Preconditioning via a Schur Complement Method: an Application in State Estimation*

Luise Blank†

published in SIAM J. Sci. Comput., Vol. 25, No. 3, 942-960, (2003).

Abstract

In this work the Schur complement is analyzed and applied for preconditioning symmetric matrices. In particular we eliminate the in modulo extreme eigenvalues by computing the Schur complement with respect to a coarse subspace, which captures the corresponding eigenvectors. The potential of this method is presented as an application in state estimation, where we employ a multiscale approach based on a wavelet-Galerkin discretization. The hierarchy of arising optimization problems is solved in the framework of nested iterations. Due to the wavelet approach, for all scales we can use one coarse adapted subspace in the Schur complement method. Typically this provides not only a reduced condition number but also large gaps in the arising eigenvalue distribution, which is favorable for several iterative methods. Numerical investigations show the enormous reduction in the condition number, the necessary cg-iteration steps and, consequently, in the computation time.

Key words: Schur complement, preconditioning, deflation, state estimation, multiscale approach, wavelets.
MSC: 15A12, 65F, 65K, 49M, 93C

1 Introduction

The most common Schur complement technique applies the Schur complement reduction repeatedly until the dimension of the Schur complement is small enough for fast inversion. The reduction in dimension is a key feature. However, it requires that the inversions in each reduction step to be performed efficiently. Typically one searches for a subspace in which the corresponding matrix is block diagonal, well conditioned and not too small.

In contrast, in this paper we would like to present an application of the Schur complement method in which we apply the reduction step only once. We look for a subspace in which the corresponding matrix is probably dense, as small as possible, but captures the ill-conditioning of the matrix. Consequently, the dimension of the Schur complement remains large, but the goal is to reduce the condition number drastically in comparison to the original matrix. Hence the Schur complement is used for eigenvalue deflation.

The potential of the method is presented as an application in state estimation. For better motivation let us briefly describe the general framework in which we apply the method. In chemical engineering complex dynamic optimization problems have to be solved online. In [4, 8] a multiscale approach based on wavelets is presented in which a hierarchy of successively, adaptively refined problems is constructed. These

*This work has been supported by the “Deutsche Forschungsgemeinschaft” under grant MA1188/6.
†NWF I – Mathematik, Universität Regensburg, D–93040 Regensburg, Germany,
e-mail: luise.blank@mathematik.uni-regensburg.de

problems are solved in the framework of nested iteration as long as the real-time restrictions are fulfilled. Each discrete problem has to be solved only with an accuracy comparable to the current approximation error. Hence the use of an iterative solver also suggests itself for the arising systems of linear equations. Obviously the eigenvalue distribution plays an important role in the efficiency. In [6] it is shown that by using properly scaled wavelets the condition numbers can be bounded independently of the refinement level. However, the in modulo smallest eigenvalues are still close to zero due to inherent features of the estimation problem, i.e., the original ill-posedness, the low measure of observability, and the stiff involved differential equations.

A number of articles have established the benefits of eigenvalue deflation when using Krylov subspace methods. Most common is the deflation of nearly invariant subspaces. Some authors augment the Krylov subspace with approximate eigenvectors that either are obtained during the iteration [24, 26] or are known a priori [29]. Others use subdomain deflation, which is related to domain decomposition [16, 22, 25]. The Schur complement fits into the framework of the projection used in this technique of eigenvalue deflation. However, in general it does not employ invariant subspaces. To distinguish mesh dependency of eigenvalues, Schur complement techniques are used, for example, as saddle point preconditioners for discrete Navier–Stokes equations [15] or elliptic problems with Neumann boundaries (see [3] and references therein). Using the Schur complement to eliminate negative eigenvalues of indefinite but symmetric systems arising in nonlinear elasticity is considered in [27] and is based on the Sylvester law of inertia.

We apply the above Schur complement preconditioning in order to eliminate the smallest eigenvalues. The idea is based on the following: roughly speaking, if we compute the Schur complement with respect to a subspace containing the eigenvectors corresponding to the smallest eigenvalues, then these eigenvalues are eliminated in the Schur complement. To approximate such a subspace we consider the underlying infinite dimensional, but wavelet discretized, system. We try to find a very coarse, but adapted, set of wavelet basis functions that captures the eigenfunctions (and with them the corresponding eigenvectors) to the smallest eigenvalues. Due to the multiscale representation, one expects that this coarse wavelet space will also capture good approximations of the eigenvectors to the smallest eigenvalues of the truncated systems. Then, for the whole hierarchy of systems, the Schur complements with respect to this coarse, fixed subspace will give rise to well-conditioned systems.

The outline of the paper is as follows. In section 2 the formulation of the dynamic state estimation is briefly introduced. Furthermore wavelets and their basic properties are described in order to motivate their application in this context. Then we present the wavelet-Galerkin discretization for the multiscale approach, which leads to the sequence of linear equation systems that we will study. Wavelet-based preconditioning eliminates the scale-to-scale increase of the large eigenvalues. We conclude this section by commenting on the (in our application) typically arising eigenvalue distributions and their origins. The analytical results presented in section 3 provide the foundation of the application of the Schur complement method as a preconditioner. As a first step, in section 3.1 we deduce estimations for eigenvalues of symmetric preconditioned matrices with respect to the eigenvalues of the original matrix. Using the fact that the Schur complement arises by a particular symmetric preconditioning, we prove under certain assumptions the elimination of the extreme eigenvalues. The choice of an appropriate subspace is then motivated in section 3.2. In section 4 the nested iteration algorithm based on the Schur complement technique is sketched and its application is numerically investigated. We illustrate the drastic effect on the eigenvalue distribution and, hence, on the condition number and iteration number. In section 4.4 we analyze the influence of the choice of the subspace in the numerical results. Finally, in section 5 we summarize the conclusions.

