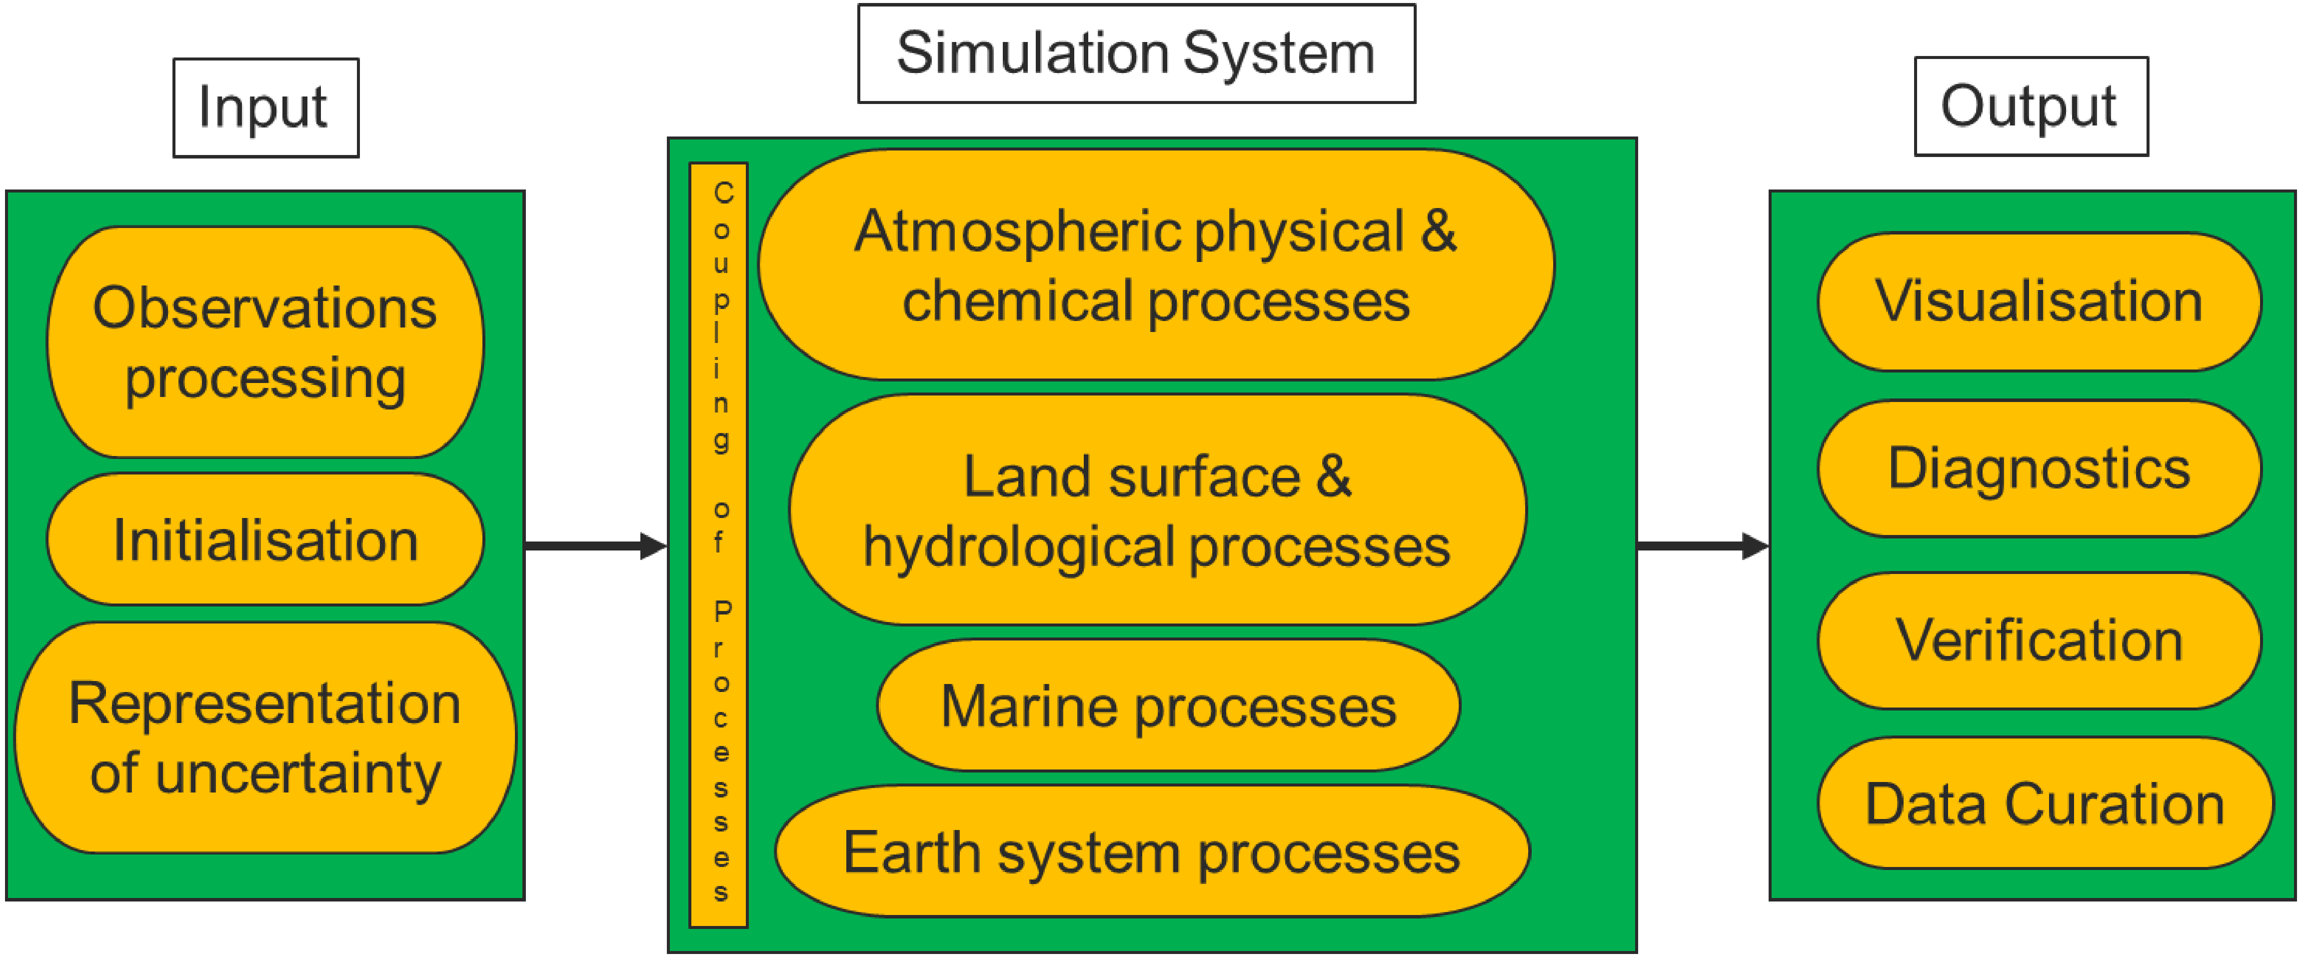

(Note: bulk of work lies with Met Office's Next Generation Modelling Systems Programme
ExCALIBUR activities augment & accelerate various aspects)

Use Case Activities

Work package 1 Component Model Co-design

- UK Chemistry & Aerosol (UKCA) model Design
- Marine Systems (NEMO) design
- Atmospheric Model data layout and memory access design
- Atmospheric observation pre-processing and assimilation
- Marine Systems (WAVE WATCH III) design
- Verification system components design, integration and testing
- Data assimilation algorithm research
- Efficient application of atmospheric model parametrisation schemes
- Joint UK Land Environment Simulator (JULES) model design

Work package 2 System Co-design

- Design and implementation of new weather and climate diagnostic system
- Spatial decoupling of dynamics, physics and chemistry
- Coupling multiple components as a single executable
- Framework for application input configuration and validation

Work package 3 System Integration

- Adaptation of Marine Systems for GPU-based and other HPC systems
- Pull through of LFRic multi-mesh work
- Ensemble-based Next Generation Weather & Climate Operational Suites
- Adaptation of atmospheric physics kernels to GPUs
- Pull through of UKCA work
- JEDI workflow

