41

An Efficient FGMRES Solver for the Shallow Water Equations based on Domain Decomposition

Serge Goossens, Kian Tan, and Dirk Roose

1 Shallow Water Equations

The *Shallow Water Equations* (SWE) are a set of nonlinear hyperbolic equations, describing long waves relative to the water depth. Physical phenomena such as tidal waves in rivers and seas, breaking of waves on shallow beaches and even harbour oscillations can be modelled successfully with the SWE. The 3D SWE (1.1)–(1.3) given below for Cartesian (ξ, η) coordinates are based on the *hydrostatic assumption*, that the influence of the vertical component of the acceleration of the water particles on the pressure can be neglected.

$$\frac{\partial u}{\partial t} + u\frac{\partial u}{\partial \xi} + v\frac{\partial u}{\partial \eta} + \frac{\omega}{H}\frac{\partial u}{\partial \sigma} - fv + g\frac{\partial \zeta}{\partial \xi} - \nu_H\left(\frac{\partial^2 u}{\partial \xi^2} + \frac{\partial^2 u}{\partial \eta^2}\right) - \frac{1}{H^2}\frac{\partial}{\partial \sigma}\left(\nu_V\frac{\partial u}{\partial \sigma}\right) = 0 \tag{1.1}$$

$$\frac{\partial v}{\partial t} + u\frac{\partial v}{\partial \xi} + v\frac{\partial v}{\partial \eta} + \frac{\omega}{H}\frac{\partial v}{\partial \sigma} + fu + g\frac{\partial \zeta}{\partial \eta} - \nu_H\left(\frac{\partial^2 v}{\partial \xi^2} + \frac{\partial^2 v}{\partial \eta^2}\right) - \frac{1}{H^2}\frac{\partial}{\partial \sigma}\left(\nu_V\frac{\partial v}{\partial \sigma}\right) = 0 \tag{1.2}$$

$$\frac{\partial \zeta}{\partial t} + \frac{\partial (Hu)}{\partial \xi} + \frac{\partial (Hv)}{\partial \eta} + \frac{\partial \omega}{\partial \sigma} = 0 \tag{1.3}$$

We denote by ζ the water elevation above some plane of reference, hence the total water depth is given by $H = d + \zeta$, where d is the depth below this plane of reference. The scaled vertical coordinate $\sigma = \dfrac{z - \zeta}{d + \zeta}$ varies between -1 at the bottom and 0 at the free surface. The velocities in the ξ- and η-directions are denoted by u and v respectively, while ω represents the transformed vertical velocity. The parameter f accounts for the Coriolis force due to the rotation of the Earth. The viscosity is modelled using ν_H and ν_V. In each σ-plane ν_H models the "horizontal" viscosity, while ν_V describes the viscosity in the vertical (σ) direction.

Ninth International Conference on Domain Decomposition Methods
Editor Petter E. Bjørstad, Magne S. Espedal and David E. Keyes

2 Alternating Operator Implicit Method

For the time integration we use the two-stage Alternating Operator Implicit (AOI) time splitting method, which has been developed at Delft Hydraulics [dG93]. This method is unconditionally stable and second order accurate in time. In the first stage (most of) the advection and diffusion terms in the momentum equations are handled implicitly, while the continuity equation is integrated explicitly. The resulting two linear systems for the intermediate u and v are solved by Red-Black Gauss-Seidel iterations. During the second stage the continuity equation is treated implicitly. Substitution into the continuity equation of the momentum equations, in which the velocity components are now handled explicitly, leads to a nonlinear system for the water elevation ζ. For each time step n, we perform Q fixed point iterations to solve this nonlinear system. Introducing an iteration counter q ($q = 1, 2, \ldots, Q$) and multiplying the pressure terms in the momentum equations with $H^{(n,q)}/H^{(n,q+1)}$, we obtain

$$\left(I - \nu_\xi \frac{\partial^2}{\partial \xi^2} - \nu_\eta \frac{\partial^2}{\partial \eta^2}\right) \zeta^{(n,q)} = f^*, \tag{2.4}$$

where the right-hand side f^* involves previously computed values and where $\zeta^{(n,q)}$ denotes the water elevation at iteration q of time step n. In the remainder of the paper we drop the superscripts. The imposed boundary conditions might be of Neumann type (e.g. closed wall) which could lead to a nonsymmetric linear system after discretisation. The pseudo viscosities ν_ξ and ν_η mainly depend on the time step and the total depth, which makes the linear system *nonsymmetric*.

