Geosci. Model Dev. Discuss.,
https://doi.org/10.5194/gmd-2019-333-RC2, 2020

Geoscientific Model Development Discussions

Interactive comment on "WRF-GC: online coupling of WRF and GEOS-Chem for regional atmospheric chemistry modeling, Part 1: description of the one-way model (v1.0)" *by* Haipeng Lin et al.

Anonymous Referee #2

Received and published: 13 March 2020

This paper describes how GEOS-Chem has been implemented in the WRF model. The topic and the description of its implementation is appropriate for the GMD audience. There are certain details that need elaboration for future users of the code and I have made several suggestions. In addition, some the statements are clearly biased that overstate the capabilities of the model without actual proof. At times, the text sounds more salesmanship than scientific. Both of these concerns should be relatively easy to address.

Major Comments:

1) Boundary layer mixing is handled separately by GEOS-Chem. I am a concerned

about this approach which is glossed over. Most, but not, all of the meteorology from WRF is used and that point is passed over in other parts of the text. WRF-Chem uses parameters from the WRF boundary layer parameterizations so that the vertical mixing is treated in as similar as possible. In this way, mixing of chemical species occurs within the same boundary layer depth as the meteorology. It is not clear whether vertical mixing in GEOS-Chem and WRF are consistent. If not, it may be possible for GEOS-Chem to have a deeper or shallower boundary layer than in WRF. If deeper, than more chemistry variables will be transported by free tropospheric air and boundary layer air. The authors should delve into this in more detail to let users know what the approach for vertical mixing actually is and the potential consequences. This applies to the results shown in Figure 4.

2) At several places in the manuscript, the authors point out the reasons for the advantages of coupling GEOS-Chem within WRF. For the GEOS-Chem community the advantages are obvious. But it is not as clear what the WRF-Chem community gains. The WRF-Chem community would have another chemistry option, but this paper does not explore the types of chemical processes that might be missing (perhaps halogen chemistry) already in the model. The current aerosol treatment is rather simple (is it even "state-of the science"?) and similar in many respects to existing simple aerosol options in WRF-Chem. Users choose particular treatments in a community model, such as WRF, for various reasons such as overall performance (i.e. how well the model represents reality), physics complexity, and computational considerations. The authors have failed to articulate what the advantages are to the WRF community. The future developments beyond v1.0 may change this picture and make that argument more apparent. But if the authors want to frame their arguments in this paper, then a few more concrete points of the advantages to the WRF community are needed.

Specific Comments:

Lines 7-8: "is designed to by easy to use, ... extendable, and easy to update" are relative terms and depends on an individual's point of view. This is something I'm

struggling with here. “is designed” is probably the key phrase. It is hard for a reviewer to verify the last part of the sentence without using the code itself. WRF itself is a rather complex model, although users with “sufficient” expertise in atmospheric models and computational hardware can learn how to run the model in a short period (i.e. days) of time. Those without “sufficient” expertise, might not describe it as easy to use. Nor am I sure what “extendable” means here. All computational models by their very nature are extendable by modifying code.

Lines 15-17. I have some concerns regarding the statement on PBL heights, in which I will comment on later in the appropriate section.

Lines 18-19: The sentence “Both parent models . . .” is redundant with an earlier statement. It probably can be deleted and the thought merged with the earlier statement.

Lines 27-28: “regional” is used twice in this sentence and is redundant and is self-evident.

Lines 33-34: “to better serve the public, inform policy makers, and advance science” is a laudable goal. However, I think these types of models are used primarily to “advance science”. WRF has been used to provide short-range operational forecasts (e.g. HRRR model), but NOAA is phasing out the use of WRF. WRF-CMAQ has been serving EPA regulatory interest in addition to advancing science, but I am not sure GEOS-Chem serves that purpose.

Lines 73-86: The authors state that online models are more difficult to keep up to date than offline models. I disagree. There are many factors that contribute to how quickly and the frequency of updates made in models (some described by the authors) that have nothing to do with whether the model is online or offline. The authors first mention resources used to provide benchmarking, validating, and documentation, which is probably a good thing for both offline and online models. This process does take longer for more complex models, but offline models can be complex too. It is also not a good idea to translate research findings into publicly available without due diligence. Some-

times new scientific findings are proven to be incorrect and/or the findings are based on a limited case study; therefore, it is not wise or appropriate for use by a broad community. The authors second point on expertise residing in different communities. I agree this is an issue, but this is largely governed by how the organization wishes to develop the code, and not whether the model is offline or online. In the case of WRF, it has always been a community model so that contributions originate from various organizations and universities. Nor has NCAR insisted that all physics schemes be compatible with each other, and sometimes there are good reasons why some physics schemes should not be made compatible. Had NCAR decided to be the sole developer, the code would no doubt be more streamlined but would probably lack scientific options many users want. It seems strange to me that the authors are arguing that bring communities together introduces problems, when the purpose of WRF-GC is yet another iteration of bringing very different communities together. The authors correctly point out that not everything in WRF-Chem (and WRF for that matter) is compatible, but it is not clear whether the same holds for the various treatments in GEOS-Chem.