Weather & Climate prediction system

Activities

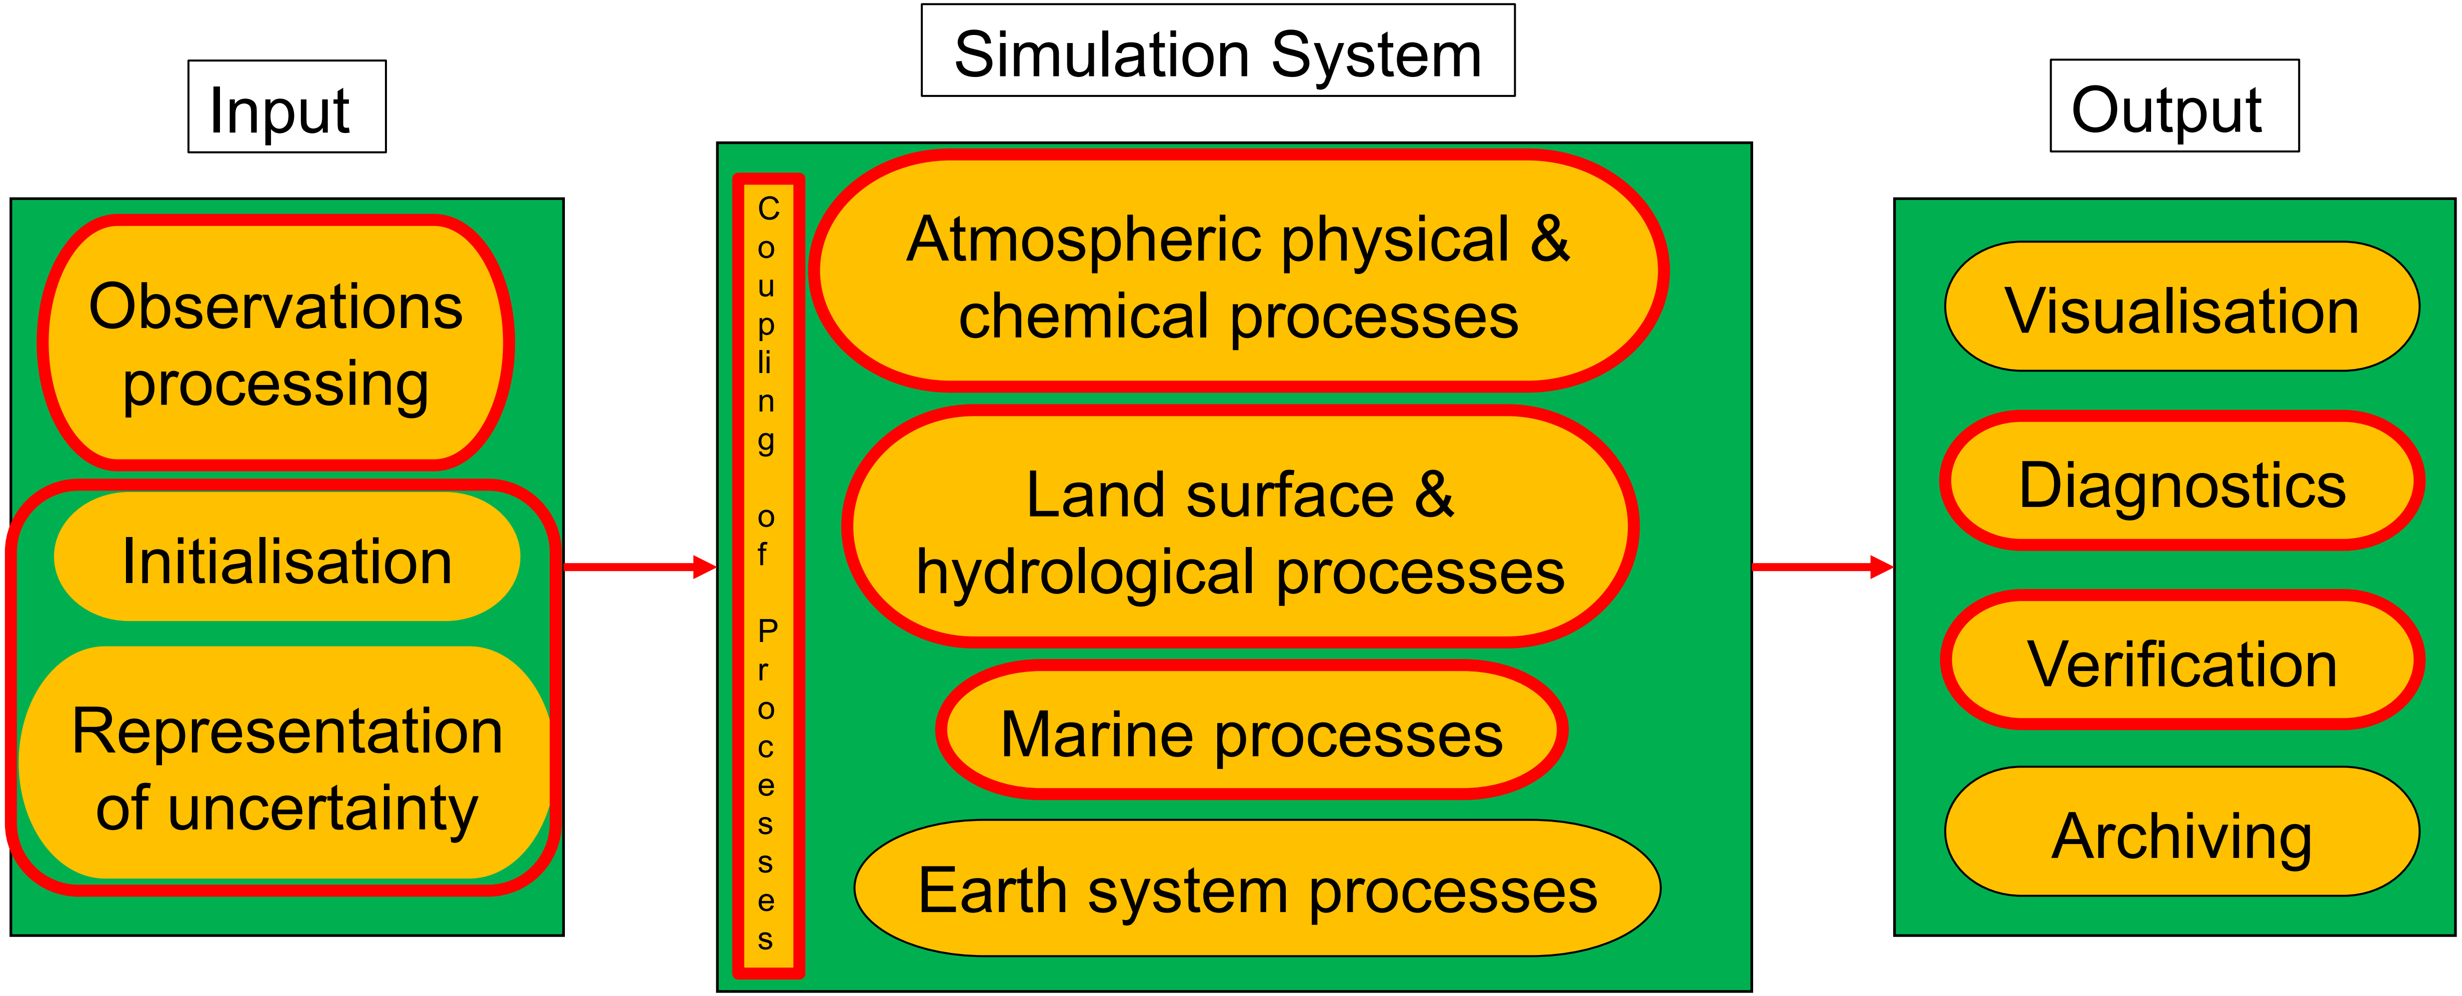

Use Case Activities

Work package 1 Component Model Co-design

- UK Chemistry & Aerosol (UKCA) model Design
- Marine Systems (NEMO) design
- Atmospheric Model data layout and memory access design
- Atmospheric observation pre-processing and assimilation
- Marine Systems (WAVE WATCH III) design
- Verification system components design, integration and testing
- Data assimilation algorithm research
- Efficient application of atmospheric model parametrisation schemes
- Joint UK Land Environment Simulator (JULES) model design

Work package 2 System Co-design

- Design and implementation of new weather and climate diagnostic system
- Spatial decoupling of dynamics, physics and chemistry
- Coupling multiple components as a single executable
- Framework for application input configuration and validation

Work package 3 System Integration

- Adaptation of Marine Systems for GPU-based and other HPC systems
- Pull through of LFRic multi-mesh work
- Pull through of UKCA work
- JEDI workflow
- Ensemble-based Next Generation Weather & Climate Operational Suites
- Adaptation of atmospheric physics kernels to GPUs

ExCALIBUR

Physics-Dynamics-Chemistry Coupling with components of different resolutions in LFRic-atmosphere

Tom Bendall, Alex Brown, **Ben Shipway** (Met Office)

Atmospheric Model data layout, memory access design system, and spatial decoupling of processes

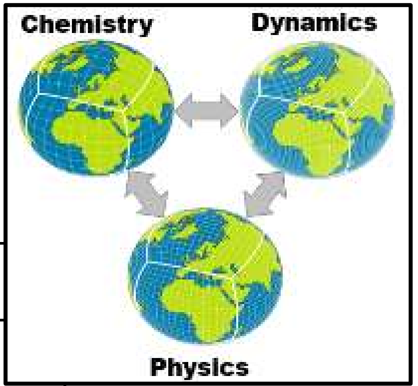

- **Only do it to the accuracy needed:** Framework in place for mixed-precision.

- **Do it using the optimal data layout:** Implementation of the “i-first” data transpose to the microphysics code in the LFRic basic-gal model shows a **4x speed up** in that part of the model.

- **Only do what is needed:** Implementation of capability to run transport, dynamics, physical parametrizations, … at different resolutions within same model.

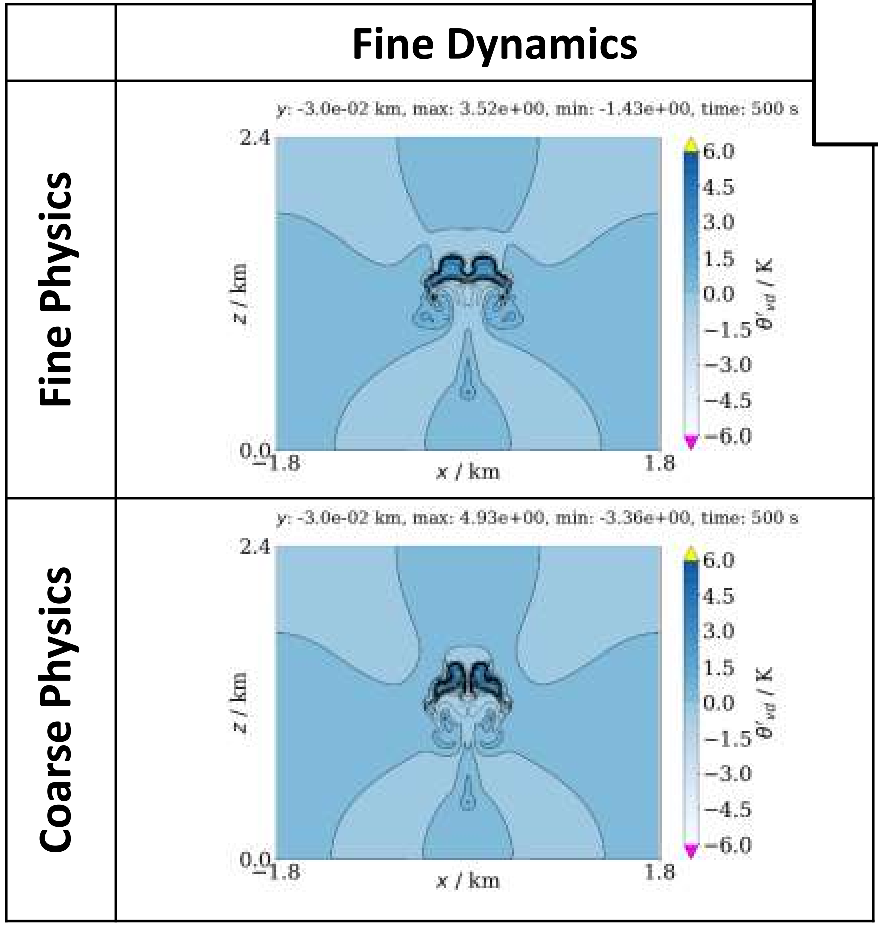

A moist bubble test showing the impact of running the physics (condensation/cloud scheme) at a coarser resolution.

Spatial decoupling of processes

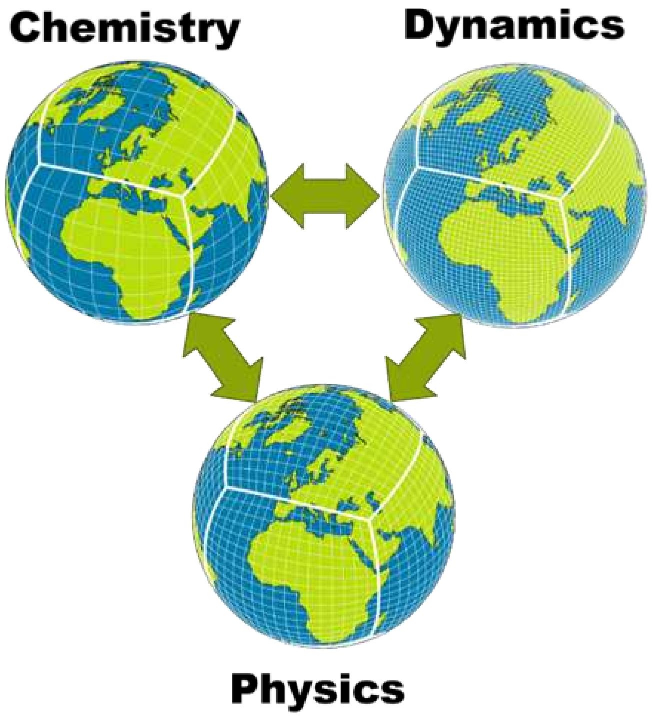

Developed capability to run *different components* of LFRic atmospheric model *at different horizontal resolutions*

Resolution of some parts of model can be increased/decreased independently of others, so computational resources can be targeted/saved

The formulation is designed to preserve key **physical properties**.

