Bridging Communities: Technical Concerns for Building Integrated Environmental Models

Bryan Lawrence (@bnlawrence)
(and others, see next page)

National Centre for Atmospheric Science
University of Reading
Science and Technology Facilities Council

National Centre for
Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk

This presentation based on a draft paper in preparation

Bridging Communities: Technical Concerns for Integrating Environmental Models

B.N. Lawrence[1,2,3], V.Balaji[4,5], V. Bell[6], M. Carter[7], C. DeLuca[8], S. Easterbrook[9], R. Ford[10], and A. Hughes[11]

[1] Department of Meteorology, University of Reading, U.K.
[2] National Centre for Atmospheric Science (NCAS), Natural Environment Research Council, U.K.
[3] Centre for Environmental Data Archival, STFC Rutherford Appleton Laboratory, U.K.
[4] NOAA Geophysical Fluid Dynamics Laboratory
[5] Princeton University, Princeton, NJ
[6] Centre of Ecology and Hydrology, Natural Environment Research Council U.K.
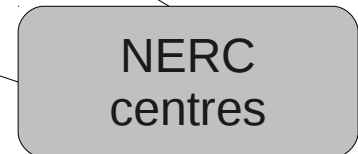[7] Met Office, Hadley Centre, Exeter, UK
[8] NOAA Cooperative Institute for Research in Environmental Sciences, Boulder, CO, USA
[9] Department of Computer Science, University of Toronto, Canada
[10] Scientific Computing Department, STFC Daresbury Laboratory, U.K.
[11] British Geological Survey, Natural Environment Research Council, U.K.

NERC centres

informed by Annette Osprey and Luke Abrahams

VERY MUCH WORK IN PROGRESS

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk

National Centre for Atmospheric Science

A "virtual centre" with five divisions and a directorate

Uni of York
NCAS Air Quality Lead

Two Facility Divisions

Models and Data (M&D)
Observational Facilities

M&D consists of
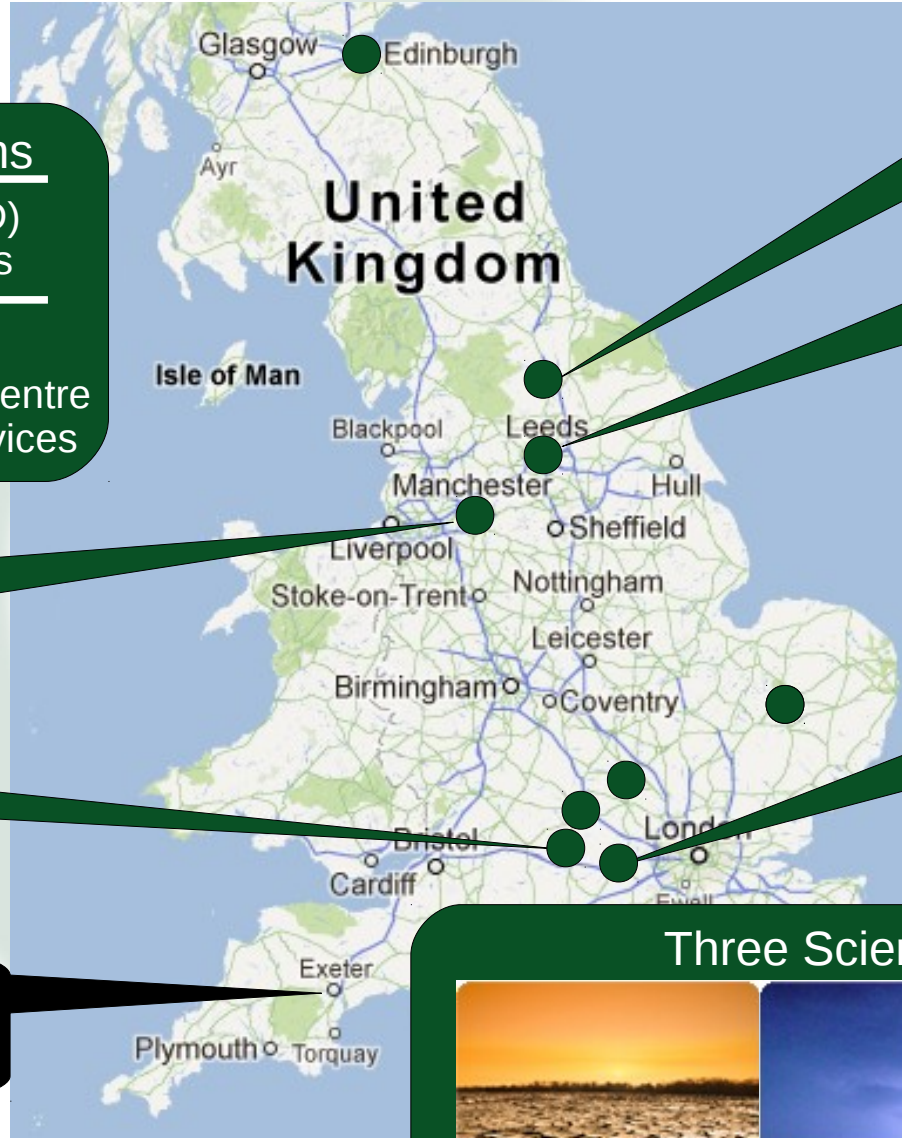British Atmospheric Data Centre
Computational Model Services

Uni of Leeds
NCAS Directorate
NCAS Weather Lead

NATURAL ENVIRONMENT RESEARCH COUNCIL

Uni of Manchester
NCAS Observations Lead

Uni of Reading
NCAS Climate Lead
NCAS M&D Lead

STFC Rutherford Lab
NCAS BADC

Three Science Divisions

Partner Institution
Met Office

CLIMATE   WEATHER   AIR QUALITY

Not all NCAS sites shown!

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk

Outline

Part One: Drivers and Context:

- Science

- Communities

- ~~Exascale~~

Segue to Part 2 via Coupling Technology

Part Two: Use Cases

- UKCA: the UK Chemistry Atmosphere

- (Modelling water flow in the Thames Basin)

There are no conclusions. That's why I'm here.

