First Applications of the New Parallel Krylov Solver for MODFLOW on a National and Global Scale

Jarno Verkaik & Joseph D. Hughes & Edwin H. Sutanudjaja & Paul E.V. van Walsum

jarno.verkaik@deltares.nl

Overview

Integrated high-resolution hydrologic models are increasingly being used for evaluating water management measures at the field scale. Unfortunately, their drawbacks are large memory requirements and long run times.

An efficient technique for obtaining realistic run times and memory requirements is parallel computing, where the problem is divided over multiple processor cores. The new Parallel Krylov Solver (PKS) for MODFLOW is applied to the Netherlands Hydrological Instrument (NHI) model and the PCRaster Global Water Balance (PCR-GLOBWB) model. PKS uses the Message Passing Interface (MPI) and can be easily combined with Open Multi-Processing (OpenMP).

Mathematical model

PKS includes the Preconditioned Conjugate Gradient (PCG), BiConjugate Gradient Stabilized and Generalized Conjugate Residual linear Krylov solvers. A Restricted Additive Schwarz parallel preconditioner is applied:

$$M^{-1} = \sum_i \tilde{R}_i^T A_i^{-1} R_i$$

where A_i correspond to the interior coefficients of the linear system for the (overlapping) subdomain, $\tilde{R}_i$ and R_i are the restriction operators for the non-overlapping and overlapping subdomain, respectively. For the subdomain solve, an incomplete LU-factorization is used as a preconditioner.

For the preliminary results, we restrict ourselves to non-overlapping PCG exclusively. Furthermore, for the non-linear solution, we restrict ourselves to Picard iteration.

Implementation

PKS is largely based on the unstructured PCGU-solver for MODFLOW-USG and supports both structured and unstructured grids. For unstructured grids, PKS is fully integrated in MODFLOW-USG using METIS graph partitioning. For structured grids, PKS is integrated in iMODFLOW, an accelerated version of MODFLOW-2005 that is implicitly and online coupled to the MetaSWAP unsaturated zone model. In this code, both the MODFLOW and MetaSWAP cells share the same partitions, which can be generated uniformly or with the static Recursive Coordinate Bisection (RCB) method. Both codes with PKS support parallel input/output to separate files for each grid.

