Geophysical Research Abstracts
Vol. 14, EGU2012-2642, 2012
EGU General Assembly 2012

OpenGeoSys-GEMS: Hybrid parallelization of a reactive transport code with MPI and threads

G. Kosakowski (1), D.A. Kulik (1), and H. Shao (2)

(1) Paul Scherrer Institut, Department of Nuclear Energy and Safety Research, Villigen, Switzerland (georg.kosakowski@psi.ch), (2) Helmholtz Centre for Environmental Research - UFZ, Environmental Informatics, Leipzig, Germany

OpenGeoSys-GEMS is a generic purpose reactive transport code based on the operator splitting approach. The code couples the Finite-Element groundwater flow and multi-species transport modules of the OpenGeoSys (OGS) project (http://www.ufz.de/index.php?en=18345) with the GEM-Selektor research package to model thermodynamic equilibrium of aquatic (geo)chemical systems utilizing the Gibbs Energy Minimization approach (http://gems.web.psi.ch/). The combination of OGS and the GEM-Selektor kernel (GEMS3K) is highly flexible due to the object-oriented modular code structures and the well defined (memory based) data exchange modules.

Like other reactive transport codes, the practical applicability of OGS-GEMS is often hampered by the long calculation time and large memory requirements.
• For realistic geochemical systems which might include dozens of mineral phases and several (non-ideal) solid solutions the time needed to solve the chemical system with GEMS3K may increase exceptionally.
• The codes are coupled in a sequential non-iterative loop. In order to keep the accuracy, the time step size is restricted. In combination with a fine spatial discretization the time step size may become very small which increases calculation times drastically even for small 1D problems.
• The current version of OGS is not optimized for memory use and the MPI version of OGS does not distribute data between nodes. Even for moderately small 2D problems the number of MPI processes that fit into memory of up-to-date workstations or HPC hardware is limited.

One strategy to overcome the above mentioned restrictions of OGS-GEMS is to parallelize the coupled code. For OGS a parallelized version already exists. It is based on a domain decomposition method implemented with MPI and provides a parallel solver for fluid and mass transport processes.
In the coupled code, after solving fluid flow and solute transport, geochemical calculations are done in form of a central loop over all finite element nodes with calls to GEMS3K and consecutive calculations of changed material parameters. In a first step the existing MPI implementation was utilized to parallelize this loop. Calculations were split between the MPI processes and afterwards data was synchronized by using MPI communication routines.

Furthermore, multi-threaded calculation of the loop was implemented with help of the boost thread library (http://www.boost.org). This implementation provides a flexible environment to distribute calculations between several threads. For each MPI process at least one and up to several dozens of worker threads are spawned. These threads do not replicate the complete OGS-GEM data structure and use only a limited amount of memory. Calculation of the central geochemical loop is shared between all threads. Synchronization between the threads is done by barrier commands. The overall number of local threads times MPI processes should match the number of available computing nodes.

The combination of multi-threading and MPI provides an effective and flexible environment to speed up OGS-GEMS calculations while limiting the required memory use. Test calculations on different hardware show that for certain types of applications tremendous speedups are possible.