– There is a bit of a summary :-)

National Centre for
Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk

The first and second problems of climate change science (from Rowan Sutton)

Pre AR4: *Are human activities changing climate?*
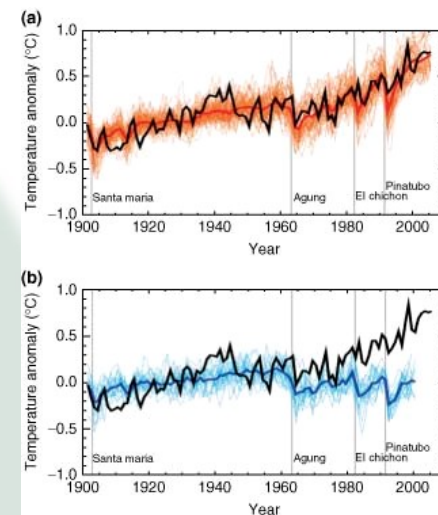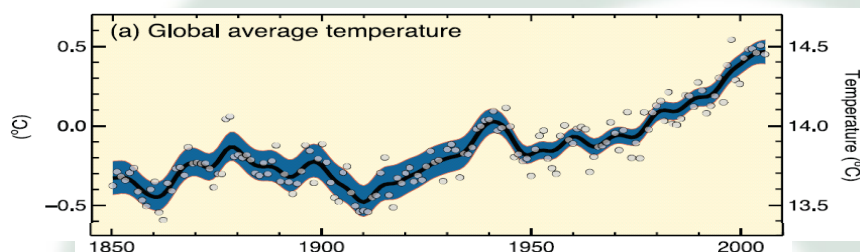

Post AR4:

- *What is the signal of anthropogenic climate change on the regional and local scales that really matter to individuals, economies and societies?*

- *What does/will climate change look like where I live? Temperature not the be-all and end-all of answering that question!*

