Comparison of Preconditioning Techniques for Optimizing a Nonhydrostatic, Parallel Tsunami Simulation Model

Annika Fuchs, Stephan Frickenhaus, Wolfgang Hiller and Natalja Rakowsky

Alfred Wegener Institute for Polar & Marine Research, Bremerhaven, Germany

Nonhydrostatic Simulation Model

The Tsunami Modelling Group at Alfred Wegener Institute developed an unstructured mesh Finite Element code named *TsunAWI* [1], based on the Shallow Water Model with a hydrostatic approach. For higher accuracy there is an extended, nonhydrostatic TsunAWI branch adapted from the idea to correct the hydrostatic velocity components after each time step [3]. This is expensive in memory requirement and computing time since the computation of three additional unknowns per grid node is required, just like the solution of several large, sparse systems of linear equations in each time step.

Optimization

Since the major percentage of the additional work is claimed by the systems of linear equations, it is reasonable to accelerate these for optimization. Firstly the mass matrices of two systems are approximated by *lumped matrices*, so they can be solved explicitly. The remaining system

$$Ax = b \qquad (1)$$

is solved by the Krylov Subspace Solution Method *GMRES*. As the convergence behaviour of this method depends on the properties of the matrix A, it is usefull to construct a *preconditioning matrix* K^{-1} that approximates A^{-1} and is cheap to compute.

Domain Decomposition

The nodes of the grid are distributed to several partitions / processors, fig. 1. After resorting the global indices block by block, the pattern of A looks like fig. 2 illustrates.

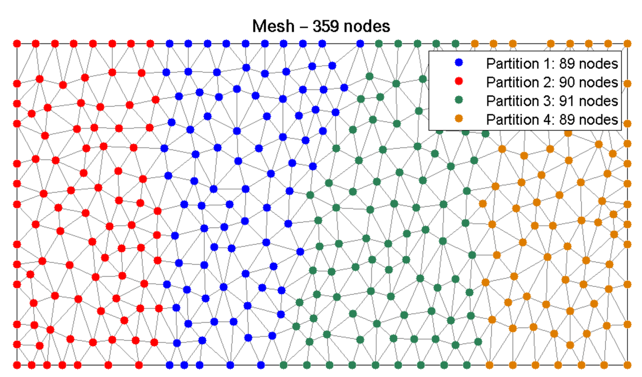

Figure 1: Decomposited grid. The colors indicate to which partition the grid nodes belong.

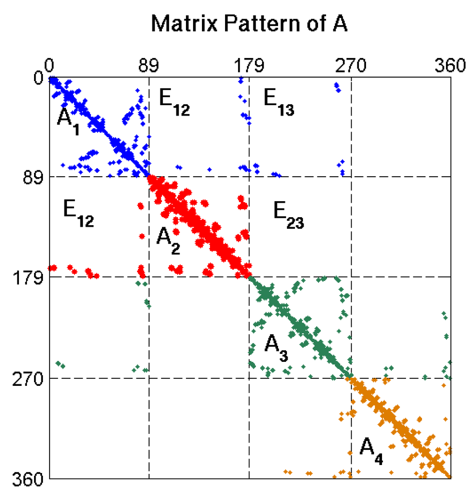

Figure 2: Corresponding distributed matrix A.

A single partition P_i has to solve a local part of the global system:

$$A_i x_i + \Sigma_j E_{ij} x_j = b_i \qquad (2)$$

in which the values of x_j are contained at other partitions.

Used software packages are METIS *4.0* from G. Karypis and V. Kumar and *pARMS, version 3.2* from Y. Saad et al with embedded own implementations.

Preconditioners

All following preconditioning techniques use the *Incomplete LU Factorization (ILU)*. Unfortunately, the number of nonzero entries of a complete LU decomposition of a sparse matrix calls for too much memory capacity. So the approximation $\tilde{L}\tilde{U} \approx A$ must be enough. There are several approaches how to force the sparsity of $\tilde{L}$ and $\tilde{U}$. Here, *ILU(2)*, *ILU(3)* and *ILUT* are tested.

The *Block Jacobi* preconditioning technique operates on the local diagonal block A_i, the offdiagonal blocks E_{ij} (compare fig. 2) are ignored. A_i is subject to an ILU factorization and the local preconditioning matrix is $K_i^{-1} = \{\tilde{U}_i^{-1}\tilde{L}_i^{-1}\}$. A better approximation of A^{-1} offers the *restricted additive Schwarz method* by communication of the interface nodes of neighbouring partitions. The extended local block A_i^{ext} (fig. 5) is submitted to an ILU factorization.

