

260629-GMD-Review

The authors have done a good job of addressing the previous referees' comments, and I have only two minor things that they might choose to address in their final submission:

1. P6, L150. Figure 2 shows the run time, but not the memory. It'd be rewrite this to be clear about that (it is not necessary to have a figure for the memory, but the text should not imply that figure 2 does show that).
2. I am confused about the 200 core concurrency limit, what drives that, is it just the available compute/scheduler architecture? Is this somehow related to P13L290? There are several things here I don't understand. In slurm you can request the memory resources you need, there are other compute types on JASMIN/LOTUS I believe ... and why can't you run batches in parallel as well? (In which case you have much more memory available, but still not enough to have 16 GB per task for 200 tasks which is what they say they need). If this was a "production configuration choice" the fine, but please say so. I've missed something here, but maybe that's my problem and I just missed the key sentence.