Project: **993**
Project title: **PalMod 4.4 Algorithmic and Implementation Performance Optimization**
Project lead: **Hendryk Bockelmann**
Report period: **2016-01-01 to 2016-12-31**

The work package 4.4 of the PalMod project (Paleo Modelling: A national paleo climate modelling initiative) is devoted to continuous performance optimization provision of the global fully coupled comprehensive Earth System Model MPI-ESM suitable for long-term transient paleoclimate simulations. The paleo-versions of the MPI-ESM under investigation are built upon the existing coarse and low resolution configurations MPI-ESM-CR and MPI-ESM-LR.

In the 2016 time frame we performed several in depth analysis of the aforementioned configurations on mistral to identify most time consuming portions of the code where optimizations should reveal most benefit. First the classical MPI parallel version of the ocean and atmosphere component of MPI-ESM was profiled per timestep wise – see Figure 1 and 2. It turned out that for the ocean component only the coupling timesteps cause a noticeable overhead in runtime compared to the other timesteps. While for the atmospheric component the radiation module was found as a major bottleneck in runtime. Therefore, the component-based parallelism for ECHAM was chosen to be implemented in 2017 – refer to the request of resources.

As a second step, we started analysing and optimising the hybrid MPI-OpenMP version and tried to figure out to what extend the Hyperthreading capabilities of Intel CPUs lead to better load balancing between the ocean and atmospheric component. This work is still ongoing, first results are show in Figure 3 where the MPI overhead in the ocean is still present while the coupling is almost negligible for ECHAM.

Finally, we used Mistral to investigate OpenACC performance portability between two different (hardware and software) platforms. We used one computing node from Mistral gpu partition and applied a benchmark to study OpenACC capability to maintain performance of the same code compiled with the PGI compiler for NVIDIA OpenACC Kepler K80 and compare it with the previous results obtained from the CRAY compiler and the NVIDIA Kepler K20x. We also used the same platform to study CUDA and OpenACC performance. In addition, we would like to apply some profiling and performance analyses of ECHAM6 on Mistral shortly.

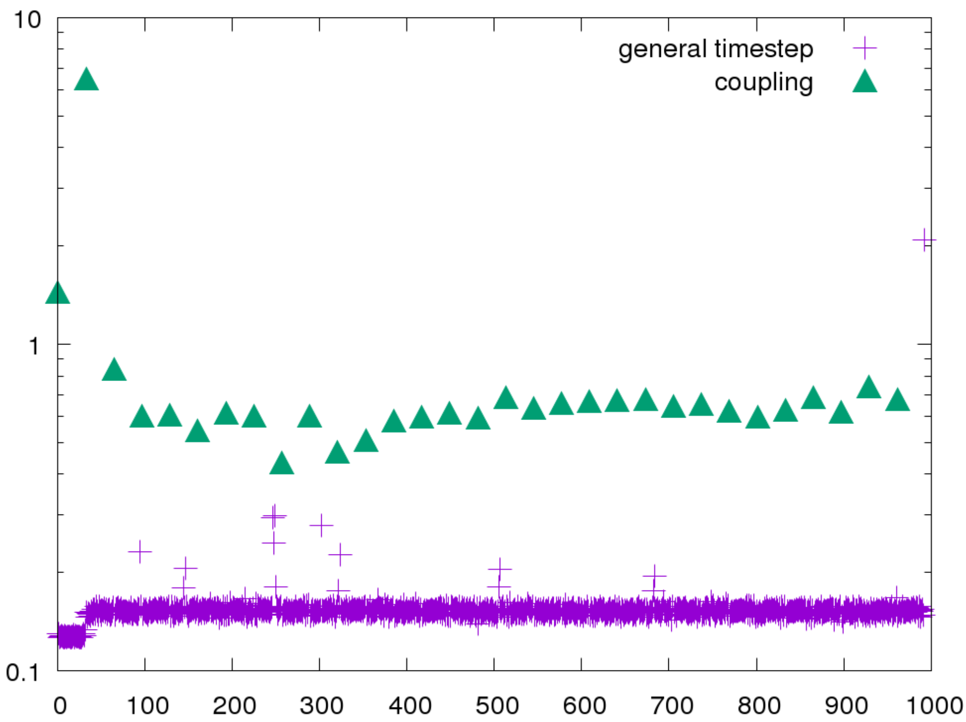

Figure 1: Analysis of runtime per timestep for MPIOM.

