Geosci. Model Dev. Discuss.,
https://doi.org/10.5194/gmd-2018-42-AC2, 2018

Geoscientific Model Development Discussions

Interactive comment on "MP CBM-Z V1.0: design for a new CBM-Z gas-phase chemical mechanism architecture for next generation processors" *by* Hui Wang et al.

Hui Wang et al.

wqizhong@bnu.edu.cn

Received and published: 18 August 2018

This paper describes improvements in vectorisation of the Carbon Bond Mechanism Z. While interesting, significant problems need to be addressed before I can recommend its publication in GMD.

Response: The authors thank you for encourage and sharing your time in this manuscript. We have merged your two files of comments and the detailed responses to the comments are given point to point as followed, and the modifications followed the comments will be added into the revised manuscript.

General Comments

1. The manuscript contains numerous grammatical and spelling mistakes, and I will attempt to highlight these in detail in my specific comments below. I would recommend the authors check any resubmission carefully.

Response: We have modified these grammatical and spelling mistakes of the current manuscript following your guides, and we will check the revised manuscript carefully. And a native speaker will be invited to help us to improve the language before the resubmission.

2. Throughout the manuscript the authors make reference to "the chemistry transport model", "the CTM", and "the air quality model" without first describing which specific CTM or AQ model they are referring to, if any. I believe from the context that the model the authors use is in fact the GNAQPMS model, first mentioned on P2L13. If the authors mean a specific CTM it should be discussed in this context, although in some places (e.g. P3L15) the authors mean CTMs in general. If the authors mean any CTM then the phrase "Chemistry Transport Models" or "CTMs" would be appropriate.

Response: The authors agree with your comments. We would follow your advice and use the name of the GNAQPMS model when we referring to the specific model, meanwhile, we would use the phrase "Chemistry Transport Models" or "CTMs" when referring to CTMs in general.

3. The paper itself only covers improvements to the CBM-Z module, which I assume is included in some way into GNAQPMS, although this is not discussed by the authors. The CBM-Z code is provided by the authors on Zenodo, which is great to see, although I find it difficult to understand the improvements made in the code by examining Figures 2 and 3 in the manuscript. While graphical representations of these optimisations are useful, I would like to see how these were implemented in practice by the authors giving specific code examples within the manuscript.

Response: The authors appreciate your comments. Since we would only focus on the CBM-Z module and isolate the impact of other modules in this manuscript, we didn't show the results of GNAQPMS model and only discussed CBM-Z module. We will follow your advice and add the sample code to illustrate the implement of codes.

4. The authors only give results from a CBM-Z standalone model, rather than having incorporated these improvements back into their CTM and presenting results from there. If available, I would certainly like to see what this does to model performance, as I feel it would strengthen this work greatly. As it is, they present only two case studies, where emissions are zero and meteorological conditions were constant. I would expect to see simulations of a number of different environments similar to those seen in simulations, i.e. free troposphere, boundary layer, urban, rural etc.

Response: The authors thank for your comments. We accept your advices and will conduct more simulation experiments by using the CTM incorporated with the optimized CBM-Z model. We will test the performance of model under diverse scenarios, and corresponding results will be provided in the revised manuscript.

5. CBM-Z output is plotted for 10 model hours, and within this time the relative error introduced by the optimisations is less than 0.05% (just), as seen in Figure 4 and mentioned throughout the manuscript. However, it is clear that this error is increasing for some species (e.g. H2O2, SO2, and H2SO4 as presented, and possibly others not shown). If these simulations were run for longer than 10 hours, would these errors still be below 0.05%? For confidence in the improvements described I would expect to see that the errors remain low for the length of a typical CTM simulation, which could be years depending on how the author's CTM is used. If the conditions used are more realistic (see point 4 above) do these errors increase?

Response: The authors appreciate your constructive comments. We agree with your comments. Indeed, the error of some species is increasing with time, and we also concern about this issue. Therefore, we would further investigate the source of the

error by conducting more tests and try to find some solutions like using more reliable compiling flags to constrain the error. We will use the long-time simulation to see the trend of this error, and the CTM's simulation would be used to investigate its impact in real situations. In addition, more species with diverse chemical properties will be analyzed and results will be presented in the revised manuscript.