A far more difficult ***grand challenge research and development*** problem

Greatly enhanced national and international collaboration *and strategy* essential
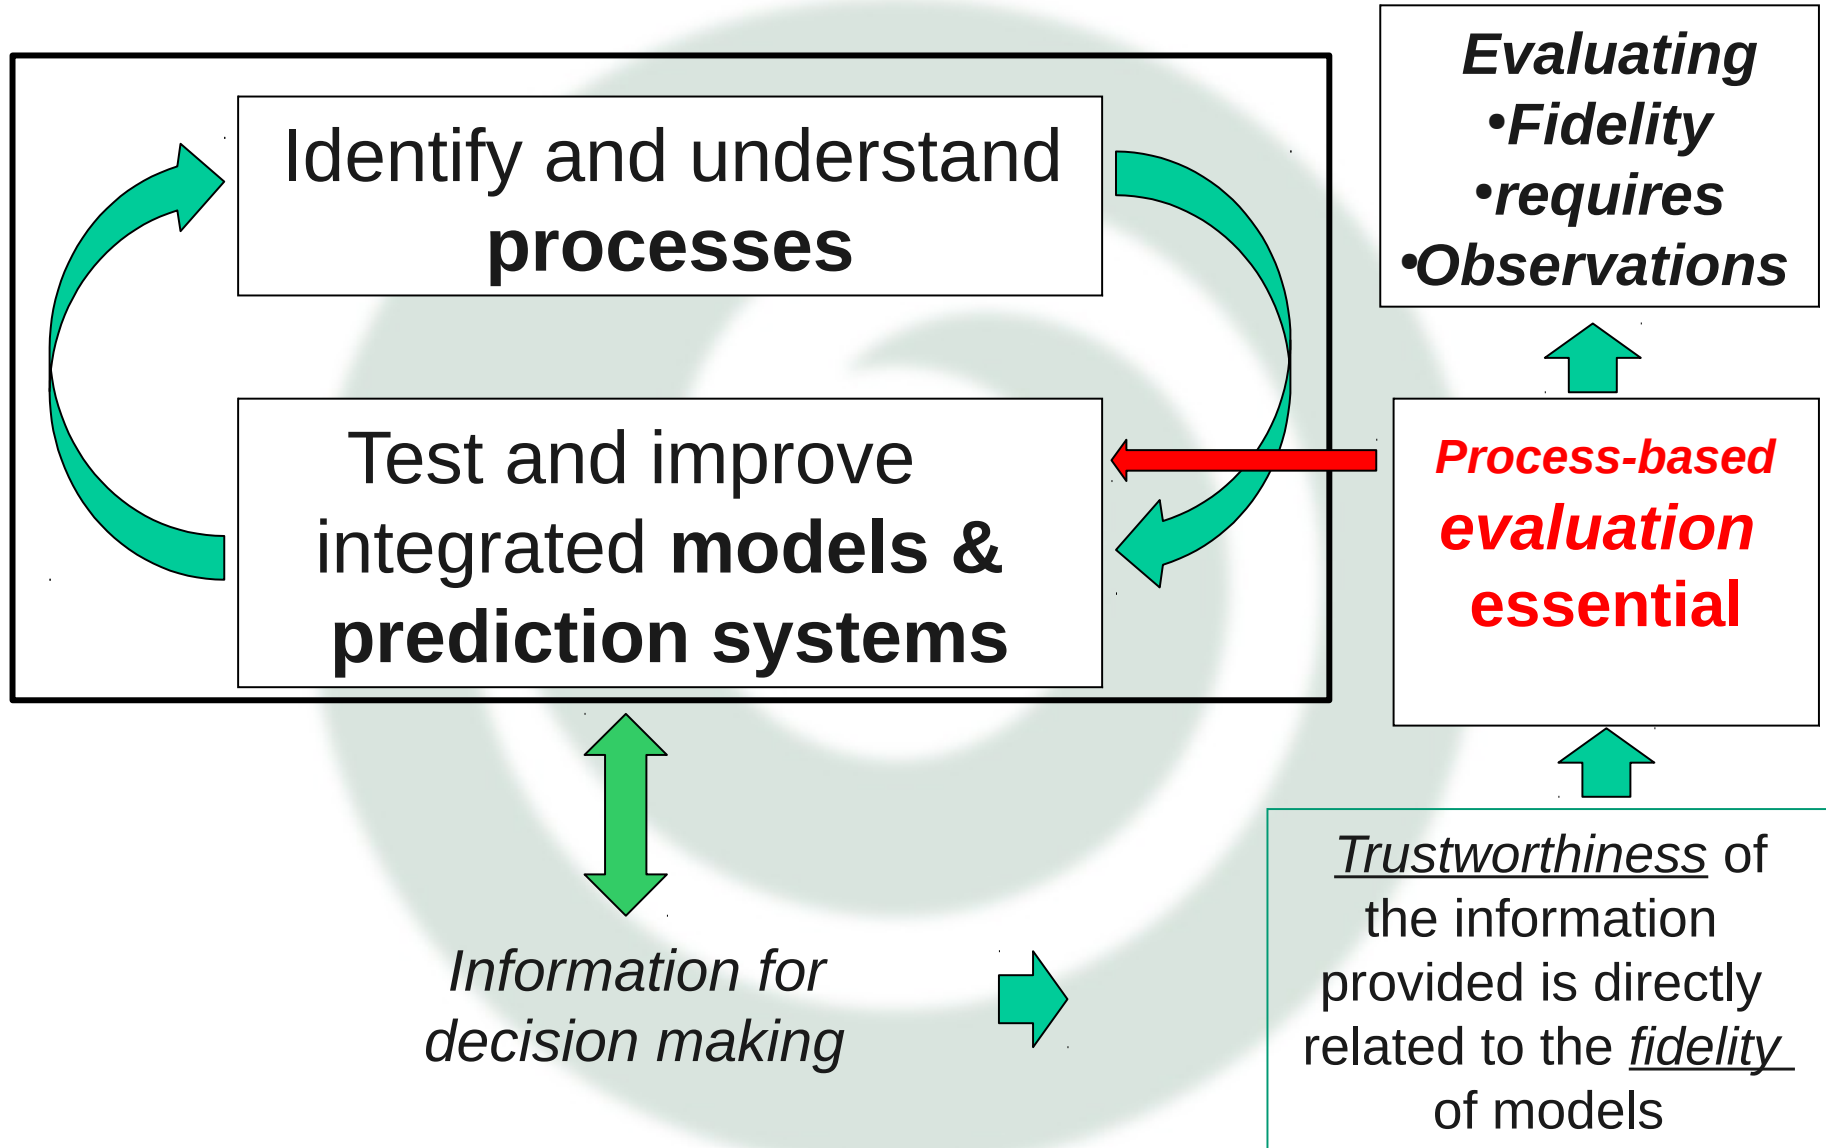
Many, many processes, so many, many communities *interacting.*
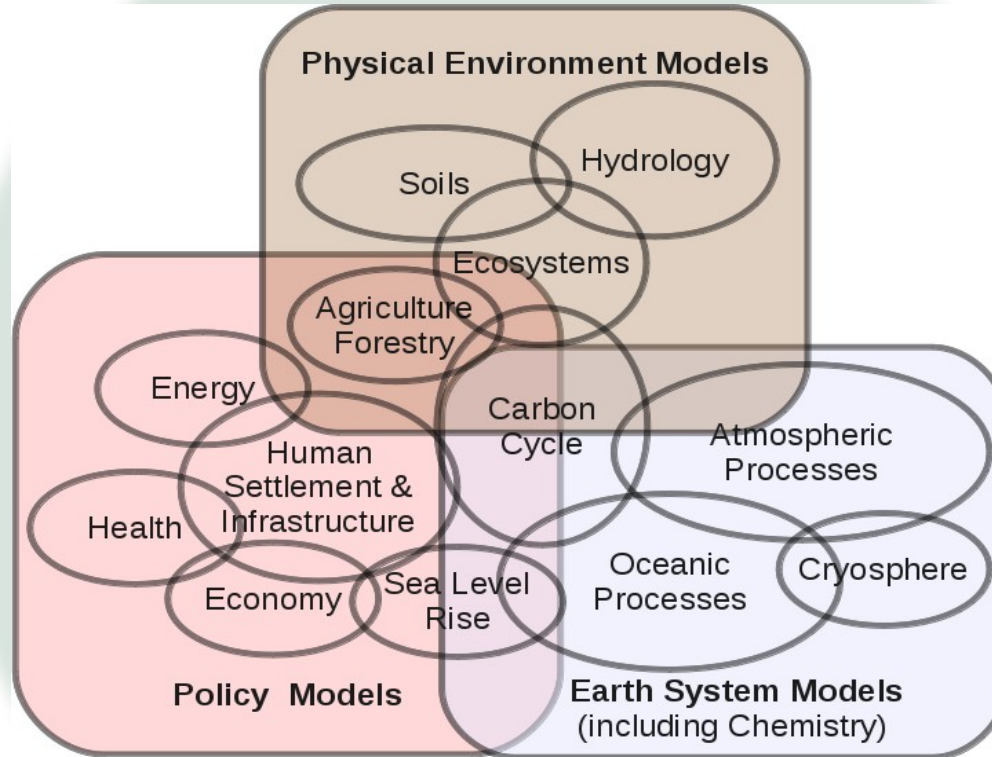
We can't add all these processes into our models, but these communities will interact with common models *and* via data coupling.

(Figure adapted from Moss et al., 2010).

Model Validation is Crucial to the Process

Identify and understand **processes**

Test and improve integrated **models & prediction systems**

Evaluating
- *Fidelity*
- *requires*
- *Observations*

Process-based **evaluation essential**

Information for decision making

Trustworthiness of the information provided is directly related to the _fidelity_ of models

Adapted from R.Sutton, June 2012

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk

Process-Based Evaluation and Coupling

Evaluation at the process level is hard to do:

➜ Getting a "host model" to follow a real situational trajectory for comparison is difficult (even with nudging, assimilation reanalysis, et al).

➜ Using a hierarchy of models is a key part of the process, understanding the impacts of lack of resolution, and removal of processes as we go to larger scale.

➜ Interactions between scales and between model components can be important.

➜ Interactions between the communities responsible for understanding these processes VERY important.

➜ Understanding, and using other peoples code, is an important part of the scientific understanding!

➜ Divergence of models which share the same name hinders this methodology!