Since the classical five point star stencil is used, a discrete equation of the form

$$(b_{i,j} + b_{i,j}^{(x)} + b_{i,j}^{(y)})\zeta_{i,j} + a_{i,j}\zeta_{i-1,j} + c_{i,j}\zeta_{i+1,j} + d_{i,j}\zeta_{i,j-1} + e_{i,j}\zeta_{i,j+1} = f_{i,j} \tag{2.5}$$

is obtained for each grid point (i, j) and the resulting linear system has a pentadiagonal structure. In practice it often suffices to take $Q = 2$. Until recently an ADI iteration was used for solving system (2.4).

The main topic addressed in this paper is the application of a *Domain Decomposition Preconditioner* in combination with the *Flexible GMRES* (FGMRES) method to solve (2.4). The original ADI method is used as a preconditioner in the subdomain solver only.

3 Generalised Additive Schwarz Preconditioner

The domain decomposition preconditioner which is employed in accelerating the FGMRES method to solve (2.4) on the entire domain is based on a *Generalised Additive Schwarz Preconditioner* (GASP). Let $R_i : \Omega \mapsto \Omega_i$ denote the (linear) restriction operator that maps onto subdomain i by selecting the components corresponding to this subdomain. The matrix $M_i = R_i A R_i^T$ denotes the principal submatrix of the matrix A associated with subdomain Ω_i. The result of applying the GASP can be written as a sum of the extensions of the solutions of independent

subdomain problems, which can be solved in parallel.

$$M^{-1} = \sum_{i=1}^{p} R_i^T M_i^{-1} R_i \tag{3.6}$$

We elaborate on this GASP for the case of two subdomains separated by the interface Γ as shown in Fig. 1. Extension to more subdomains is straightforward. At the heart of our GASP lies an *extension* of the subdomains to (physically) slightly overlapping grids. With a proper definition of the overlap, the restrictions R_i can be defined in such a way that the original discretisation is "distributed" across the subdomain operators M_i. Since the classical five point star stencil is used an overlap of two grid lines is sufficient. Figure 2 illustrates the extension process. In the discretisation, points

Figure 1 Grid before partitioning

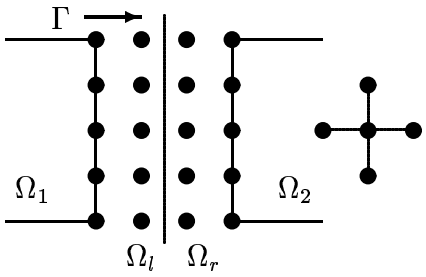

Figure 2 Grid after partitioning

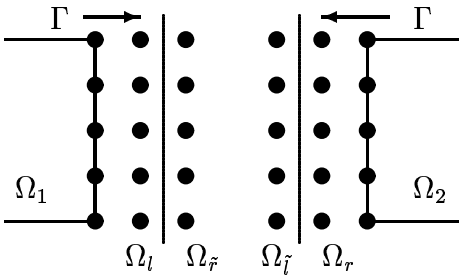

in subdomain Ω_1 are only connected to points in Ω_1 or in Ω_l. Similar statements can be made about the points in Ω_l, Ω_r and Ω_2. This leads to the following block structured linear system.

$$\begin{pmatrix} A_{11} & A_{1l} & 0 & 0 \\ A_{l1} & A_{ll} & A_{lr} & 0 \\ 0 & A_{rl} & A_{rr} & A_{r2} \\ 0 & 0 & A_{2r} & A_{22} \end{pmatrix} \begin{pmatrix} \zeta_1 \\ \zeta_l \\ \zeta_r \\ \zeta_2 \end{pmatrix} = \begin{pmatrix} f_1 \\ f_l \\ f_r \\ f_2 \end{pmatrix} \tag{3.7}$$

After extension towards overlap, and thus duplication of Ω_l and Ω_r into $\Omega_{\tilde{l}}$ and $\Omega_{\tilde{r}}$, we obtain an enhanced system of equations in which we still have to specify the relation between the "overlapping" unknowns. The obvious way is just to state that the values in the duplicated subdomains $\Omega_{\tilde{l}}$ and $\Omega_{\tilde{r}}$ should be copied from the values in the original subdomains Ω_l and Ω_r respectively. This is known as the *Dirichlet-Dirichlet* (DD) coupling. The enhanced system of equations with this DD coupling can be written as follows.