6. The authors describe running the CBM-Z model using a single point and over a number of grids for testing, which I believe to be a spatial grid from the context. In this case, is there any transport or mixing between grid-points, or any differences in meteorological variables? If not, how was this configured and set-up is each point solving the same conditions at all times? If so, this will not be representative of real-world usage where there will be a large amount of variation across the domain. For the single-point model, is it integrated in time? If not, how is it configured? I again question whether it would be better to perform these 3D simulations with a CTM instead of a stand-alone CBM-Z model, as it would then allow for a better understanding of how these improvements impact their model in practice. If this is not possible, the authors should give reasons as to why this is.

Response: The authors appreciate your comments. The meteorological variables, e.g. temperature, are varied with grids in the multiple grid test. There is also no transport as well as diffusion but only gas-chemistry process in that test to isolate the impact of other processes. We agree that this test is insufficient to illustrate the improvement in real world, and we will test results by using the CTM to further present the improvement of our optimized scheme. In addition, the single-point model is integrated in time as a box model to show the impact of optimization on results, and in the revised manuscript, we will also present the results of CTM to supplementarily its effect on model results.

My major concerns with the manuscript as presented are:

A. Insufficient conditions used for testing

Response: We fully understand your concern and we would use the CTM with the

optimized CBM-Z scheme to test the performance and validate the output of the model.

B. Insufficient analysis of the errors introduced by the optimisations.

Response: In the current manuscript, we have stated that using -O0 compiling flag could diminish the difference of the results, which demonstrate there is no logical and artificial errors of the optimized codes, but we didn't further investigate the impact of optimizing codes on results and how to constrain the error by using some specific compiling flags. Therefore, we will try to investigate the source of error and potential way to constrain the errors.

Specific Comments and Technical Corrections

P1L11 - computational

Response: We have accepted your advice and modified in the revised manuscript as following:

"Precise and rapid air quality simulation and forecasting are limited by the computational performance of the air quality model used, and the gas-phase chemistry module is the most time-consuming function in the air quality model."

P1L12 - model used.

Response: The authors appreciate your comments. We have accepted your advice and modified in the revised manuscript as following:

"Precise and rapid air quality simulation and forecasting are limited by the computational performance of the air quality model used, and the gas-phase chemistry module is the most time-consuming function in the air quality model."

P1L18 - Knights Landing

Response: We have accepted your advice and modified in the revised manuscript as:

"The Intel Xeon E5-2697 V4 CPU and Intel Xeon Phi 7250 Knights Landing (KNL) are

used as the benchmark processors."

P1L19 - I question whether the <0.05

Response: The authors appreciate your comments. The criteria of 0.05% comes from the results of single point test, and as mentioned above, we would further investigate the source of error by conducting more simulation experiments.

P1L29 - Chemistry Transport Models

Response: The authors appreciate your comments. We have accepted your advice and modified in the revised manuscript as following:

"As a useful tool for air quality problems, Chemistry Transport Models (CTMs) are widely used in studies of air quality"

P2L1 - a CTM

Response: The authors appreciate your comments. We have accepted your advice and modified in the revised manuscript as following:

"As the core of the AQF system, a CTM requires a large number of computational resources to simulate the complex chemical and physical processes."

P2L1 - computational

Response: The authors appreciate your comments. We have accepted your advice and modified in the revised manuscript.

"As the core of the AQF system, a CTM requires a large number of computational resources to simulate the complex chemical and physical processes."

P2L3 - relatively simple processes are adopted in CTMs to minimize

Response: The authors appreciate your comments. We have accepted your advice and modified in the revised manuscript as:

“To satisfy the demand of routine air quality forecasting in a timely manner, coarse spatial resolution and relatively simple processes are adopted in CTMs to minimize the use of computational resources.”

P2L4 - computational

Response: The authors appreciate your precious comments. We have accepted your advice and modified in the revised manuscript as following

“To satisfy the demand of routine air quality forecasting in a timely manner, coarse spatial resolution and relatively simple processes are adopted in CTMs to minimize the use of computational resources.”

P2L5-6 - I would suggest this: "Therefore, air quality simulation studies can benefit significantly by improving the performance of the CTM used."

Response: The authors appreciate your constructive suggestion. We have accepted your advice and modified in the revised manuscript as following:

“Therefore, air quality simulation studies can benefit significantly by improving the performance of the CTM used.”