Towards a national strategy for Earth System Modelling (adapted from Rowan Sutton version)

Global Earth System Modelling

UKESM1

Major focus on Biogeochemical cycles and Ice sheets

Global & Regional Modelling of the Physical Climate System ("Atmosphere-Ocean-Sea Ice-Land")

Includes pushing the frontiers of resolution & physical parameterisations

Integrated Environmental Modelling

Focus on local-regional scales

Strategy under development

Coupling technologies, plural

Sea & Land Ice (CICE + GLIMMER)	**Ocean** (NEMO + Biogeochem)	**Atmosphere** (HadGEM3 + **UKCA**)	**Land** (JULES inc veg., C & N cycles)	**local-regional environmental process models**

Process-based evaluation (across scales)

Technical community support (for communities, plural)

Segue from the Science to the Technology

We begin with communities, and their models.

We progress to interacting communities, and interacting models.

- Generally one community modelling paradigm dominates how that is done! The "top-model", often an atmosphere dynamical core (or it's driver) …
- Almost immediately we start to see a code divergence, as the coupled version differs from the standalone version.

We know that not all communities are going to be able to interact by direct two-way coupling via a "top-model".

- This simply doesn't scale, socially, or technically.
- But we don't always know what things we can neglect in terms of feedback. We need to experiment.

Two use cases to consider:

- Can we mitigate against that code divergence?
- Can we simplify the interfaces to support experimentation?

From Evaluation, via Versions, to Confidence

Most "sub-"models exist as standalone models first.

Most "sub-"models still exist as standalone models.

Most "sub-" models exist with multiple variants/versions in use, with **different variants** "coupled" into ESMs than are used standalone.

Different levels of evaluation and confidence may exist for those different variants.

How confident are we in the scientific traceability?

A representative sample of coupling technologies

1) Direct/Bespoke

2) ESMF
- – A framework

3) OASIS
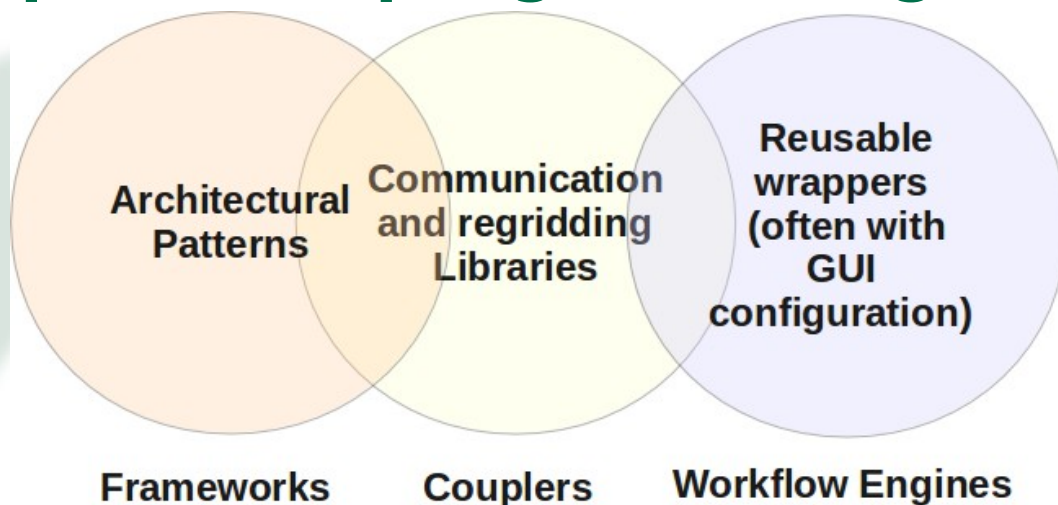- – A coupler

4) OpenMI
- – A limited framework

5) CSDMS
- – minimally intrusive framework (Basic Model Interface) + library implementation of the Common Component Architecture (CCA)

6) Kepler
- – Workflow Management (coupling via files)

7) *BFG (Bespoke Framework Generator)*
- – *Metadata driven coupling a la carte*

Architectural Patterns

Communication and regridding Libraries

Reusable wrappers (often with GUI configuration)

Frameworks   Couplers   Workflow Engines

Different implications for code and workflow

Questions to consider: How usable (and reusable) are these approaches? In particular, how intrusive/invasive is the approach?

If the methodology is difficult to approach, intellectually, or in terms of the implementation, it can be difficult for all communities involved in coupling to have equal knowledge & that's not good for the science!

If the methodology is intrusive, this might have real implications for the necessity for having multiple versions of the component models.

- Two forms of intrusiveness to consider:
  - the need for refactoring (changing and/or reordering code), and
  - Sheer volume of code inserted/needed/to-be-comprehended
- (Hidden dependencies on other component code and behaviour … can't be avoided, but can it be minimised?)

All of these apply to the entire work flow, not just the running model! Need to consider debugging, evaluation, post-processing data formats etc.

Coupling Requirements

1) **Two sides of the interface need to provide the right variables.**

 – And have they been modelled "sufficiently" well?