$$\begin{pmatrix} A_{11} & A_{1l} & 0 & 0 & 0 & 0 \\ A_{l1} & A_{ll} & A_{lr} & 0 & 0 & 0 \\ 0 & 0 & I & 0 & -I & 0 \\ 0 & -I & 0 & I & 0 & 0 \\ 0 & 0 & 0 & A_{rl} & A_{rr} & A_{r2} \\ 0 & 0 & 0 & 0 & A_{2r} & A_{22} \end{pmatrix} \begin{pmatrix} \zeta_1 \\ \zeta_l \\ \tilde{\zeta}_r \\ \tilde{\zeta}_l \\ \zeta_r \\ \zeta_2 \end{pmatrix} = \begin{pmatrix} f_1 \\ f_l \\ 0 \\ 0 \\ f_r \\ f_2 \end{pmatrix} \tag{3.8}$$

Tan [TD88] showed that the spectral radius of the preconditioned operator AM^{-1} and thus the convergence properties of a *Krylov Subspace Method* preconditioned by a GASP as given by (3.6), are improved by pre-multiplying the linear system with a properly chosen nonsingular matrix P of the form.

$$P = \begin{pmatrix} I & 0 & 0 & 0 & 0 & 0 \\ 0 & I & 0 & 0 & 0 & 0 \\ 0 & 0 & C_{lr} & -C_{ll} & 0 & 0 \\ 0 & 0 & -C_{rr} & C_{rl} & 0 & 0 \\ 0 & 0 & 0 & 0 & I & 0 \\ 0 & 0 & 0 & 0 & 0 & I \end{pmatrix} \tag{3.9}$$

This can also be interpreted in terms of imposing more general conditions at the subdomain interfaces. This approach was originally introduced by Lions [Lio90] and subsequently used by e.g. Hagstrom et. al. [HTJ88] and Nataf and Rogier [NR95]. The submatrices C_{lr}, C_{ll}, C_{rr} and C_{rl} are chosen to achieve a clustering of the eigenvalues of the preconditioned operator, subject to the condition that P remains nonsingular. This gives rise to the *Locally Optimised Block Jacobi* preconditioners which are thus based on the enhanced system of equations $A\zeta = f$:

$$\begin{pmatrix} A_{11} & A_{1l} & 0 & 0 & 0 & 0 \\ A_{l1} & A_{ll} & A_{lr} & 0 & 0 & 0 \\ 0 & C_{ll} & C_{lr} & -C_{ll} & -C_{lr} & 0 \\ 0 & -C_{rl} & -C_{rr} & C_{rl} & C_{rr} & 0 \\ 0 & 0 & 0 & A_{rl} & A_{rr} & A_{r2} \\ 0 & 0 & 0 & 0 & A_{2r} & A_{22} \end{pmatrix} \begin{pmatrix} \zeta_1 \\ \zeta_l \\ \tilde{\zeta}_r \\ \tilde{\zeta}_l \\ \zeta_r \\ \zeta_2 \end{pmatrix} = \begin{pmatrix} f_1 \\ f_l \\ 0 \\ 0 \\ f_r \\ f_2 \end{pmatrix} \tag{3.10}$$

This enhanced system of equations can be written in terms of the 3×3 blocks. Defining the restriction operators R_1 and R_2 in terms of the index sets corresponding to ζ_1, ζ_l and $\tilde{\zeta}_r$ on the one hand and $\tilde{\zeta}_l$, ζ_r and ζ_2 on the other hand, the GASP can be written as the block diagonal matrix M with

$$M = \begin{pmatrix} R_1 A R_1^T & 0 \\ 0 & R_2 A R_2^T \end{pmatrix}. \tag{3.11}$$

4 Flexible GMRES

Applying FGMRES in combination with the GASP described above to solve (2.4) is straightforward. The FGMRES method developed by Saad [Saa93] is a Krylov subspace method which allows the introduction of a set of well-chosen vectors in the search space. We assume for convenience that AM^{-1} is normal. The FGMRES algorithm computes the fundamental relation

$$AZ_m = V_m H_m + h_{m+1,m} v_{m+1} e_m^T \tag{4.12}$$

where $Z_m = \begin{pmatrix} z_1 & z_2 & \dots & z_m \end{pmatrix}$ is the matrix containing the search directions and the matrix $V_m = \begin{pmatrix} v_1 & v_2 & \dots & v_m \end{pmatrix}$ is defined by its columns. The matrix H_m is

a square $m \times m$ upper Hessenberg matrix whose elements are computed during the orthogonalisation process of the v-vectors, consequently $V_m^T V_m = I$. Using a fixed preconditioner with FGMRES is equivalent to using right preconditioned GMRES with this preconditioner. In this case, the matrix Z_m can be computed by applying the fixed preconditioning operator M^{-1} to the matrix V_m.