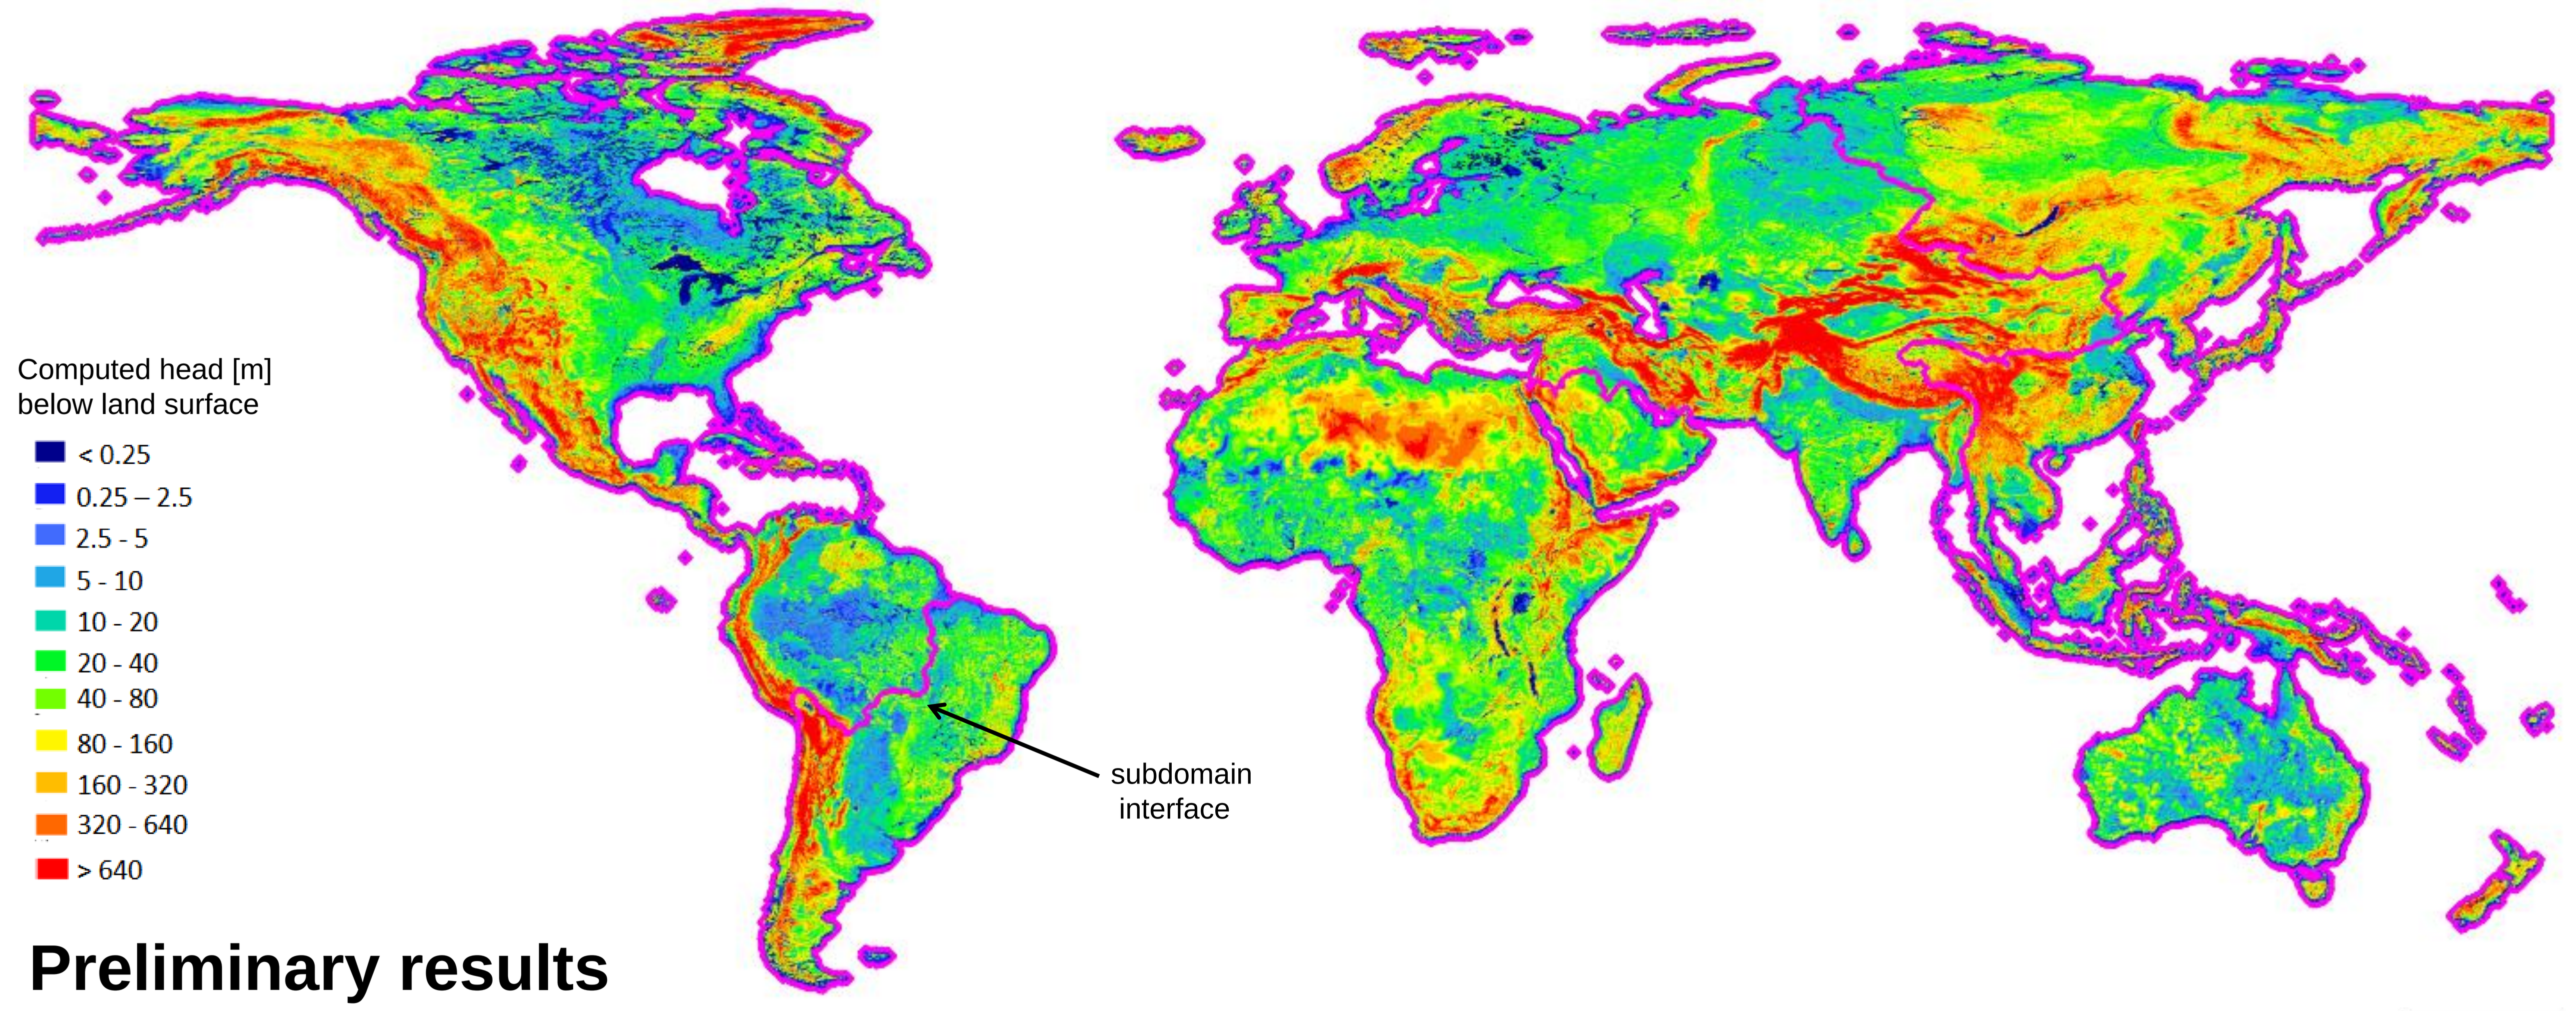

Preliminary results

Numerical experiments for both the NHI and PCR-GLOBWB model were carried out on the Cartesius Dutch National supercomputer. The Cartesius has ±40,000 computational cores and a fast InfiniBand interconnect. Experiments were done on nodes consisting of 2 Haswell 12-core CPUs (E5-2690 v3) with 64 GB RAM. Experiments for the NHI model were also carried out on a Windows machine, consisting of 2 Haswell 16-cores CPUs (E5-2698 v3) with 128 GB RAM.

The NHI is a state-of-the-art coupling of several models: a 7-layer confined groundwater iMODFLOW model (±10M 250m×250m cells), a MetaSWAP model for the unsaturated zone (Richards emulator of ±0.5M cells), and a surface water model (MOZART-DM). Transient groundwater flow is simulated for 2006 with daily timesteps for a) including surface water simulation on the Windows machine, b) excluding surface water simulation on the Cartesius. For a), surface water computes in serial and is off-line coupled. Stopping criteria are HCLOSE = 0.001m and RCLOSE = 100m3.

The PCR-GLOBWB model provides a grid-based representation of global terrestrial hydrology and this work uses the version that includes a 2-layer confined MODFLOW groundwater model (±4.5M 10km×10km cells). The benchmarks are for the stand-alone MODFLOW-USG model, simulating January 1992 with daily timesteps. Two variants were considered: a) the SIO variant (Structured Input/Out) were the model input/output is set-up in a structured way for 53 global catchments and each partition contains one or more catchments, and b) the NO-SIO variant, where METIS is used as a pre-processing and all input data are clipped to the generated partitions initially. Stopping criteria are HCLOSE = 0.001m and RCLOSE = 1.37e+06m3.

The scaling results show that the new PKS solver can result in significant speedups. We expect that for future research this can be improved by optimizing the load balance, reducing iterations (by overlap and/or two-level preconditioning), and applying OpenMP.