 – (This is about our scientific confidence in the individual models.)
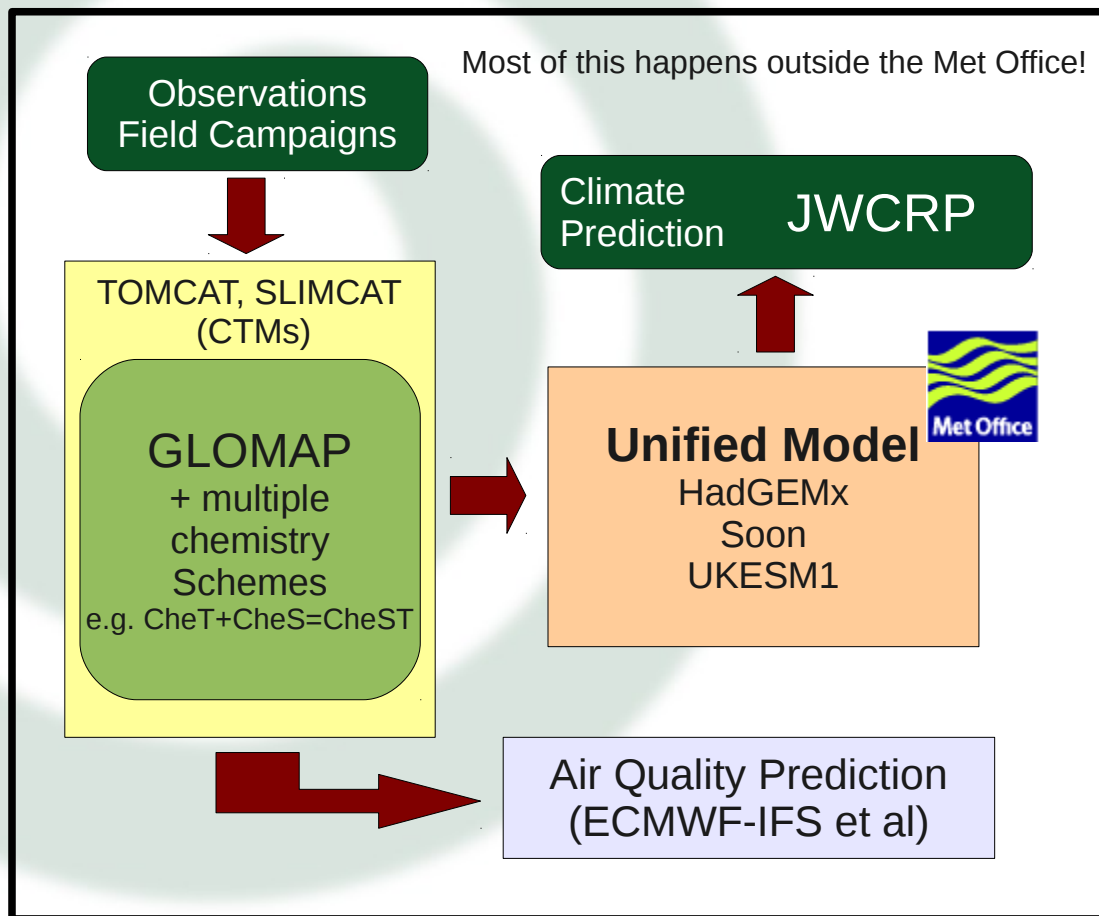
2) **Can the exchange be modified explicitly**
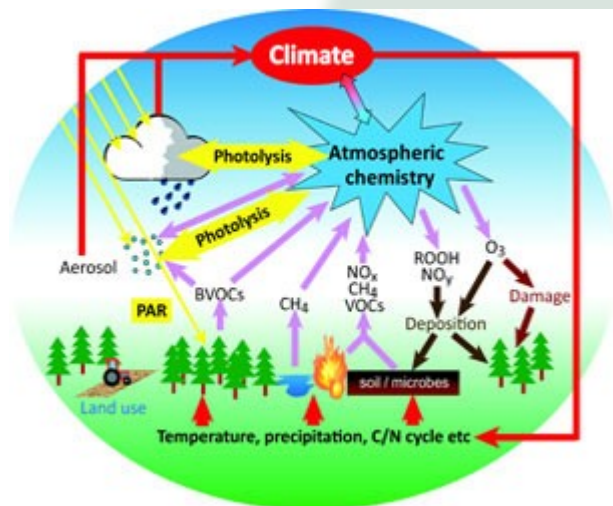
 – Solution is stable if the future state of variables in either component can be calculated from past states in the other.

 – (This is going to be problematic in the use cases I'm going to discuss! New approaches: service models!)

3) **Are the variables on the same grid.**

 – Or can they be made to be so.

 – (We know we can solve this one, but maybe not at exascale)

UKCA: The Programmatic View

(I should say that what follows is my personal opinions of UKCA, from the outside, looking in … some of these ideas have been discussed with individuals in UKCA … but I'm no expert on it … yet.)

Most of this happens outside the Met Office!

Observations Field Campaigns

Climate Prediction **JWCRP**

TOMCAT, SLIMCAT (CTMs)

GLOMAP + multiple chemistry Schemes e.g. CheT+CheS=CheST

Unified Model HadGEMx Soon UKESM1

Air Quality Prediction (ECMWF-IFS et al)