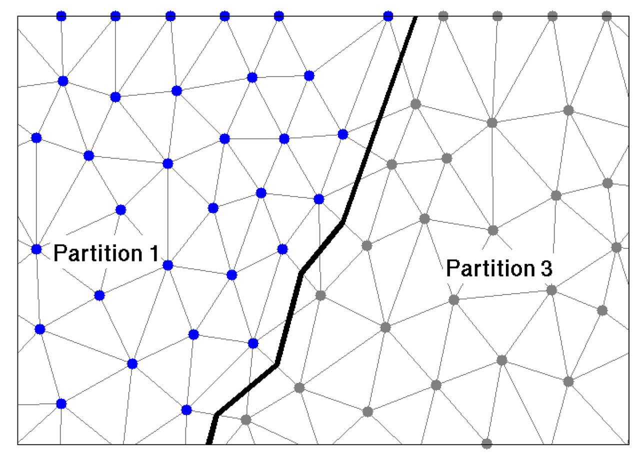

Figure 3: Strict isolation of different partitions.

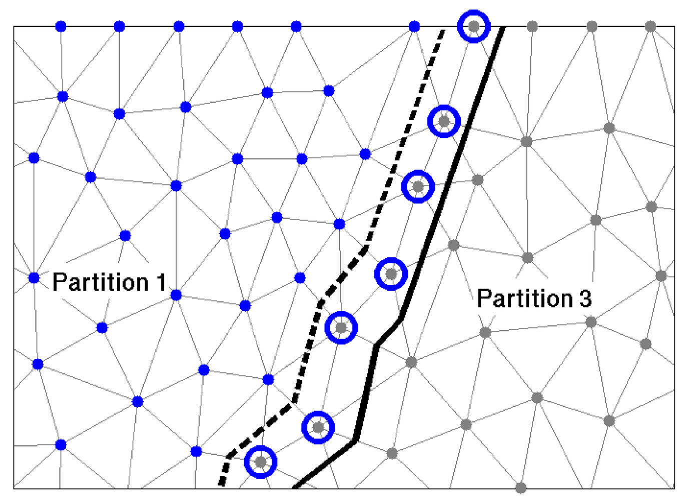

Figure 4: Partitions with overlap 1.

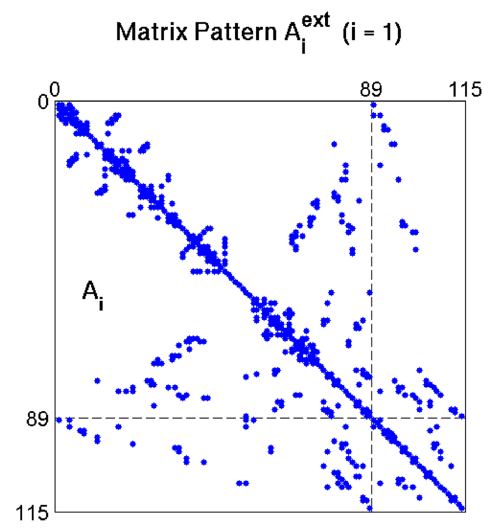

Figure 5: An extended local block A_i^{ext}.

By grouping local grid nodes in two categories, interior ones and interface nodes, (2) can be written as

$$\underbrace{\begin{pmatrix} B_i & F_i \\ E_i & C_i \end{pmatrix}}_{A_i} \underbrace{\begin{pmatrix} u_i \\ v_i \end{pmatrix}}_{x_i} + \begin{pmatrix} 0 \\ \Sigma_j E_{ij} v_j \end{pmatrix} = \underbrace{\begin{pmatrix} f_i \\ g_i \end{pmatrix}}_{b_i}. \qquad (3)$$

and this leads to

$$u_i = B_i^{-1}(f_i - F_i v_i), \qquad (4)$$

$$S_i v_i + \sum_j E_{ij} v_j = g_i - E_i B_i^{-1} f_i, \qquad (5)$$

with $S_i = C_i - E_i B_i^{-1} F_i$.