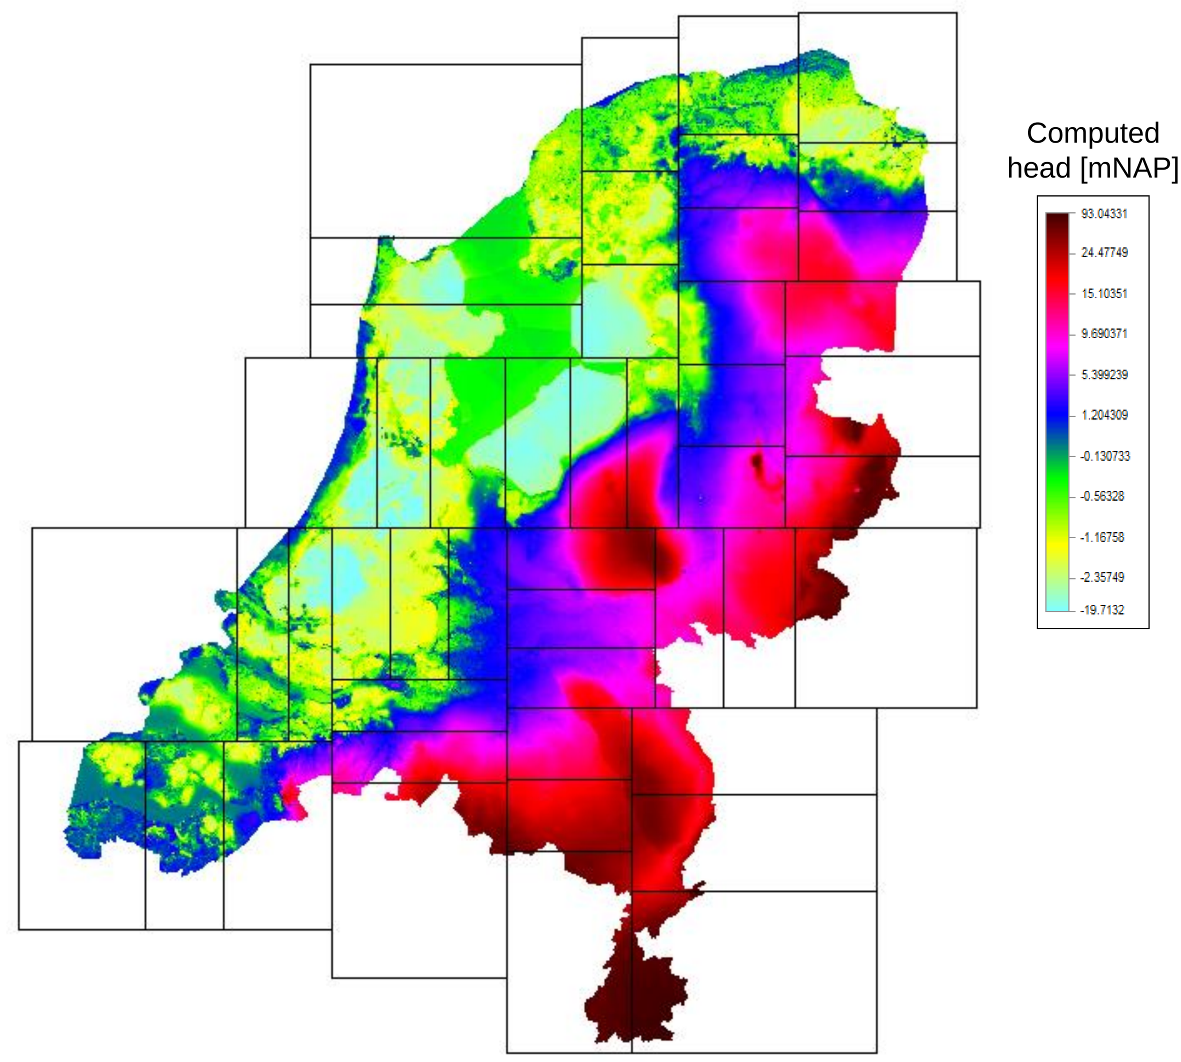

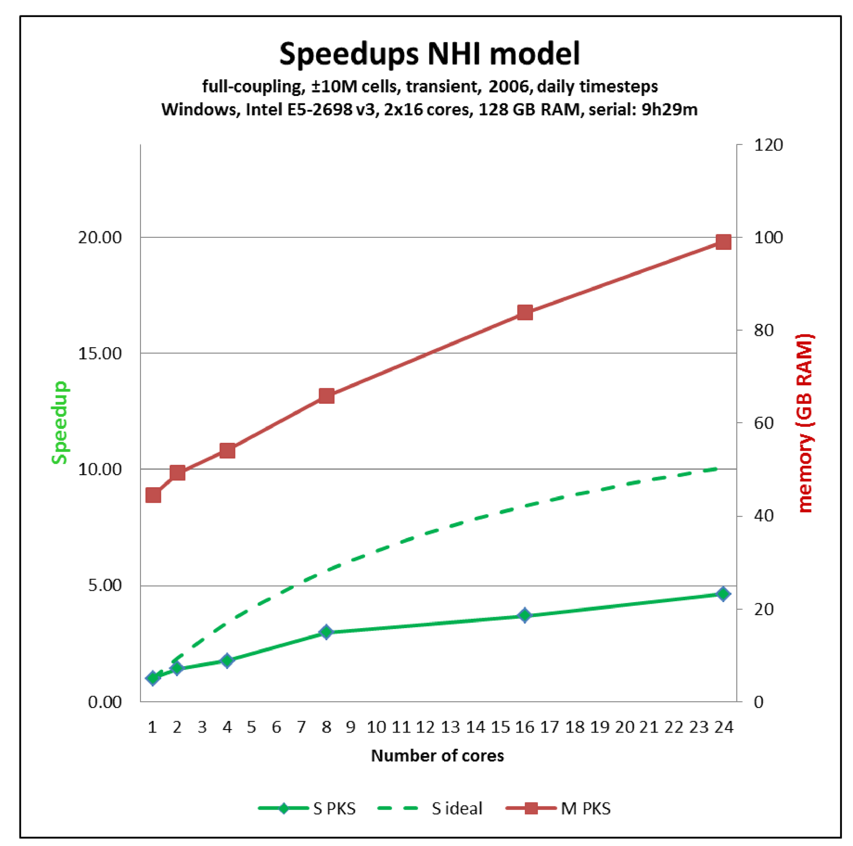

Measured speedups (green) and memory usage (red) for the full-coupled NHI model on the Windows machine. Serial computation takes 9h29m and requires 44.5 GB RAM.