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk

UKCA: Aerosol from a science perspective
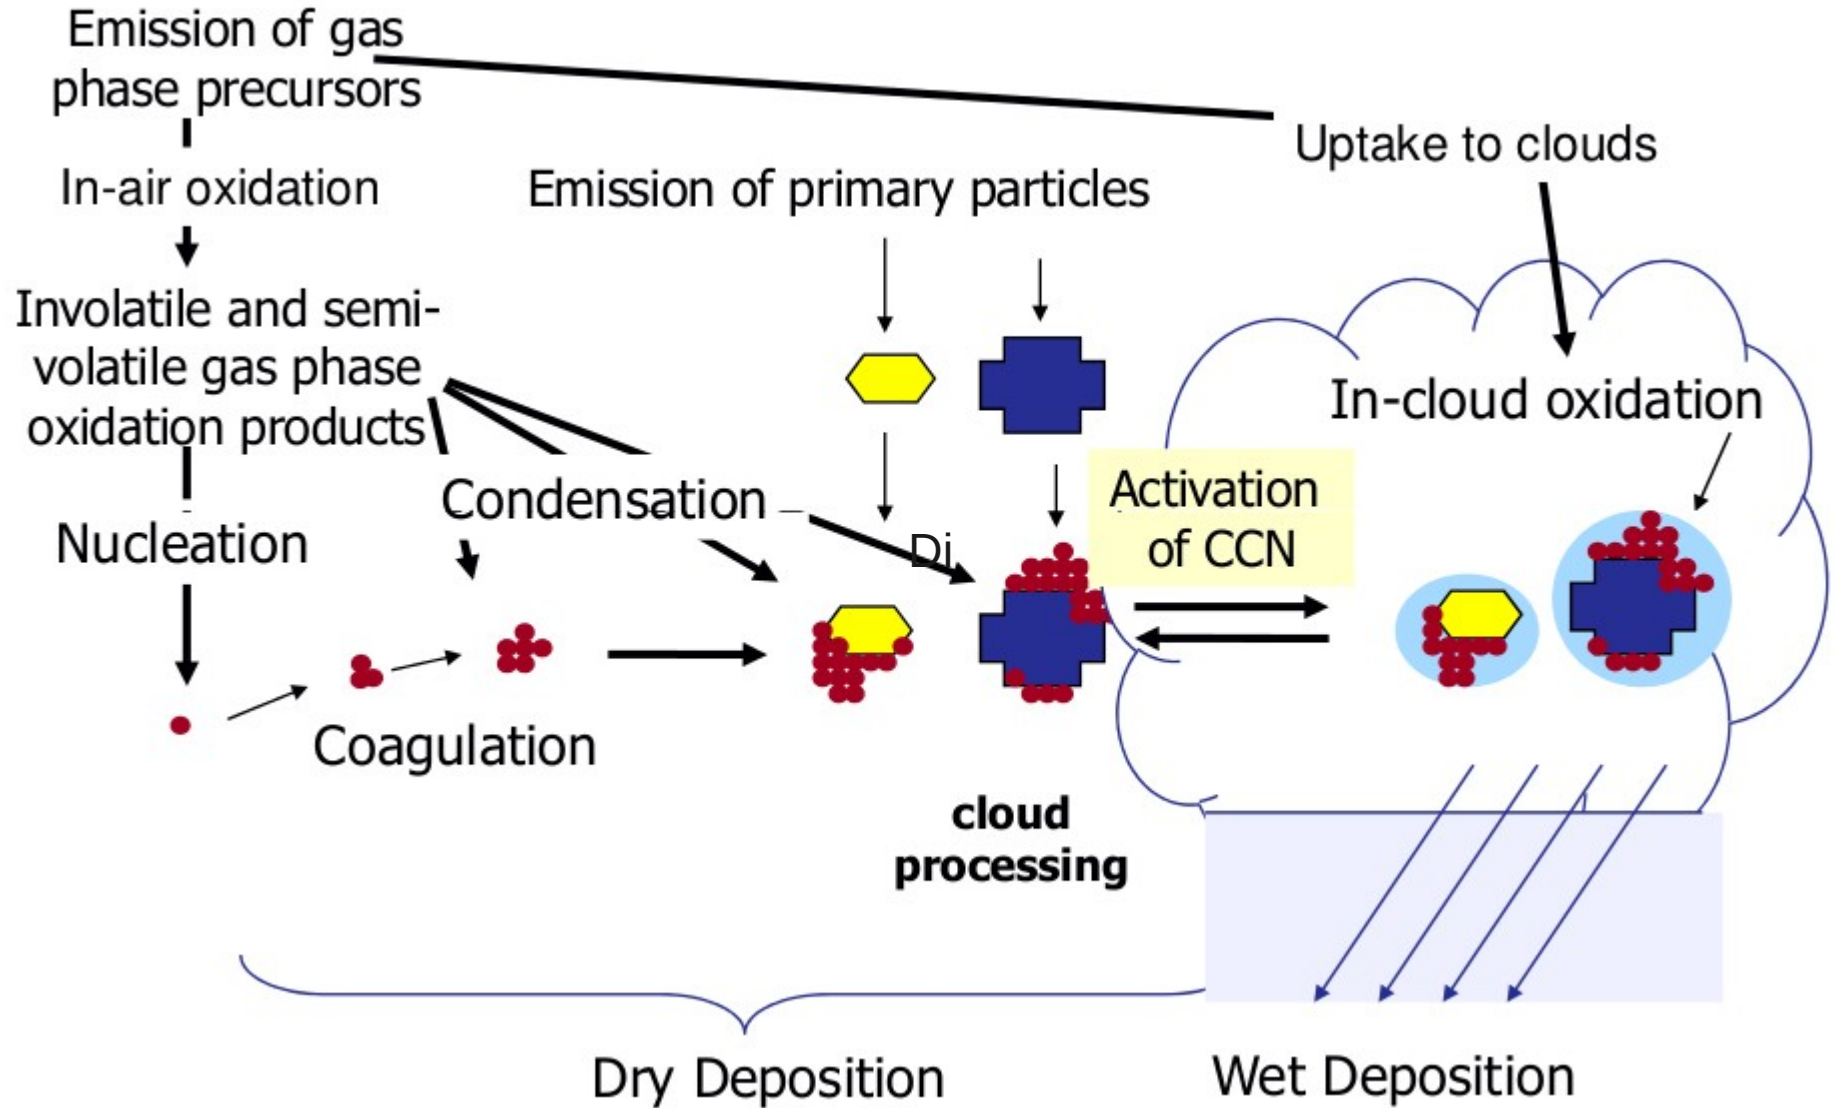
Diagram from Graham Mann (NCAS, University of Leeds)

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk

UKCA from a code perspective

Typical UM Timestep

ATM_STEP
- Init + Polar Filtering
- atm_phys1: slow physics
- SLT: advection
- atm_phys2: fast physics
- diffusion
- Helmholtz pressure solver
- updates
- classic aerosol
- energy_correction
- diagnostics

UKCA_MAIN
- Init
- Emissions
- Chemistry (Dry, Wet)
- Aerosol (inc chem tendencies)

Write Restarts

Diagnostic Meaning

Update forcing via ancillary

Boundary Layer · Dynamics Diagnostics · Wetlands (Hydrology) · Large Scale Precipitation

One-Way
D1
Two-Way

UKCA Chemistry And Aerosol

q

Prognostic Variables · Tracer Advection · Convection · Classic Aerosol · Radiation

UM Atmosphere

UKCA uses "embedded" "coupling" (ie, it's not "coupled")

UKCA called sequentially after rest of "atmosphere" step.

Depends on a range of processes, including diagnostics, calculated previously (!?) in that step!

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk

UKCA the community context

The Met Office

- Clearly the UM is the "top model", which really means the UM atmosphere!
- MO controls the integration (when, what) into new UM versions, and is constrained by many things in doing so.
- MO versions not always trivially ported to academic computing environments.

The ROW (rest of the world, primarily NCAS staff in Cambridge, Leeds & Oxford, but Met Office and NIWA too.)

- Develops UKCA
- Uses the components in other models (e.g. CTMs)
- Uses older versions of the ESM with newer versions of UKCA …

www.ncas.ac.uk

UKCA: What's wrong?

Until the advent of MONSooN (a common supercomputer development environment) *progress was painfully slow, and it's still difficult.*

Differing versions.

Poor performance.

Worries about migrating UKCA into new UM environment (ROSE).

The bottom line is that the UKCA process has not optimised the model runtime performance, OR the model development time OR the scientific confidence in the latest ESM version (since it's never the latest scientific version).

I would assert that the coupling environment (or lack thereof) is part of the problem!