1. **Reversibility**: mapping from a coarse mesh to a fine mesh and back should leave a field unchanged
2. **Preservation of a steady-state**: if a physics parametrisation on one mesh results in no change, it should not change the state on the other mesh
3. **Local conservation of mass of moisture, chemicals and aerosols mixing ratios**
4. **Preservation of constant moisture, chemical and aerosol mixing ratios**
5. **Avoid generation of negative moisture mixing ratios**

ExCALIBUR

Input Data
Other Ancillary
Files

Dynamics / Physics

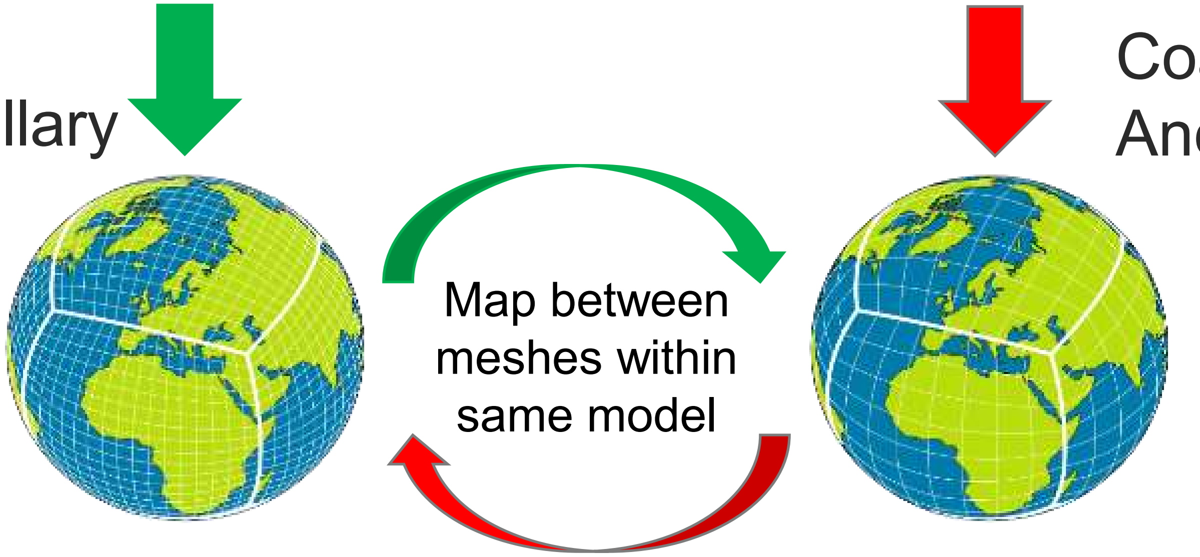

Coarse Aerosol
Ancillary Files

RadAer = Radiation/Aerosol interaction

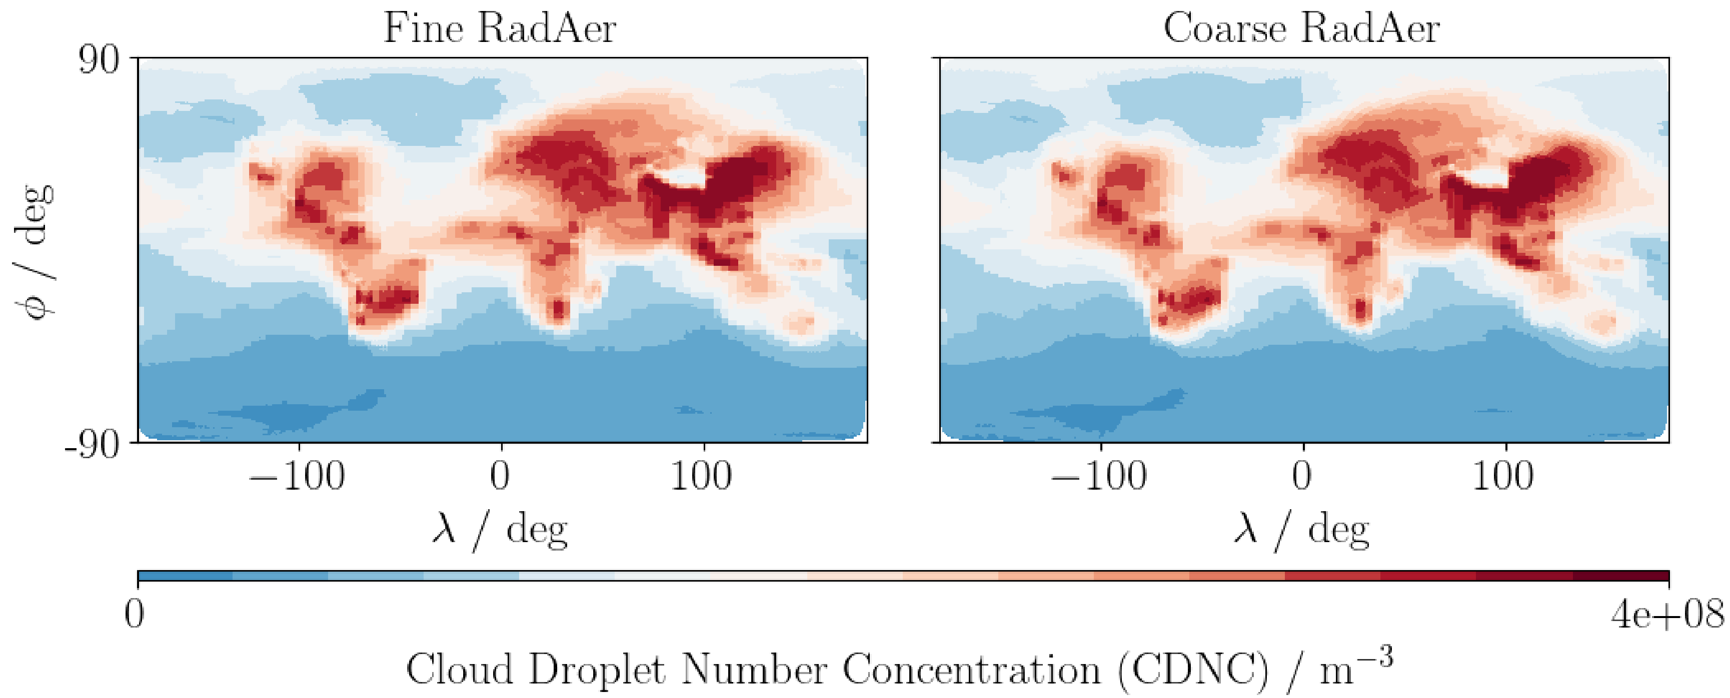

Joint UK Land Environment Simulator (JULES)

Doug Clark, Rich Ellis, **Emma Robinson** (CEH)

Dave Case, Bryan Lawrence, Grenville Lister, David Livings, Simon Wilson (Reading)

ExaJULES: Model development for JULES in LFRic

- JULES = land surface component of *coupled model* (UM and LFRic)
- Also exists as a *standalone model* driven by input meteorology
- Historically these have had two different implementations:
  - ➢ Not a good separation of concerns
  - ➢ Leads to difficult pull-through of science from standalone to coupled system
  - ➢ Existing JULES standalone can't take advantage of exascale developments

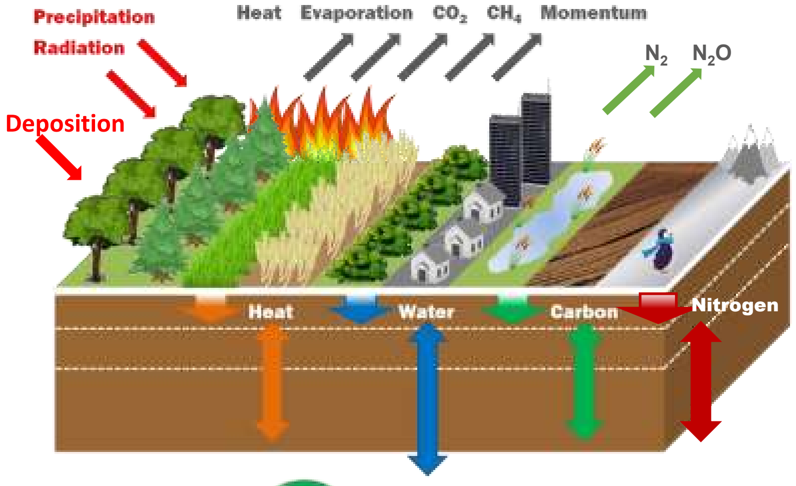

JULES
Joint UK Land Environment Simulator

ExaJULES aim:
build a prototype JULES application in the LFRic system

ExaJULES: Outcomes

- ✓ Designed and implemented a JULES LFRic application
  - Bit comparison with existing benchmarks
  - Infrastructure to be included in LFRic trunk
- ✓ Assessment of performance of JULES science code
  - Recommendations for performance optimisation
- ✓ Analysis and recommendations for coupling land and atmosphere on different meshes

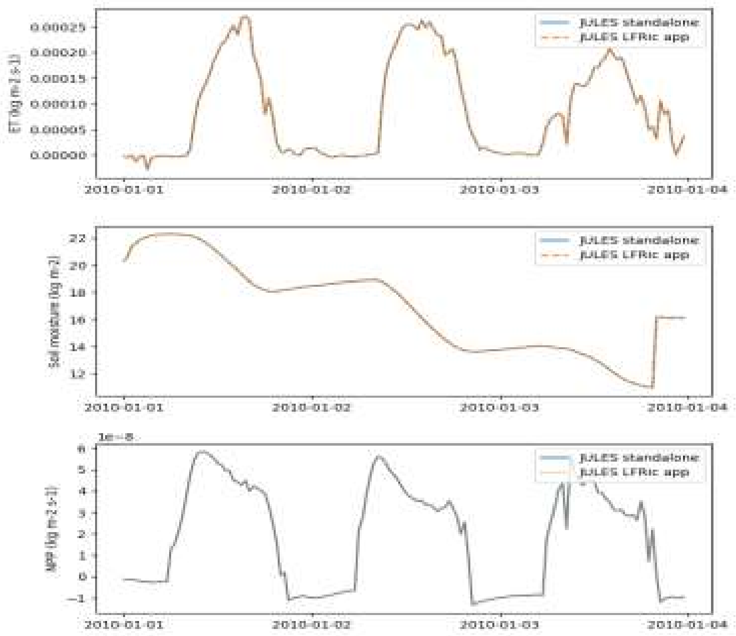

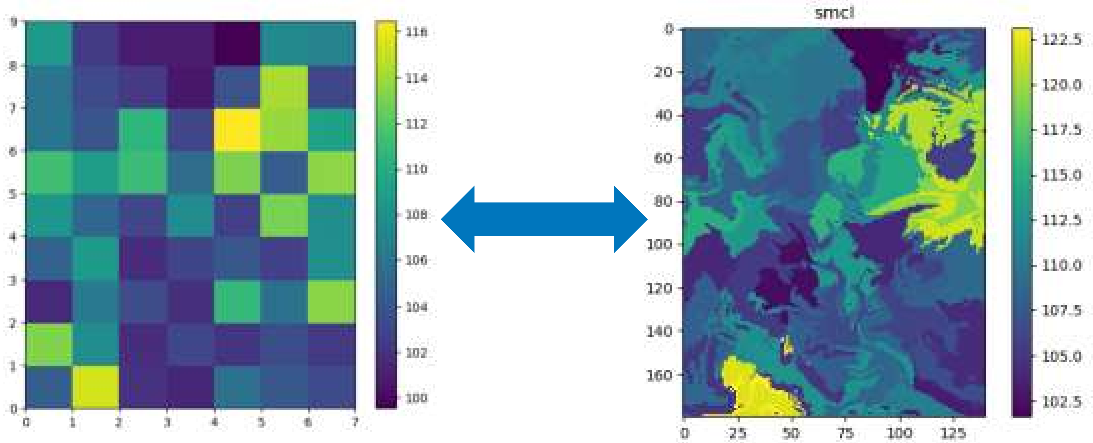

UK Centre for Ecology & Hydrology

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

ExCALIBUR

Marine Systems

Chris Dearden, Rupert Ford, **Andy Porter**, Sergi Siso (STFC)