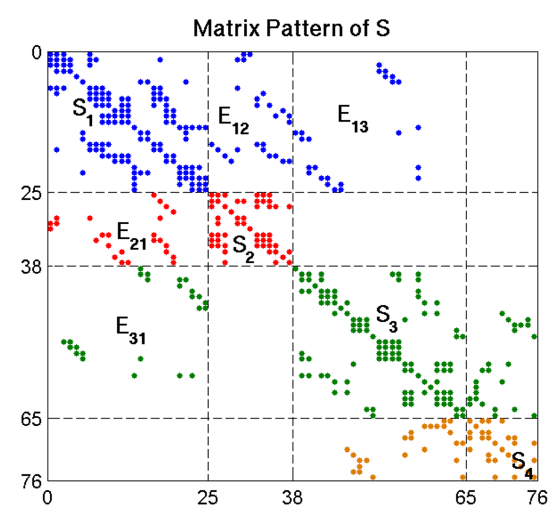

Figure 6: The global Schur matrix.

Via an ILU factorization of the local *Schur complement* S_i, (5) can be solved approximately by an inner GMRES solver. Another possibility is to apply the restricted additive Schwarz technique to the global Schur matrix (fig. 6) and factorize S_i^{ext} (*SchurRAS* [2]). After this, equation (4) can be solved by the results at the interface nodes.

Results

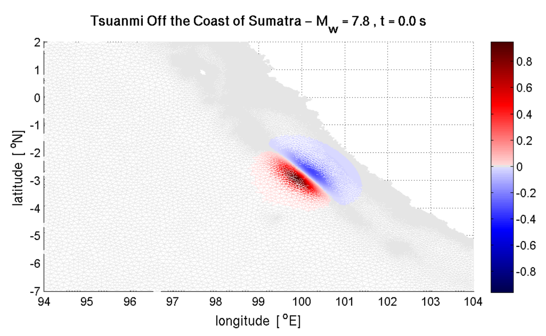

Figure 7: Initial surface elevation.

As test case serves a tsunami simulation off the coast of Sumatra. The used mesh consists of 692061 nodes and 1256019 elements. An earthquake with magnitude 7.8 initiates the scenario, fig. 7. Each combination of preconditioner (BJ, rAS, Schur+GMRES, SchurRAS) and ILU factorization (ILU(2), ILU(3) and ILUT) is tested with 2, 4, 8 and 16 POWER6 processors. Alltogether there are 48 runs with measurements of setup time t^i_{setup} of the preconditioner and the time required to solve system (1) t^i_{apply} in each time step i. Fig. 8 presents the average values of the first 1800 time steps ($t_{end} = 30$ min).