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk

UKCA: What to do?

Relatively straightforward to "un-embed the code" and couple (quickly) using OASIS3-MCT.

- It's primarily the way it is for historical reasons, but some have asserted that changing this is a bad idea (losing the "efficiency" of common memory on the processor).

- I would argue that with a runtime that is five times longer with UKCA, that's the wrong efficiency to prioritise! If coupled, then

  – Easier to get concurrency (hopefully, might have stability issues).

  – Development processes in the two different communities should block each other less frequently (but still issues with scientific dependencies on diagnostics).

National Centre for
Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk

By Contrast:
Understanding Drought in the Thames Basin

Drought in the Thames basin is a complex interplay of precipitation, river-flow, ground-water recharge, geology (sub-surface aquifer behaviour, and flow through porous rocks).
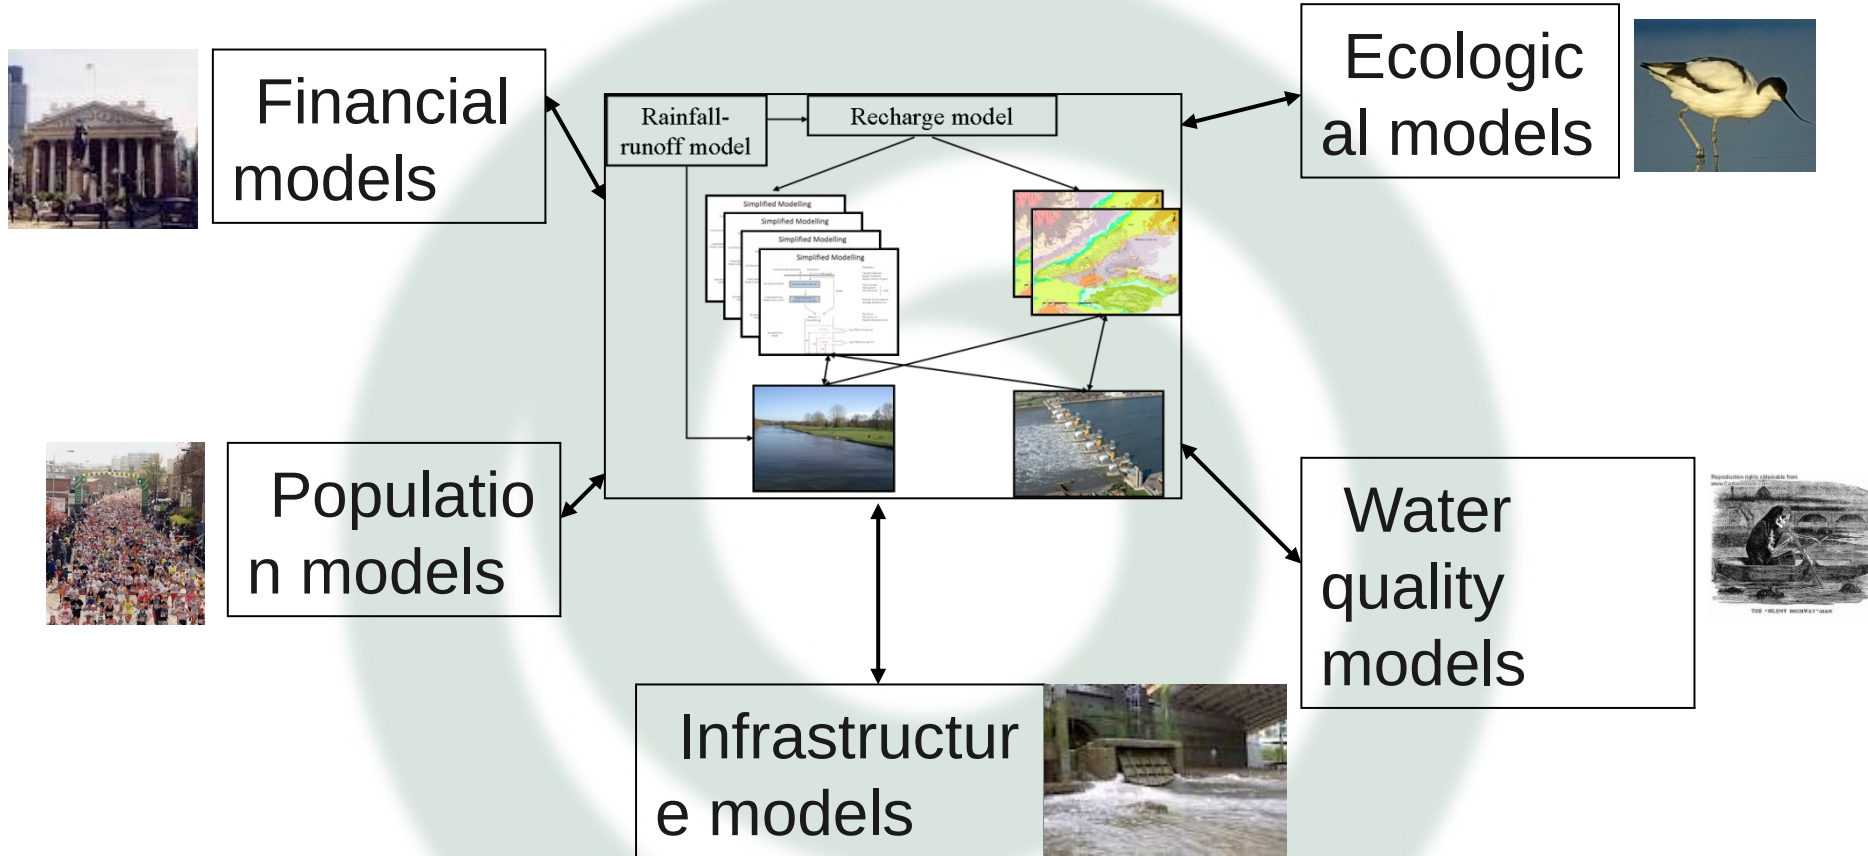
In one possible future, the UK will have a land surface "top model" suitable for applicability at the basin or regional scale.

It will run coupled in ESMs in some applications, and run as the "top model" in others (using ESM output as one model component).

This utility will bring scientific traceability … as well as supporting differing application domains.

