

Interactive comment on “Quasi-Newton Methods for Atmospheric Chemistry Simulations: Implementation in UKCA UM Vn10.8” by Emre Esenturk et al.

Anonymous Referee #2

Received and published: 4 April 2018

In this paper the authors present a methodology for improving the time-to-solution for atmospheric chemistry modules. Given the large computational burden these modules carry, any robust methodology to reduce this cost is valuable. The study is certainly within scope of the journal, but I would recommend a few points are addressed prior to publication.

Abstract:

You state that ‘each QN call costing 27% of a call to the full NR routine.’. I would recommend this is made more specific with regards to overall cost to the entire box-model simulation. At the moment the cost is presented as an individual function call. Whilst

C1

an interesting statistic, and I appreciate different solvers will use different approaches to Jacobian iterations, this would be a valuable take-home message. As the authors note in other areas of the paper, some large scale models have associated costs of ~70% total run-time from the chemical routine.

In the abstract you also state that ‘The blended QN method also improves the robustness of the solver, reducing the number of grid cells which fail to converge after 50 iterations.’ Can you put this in wider context? What are the typical number of grid cells that would otherwise fail to converge after 50 iterations? Even if the efficiency gain is small, this could be an important reason for adopting the method. I appreciate space is tight in the abstract, but do make sure this is covered elsewhere.

General points: On the general cost saving point above, do the authors feel a 3% reduction in cost is significant for the UKCA? Is this within cost variability for other implicit solvers available? This might be an interesting comparison, even within the box-model simulations. This relates to an issue with code availability, noted at the end of this review, presuming standard optimisation options can be used with standard ODE libraries.

Most Jacobian matrices in atmospheric mechanisms are sparse. Is there any perceived scaling issues with sparsity with your method?

Section points: Section 1, page 3 lines 5 onwards. Here you discuss ‘construction’ of a Jacobian. Defining the Jacobian of a gas phase solver is relatively easy and, whilst pedantic, I’m unsure if the authors are referring to explicit calculation of a Jacobian during solver iterations or some approximation via finite difference. Please clarify since you then discuss approximating the Jacobian to reduce costs, where a finite difference approach would be more expensive than an explicit definition for any given chemical mechanism.

Section 2.3, Page 7, line 25 – is the tolerance relative and defined as a percentage?

C2

Section 2.3. I'm interested in the strategy in the existing solver for testing convergence before reducing the time-step and re-calculating. Presumably this is a heterogeneous issue on any given domain, but how are 50 iterations chosen as a point to reduce the time-step and wouldn't an adaptive method work?

Section 3.2.1 page 18. The authors make an interesting comment on potential speedups for 'spatial systems modelled by partial differential equations'. This includes reaction-diffusion problems, or processes in particulates. However it is not clear how generally applicable the method could be. Is it possible the method could ever lead to a decrease in performance?

Section 3.2.2, page 22 line 7: 'If the 3D model predictions for the two species which are on the opposite sides of the lifetime spectrum are very close, then it is very likely that physical values for all other species which have intermediate lifetimes will also be close.' Is it not relatively easy to demonstrate the range of propagation errors in all species? I would suggest a demonstration of this is included rather than a qualitative statement on potential. If this is not straight forward, please state why.

Section 4 Line 11: 'We also demonstrated that, the suggested method..'. Please remove the comma

You also state that 'The differences in chemical concentrations between the control run and that using the blended QN method are negligible for longer lived species, such as ozone, ' Please quantify 'negligible'.

Section 4 Code Availability. Perhaps I have misunderstood licensing issues here, but this is a slightly disappointing end. Is there a reason why at least an example chemical box-model with the QN method could not be supplied? Even if this was hard-wired, without using a package such as KPP, it fits within the clear procedures and ethos now pervading all GMD papers. I'm not sure if I could reproduce your results and check the potential exciting co-benefit for other models. This could be a simple oversight, but I would suggest the authors check with the paper and editor on providing a minimum

C3

statement on the availability of this component. Will the optional QN methods be part of a new release? Maybe I have missed this within the main body of text, or it is implicit with the paper. If this is the subject on on-going work for which the group wishes to retain IP, which is absolutely fine, then I would simply state we could all wait for some exciting follow up studies. If there is a perceived issue with general applicability, this should also be stated.

Interactive comment on Geosci. Model Dev. Discuss., <https://doi.org/10.5194/gmd-2018-32>, 2018.

C4