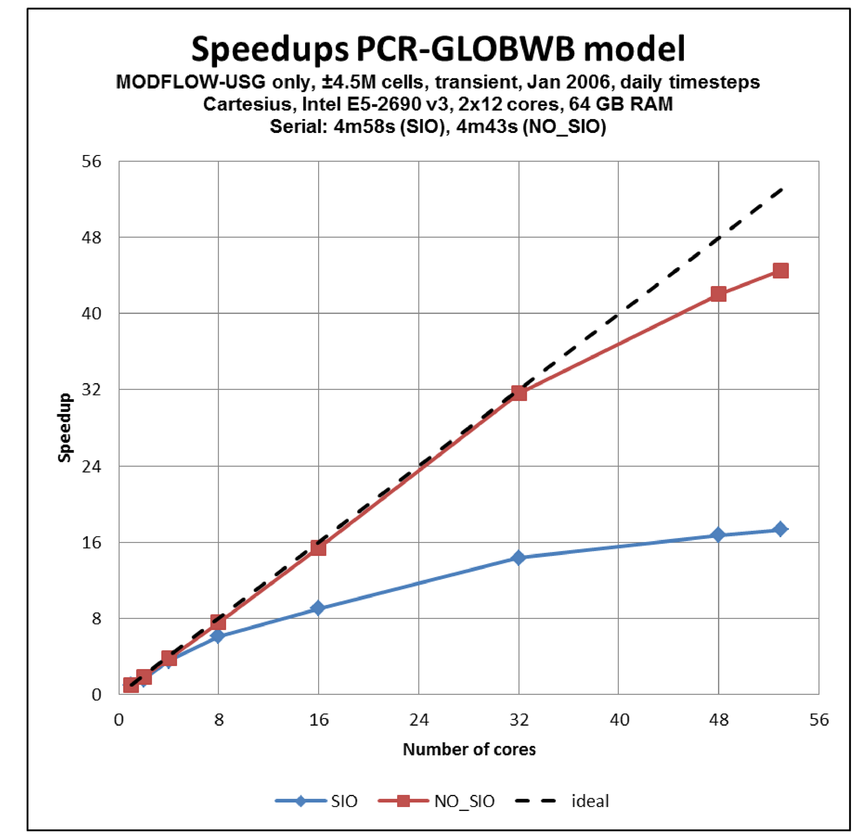

Measured speedups for the groundwater-only PCR-GLOBWB model on the Cartesius. Serial computation takes 4m58s seconds (SIO) and 4m43s seconds (NO_SIO), and requires 2.4 GB RAM.