Mike Bell, Daley Calvert, Iva Kavčič, Harry Shepherd (Met Office)

Andrew Coward (NOC)

Dave Case, Bryan Lawrence, Grenville Lister (Reading)

Seth Camp, Wayne Gaudin (Nvidia)

Suite of Marine Modelling Systems

NEMO → Nucleus for European Modelling of the Ocean

NEMO MEDUSA → Model of Ecosystem Dynamics, nutrient Utilisation, Sequestration and Acidification

NEMOVAR → Data Assimilation system for NEMO

WAVEWATCH III → NOAA's surface wave model

UKRI Science and Technology Facilities Council Hartree Centre

National Oceanography Centre

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

Met Office

NVIDIA

Performance Portability for Marine Modelling Systems

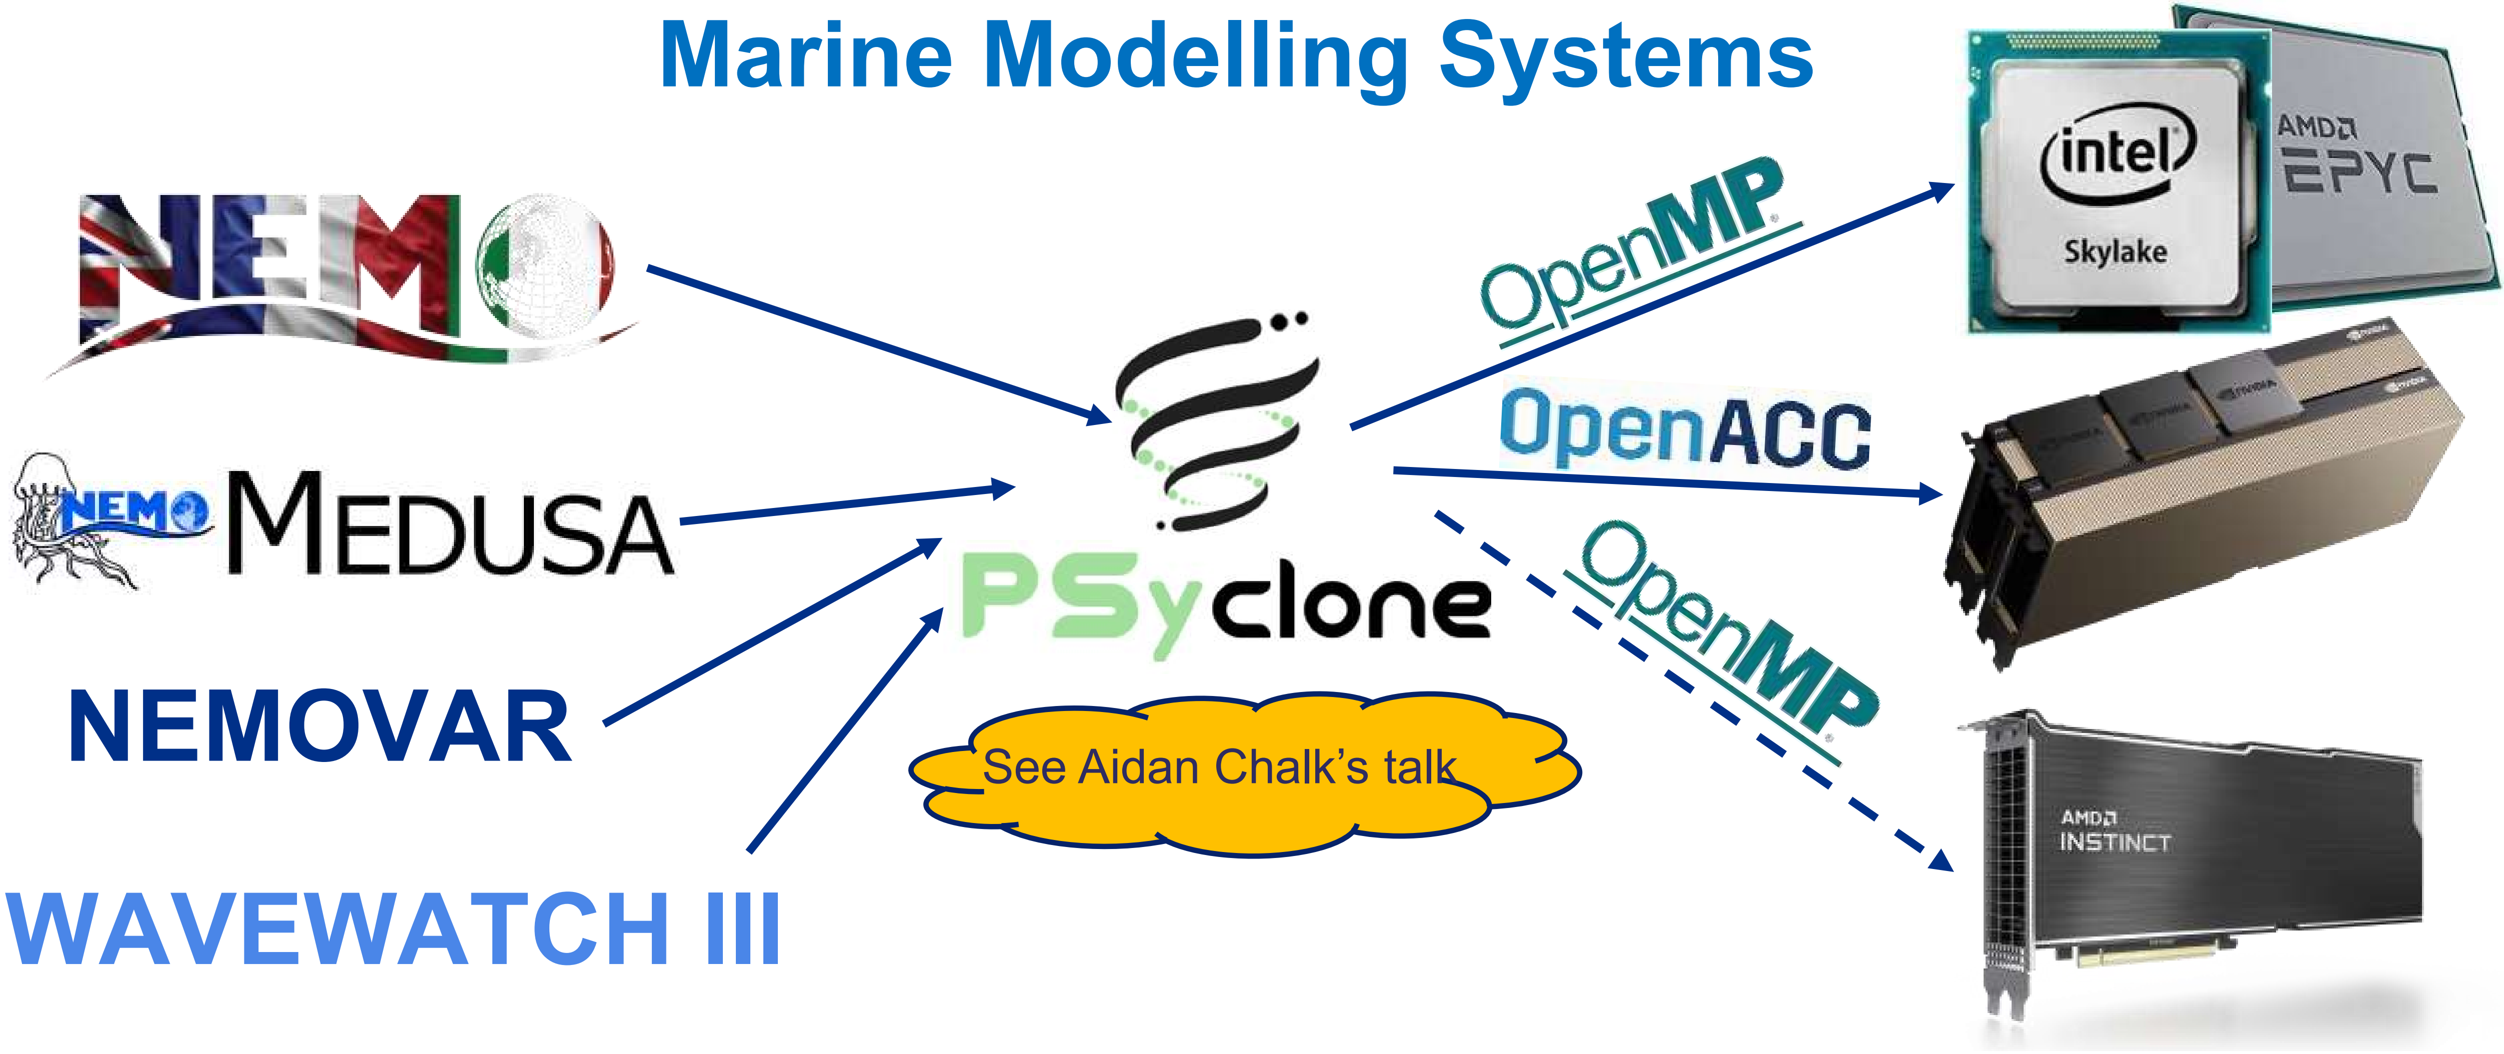

UKRI Science and Technology Facilities Council
Hartree Centre

National Oceanography Centre

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

Met Office

NVIDIA

Running NEMO on GPUs

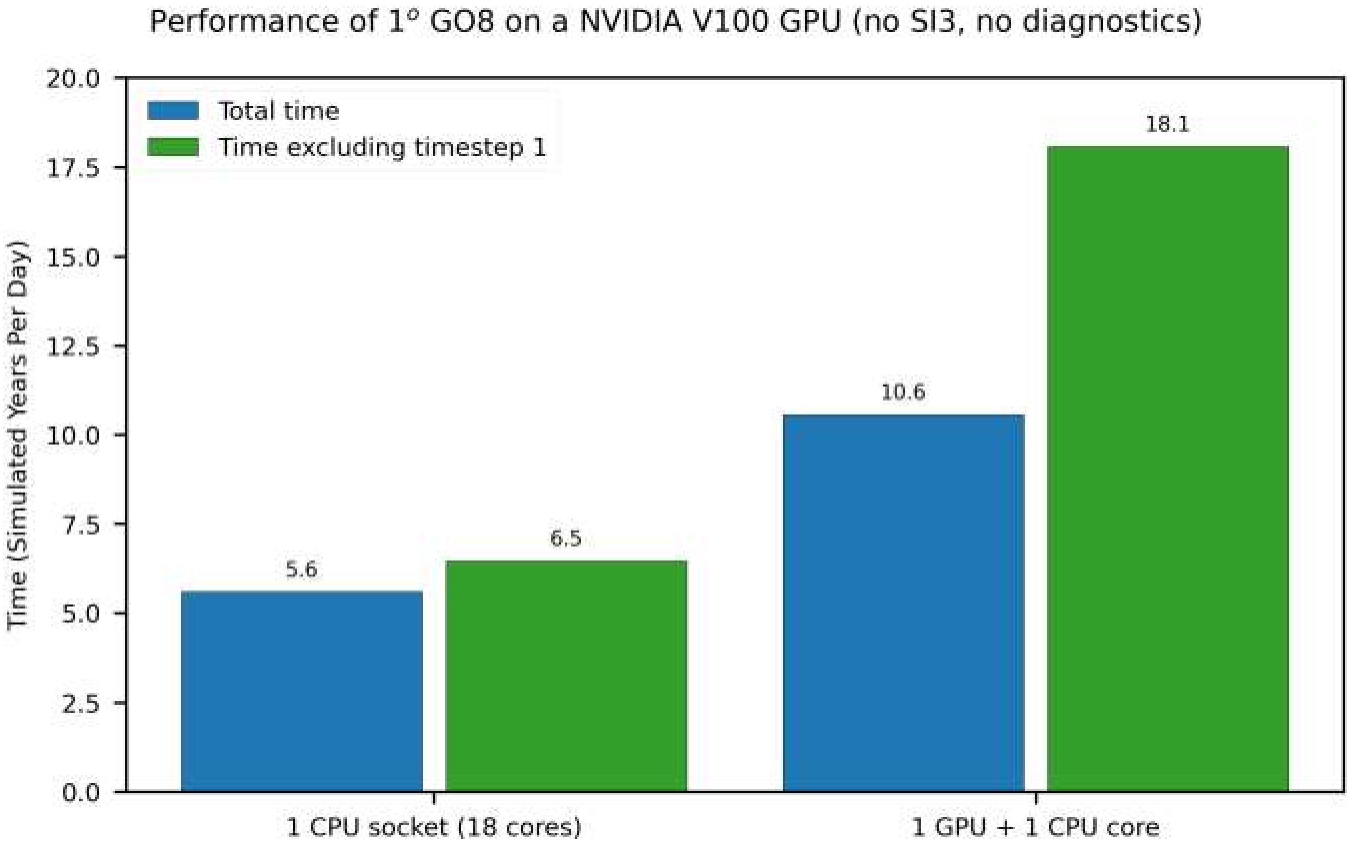

- PSyclone used to transform codes for running on GPUs
  - As of NEMO 5, PSyclone preprocessing step is supported by NEMO's native build system
- Ocean-only configuration of GO8 (NEMO 4.0.2):
  - Runs **~3.4x faster** on a GPU compared to a CPU (1° grid)
  - Runs on 90 GPUs with performance equivalent to 270 CPUs (1/12° grid)
- For reference configurations (GOSI9, GOSI10)
  - 1° GO8 configuration of NEMO 4.0.4 currently runs **~2.8x faster** on a GPU than on a CPU

Pull-through of UKCA stand-alone capability and adaptation of code for GPU-based and other HPC systems

Joseph Abram, Neal Carr, Mohit Dalvi, John Hemmings, Alan Hewitt, Adrian Hill, **Fiona O'Connor**, Joseph Wallwork (Met Office)

Luke Abraham, Scott Archer-Nicholls, Alex Archibald, Sue Cowen, Kiril Dichev, Chris Edsall, Konstantinos Kokalis, Catherine Pitt, Arjen Tamerus (Cambridge)

Chris Dearden, Rupert Ford, Andy Porter, Sergi Siso (STFC)

Seth Camp, Wayne Gaudin (Nvidia)

The UKCA sub-model

United Kingdom Chemistry and Aerosols (www.ukca.ac.uk)

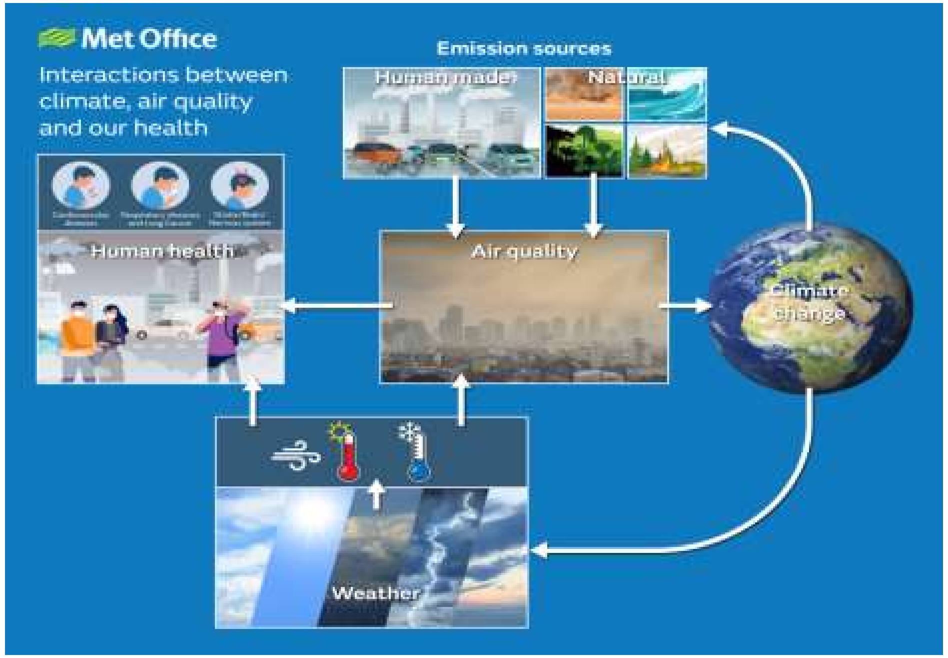

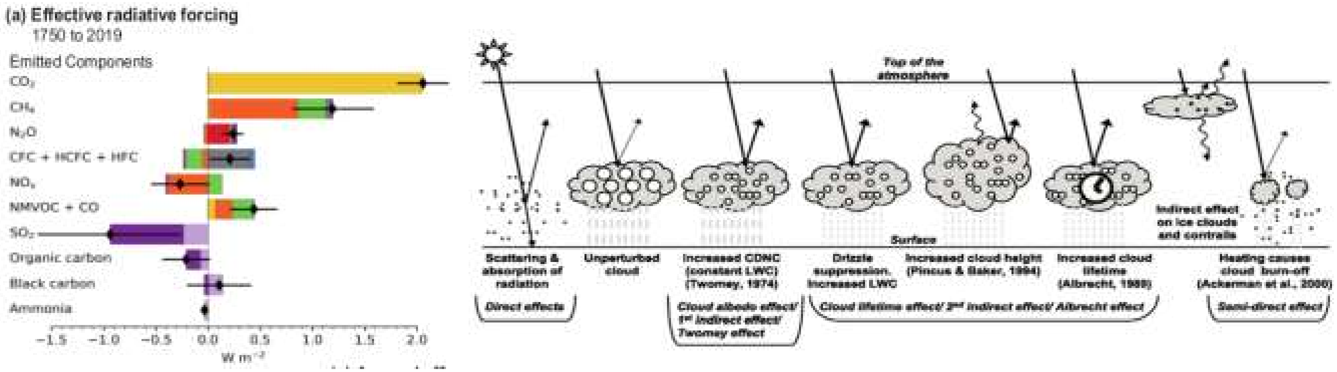

Community Modeling Framework

National Centre for Atmospheric Science · University of Cambridge · University of Leeds · University of Reading · University of Oxford · NIWA · CSIRO · University of Exeter · The University of Edinburgh

and many others ....

- More than 100 users/ developers