Lines 98-98. This sentence is kind of a put-down to WRF (in the context of this paragraph) and other models listed in Table 1. As with another statement in the introduction, this depends on one’s point-of-view and difficult to prove. What is “state-of-the-science” anyway? This overused phrase is almost meaningless at this point. For examples, some chemistry and aerosol disciplines are changing very rapidly (weekly) and it is very difficult to argue which model has the most “up-to-date” science. I suggest the authors rephrase this sentence to be a bit more fair and unbiased.

Line 100: The authors state that that GEOS-Chem is driven by the meteorological fields simulated by WRF. But after reading material later in the paper, this may not be entirely true. It seems that GEOS-Chem still uses its own turbulent vertical mixing which could be different from WRF. I have some other comments on this point later in the paper.

Lines 168-169: Based on this sentence, it sounds like the bulk treatment is similar to GOCART. But the last sentence in the paragraph implies a modal treatment for dust

and sea-salt. So is dust and sea-salt in bulk bins or actually prescribed using a size distribution? Just need to be consistent in the text here.

Lines 169-172: These two sentences might be better at the end of the paragraph. It would be better to have all the discussion on the present aerosol model together, rather than broken up with what is not in the code in the middle.

Lines 181-185: I think some discussion is needed regarding higher-resolution emission datasets. GEOS-Chem has been used traditionally at global scales and there are several global emission datasets available. But if the point of WRF-GC is to run at higher resolution, this would be defeated in part by an emission inventory that is much coarser than what could be simulated by the model. What is the strategy for that? Will other inventories used by EPA and the WRF-Chem community be used? Or are users expected to generate emissions on their own? Also, emissions are discussed in more detail later in Section 3.3.1. So I am wondering if it is even necessary to talk about emissions here. This material can be merged into Section 3.3.1

Line 198: I am not sure what "greater modularity" means here. WRF-Chem is structured in a way to have modularity for important chemistry and aerosol processes. For example, there are several option for online emissions, deposition, etc. The modularity also permits users to add their own treatments if they wish. In some cases, there are treatments that can be used for multiple chemistry or aerosol options. For example, wet scavenging can be handled similarly for MADE-SORGAM (modal aerosols) and MOSAIC (sectional aerosols). But I doubt if the modules in GEOS-Chem can be used or accessed by other treatments in the code, given the structure on how it was implemented. So, if the authors want to use the phrase "greater modularity" some description is need to know exactly what this means.

Line 206: What is the "WRF-to-chemistry interface". Is GEOS-Chem handled in the registry.chem file, much like the other chemistry options? Given Figure 1, I think this will be discussed later and it looks like the material is in Section 3.2.2. I just found this

a bit vague at this point.

Line 213: See the first major comment above.

Line 220: In Section 3.2, the text implies that GEOS-chem is compiled and run on top of the host WRF model. The text implies but does not explicitly state that the addition of GEOS-Chem has not impaired the use of WRF-Chem. One of the rules in the WRF community is that any new additions must be shown to not "break" other parts of the code. So have the authors performed tests with the new code to ensure that other chemistry options produce the same results went the code is compiled with GEOS-chem?

Lines 312-335: As the authors imply, subgrid and removal processes, will depend on both "resolved" and "unresolved" clouds. In contrast with global models, simulations at fine enough resolutions may be best run with any "unresolved" cloud parameterizations. Ideally, a cumulus parameterization would have a progressively smaller and smaller impact at higher and higher resolutions. The subgrid and removal processes depend on the current behavior of what is in WRF. It would be useful to provide some discussion for users regarding the implications of these assumptions. Since this paper describes a new modeling framework, it would have been interesting to see some simulations at coarse and fine resolutions to demonstrate the differences on the vertical transport and wet scavenging. WRF-GC users will be able to run at smaller grid spacings, which is a good thing, but will they understand the subtleties of these assumptions that are not normally encountered at global scales?

