

A Separate Repository for UKCA

N. Luke Abraham

Following a JWCRP initiative, UKCA is refreshing its governance structure. As part of this refresh, moving the UKCA code-base out of the UM trunk has been discussed and was considered a positive step for the UKCA project, for several reasons, including:

1. It would better facilitate running UKCA in a standalone mode (sometimes called a *box-model*). There have been at least three previous attempts to make a standalone UKCA executable of the chemical solver, which have been shown to be useful for model development and testing.
 - These previous attempts have taken time to produce and have not been kept in step with the UM trunk, meaning that their usefulness has become less and less with time, as UM developments supersede the standalone model. Older box-models have eventually been abandoned as too expensive to maintain.
 - As well as a single grid-cell model, there is also interest in building column and trajectory-following models
 - This standalone model would be incredibly useful for the testing and development of new schemes, especially ones which would require hundreds of tracers and thousands of reactions that are unable to be used within the UM at present.
 - A standalone model would also be useful scientifically for many studies that do not require a global model. This model would improve community engagement as it could be provided under a different license arrangement, enabling a wider uptake of UKCA.
2. It would, by necessity, involve the cleaning-up of the UKCA interface. Currently the UKCA API is not well defined and involves a mixture of tagged STASH requests and subroutine calls. Here UKCA would benefit from examining how JULES and SOCRATES interface with the UM and developing a best practice from this in collaboration with the UM Systems Team. This will also aid when coupling to LFRic.
3. It should make UKCA more flexible in terms of how it could couple with other models. There are currently no plans to fundamentally change how UKCA interacts with the UM (i.e. built within the same executable using subroutine calls). However, being in a separate code base would more easily allow for making a separate UKCA binary and using a coupler such as OASIS, which may be beneficial in terms of improving UKCA load-balancing, for example.

Current plans would require the creation of a separate repository that would be used to create this standalone model, ingesting the UKCA code from the UM trunk. The refresh of the UKCA API would be planned in discussion with the UM Systems Team prior to any code being written. This work should have minimal impact on the UM, as UKCA would not be considered for removal to this separate repository until this standalone model is shown to be working. As this work is ongoing the UM should also benefit from the improvements to the UKCA interface being made as part of this work.

A detailed workplan would need to be agreed between the Met Office and NCAS, covering who would do this work and in what order the various UKCA subroutines should be worked on. Advice and support from the UM Systems Team would be required during this project, and approval from the UMPB will be required due to the additional overhead this would place on UM maintenance and its build system. Licensing is also an issue that will need to be addressed, as UKCA makes use of several routines from external sources, in addition to the code supplied by the original UKCA partners of the universities of Cambridge, Leeds, and Oxford.