CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

Dr. Antigoni Voudouri
Hellenic National Meteorological Service
(HNMS)
14 El. Venizelou,
16777 Elliniko
Greece

Lugano, 22 September 2017

Ref: Your allocation request for a Production Project at CSCS

Dear Dr. Voudouri,

We have finished the review process for research proposals submitted until May 19th, 2017. Based on the technical and scientific reviews and the final assessment of the Scientific Advisory Committee (SAC), we shall support your proposal entitled "Optimization of a calibration procedure for weather prediction model" with an allocation of:

- 400'000 node hours over 1 year on the Cray XC50 – Hybrid
- 35 TB disk space

and a duration of one year, that is until September 30th, 2018. The allocation is available from October 1st, 2017 until September 30th, 2018. You will receive information about your new project separately.

Please find enclosed the technical and scientific reviews and the recommendations of the SAC.

Sincerely yours

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

Dr. Maria Grazia Giuffreda
Associate Director
& Head of User Engagement and Support

Tel: +41 91 610 8223
Email: mgg@cscs.ch

CSCS

Scientific Advisory Committee Recommendation

Scientific Reviews:

Outstanding ☐ Excellent ☒ Good ☐ Fair ☐ Poor ☐

SAC

The SAC recognizes the potential scientific value of the proposal and agrees with the scientific and technical reviews. The SAC agrees to grant the full allocation if concerns by scientific and technical reviews are properly addressed. The SAC recommends that the project should be carefully monitoring to prevent the usage pattern of the past project granted at CSCS.

The Committee thinks the proposal should be

(A) accepted as is with … TB of storage and a duration of one year
(B) accepted, but requested allocation cut to 400'000 node hrs, 35 TB and the duration of one year
(C) rejected

CSCS

Scientific Reviews

Scientific Review 1:

Significance of the proposed research:

This is a nice project that builds on work that has already shown that an objective tuning methodology for a numerical weather prediction model performs at least as well as manual "expert tuning". This work is not entirely original, in that automated tuning schemes have been available for some years. Even so, the applicants propose to extend this work by examining the impact of soil moisture initialization, which is well known to be an important factor in the skill of numerical weather prediction models. Objective tuning has a number of advantages over expert tuning and could save a lot of time when a model is implemented over a new region or used in a different climate. If the proposed project has any weaknesses, it is that it appears to be proposing the incremental development of a tool rather than a fundamental aspect of science. Of course, this can also be a strength because of a likely higher probability of success.

Soundness of research methods and tools:

The project appears feasible scientifically and the tools proposed to be used are appropriate for the completion of the scientific goals of the project.

A couple of comments on the proposal in its current form. It would have been good to see further details on the methodology of the soil moisture initialization, as it is not clear from the proposal exactly how this is being done, and that is important for evaluating the feasibility of this proposal. Secondly, the authors themselves state that the objective tuning method may turn out to be too expensive to be performed routinely. This is indeed a risk but I suppose it is not as crucial a consideration when evaluating a research project of this kind at this stage, as computing power will increase in the future, making such an objective tuning methods more practical.

The authors also mention the possibility of constructing an objective tuning scheme for extreme events. It would be useful if the proposal had stated what extreme events the applicants had in mind, and given some explanation of why they believe that this might be feasible.

Appropriateness of project timeline and resources:

The technical justification for this project is quite impressive and it is obvious that the applicants are experienced users of similar computing systems and have thought a great deal about how well their code is likely to run on this system.

A significant concern that I have is that computer time granted to this group for a previous project on this computer system was not used particularly well, with less than 50% of the grant for computer time used. This is perhaps less of a concern in the proposed project, as apparently most of the scientific groundwork has been completed and so issues that are likely to decrease utilisation of computer time have been addressed. Nevertheless, there is still the potential for the implementation and testing of the soil moisture initialization to cause delays, which is why it would have been good to see more convincing details of how that is going to be done, and by who.

Another potential issue is that this proposal assumes the simultaneous running of six ensemble members and only 20% downtime, which would include periods of time when the ensemble members are not running due to crashes that had not yet been fixed by the project team. This seems to me to be quite optimistic, but perhaps the managers of this computing system would have a better feel for this issue. As a result, to accomplish this project, the applicants will need to be awarded the proposed resources in full, as they state.

PI's qualification:

The team appears clearly suited to completing this project. PI Voudouri has published work on the objective calibration scheme in this model and is very familiar with it, and has already participated in a closely related project. This multi-institutional team appears to have a wide range of relevant expertise. I have no real concerns about the team.

Request for clarifications:

Additional comments:

Scientific Merit:

Outstanding	Excellent	Good	Fair	Poor
☐	☒	☐	☐	☐

CSCS

Scientific Reviews

Scientific Review 2:

Significance of the proposed research:

Using optimization methods to calibrate model parameters objectively is relatively new in numerical weather prediction modeling community, as compared to hydrological, financial and communications communities. This is because calibration of NWP models poses special challenges, including dealing with high model complexity (e.g., a large of model parameters, needs to calibrate against to the observations of many different meteorological variables) and computational constraints. However, the potential gain in NWP model performance through model calibration is unmistakable. The needs for model calibration of atmospheric models including NWP models and climate models have been recognized more clearly, as shown in a recent Bulletin of American Meteorological Society (BAMS) paper by Hourdin et al. (2016), who elaborated on the issues, challenges, and potential approaches to model calibration of climate models. Another paper by Duan et al. (2017) described how automatic model calibration can be a new way to improve NWP model performance, which clearly demonstrated the significant improvement in 5-day precipitation forecasts in the Greater Beijing area. The PI and team of this proposal obviously understand the issues and have been working on them for a number of years. They have developed some proto-type, workable tools and have achieved significant understanding and promising results from prior related project. The tools may need further improvement and be tested more rigorously. I totally support their efforts and think that continuation of the research into objective calibration methods of NWP models is of high importance and should contribute to improved NWP model performance.