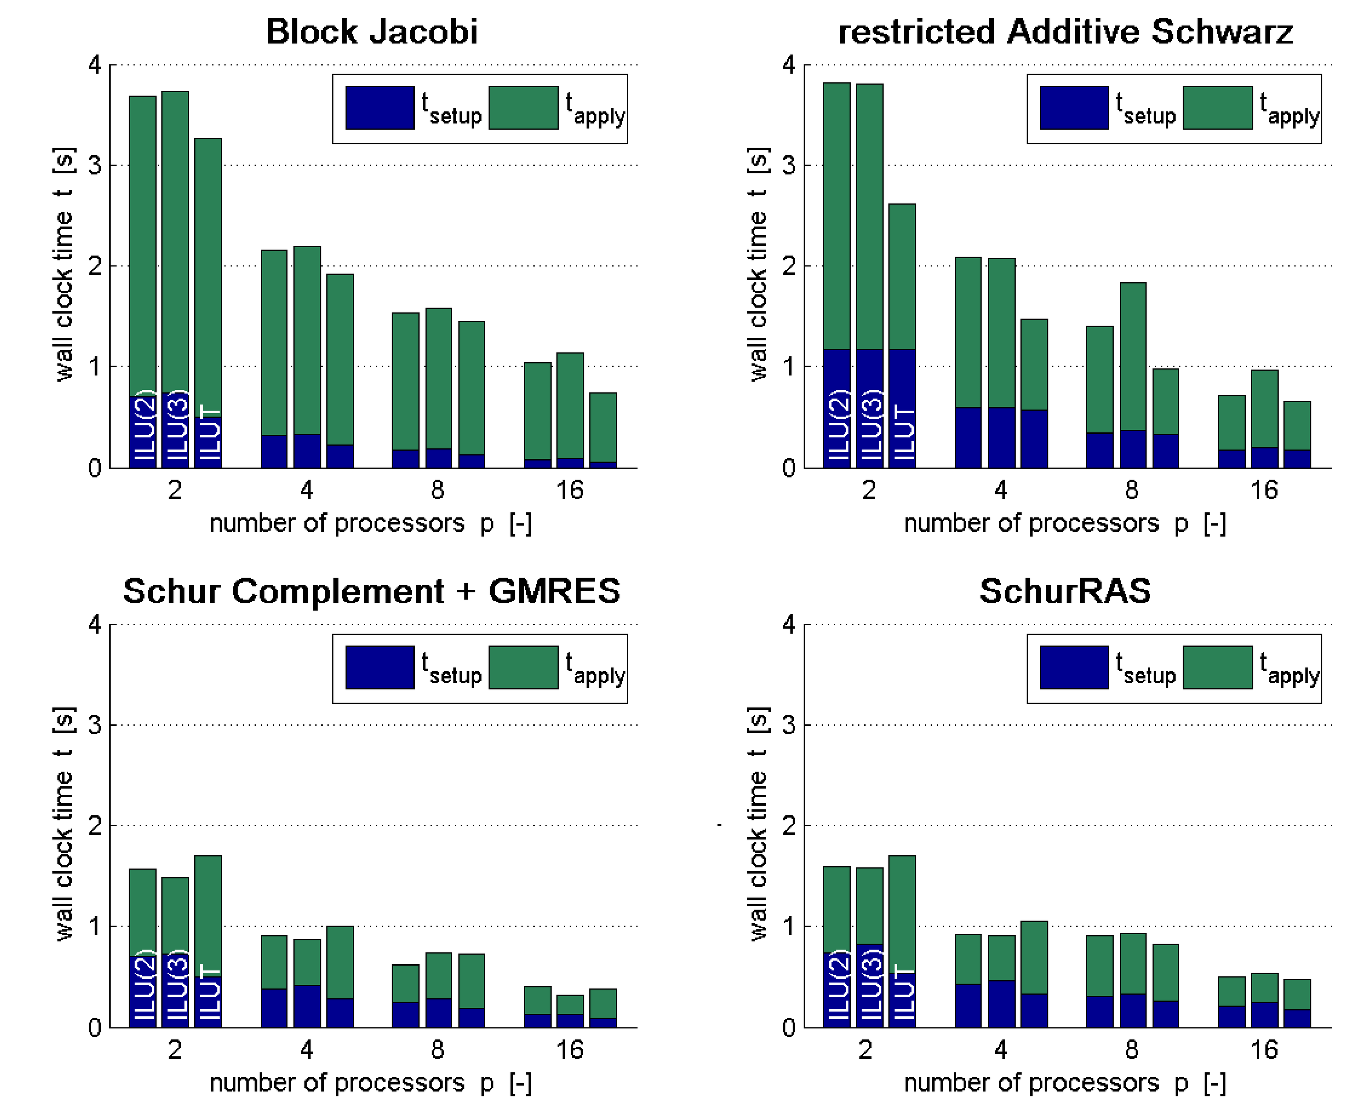

Figure 8: Average time duration per time step.

Due to the communication of interface values (what costs setup time) a better approximation of A^{-1} is available which accelerates the convergence rate. By factorization of the local or extended Schur matrix the system is reduced to a smaller one what can be solved very fast.

References

[1] Harig, S., Chaeroni, C., Pranowo, W.S., Behrens, J.: *Tsunami simulations on several scales: Comparison of approaches with unstructured meshes and nested grids.* Journal of Ocean Dynamics, 2008, 58:429-440.

[2] Li, Z., Saad, Y.: *SchurRAS: A restricted version of the overlapping Schur complement preconditioner.* Report umsi-2004-76, Univ. of Minnesota, Minnesota Supercomputer Institute, 2004. Available at: http://www-users.cs.umn.edu/ saad/PDF/umsi-2004-76.pdf

[3] Walters, R. A.: *A semi-implicit finite element model for non-hydrostatic (dispersive) surface waves.* Int. J. Numer. Meth. Fluids, 2005, 49:721-737.

Annika Fuchs • Alfred Wegener Institute, Am Handelshafen 12, 27570 Bremerhaven • Ph: +49(471)4831-2904 • E-mail: Annika.Fuchs@awi.de • Web: http://www.awi.de