Response for:

"Implementation and scaling of the fully coupled Terrestrial Systems Modeling Platform (TerrSysMP v1.0) in a massively parallel supercomputing environment – a case study on JUQUEEN (IBM BlueGene/Q)"

F. Gasper, K. Goergen, P. Shrestha, M. Sulis, J. Rihani, M. Geimer, and S. Kollet

Dear Dr. Simon Unterstrasser

Thank you very much for your feedback. In the following response we address all of your comments.

Your original comments are formatted in blue, while our response is formatted in black font.

Following the GMD Editorial (http://www.geosci-model-dev.net/6/1233/2013/gmd-6-1233-2013.pdf) a section on "Code availability" must be included, preferably at the end.

We added a new section in the manuscript with following content:

"Code availability

TerrSysMP is freely available for academic, non-profit research and application after registration at https://git.meteo.uni-bonn.de.*"*

It is common GMD practice to include the version number in the title and the abstract, not only in line 105.

We added the version number in the title and abstract.

Abstract, l24-25: the statement is too general. Please add something like "for a specific test case".

This sentence was changed in the manuscript as follows:

"In massively parallel supercomputer environments, the coupler OASIS-MCT is recommended, which resolves memory limitations that may be significant in case of very large computational

domains and exchange fields as they occur in these specific test cases and in many applications in terrestrial research."

Line 65: Could you add definitions of the terms high- and hyper-resolution. Or did I miss it?

In the manuscript the term is now explained as follows:

"The need for high- or hyper-resolution (e.g., 1km lateral grid spacing over continental computational domains), coupled simulations of the terrestrial system originates […]"

line 148: delete "currently" or update the numbers.

We edited the manuscript as suggested.

line148-150: trade-off or drawback?

It is a trade-off: Less CPU clock rate for higher reliability and smaller power/cooling requirements.

l171-173: sentence structure complicated!
Scalasca 1.4.3. was used as a profiling .....

We edited the sentence as follows:

"Scalasca 1.4.3 was used as a profiling and tracing tool to analyze the runtime behavior of TerrSysMP, identify performance bottlenecks and determine the optimum (static) load balance (i.e. resources allocation for each experimental setup)."

lines232-233: This would be clearer if you expand your explanation as done in your Final Author Comments (75%/12.5%/12.5%)

We edited the manuscript as follows:

"[…] since the load of the component models is roughly distributed like: 75% / 12.5% / 12.5%, which does not follow powers of two."

line240: T is not defined formally/mathematically correct. Function time T(x,y), what is x and y?

The equation (1) and its explanation were revised as follows:

"In both setups, the parallel efficiency $E_{nb}(n)$ *[%] in our study is defined as:*

$$E_{nb}(n) = \frac{1}{nb} \cdot \frac{T_1(n)}{T_{nb}(n)} \cdot 100 \qquad (1)$$

Where $T_1(n)$ is the runtime with one nodeboard and $T_{nb}(n)$ the runtime with nb nodeboards and the problem size n. Thus, in case of perfect parallel weak scaling without communication overhead, the simulation platform would exhibit an efficiency of $E_{nb}(n)$=100% for nb nodeboards and the problem size n."

Typos (mostly commas) in lines 65, 71, 201 and 241

All listed typos were corrected in the manuscript.