$$Z_m = M^{-1} V_m \tag{4.13}$$

The GASP described above is a fixed preconditioner if and only if the linear systems in the subdomains are solved to full precision. The convergence of FGMRES is mainly governed by the eigenvalue distribution of the preconditioned operator AM^{-1}. In particular, convergence acceleration can be expected for well-separated extreme eigenvalues. Also, the stage of the FGMRES process in which acceleration occurs is related to the convergence of the Ritz values to extreme eigenvalues [VdSVdV86, VdVV93]. This phenomenon can be made visible by explicitly computing the Ritz

Figure 3 Spectrum of H_m: Dirichlet-Dirichlet Coupling

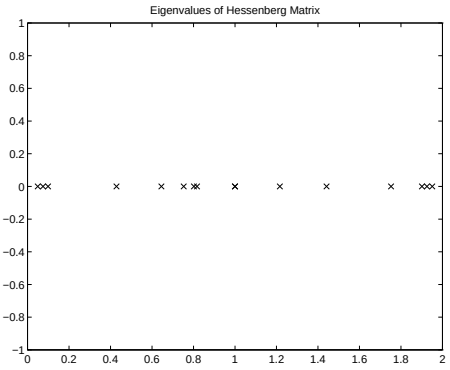

Figure 4 Spectrum of H_m: Locally Optimised Coupling

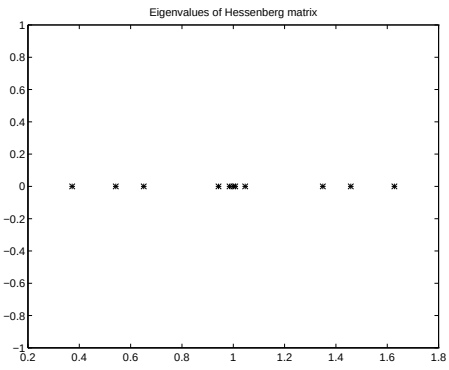

values, i.e. the eigenvalues of $H_m = V_m^T (AM^{-1}) V_m$, in the course of the FGMRES process. The eigenvectors corresponding to the outliers of this spectrum represent the eigenvector components to be removed from the initial residual which "uphold" the convergence. Due to the construction of the GASP considered here, this Ritz spectrum, at least for meshes not too fine, typically resembles the spectrum as depicted in Fig. 4, i.e. a few well-separated outliers and a cluster around 1. For comparison we show in Fig. 3 the Ritz spectrum of the domain decomposition preconditioner for the same problem when Dirichlet-Dirichlet coupling is used. This spectrum does not show a clear separation of a cluster of eigenvalues around 1 and some outliers. On the contrary, the eigenvalues are spread out over the open interval $(0, 2)$ and a lot of the eigenvalues are either close to 0 or to 2. The eigenvalue distribution explains the slow convergence of this domain decomposition method with DD coupling when it is used as a solver, because the spectral radius of the matrix $(I - AM^{-1})$ is close to 1. In the next section we try to exploit the nice spectral properties of the GASP.

Table 1 Number of iterations needed to solve the second linear system when the reuse of vectors in the subspace is done by truncation, assembling (rank-k) or assembling of preconditioned Ritz vectors (k outliers) for the rectangular basin partitioned in 4 strips.

truncation	GASP appl.	rank-k	GASP appl.	k outliers	GASP appl.
z_1, z_2	9 or 10	2	9	2	8
z_1, z_2, z_3	9 or 10	3	8 or 9	3	7 or 8
$z_1, \ldots, z_4$	8 or 9	4	8	4	7
$z_1, \ldots, z_5$	8	5	7	5	6
$z_1, \ldots, z_6$	7 or 8	6	6	6	5
$z_1, \ldots, z_7$	6 or 7	7	6		
$z_1, \ldots, z_8$	6	8	6		
$z_1, \ldots, z_9$	5 or 6	9	5 or 6		
$z_1, \ldots, z_{10}$	5	10	5 or 6		