- Core component of the UK Earth System Model → IPCC AR6 report
- Useful for assessing impacts of Climate Policies (via emissions)
- Essential to continue onto the Next Generation Modeling system (LFRic/Momentum)
- Around 165,000 lines of Fortran across 297 files within `src/atmosphere/UKCA`

ExCALIBUR

ExCALIBUR UKCA Projects

Phase 1 (2020-2022)

- Development of UKCA Box model
- (with Nvidia): offloading core ASAD solver ‘mini-app’ to GPU -> 330x faster

Phase 2 (2022-2024)

- Integrating ASAD mini-app GPU developments
- Feasibility of extending to other UKCA components
- Feasibility of using Psyclone code transformation utility for automated GPU offloading

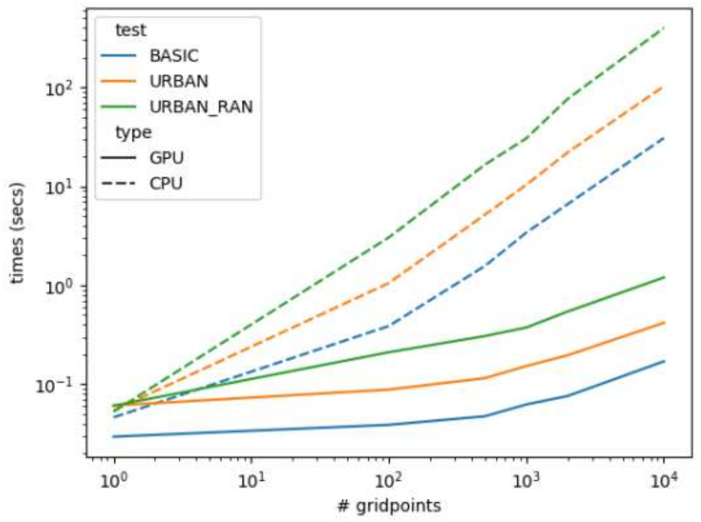

Figure 2: Runtime of ASAD with or without OpenACC GPU offloading

ExCALIBUR UKCA Projects

Science and Technology Facilities Council

PSyclone

Source-to-source Fortran compiler designed to programmatically optimise, parallelise and instrument HPC applications via user-provided transformation scripts

Manual Port

```
!$acc kernels if(ln_gpu) async(3)
!$acc loop gang vector independent
DO jl = 1, kl
  IF (l_convergence_mask(jl)) THEN
    !$acc loop seq
    DO js = 1, jpspec
      prod(jl,js) = 0.0
      slos(jl,js) = 0.0
    END DO
  END IF
END DO
!$acc end kernels

! 2. Compute reaction rate of change (kA or kAB
!    ------- -------- ---- -- ------ --- -- ---

! 2.1 Reactions with one reactant.
!$acc kernels if(ln_gpu) async(4)
!$acc loop gang vector independent
DO jl = 1, kl
  IF (l_convergence_mask(jl)) THEN
    !$acc loop seq
    DO jr = 1, nuni ! Reversed loop order
      ir1 = nspi(jr,1)
      prk(jl,jr) = rk(jl,jr) * y(jl,ir1)
    END DO
  END IF
END DO
!$acc end kernels
```

Psyclone transformed

```
!$acc kernels
!$acc loop gang vector independent
do jl = 1, kl, 1
  if (l_convergence_mask(jl)) then
    !$acc loop independent
    do js = 1, jpspec, 1
      prod(jl,js) = 0.0
      slos(jl,js) = 0.0
    enddo
  end if
enddo
!$acc loop gang vector independent
do jl = 1, kl, 1
  if (l_convergence_mask(jl)) then
    !$acc loop seq
    do jr = 1, nuni, 1
      ir1 = nspi(jr,1)
      prk(jl,jr) = rk(jl,jr) * y(jl,ir1)
    enddo
  end if
enddo
!$acc end kernels
iunip1 = nuni + 1
!$acc kernels
!$acc loop gang vector independent
do jl = 1, kl, 1
  if (l_convergence_mask(jl)) then
```