P2L16 - what is meant by "improving the frequency of air quality forecasting" in this context?

Response: The authors appreciate your comments. Yes, it means “improving the frequency of air quality forecasting” in the context and we have revised this sentence as following:

“The AQF system can also benefit from the performance improvement by adopting a higher model resolution and improving the frequency of air quality forecasting.”

P2L19-21 - do you have any references for the trend in changes to computing architecture that can be quoted here?

Response: The authors appreciate your comments. Yes, Xu et al. (2015) and Lawrence et al. (2018) could be the appropriate reference here.

P3L4 - what is "the air quality model" in this context? Do you mean "in several air quality models"?

Response: The authors appreciate your comments. The “air quality model” here refers to the “CTMs”. Following the second general comments of the reviewer, we would use the identical phase “CTMs” in the revised manuscript.

P3L9 - what do you mean by "architecture" in this context?

Response: The authors appreciate your comments. The phase “structure” may be more appropriate in this place, and it means we modified the codes structure of CBM-Z scheme, as we describe in the context, to improve the vectorization of codes.

P3L14 – tropospheric

Response: We have accepted your advice and modified in the revised manuscript as:

“CBM-Z is a lumped-structure photochemical mechanism that was developed to meet the needs of city-scale to global-scale tropospheric chemical simulation”

P3L15 - in CTMs

Response: We have accepted your advice and modified in the revised manuscript as following:

“The original scheme contains 67 species and 132 reactions. CBM-Z has been widely used in CTMs”

P4L1 - delete "The" and just start the sentence "CBM-Z also. . ."

Response: The authors appreciate your comments. We have accepted your advice and revised in the manuscript as:

“CBM-Z module still contains many scalar operations.”

P4L6 - simulations

Response: The authors appreciate your comments. We have accepted your advice and modified in the revised manuscript as:

“Fortunately, contiguous model grids may have similar chemical processes in air quality simulations”

P4L14 - space after AVX-512

Response: The authors appreciate your comments. We have accepted your advice and modified in the revised manuscript.

P4L19 - do you mean "implement fine-grained parallelization"?

Response: The authors appreciate your comments. Yes, it is exactly what we mean and we have modified this part.

P4L21 - fine-grained

Response: The authors appreciate your comments, and we have revised in the manuscript as:

“Our goal is to implement fine-grained parallelization based on the SIMD and the grids that are distributed to a specific processor operate in parallel using the VPUs on each core.”

P4L21 - delete the comma after SIMD

Response: The authors appreciate your comments. We have accepted your advice and modified in the revised manuscript.

P5L5 - does the solver use a fixed number of iterations, or does it integrate to convergence?

Response: The authors appreciate your comments. The solver uses the explicit algorithm and it didn’t need iterations. The solver in this test is Modified-Backward-Euler

(MBE) method. The detailed description of the solver could be found in Feng et al. (2015) and Feng et al. (2017).

P5L13-14 - does the fact that calculations are performed on all grids but not all grids are copied back introduce a possible inefficiency? Are these grids taking time to solve that could be better spent doing something else?

Response: The authors appreciate your comments. No, this process can keep the high computational efficiency of Vector Processing Union (VPU) considering that the chemistry processes among grids are different. It avoids the logical judgements to assort the grids so that corresponding computations can be finished on VPU simultaneously, which means that the computations of multiple grids could be finished at the same time without doing extra logical judgment processes.

P5L22-23 - as I have mentioned in the General Comments, how representative are these examples? How exactly were they configured?

Response: The authors appreciate your comments. As we mentioned in the manuscript, the goal of the single-point case is to validate the model’s output for debugging logical or artificial errors, and the 3-dimension (3D) case is used for the performance testing. The 3D case is derived from the real CTM simulation, and the meteorological conditions like relative humility of this case is diverse among grids but constant during the simulation. The 3D case does not contain the emission, transport and diffusion process but only gas-chemistry process.

Following your advice, we would conduct some real-scenario simulation by using CTM. In addition, we will provide more clear descriptions of the test conditions as well as configuration in the revised manuscript.

P6L17 - as I have mentioned in the General Comments, how robust is the error of range <0.05

Response: The authors appreciate your comments. The criteria of 0.05% is based

on our current test and we can't answer how robust of these criteria is without further test and investigation. So we will use the long-time simulation as well as the CTM's simulation to investigate the source and possible solutions for the errors.

