

LFRic and UKCA – some observations

Nick Savage 2015-01-30

Background

Since 2012 a project called Gung Ho has been underway to design a new dynamical core to replace ENDGAME. This will be on a grid designed to avoid the problems lat-lon grids have at the poles. LFRic is the name of the new infrastructure which will be used to parallelise Gung Ho and couple it to the rest of the physics code needed to turn a dynamic core into a complete NWP or climate model.

There is now a very strong driver to do this work due to changes in the nature of supercomputing. Processors are not getting any faster (they may even become slower) but we are getting more and more processors. On the next but one supercomputer (due around 2020) it is likely (according to Chris Maynard) that the current codes will actually run slower rather than faster and therefore we have to have the new model ready in time for that computer.

The timeline for the new model is as follows:

- LFRic Phase 1: 2012-2016. Design LFRic architecture, version 0.1 of the LFRic infrastructure, Develop LFRic implementation plan
- LFRic Phase 2: 2016-2019. Delivers version 1.0 of the LFRic infrastructure incorporating version 0.1 of the LFRic atmosphere. Coding “sprint” to re-write physics and other codes to LFRic design. Ends when LFRic atmosphere passes a certain level of scientific and technical performance.
- LFRic Phase 3: 2019-2022: Operational delivery planned and implemented. Delivers version 1.x of the LFRic atmosphere suitable for operational use and climate evaluation

The aim of this document is to take a preliminary look at the implications of this for composition modelling. I have some more general concerns as well which are in Appendix A

Why UKCA is special

There are 4 main reasons why UKCA is different from other physics components such as radiation and convection.

- Community developed. While most other schemes have been developed by the Met Office with a single scientist leading the development at any one time (or a small group working closely together) the UKCA code development process has been very different. UKCA has a set of chemistry routines adapted from the p-TOMCAT model at Cambridge, an aerosol scheme adapted from the offline GLOMAP-mode and also uses a photolysis scheme developed at Harvard. The Met Office has had the main responsibility for coupling these components with each other and the rest of the Unified Model.

- **Cost.** In air quality modelling or an Earth system model, the cost of adding chemistry is the largest cost. For example at 12 km, AQUM is about 5 times slower than the same resolution model without chemistry or physics. In spite of this, only small amount of time have been spent on the optimisation of UKCA. The main piece of work into the ASAD solver by the UM optimisation team failed to speed up the code and concluded that it needed a complete rewrite so that Open MP parallelism could be exploited.
- **Complexity.** Unlike say convection where a single scheme with a few modifications meet all the user requirements, UKCA has been developed with the aim of the of being applied to a wide range of scientific requirements from earth system modelling to ozone hole recovery via air quality forecasting. It can be used with aerosols off, with aerosols coupled to a simple chemistry scheme which uses offline oxidants, or with aerosols and chemistry fully coupled. This all makes the code much more complex to write and maintain. Finally, there is a desire to try and maintain the aerosol code in a consistent form across 3 models: TOMCAT, UKCA and the IFS.
- **Coupling.** The chemistry and aerosol schemes have very large number of inputs from other physics sections. These are mostly fields that are normally thought of as diagnostics and so are not passed up to the top level of the calling tree routinely. This has two consequences. Firstly a special interface for these fields must be written so that UKCA can get the inputs it needs without having to modify large parts of the other UM sections. This interface is not easy to work with, if one has to add additional fields, and requires the user to add additional STASH items so that these diagnostics are written to the correct part of memory for UKCA to get. This large number of input fields also makes UKCA vulnerable to changes elsewhere in the UM which can either prevent UKCA from running or cause unacceptable scientific deterioration in model skill. There are also a large number of UKCA fields which can be fed into the radiation and cloud schemes of the model but not for all model configurations which use UKCA (if the aerosols are off, they cannot be used in the radiation for example).