5 Reuse strategies

The main motivation for using FGMRES instead of GMRES is that the former — in contrast to the latter — accommodates variable preconditioning; any vector z can be put into the search space Z_m as long as its image Az is known in order to be able to compute the correction to the residual. This property in combination with the observation that our specific time integration method results for each time step in a sequence of systems (2.4) has raised the question whether it is possible to reuse previously computed search vectors during the solution of the next systems by FGMRES. Obviously, one advantage of reusing vectors is that it is a lot cheaper than applying the (expensive) GASP which after all requires the solution of a linear system in each subdomain. Also, when (approximations of) the preconditioned eigenvectors that uphold the convergence are collected in the search space, accelerated convergence might be achieved from the first newly computed z-vector on. Several strategies to reuse vectors from an already generated subspace have been tested. In practice, we always use $Q = 2$, i.e. two systems must be solved in each time step. We focus on the solution of the second system in each time step ($q = 2$), possibly reusing information from the search space Z_m built during the solution of the first system ($q = 1$). We formulate the following reuse strategies:

1. **truncation:** introduce the first k z-vectors $z_i^{(R)} = z_i$ $(i = 1, \ldots, k)$.
2. **assembling**: introduce the best rank-k approximation of $\text{SPAN}\{z_1, z_2, \ldots, z_m\}$.
3. **assembling of preconditioned Ritz vectors:** introduce k preconditioned approximate eigenvectors corresponding to outliers: $z_i^{(R)} = Z_m y_i$ $(i = 1, \ldots, k)$, where y_i is the eigenvector of H_m corresponding to the eigenvalue λ_i, i.e. $H_m y_i = \lambda y_i$.

The trivial truncation strategy 1 gives an indication of what can be expected from more sophisticated reuse strategies. The case in which all z-vectors are reused, allows us to verify whether the Arnoldi process is able to quickly generate a reasonably approximate eigenspectrum. The second strategy requires the computation of the

singular value decomposition of $Z_m = U\Sigma V^T$. When the singular values are ordered $\sigma_1 \geq \sigma_2 \geq \ldots \geq \sigma_m \geq 0$, the best rank-$k$ approximation of $\text{SPAN}\{z_1, z_2, \ldots, z_m\}$ is given by $\text{SPAN}\{u_1, u_2, \ldots, u_k\}$. The use of a best rank-k approximation is motivated by the fact that the column space of Z_m contains preconditioned approximate eigenvectors corresponding to the outliers. The hope is that a lower dimensional approximation still contains an approximation of these eigenvectors. The results seem to indicate that this strategy is not entirely capable of filtering out the preconditioned eigenspace. The third reuse strategy relies on the observation of clustered eigenspectra in combination with a (small) number of clearly distinguishable outliers as is the case in Fig. 4. An explicit construction of the preconditioned eigenspace corresponding to the outliers is then possible. Note that the construction is also based on (4.13). Based on the results in Table 1 the assembling strategy of preconditioned Ritz vectors has been chosen for further experiments.

6 Test Case and Results

The test case is concerned with the flow in a 8000m by 1200m rectangular basin which is 8m deep. The uniform grid has one layer in the vertical (σ) direction which contains 80×12 grid points in the horizontal direction. The prescribed boundary conditions are as follows. The north and south boundaries are closed, leading to Neumann boundary conditions in (2.4). At the east boundary the water elevation is kept constant at $\zeta = 2$m. At the west boundary the water elevation is prescribed to model the tide, yielding $\zeta = 2 + \sin \frac{2\pi t}{3600}$ m. Tests have been carried out with stripwise partitionings of the rectangular domain into 4 and 8 subdomains. A test with 4 subdomains and a mesh width of 50m instead of 100m was also done to see the effect of refinement on the number of outliers and the separation between the outliers and the cluster.

The convergence histories showing the scaled (with $\|f\|_2$) residual norm as a function of the dimension of the search space for the FGMRES algorithm applied to the two linear systems arising each time step are shown in Fig. 5. The convergence history for the first linear system starts at about 0.2 and drops below the adopted threshold of 10^{-9} after 11 iterations. This requires 11 applications of the GASP. Figure 4 shows the eigenvalues of the Hessenberg matrix H_m constructed by FGMRES during the solution of this first linear system. The six eigenvalues that are not close to 1 are the outliers which "hamper" fast convergence of FGMRES. The convergence history for the second linear system starts off with a plateau at about 0.001, dropping sharply to reach the tolerance criterion after 11 iterations as well. The plateau corresponds to the reuse of assembled preconditioned Ritz vectors. This corresponds to the removal of the approximate eigenvectors associated with the outliers from the initial residual, a process which hardly reduces the norm of it. However it makes FGMRES converge as if the outliers were not present at all; starting from the first newly computed search vector the residuals decrease rapidly, at the same speed as in the end stage of the solution of the first system. Because of the reuse, solving the second linear system requires only 5 applications of the GASP.

