



How to use MPI communication in highly parallel climate simulations more easily and more efficiently.

Jörg Behrens, Moritz Hanke, and Thomas Jahns
DKRZ, Application Support, Hamburg, Germany (behrens@dkrz.de)

In this talk we present a way to facilitate efficient use of MPI communication for developers of climate models. Exploitation of the performance potential of today's highly parallel supercomputers with real world simulations is a complex task. This is partly caused by the low level nature of the MPI communication library which is the dominant communication tool at least for inter-node communication. In order to manage the complexity of the task, climate simulations with non-trivial communication patterns often use an internal abstraction layer above MPI without exploiting the benefits of communication aggregation or MPI-datatypes. The solution for the complexity and performance problem we propose is the communication library YAXT. This library is built on top of MPI and takes high level descriptions of arbitrary domain decompositions and automatically derives an efficient collective data exchange. Several exchanges can be aggregated in order to reduce latency costs. Examples are given which demonstrate the simplicity and the performance gains for selected climate applications.