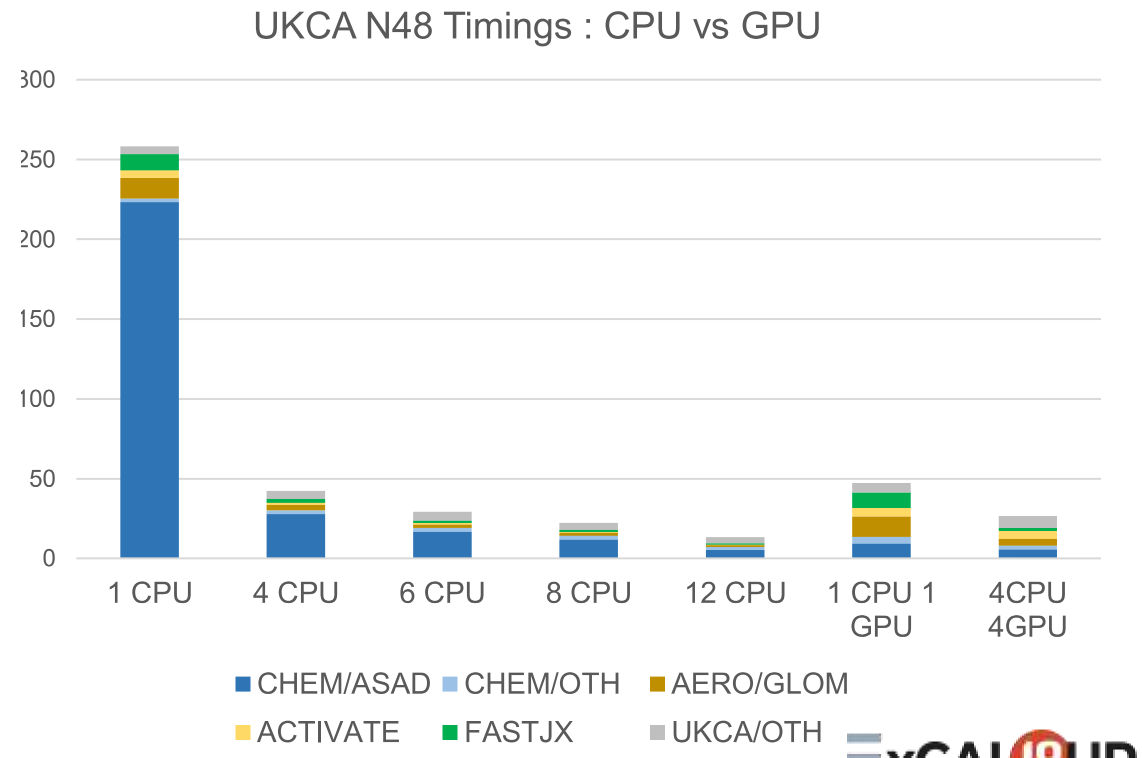

Cross-Cutting Work Packages

Common Approaches and Solutions

Potential disruptors

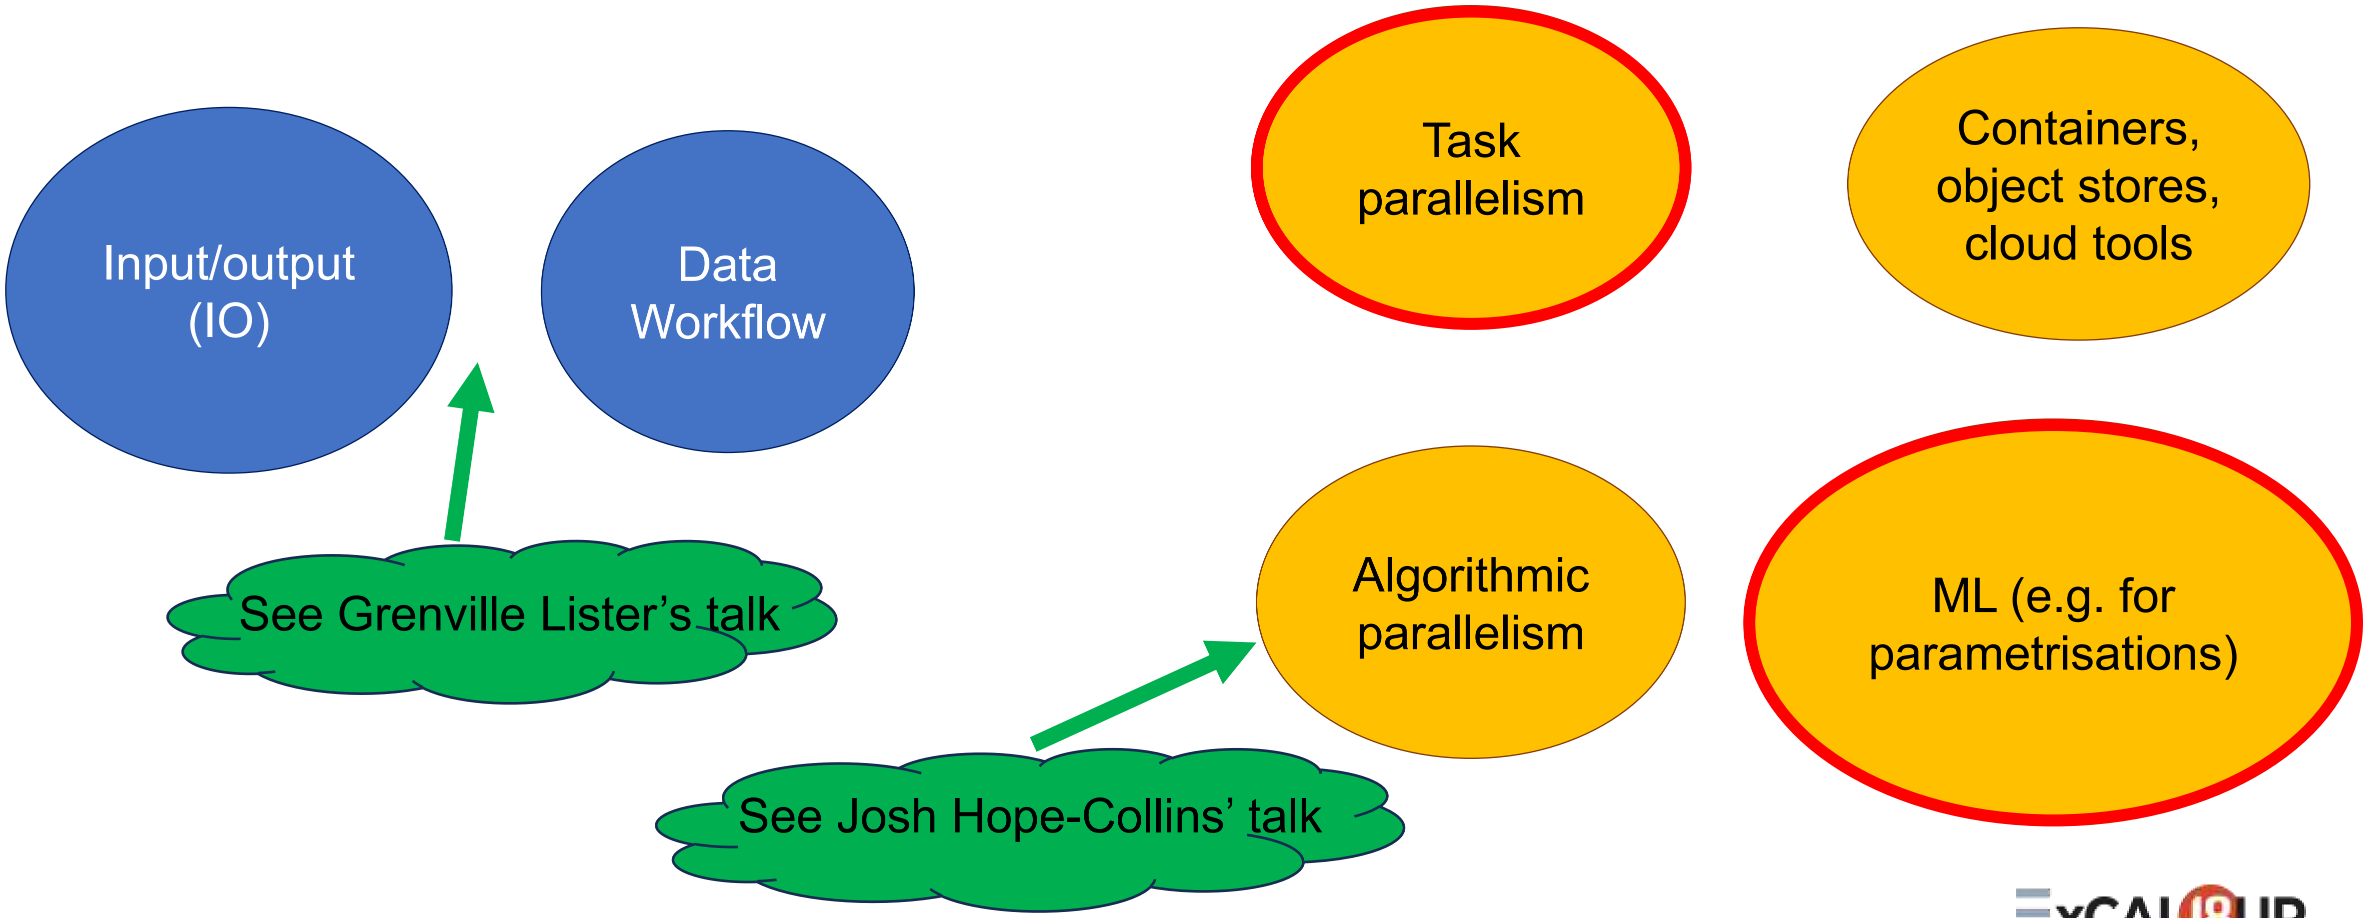

Task Parallelism

Mladen Ivkovic, Alan Real, Adam Tuft, Marion Weinzierl, **Tobias Weinzierl** (Durham)

Aidan Chalk, Rupert Ford (STFC)

Task parallelism – ‘taskification’

What constitutes a task?