Instead of computing the approximate eigenvectors at each time step from the matrix Z_m constructed during the solution of the first linear system ($q = 1$), we

Figure 5 Convergence histories of the preconditioned FGMRES method

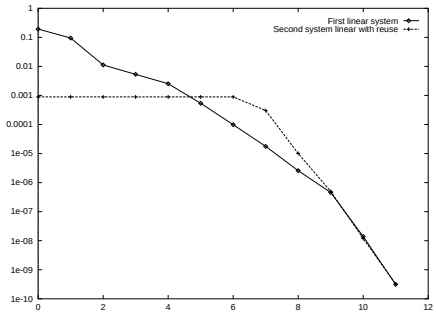

Figure 6 Convergence histories of the preconditioned FGMRES method

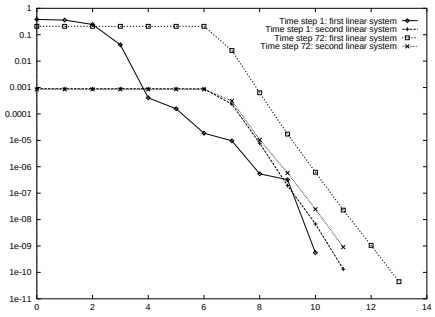

can also construct the approximate eigenspace only once, i.e. from the first linear system arising in the first time step. Moreover, the eigenspace is now reused in the solution process for the first linear systems ($q = 1$) of each time step as well. The convergence histories are shown in Fig. 6. As can be seen from this Fig. we save on preconditioning steps by introducing these vectors also in the search space when solving the first linear system. The convergence histories for the second linear system show that it is not necessary to compute the approximate eigenvectors at each time step, since the results with the vectors from the first time step are sufficiently close to those with the vectors from the current time step.

7 Conclusion

We have developed a Generalised Additive Schwarz Preconditioner for use within FGMRES to solve linear systems arising in the solution of the time-dependent shallow water equations. The preconditioned operator AM^{-1} has a clustered eigenspectrum with only a few outlying eigenvalues, at least for meshes not too fine. This property together with the specific time integration method enables the reuse of search vectors in the FGMRES process which leads to reductions in computation time.

Acknowledgement

This paper presents results obtained during the stay of S. Goossens at Delft Hydraulics. The authors thank E. de Goede for his help and also acknowledge the financial support by the E.C. MAST II Concerted Action "Application of High Performance Computing Techniques for the Modeling of Marine Ecosystems".

REFERENCES

[dG93] de Goede E. D. (September 1993) Een AOI methode voor TRISULA. Technical Report VR595.93/Z642, Delft Hydraulics. in dutch.

[HTJ88] Hagstrom T., Tewarson R. P., and Jazcilevich A. (1988) Numerical experiments on a domain decomposition algorithm for nonlinear elliptic boundary value problems. *Appl. Math. Lett.* 1(3): 299–302.

[Lio90] Lions P. L. (1990) On the Schwarz alternating method III: A variant for nonoverlapping subdomains. In Chan T. F., Glowinski R., Périaux J., and Widlund O. B. (eds) *Proc. Third Int. Conf. on Domain Decomposition Meths.*, pages 202–223. SIAM, Philadelphia.

[NR95] Nataf F. and Rogier F. (1995) Factorization of the convection-diffusion operator and the Schwarz algorithm. *Mathematical Models and Methods in Applied Sciences* 5(1): 67–93.

[Saa93] Saad Y. (1993) A flexible inner-outer preconditioned GMRES algorithm. *SIAM J. Sci. Comput.* 14(2): 461–469.

[Tan95] Tan K. H. (1995) *Local Coupling in Domain Decomposition.* PhD thesis, Universiteit Utrecht.

[VdSVdV86] Van der Sluis A. and Van der Vorst H. A. (1986) The rate of convergence of conjugate gradients. *Numerische Mathematik* 48: 543–560.

[VdVV93] Van der Vorst H. A. and Vuik C. (1993) The superlinear convergence behaviour of GMRES. *Journal of Computational and Applied Mathematics* 48: 327–341.