P7L4-9 - Please provide more details of these 3D simulations, as I am unclear exactly how they were set-up.

Response: The authors appreciate your comments. The 3-dimension case is derived from the real CTM simulation, and the meteorological conditions like relative humility of this case is diverse among grids but constant during the simulation. The 3D case does not contain the emission, transport and diffusion process but only gas-chemistry process. We will provide more detail about the test cases in the revised manuscript

P8L10 - It's great to see that the code is provided, but I found the structure provided confusing and the README provided lacking. However, this isn't part of this manuscript, but I would urge the authors to improve the documentation provided, include a directory listing with what the files do etc.

Response: The authors appreciate your comments. We would follow your advice to provide the codes with detailed document as well as directory structure for readers' or users' convenient.

P10 - Caption to Table 2: versions

Response: The authors appreciate your comments, and we have revised caption to Table 2 in the manuscript as:

"Compile flags of the different versions of CBM-Z."

P11 - Table 3: How many times were these tests run? Is it possible to provide an error estimate for these numbers?

Response: The authors appreciate your comments. These tests were done only one time currently on the relative stable testing platforms. It is possible that some unpre-

dictable situations would affect the results, but we will repeat the following tests more times to ensure the stability of results.

P12 - Caption to Figure 2: vectors

Response: The authors appreciate your comments, and we have accepted your advice in the revised manuscript as:

"The i and j loops, equaled latitude and longitude loops, were merged to construct one vector to reduce the number of unfilled vectors."

P12 - Figures 2: As mentioned in the General Comments, I believe that this manuscript would benefit from seeing how the code is altered with these optimizations.

Response: The authors appreciate your comments. We would provide the figures of sample codes to show the optimization processes.

P12 - Caption to Figure 3: grids

Response: The authors appreciate your comments and we have accepted your advice in the revised manuscript as:

"The flowchart shows the way to mask the heterogeneous girds to integrate grids to perform the vectorization operations according to the iregime values."

P13 - Figure 4: As mentioned in the General Comments, do these trends continue? How long until they become significant?

Response: The authors appreciate your comments. We are not sure about whether this error would continue. If these trends of error continue, it could become significant and effect the results. We also concern about this issue, therefore, we would further investigate the source of the error by conducting more tests and try to find some solution to figure it out.

P13/P14 - Figures 4 5: Which simulations are these figures from, the single-point case

or the 3D simulation? I would suggest more species are analyzed, covering a range of chemical lifetimes.

Response: The authors appreciate your comments. These figures are both from the single-point case, and we will conduct more experiments and analyze more species with diverse chemical lifetimes and properties.

Reference

Feng, F., Wang, Z., Li, J., and Carmichael, G. R.: A nonnegativity preserved efficient algorithm for atmospheric chemical kinetic equations, Applied Mathematics and Computation, 271, 519-531, https://doi.org/10.1016/j.amc.2015.09.033, 2015.

Feng, F., Chi, X., Wang, Z., Li, J., Jiang, J., and Yang, W.: A nonnegativity preserved efficient chemical solver applied to the air pollution forecast, Applied Mathematics & Computation, 314, 44-57, 2017.

Lawrence, B. N., Rezny, M., Budich, R., Bauer, P., Behrens, J., Carter, M., Deconinck, W., Ford, R., Maynard, C., Mullerworth, S., Osuna, C., Porter, A., Serradell, K., Valcke, S., Wedi, N., and Wilson, S.: Crossing the chasm: how to develop weather and climate models for next generation computers?, Geosci. Model Dev., 11, 1799-1821, 10.5194/gmd-11-1799-2018, 2018.

Xu, S., Huang, X., Oey, L. Y., Xu, F., Fu, H., Zhang, Y., and Yang, G.: POM.gpu-v1.0: a GPU-based Princeton Ocean Model, Geosci. Model Dev., 8, 2815-2827, 10.5194/gmd-8-2815-2015, 2015.

Please also note the supplement to this comment:
https://www.geosci-model-dev-discuss.net/gmd-2018-42/gmd-2018-42-AC2-supplement.pdf

---

Interactive comment on Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2018-42, 2018.