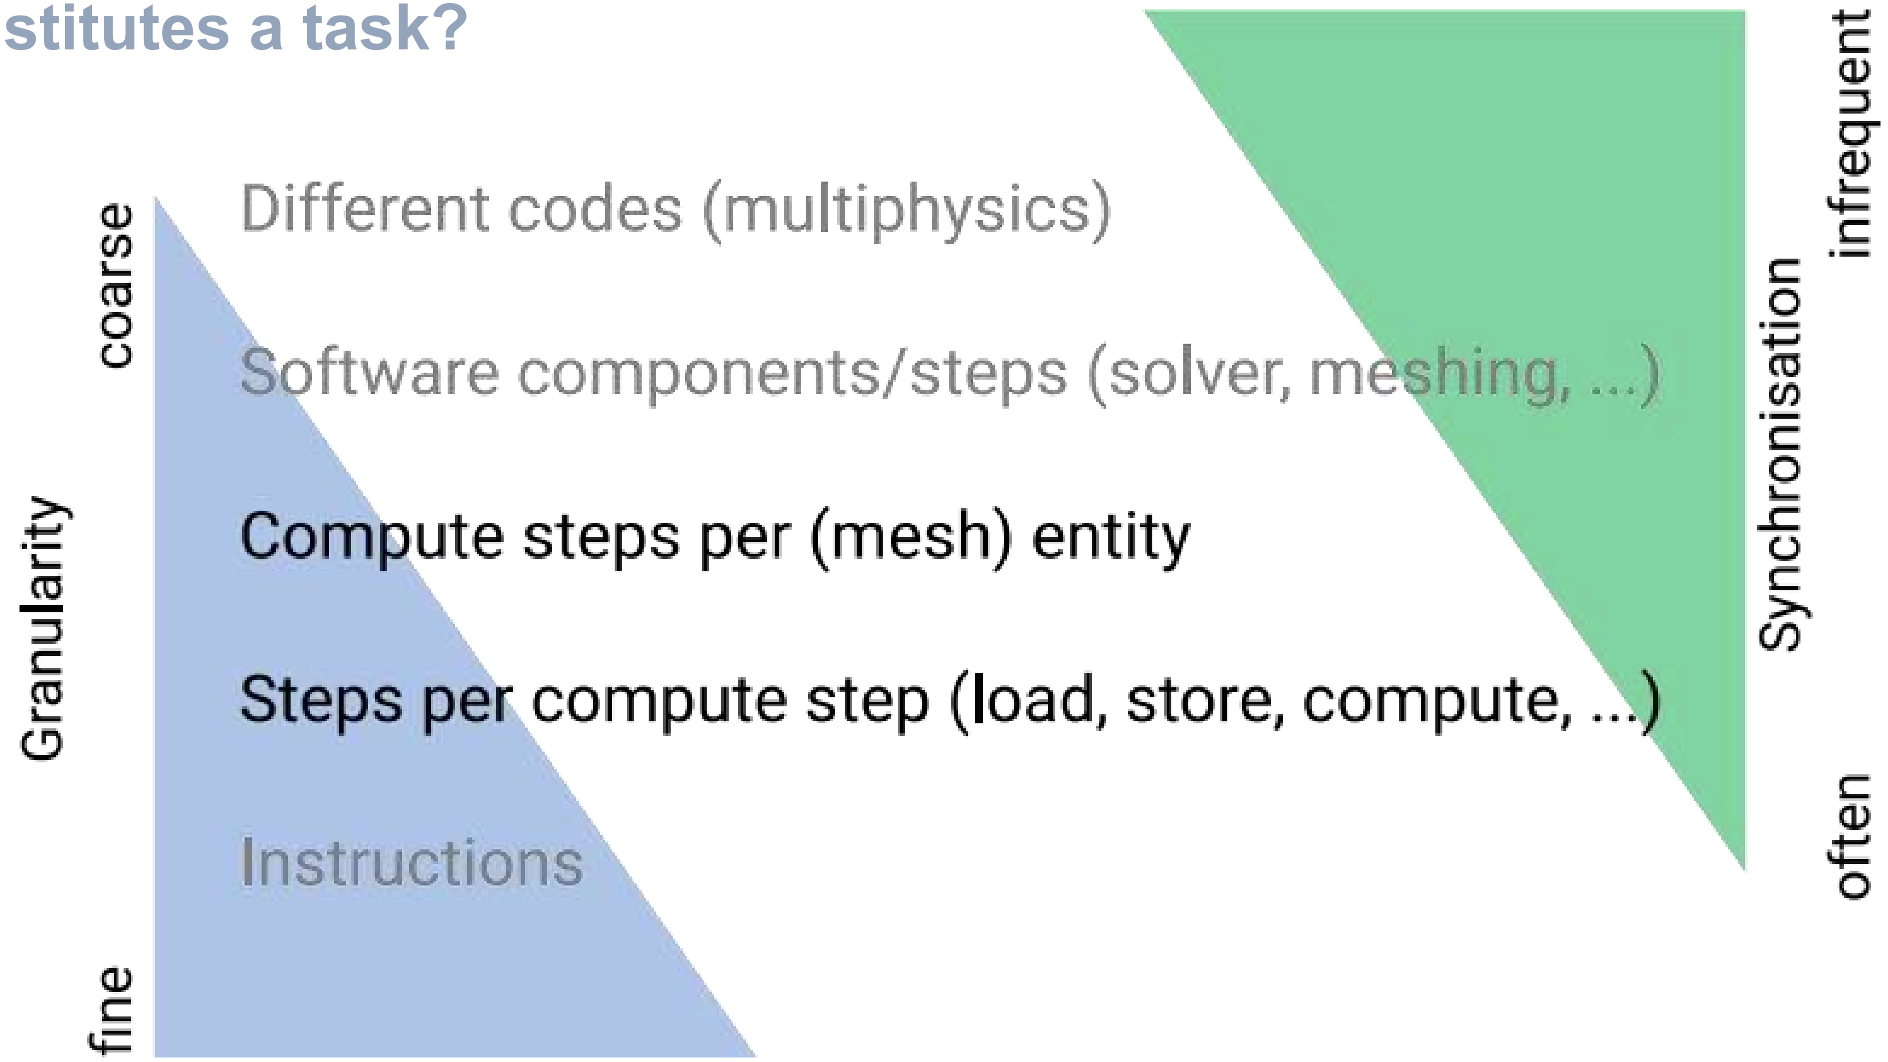

Otter Tracing

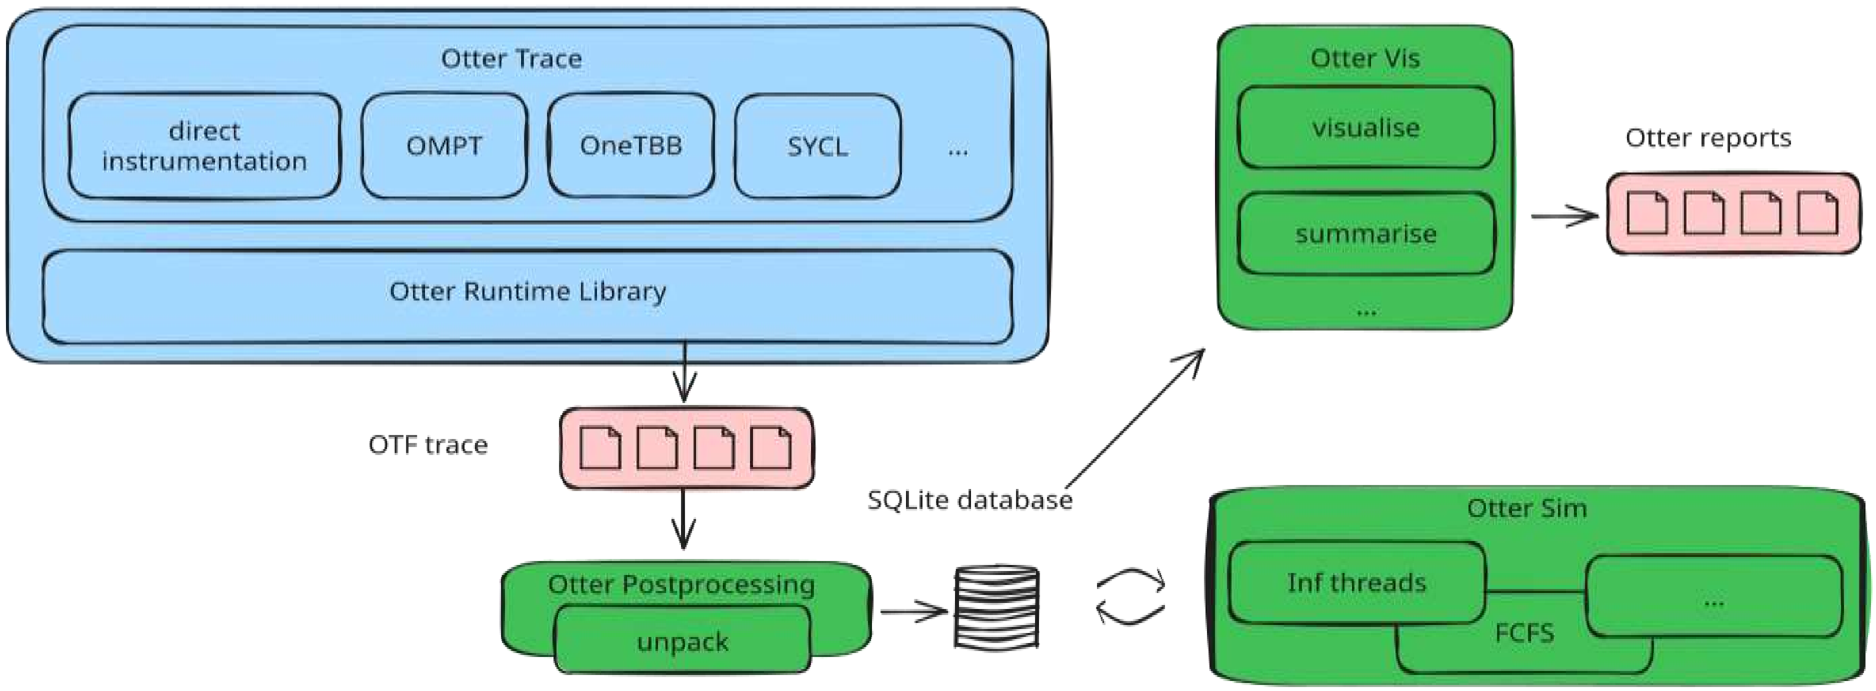

- Task execution patterns studied in Otter tool suite
- API for annotating serial or parallelised code to identify where task & loop parallelism could be introduced for performance gain
- Trace task creation, execution & synchronisation events
- All data recorded to modified OTF2 trace file
- Re-play logical task graph with various simulated schedulers
- Identify critical path
- Predict performance gain

A. Tuft. T. Weinzierl & M. Klemm, Detrimental task execution patterns in mainstream OpenMP runtimes

Machine Learning (e.g. for parametrizations)

Elena Breitmoser, Alistair Grant, Johnny Hay, **Amy Krause**, Moritz Linkmann, Jacob Page, Anna Roubickova, Iakovos Panourgias (EPCC & Edinburgh)