Implications of LFRic

Given that the whole of UKCA must be rewritten to make it consistent with the new LFRic framework (as do all physics components) and that we know the current UKCA coding is far from ideal, the question must be asked – what is the best way forward for UKCA. This offers us a chance to make the UM's replacement a model which has fewer bugs, is faster to run, is easier to extend and is easier to debug when things go wrong. The converse of this is that the risk we face if we do not engage with LFRic early

It has been suggested by various people, for different reasons that we should move UKCA completely out of the UM and use a coupler (e.g. Oasis) or a well defined interface (such as ESMF) to join together the UM and UKCA. From a code performance perspective, coupling UKCA to the UM would allow the chemistry and aerosols scheme to be run at lower resolution than the rest of the model and might enable more concurrency by allowing the chemistry to be solved while the rest of the model continues to integrate. Note however, that to achieve this, it would also be necessary to revise the way that the model deals with process splitting. At present, the chemistry is called to update gas phase and aerosol fields used on the next timestep and so blocks the rest of the model from continuing with computation. Both options would also allow UKCA to have its own release

cycle, decoupled from the UM cycle (although JULES has just moved away from this approach) and also allow the same code to be run as an offline model or a box model for process based studies (given some additional development). Given that ESMF is probably going to be the way each physics scheme in LFRic is coupled to the rest, the last of these outcomes might be achievable as a side effect of moving to LFRic.

It would also be a good time to review whether we wish to continue to use ASAD or if we want to move to another solution such as KPP and also consider any other changes we want to make. Any code we do use will need considerable rewriting if it is to be part of the LFRic code. If we move the UKCA code outside LFRic, we might be able to avoid wholesale rewriting but this might be at the expense of speed if we do not invest substantial amounts of effort on improving the code design.

Other improvements are also desirable, such as making it easier to add new tracers to the model, and more easily controlling the fields passed back and forth between chemistry and the rest of the UM.

Before any coding for UKCA in LFRic can begin these questions should be answered, or else we risk missing an opportunity to improve UKCA. Using a coupler approach raises the question of where processes like tracer transport and convective transport should be treated (inside or outside the UKCA code), how to deal with the process splitting of chemistry and aerosol process, whether it is a bad idea to have the aerosols at a different resolution to the cloud processes and probably many more.

Where next?

I would like to propose the following timeline for answering the above questions so that when coding starts in 2016, we know what approach we are attempting to implement. If we do not do this, then the risk is that the delays to starting work on coding aerosols and chemistry in LFRic will make it impossible to run the model with anything other than very basic aerosol climatologies and damage the Met Office's ability to be a leading climate centre. The process should involve negotiation (or at least consultation) with our composition modelling partners in the universities, and also seek expert advice from the LFRic team. Note that the time allowed for each aspect of this may seem long but I anticipate that any people involved in this will only be able to devote a small proportion of their time to it. I have not made any attempt at this stage to estimate the effort needed to do any of the tasks below. It is also essential to identify a single person to be accountable for this work.

- September 2014. Kick off a process of gathering requirements from all current and potential users. Also solicit advice and experiences of other model's approaches to the challenges we face. In particular ask for opinions on alternatives to ASAD and how important it is to be able to run the same code offline and online. This would initially involve sending out emails but face to face discussion and teleconferences will probably also be needed
- March 2015. Requirements gathering complete. Begin research into benefits of various solutions and
- September 2015. Draft top level design for composition modelling in Gung Ho ready. To be discussed further with all relevant people with the aim of getting agreement from the UKCA community.

- March 2016. Finalised top level design agreed with all interested parties and resources to execute implementation identified. This should be a component of the LFRic implementation plan

Appendix A. General concerns

- LFRic is heavily based on Object Oriented design and the use of Fortran 2003. How will scientists be trained in these techniques?
- LFRic is based on finite element methods which the office has little experience with. How will this impact on the physics code?
- LFRic requires all science routines to be totally rewritten. How will we find the time to rewrite code and develop the current model simultaneously?
- I fear there may be some timescale misalignment. CMIP 7 runs will have to be carried out on the next supercomputer but one and so (according to Chris Maynard) will be unable to use the current UM. However on Steve Mullerworth's timeline an ESM based on LFRic does not seem to be ready in time?
- New science developed in the period 2016 to – develop for two models or will LFRic become 5 years out of date before it is finished
- Academic take up. Some academics are still using HadCM3. What happens if our science partners don't pick up LFRic for a decade?