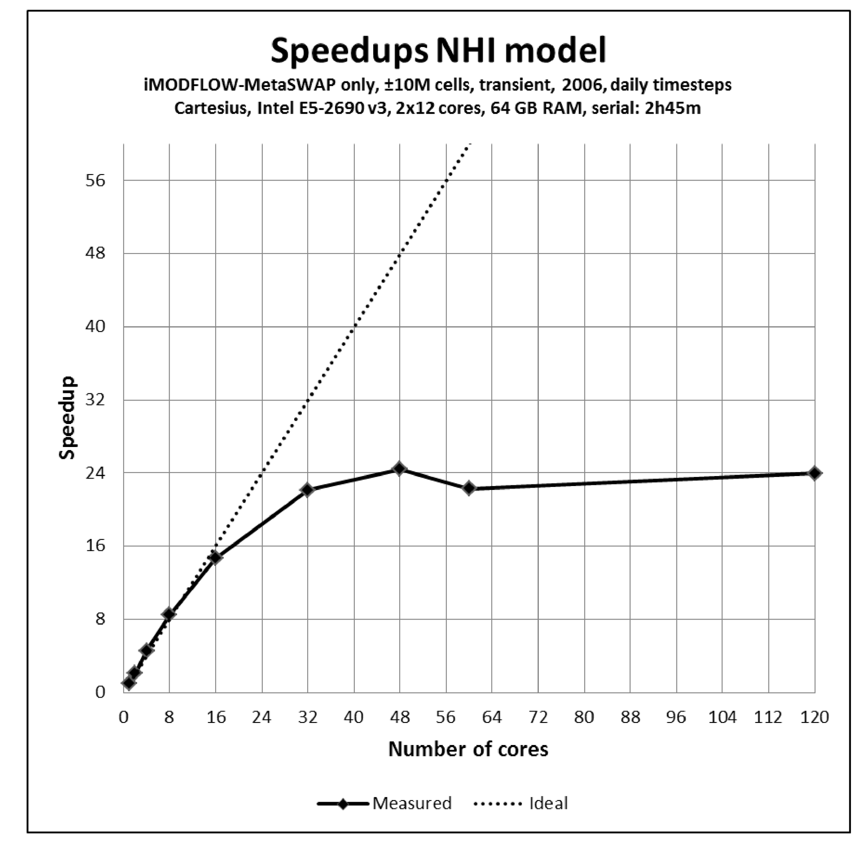

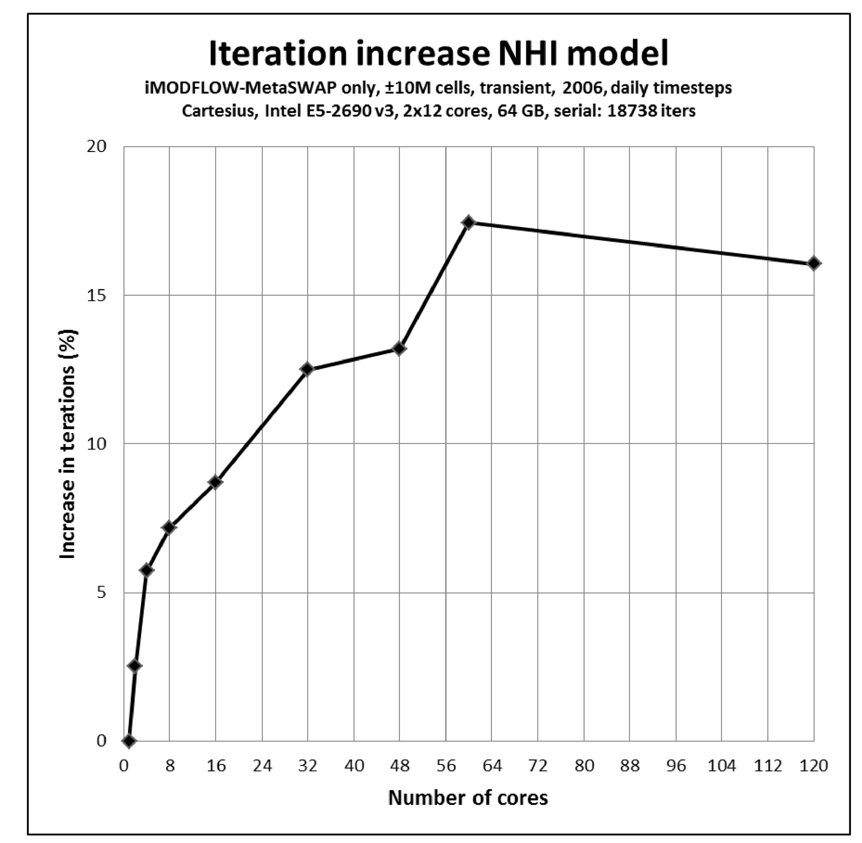

Measured speedups (top) and total inner iterations (bottom) for the groundwater-only NHI model on the Cartesius. Serial computation takes 2h45m and requires 2.4 GB RAM.