References:

Hourdin, F., T. Mauritsen, A. Gettelman, J. Golaz, V. Balaji, Q. Duan, D. Folini, D. Ji, D. Klocke, Y. Qian, F. Rauser, C. Rio, L. Tomassini, M. Watanabe, and D. Williamson, 2016: The art and science of climate model tuning. Bull. Amer. Meteor. Soc. doi:10.1175/BAMS-D-15-00135.1.

Duan, Q., Z. Di, J. Quan, C. Wang, W. Gong, Y. Gan, A. Ye, C. Miao, S. Miao, X. Liang, and S. Fan, (2017): Automatic model calibration - a new way to improve numerical weather forecasting. Bull. Amer. Meteor. Soc., 0, doi: 10.1175/BAMS-D-15-00104.1.

Soundness of research methods and tools:

Previous research results from related prior project by this team have demonstrated that using a meta-model approach to calibrate a number of critical parameters of COSMO model over a specific domain has resulted meaningful and significant improvements in NWP skill scores. This project intends to build on previous work by conducting additional experiments with an aim to build a permanent tool within the COSMO system. The new experiments include considering the effect of land soil memory, a higher spatial resolution, more diversified climatic conditions across different regions, and extreme events. Those experiments require significant amount of computational resources. Therefore, one of the key tasks is to find a balance between more reliable optimal solutions and computational resources available. Overall, I think the research methods and tools are reasonable. I believe that all the science objectives and milestones set in the proposal can be achieved. From my point of view, however, with the resources available from the piz Daint system and the expertise of the research team, the research goal can be even more ambitious than the one presented. My understanding is that the past experiments and the ones planned in the future consider only 6 model parameters. I suggest that the team considers expanding the list of tunable parameters for possible optimization. The team can refer to the works of Di et al. (2014) and Quan et al (2016), who used global sensitivity analysis (GSA) to identify the most sensitive parameters that exert significant influence over the simulation of meteorological variables of interest. They did GSA analysis using a limited number of model experiments. After the sensitive parameters are identified, the search for optimal parameters can be limited to those parameters only. Even though GSA would require additional computational resources, the computational requirement can be reduced by conducting sensitivity analysis on a limited spatial domain and with a limited simulation period. I also suggest the team to look into how Duan et al. (2017) optimized the parameters of WRF model using the Adaptive Surrogate Modeling based

Optimization (ASMO), developed by Wang et al. (2014), which may help them in the numerical experiment designs.

References:
Di, Z., Q. Duan, W. Gong, C. Wang, Y. Gan, J. Quan, J. Li, C. Miao, A. Ye, C. Tong, 2014. Assessing WRF Model Parameter Sensitivity: A Case Study with 5-day Summer Precipitation Forecasting in the Greater Beijing Area, Geophysical Research Letters, doi. 10.1002/2014GL061623
Quan, J., Z. Di., Q. Duan, W. Gong, C. Wang, Y. Gan, A. Ye and C. Miao, (2016), An evaluation of parametric sensitivities of different meteorological variables simulated by the WRF model, Q. J. R. Meteorol. Soc. 141, DOI:10.1002/qj.2885
Wang, C., Q. Duan, W. Gong, A. Ye, Z. Di, C., 2014. An evaluation of adaptive surrogate modeling based optimization with two benchmark problems, Environmental Modelling & Software, 60, P167-179, http://dx.doi.org/10.1016/j.envsoft.2014.05.026.

Appropriateness of project timeline and resources:

I have no specific comments on project timeline and resources. The amount of computational resources requested is probably enough for the tasks outlined. More resources may be helpful to ensure the successful execution of the project.

PI's qualification:

The PI has very solid meteorological background and has accumulated good experience in model calibration. The qualification is sufficient, but I suggest him to engage more with the model calibration community to keep abreast with the new development.

Request for clarifications:

None

Additional comments:

None

Scientific Merit:

Outstanding ☐ Excellent ☒ Good ☐ Fair ☐ Poor ☐

CSCS

Technical Reviews

Technical Review

Current application performance and scalability on requested CSCS platform according to the proposal:

The COSMO application has been ported and tuned to run at scale on the P100 GPUs on Piz Daint. Scalability results are shown for 96-144 nodes running the COSMO-1 model that is used in this project.

A configuration with 96 nodes is chosen, which would require 10,500 node hours to simulate one year according to the provided benchmark results.

Suitability of requested CSCS platform?

Yes [X] No [ ]

Has the applicant provided a breakdown of the amount of time requested to carry out the simulations?

Yes [X] No [ ]

Is the resource request consistent with the simulations proposed in the project plan?

The request is partially consistent. It has been padded significantly.

First, 34 simulations are required according to the formula provided, which is padded to 40 by adding "two control runs and 10% additional overhead", with no justification.

Second, 10,500 node hours per run is inflated to 13,000 by virtue of 20% for re running jobs with no justification.

Third, an additional 260,000 node hours are requested for "Different types of reduced configurations", which are not justified or discussed in the proposal.

Final Decision:

Accepted [ ] Rejected [ ] Conditionally accepted [X]

Reason for decision:

The COSMO-1 model is well suited to the requested system, and the main simulation request is sound. However, most of the padding should be removed and the unjustified post processing costs not granted, for a total request for 36 simulations of:

36*11,000 ~= 400,000 node hours