Lines 323-325: Same comment as before regarding vertical mixing. This is a different treatment than in WRF. So there could be mis-matches in how boundary layer mixing, and it impact the vertical extent of boundary layer mixing, between WRF (which is used to transport chemical species) and GEOS-Chem. Lines 350-359: The authors should include the meteorology simulation time step and chemical time step. What is not mentioned in the discussion of the model's implementation, but alluded to in the

conclusion, is that a different chemical time step can be used (which is similar to WRF-Chem). In this way, transport of chemical species are done at the meteorological time step. Chemistry is usually simulated at a coarser time step to save computer time, but some care is needed since chemical time steps that are too large could introduce uncertainties in the predictions – especially at higher spatial resolutions. Since the paper is about WRF-GC and designed to work at higher spatial resolution, some discussion is needed in the best practice for the two types of time steps. Ideally, they should be the same for consistency.

Line 353: The authors use FNL for the meteorological initial and boundary conditions. To be more consistent with the nested GEOS-Chem, why not use the meteorological conditions from that model?

Lines 356-359: It sounds like the authors are using the obs-nudging capability in WRF. Would be useful to cite a reference on that. What is not clear in the previous paragraph, for readers not familiar with GEOS-Chem, that the meteorological fields are prescribed analyses. Thus nudging in WRF would be appropriate to make the meteorology in the two simulations more compatible. The authors should be more specific on these points.

Line 363-364: Do the anthropogenic emissions vary diurnally? This seems to be an important point in simulating diurnal and peak concentrations. Here the authors only show a 6-day average. Perhaps this is okay for demonstrate WRF-GC compared to GEOS-Chem, and this distinction should be made.

Lines 394: Are these boundary layer heights the same as predicted by WRF? Or different calculations in the GEOS-Chem modules?

Lines 397-402: The 6-day averaging may be hiding some day-to-day variations when trying to assess the cause of the positive PM2.5 bias. While it is very likely that the boundary layer height issue is contributing to that, there are likely other differences in the meteorology that could also be attributing factors. The authors should at least acknowledge that if they do not wish to pursue a more detailed analysis of the simula-

tions. I assume the same boundary layer scheme is being used in both models, but the differences arise in how that scheme is driven. I presume in GEOS-Chem the meteorology is from the large-scale analyses. Whereas, in WRF the prognostic land-surface model will be controlling the evolution of surface temperature and fluxes that will drive a boundary layer parameterization. So, some additional discussion as to why these differences in the boundary layer height would be useful.

Line 417: Eliminating the need to read archived meteorology becomes even more important at high resolution offline approaches as more and more time is required for I/O.

Line 418: I do not fully understand what reading from disks means. Do the authors mean the analyses are on some long-term storage device (tapes?) in which reading is slower than conventional hard-drives? Would a fairer timing test include having meteorological fields in a place with faster access?

Line 432: Figure 6 is difficult to interpret because of the scale on the y-axis. Ideally if one doubled the number of cores, one would want the wall clock time to be reduced by a factor of 2. It is hard to see this with the current scaling. To me it looks like the total performance does not improve much greater than 100 cores. In addition to "fragmentation" the authors mention, the most CPU time in WRF is due to the advection of species. The more chemical species there are the slower the code is. The authors mention 241 species on line 163 (presumably trace gas), but there would be aerosol species on top of that. This is more than other options in WRF-Chem, but less than others. In the conclusion the paper states that the WRF-chem community would benefit from GEOS-Chem, but it is not clear what the computational cost would be compared to other approaches. A large fraction of the users in the WRF community chose "simple" chemistry and aerosol schemes to save computational cost. While the authors may not wish to perform identical simulations with different chemical options to benchmark GEOS-Chem with others, the least they could do is compare the number of species to establish some sort of computational level of complexity.

Line 440: Change “light-weight” to “efficient”.

Line 449: The conclusion would benefit (perhaps at the end) about future directions. They are alluded to elsewhere in the text, but it would be good to briefly summarize them here. That would include things like more complex aerosol treatments and fully-taking advantage of on-line coupling to include feedback effects (both of which are already in WRF-Chem). The conclusion would also be strengthened to foreshadow what would be discussed in part 2.

Lines 460-463. As I mentioned earlier, the evaluation is rather simple. This paragraph should be rephrased to indicate this is only one case study (and thus not comprehensive) and other meteorological differences may be contributing to the PM2.5 predictions.

Line 464: The first sentence is an overstatement and deceiving if a reader only looks at the conclusions. Yes, the chemistry part itself scales well but not the entire code.

Interactive comment on Geosci. Model Dev. Discuss., https://doi.org/10.5194/gmd-2019-333, 2020.