SiMLInt: Simulation and Machine Learning Integration

EPCC and The School of Mathematics, The University of Edinburgh

Motivation

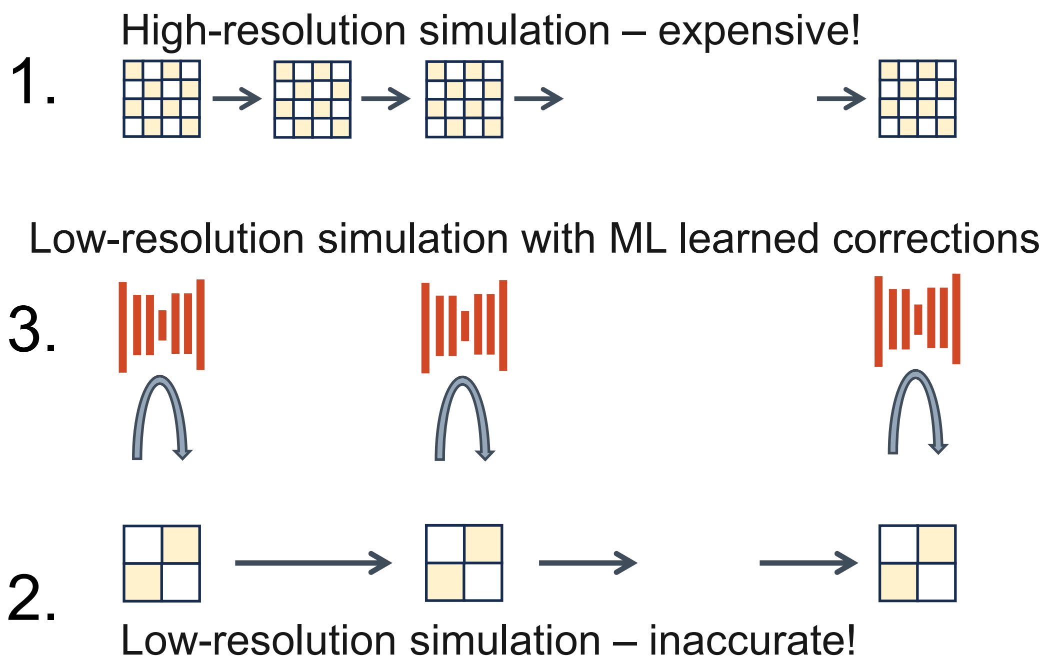

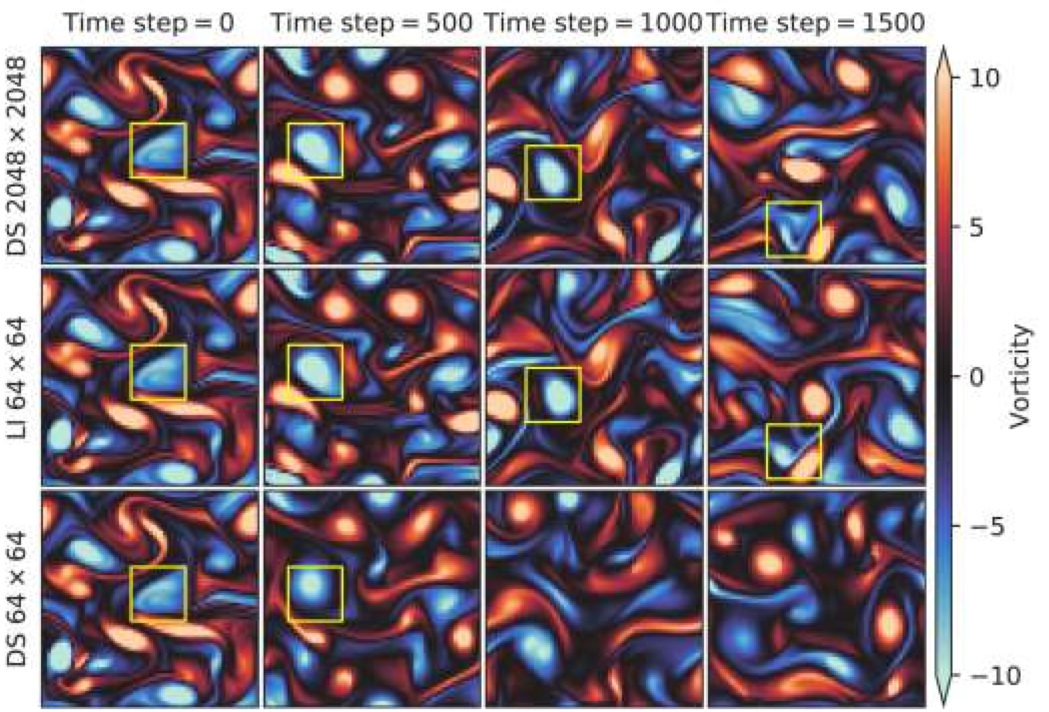

Machine learning–accelerated computational fluid dynamics, Kochkov et al., PNAS 2021 https://doi.org/10.1073/pnas.2101784118

Aim: Explore best practices for embedding Machine Learning into large scale numerical simulations

SiMLInt: Simulation and Machine Learning Integration

EPCC and The School of Mathematics, The University of Edinburgh

Outcomes (Technical Findings)

- Minimum-change ML approach: Learned Correction
- Infrastructure: SmartSim [https://github.com/CrayLabs/SmartSim]
- Use-cases: Drift-wave turbulence in tokamak plasma, Lid-driven cavity flow

Outcomes (Knowledge Exchange)

- Set of instructions to reproduce the ML-bootstrapped DNS pipeline on HPC machines [https://epcced.github.io/SiMLInt/]
- Training materials and in-person workshops
- Links with BOUT++ community and SmartSim developers
- PhD intern, MSc student

People!

A reminder of Mark Parson et al's:

UK Research and Innovation

ExCALIBUR

UK Atomic Energy Authority

Met Office

Research Software Engineer Knowledge Integration Landscape Review

The Research Software Engineer (RSE) Knowledge Integration Landscape Review was commissioned by the ExCALIBUR (**Ex**ascale **C**omputing **AL**gorithms & **I**nfrastructures **B**enefitting **U**K **R**esearch) programme to provide a comprehensive Landscape Review focussed on the concept of the Research Software Engineer and their role within UK Science in preparation for the arrival of Exascale Supercomputers.

The document written by members of the ExCALIBUR Programme community details:

- The skills required by RSEs in HPC
- The future training needs of RSEs
- Challenges faced in developing these skills and growing the number of RSEs in the UK with a specific focus on HPC
- The importance of establishing a career path for RSEs that does not rely on the conventional academic metrics.

ExCALIBUR is an exciting UK programme funded through the UK Government's Strategic Priorities Fund. This is a 5-year, £46m programme of activities led by the Met Office and UKRI.

The vital role of RSEs in the UK computational science community has grown markedly over the last decade.The need for Research Software Engineers focussed on HPC is wide ranging; from cosmology to digital archiving, HPC is used at all scales of scientific discovery. Indeed, with the field of HPC undergoing an explosion in technologies and becoming a cornerstone of research in many diverse fields, the need for RSEs with HPC experience will continue to grow.

The ExCALIBUR Research Software Engineer Knowledge Integration Landscape Review (DOI:10.5281/zenodo.4986062) can be found at https://zenodo.org/communities/excalibur-spf

Research Software Engineer Knowledge Integration Landscape Review

The report delivered the following recommendations on the future training and skills requirements for RSEs, identified a number of training and skills gaps in the current provision and discussed the issues surrounding the long-term career development of RSEs:

- UKRI should continue to invest in the development of Research Software Engineering in the UK.
- As we grow the overall number of RSE staff in Universities, National Laboratories, and other research organisations we should also grow the number of such staff with specific High-Performance Computing skills.
- An ExCALIBUR Training Programme for Research Software Engineers who want to focus on High Performance Computing should be established.
- A long-term training strategy should be developed to both train the next generation of RSEs with HPC skills and also fill the training and skills gaps identified in this report.
- A variety of different training models should be adopted – including postgraduate study, workshops, hackathons and bootcamps. Wherever possible training should be made available as online training as well as in face-to-face training opportunities.
- Clear career paths for Research Software Engineers, and funding opportunities for software development allowing them to apply and develop skills, are crucially important to ensure that, once trained, the knowledge they have gained stays in the research sector and grows over time. The contribution of software engineering needs to be recognised in university recruitment and promotion procedures.
- UKRI should ensure that it supports the message that Research Software Engineers are a highly valued resource at Universities, National Laboratories, and other research organisations by providing clear guidance for inclusion of RSEs on grants.
- Greater collaboration and transfer of skills by RSEs in both directions between the academic and industrial research sectors should be encouraged, particularly from industry to academia.
- The UKRI Exascale Supercomputer Project's software programme should ensure that it encompasses a variety of different types of software activity and ensure they are contributing to developing the RSE community with HPC skills in the UK.

Expertise

“Investment in software means investment in people.”

People!

- Met Office's equivalent of RSEs = Scientific Software Engineers (SSEs)
- We have O(100) SSEs but they are mostly 'embedded', in long established, specialist teams
- Lack rapid deployability, breadth of experience, and adaptability to new skills etc.

- ExCALIBUR enabled us to establish the 'ExCALIBUR Team' – a pool of deployable SSEs

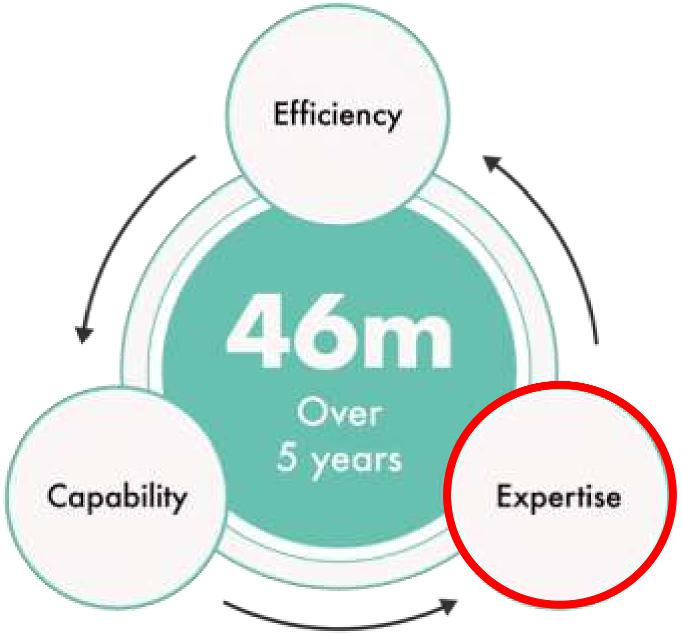

- Varied in size but ~ 6 people
- Challenging
  - Team was spun up during Covid
  - Initially not clear 'home' for the team
  - Have to actively ensure we keep the 'deployable' bit
- But it is now well established
- Keeping the name!

How are we doing with that sword?

- Huge pleasure to be involved in ExCALIBUR
  - EPSRC
  - UKAEA
  - New, reinvigorated, strengthened relationships
- Impressive breadth of scope and institutions
- Achieved huge amounts
- Timely!

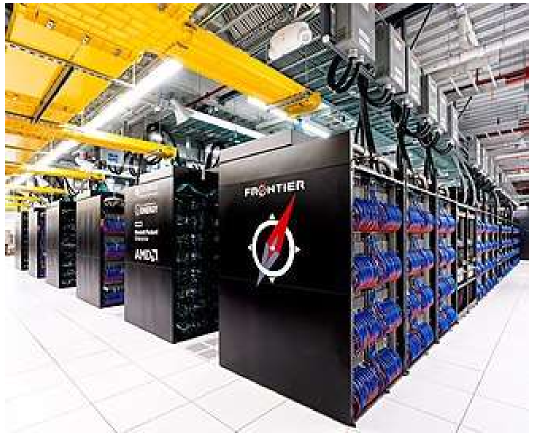

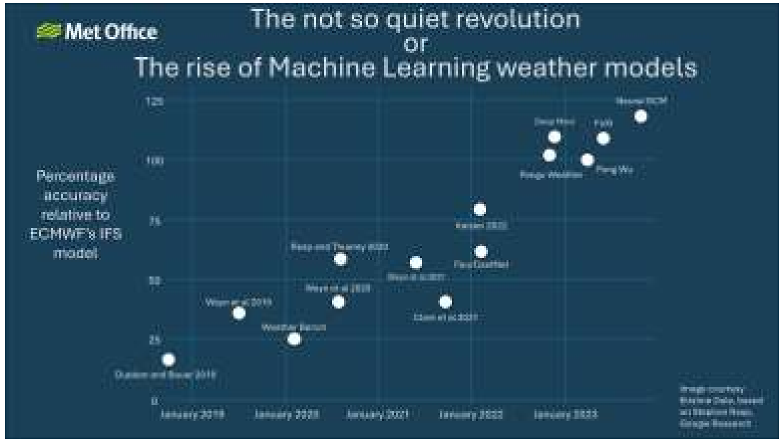

Thank you! Questions?

See https://excalibur.ac.uk/ for more