Further linkages

Financial models

Ecological models

Population models

Water quality models
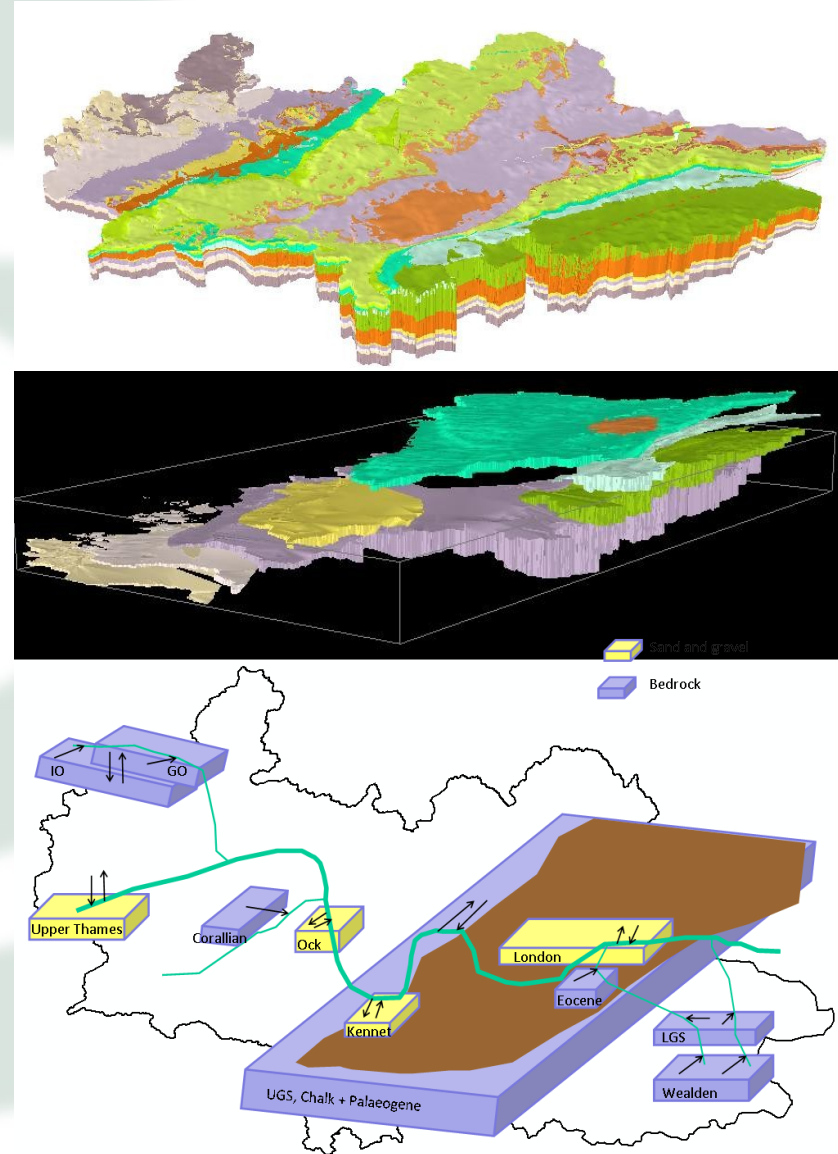
Infrastructure models

Slide courtesy of Andrew Hughes, British Geological Survey

Avoiding Top Model Coupling Paradigms: The role of generative approaches.

From a science perspective: there is no such thing as a component model! From my/your viewpoint my/your "component" model is a top model:

- Ideally I/you want to be coupling other components into my/your model.
- Consider the land surface case, running at (lower) resolution in an ESM, and at (higher) resolution being the top model (coupling precip via files) and complex ground water models …
  - Inevitably using different coupling paradigmsin those two directions!
- It simply cannot be good (efficient) science to maintain two code stacks. Much better to generate the coupling from one code stack.

Obvious role for generative tools like BFG.

- (I have yet to fully understand the possibilities of CSDMS/BMI ...)

Need: Better Workflow Tooling for "Coupling"!

It's not just about the runtime!

- – Comprehending the code!
- – Development
- – Debugging
- – Documenting
- – Validating
- – Evaluating

Comment from yesterday:
" communities interacting.. it's easy to get output, it's hard to know if it's correct … "
This talk in a sentence! (Except maybe it's not so easy to take the first step.)

All these things currently require "artisans" not "engineers" and certainly not "scientists". That has to change, and the tooling needs to facilitate all these things!

A cautionary note: On stability.

Most of our coupling between components is explicit, that is, we are updating our atmosphere using the past state of the ocean and vice versa.

It's not obvious that as we go to higher resolution and more process concurrency that this will result in good solutions in all cases!

We may need to move to more implicit coupling. That is, we introduce new models which update (some) of the variables using implicit methods and have slower explicit coupling elsewhere.

This is already an issue within model components: e.g.Wan et al 2013 in review at GMDD doi:10.5194/gmdd-6-685-2013

Their work was looking at the necessity for using an implicit solver to handle condensation, nucleation and production of sulphuric acid.

- Primarily because the solution involves finding the (small) difference of two (large) compensating terms.

Need to be careful (hence the caveat with UKCA). Also something to think about in dynamic coupling frameworks like CSDMS.

Summary (1 of 2)

Coupling is a technical solution to BOTH the science requirements and the shape of the scientific community.

Optimising for any one of those alone (or just for performance) is likely to result in short time wins at the expense of long term victory.

There is no one right solution for all communities and all problems.
- Generative techniques (e.g. BFG) or really simple framework support (e.g. CSDMS BMI) will be part of dealing with that!

Having those points in mind when we develop our coupling toolboxes should increase their utility.

Probably smart not to assume that our explicit coupling paradigms are going to survive in a higher resolution more highly concurrent exascale world.

Both Europe and the US aspire to put their scarce modelling effort into fewer fully complete and independent modelling assemblages.

It will never happen unless we address coupling as a community issue so that there is no concept of a top model (or super-framework, singular).

- We might be able to live with common frameworks: the VW model versus the Airbus model (with Boeing as an important contributor to evolutionary vigour).

- I just think it would be smarter if there were **fewer**!

National Centre for Atmospheric Science
NATURAL ENVIRONMENT RESEARCH COUNCIL

www.ncas.ac.uk