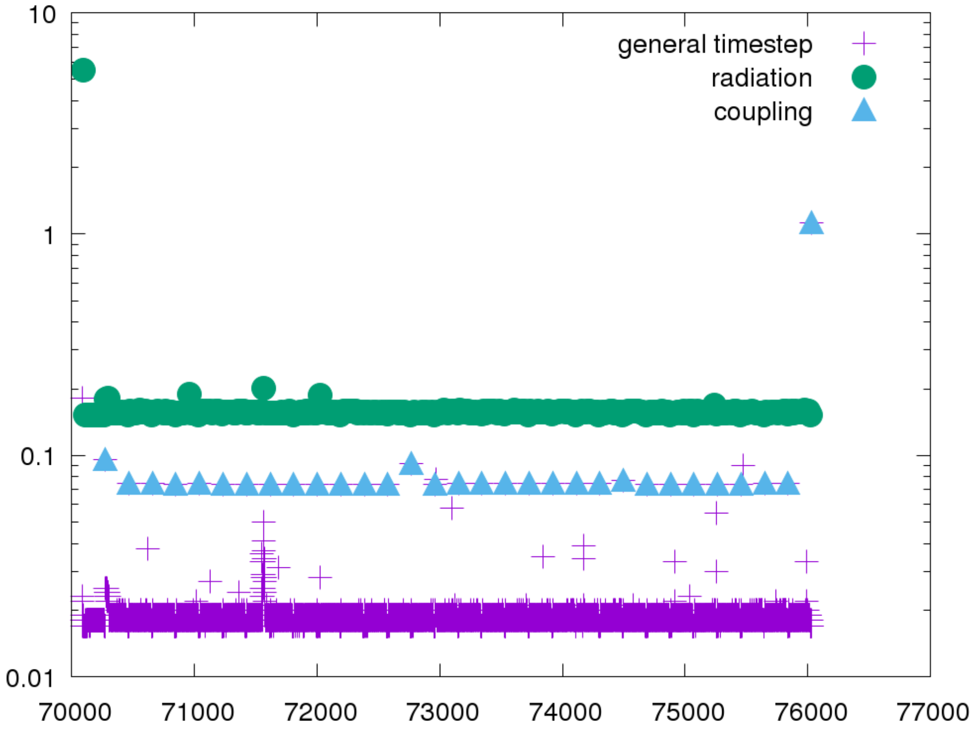

Figure 2: Analysis of runtime per timestep for ECHAM.

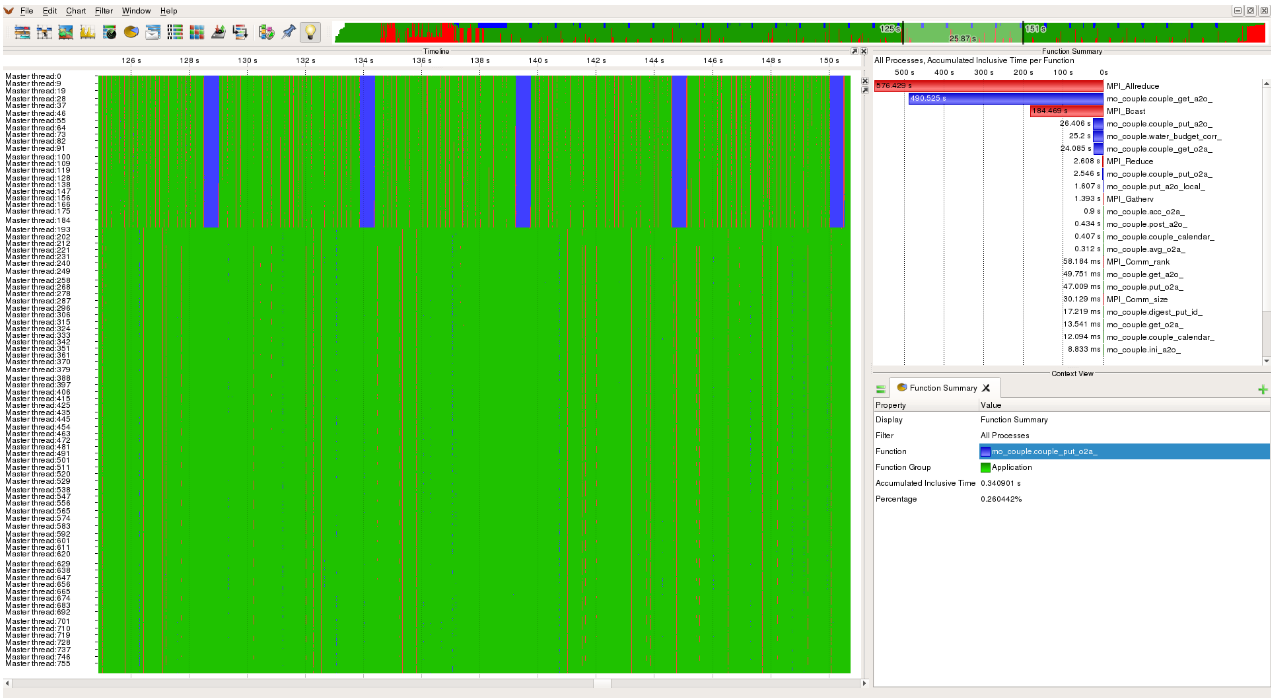

Figure 3: Tracing of MPI-ESM PalMod setting. Blue areas show the time in MPI comunication due to the coupler, while green areas are compute time for the model component