Project: **1202**
Project title: **Modelling the impact of global change on the terrestrial biosphere with LPJ-GUESS**
Principal investigator: **Matthew Forrest**
Report period: **2021-11-01 to 2022-10-31**

Migration to Levante

Following the migration, LPJ-GUESS has been configured and optimized for Levante. This required some changes to the previous setup as Levante has no local node storage whereas Mistral did. To counter this, two main optimizations have been implemented. Firstly, output files are compressed by the model before being written to disk, thus saving storage space and network traffic and reducing run time by approximately 20%. Secondly, the post-processing step of combining the individual output files from each process (LPJ-GUESS is a little archaic in the respect) has been moved from the HPC (where the standard submission framework ran it as a part of the batch job) to a local workstation, thus reducing time required on compute nodes. These optimizations, combined with the advanced capabilities of Levante have resulted in some truly excellent performance (details below).

Performance and scaling on Levante

As requested by the WLA, I have performed some performance calculations for Levante. For LPJ-GUESS spinup years, Levante needs 0.05 nodeh/year compared to my previous estimate of 2.3 nodeh/year (based on simulations at the Goethe-HLR cluster). For more computationally demanding simulations of the 20th century (which features more land use and thus more calculations) the estimate is 0.09 nodeh/year down from 3.5 nodeh/year. This is a speed increase of approximately 40 times in terms of nodeh. After allowing for the increased CPUs/node (Goethe-HKR cluster has 40 physical CPUs per node vs Levante's 128), this is still an increase of over one order of magnitude. Thus, the Levante system with the appropriately optimized run procedure provides an extremely well-suited environment for LPJ-GUESS simulations.

I also did some scaling tests by running the same simulation on 2, 4 and 8 nodes. This gave speed up factor of 1.95x for each 2x increase in CPU cores (more details in the request document). LPJ-GUESS does not include any interactions between gridcells and so no inter-core or inter-node communication required. Thus, the LPJ-GUESS is an "embarrassingly parallel" workload and near-perfect scaling is expected and indeed observed.

Simulation Status

The extensive simulation setup, input datasets and model version have prepared for the requested ISIMIP3a simulations (historical period) and these are now underway. The ISIMIP3a protocol for the combined biome and fire sectors has expanded from 8 runs per observed climate dataset to 12 runs. The number of observed climate datasets has expanded from 2 to 4. Initially I requested compute time for only the first priority observational climate datasets to reduce computation demand. However, based on initial results showing that fire models are rather sensitive to the climate datasets used and the high efficiency of Levante for LPJ-GUESS, I now intend to do the full set of runs (i.e. all four cobserved climate datasets). These will be completed in 2022 using the remaining compute allowance for 2022.

The future simulation runs (FURNACES and ISIMIP3b) have not yet commenced due to a number of factors. One is that the post-doctoral researcher from the FURNACES project has taken maternity leave for 2022 and will return February 2023 (the FURNACES project has a cost-neutral extension to allow this). Another factor is time constraints due to the extremely

time-consuming nature of the setting up the ISIMIP3a protocol simulations and other time commitments. Finally, the full set of inputs for ISIMIP3b are not yet available (although this only affects the so-called "Group 3" simulations with varying future socio-economic forcing data). The upshot is that, unfortunately, due to these delays I am requesting compute time for these simulations in 2023, although with much reduced node-hours to the effectiveness of Levante.

FirEUrisk simulations

In addition to the ISIMIP3b and FURNACES simulations, I am also requesting resources to support the FirEUrisk[1]. FirEUrisk is a large EU-funded project focusing on fire risk for Europe. Our role is to assess future fire risk with LPJ-GUESS-SPITFIRE and investigate different forest management strategies in term of reducing fire risk (see request document for more details.)

Summary and outlook

Much groundwork has been laid and progress made in 2022: optimising LPJ-GUESS for Levante, finalizing the LPJ-GUESS-SPITFIRE code version, preparing the input datasets and building the ISIMIP3a simulation framework and the commencing the ISMIP3a simulations. Consequently, the ISIMIP3a simulations should be completed by the end of the year. Two additional personnel available will be available for 2023 – one post-doc for the FURNACES project returning from maternity leave and another post-doc from the FirEUrisk project. Given the progress achieved in 2022 and the additional human power available for 2023, I feel confident that we are in a good position to take advantage of Levante to perform the requested simulations in 2023 and so fully supporting all three projects.

[1] https://fireurisk.eu/