2 Problem Formulation and Discretization

2.1 Problem Formulation

Dynamic state estimation aims to determine, from a few given, noisy measurements, all states involved in the process. The dynamic behavior of the plant is often modeled by a system of differential-algebraic equations for all states $\boldsymbol{x}$ depending on given controls $\boldsymbol{u}$. Moreover, unknown model error functions $\boldsymbol{w}$ are included in the model equations to represent unstructured model uncertainty. Typically only very few states can be measured, and these states are denoted by the output functions $\boldsymbol{y}$. Then state estimation can be formulated as a constrained optimization problem, where the difference between measured data $\boldsymbol{z}$ and corresponding modeled outputs $\boldsymbol{y}$ has to be minimized [1, 4]. Together with bounds on selected variables, the model equations form the constraints of the optimization. We stress that the problem in its original form is an ill-posed inverse problem. Therefore some regularization is necessary [8, 9].

Since we concentrate on the solution of the final linear systems, we consider here only linear model problems without inequality constraints. The estimation horizon is scaled to $[0,1]$. Hence, we focus here on the following particular setting, where we suppress for convenience explicit references to the domain $[0,1]$. For given $\boldsymbol{z} \in (L_2)^{n_y}$ and $\boldsymbol{u} \in (L_2)^{n_u}$ and unknown functions $\boldsymbol{x} \in \left(H^1\right)^{n_x}$, $\boldsymbol{w} \in (L_2)^{n_w}$, and $\boldsymbol{y} \in (L_2)^{n_y}$, the resulting optimization problem has the form

$$1/2 \min_{\boldsymbol{x}(0),\boldsymbol{w}} \left\{ ||\boldsymbol{y}-\boldsymbol{z}||^2_{L_2^{n_y}} + \varepsilon\, ||\boldsymbol{w}||^2_{L_2^{n_w}} \right\} \tag{1}$$

subject to the constraints

$$\begin{aligned} \dot{\boldsymbol{x}} - \boldsymbol{A}\boldsymbol{x} - \boldsymbol{W}\boldsymbol{w} &= \boldsymbol{u}, & (2)\\ \boldsymbol{y} - \boldsymbol{C}\boldsymbol{x} &= \boldsymbol{0}. & (3) \end{aligned}$$

Even for nonlinear model equations, (3) remains typically linear, since $\boldsymbol{C}$ selects some states as outputs. Consequently, $\boldsymbol{y}$ can be eliminated in the optimization problem. We would like to stress that in the above problem the initial data $\boldsymbol{x}(0)$ are not given but have to be estimated. They are also not included in the cost functional as is often seen in other applications. This has a major impact on the condition number. In addition, the observability measure of the model [21, 30] represented by the constraints, the stiffness of the differential equations, the incorporated model error functions, and the choice of the regularization parameter ε influences the condition of the formulated optimization problem, and hence the numerical treatment of the resulting linear systems of equations.

If we denote the Lagrange multipliers by $\boldsymbol{\lambda}$, the resulting first order necessary conditions are the following [20]:

$$\begin{aligned} \varepsilon\boldsymbol{w} - \boldsymbol{W}^T\boldsymbol{\lambda} &= \boldsymbol{0}, & (4)\\ \dot{\boldsymbol{x}} - \boldsymbol{A}\boldsymbol{x} - \boldsymbol{W}\boldsymbol{w} &= \boldsymbol{u}, & (5)\\ -\dot{\boldsymbol{\lambda}} - \boldsymbol{A}^T\boldsymbol{\lambda} + \boldsymbol{C}^T\boldsymbol{C}\boldsymbol{x} &= \boldsymbol{C}^T\boldsymbol{z}, & (6)\\ \boldsymbol{\lambda}(0) = \boldsymbol{0} \quad \text{and} \quad & \boldsymbol{\lambda}(1) = \boldsymbol{0}. & (7) \end{aligned}$$

Obviously $\boldsymbol{w}$ can be eliminated in (5) by the Lagrange multipliers $\boldsymbol{\lambda}$ using (4). Defining $\boldsymbol{R} := 1/\varepsilon \boldsymbol{W}\boldsymbol{W}^T$, we see that the weak formulation of the remaining equations is given by

$$\langle \dot{\boldsymbol{x}} - \boldsymbol{A}\boldsymbol{x} - \boldsymbol{R}\boldsymbol{\lambda}, \boldsymbol{\zeta}_1 \rangle = \langle \boldsymbol{u}, \boldsymbol{\zeta}_1 \rangle \text{ for all } \boldsymbol{\zeta}_1 \in L_2^{n_x}, \tag{8}$$

$$\langle \boldsymbol{\lambda}, \dot{\boldsymbol{\zeta}}_2 \rangle + \langle -\boldsymbol{A}^T\boldsymbol{\lambda} + \boldsymbol{C}^T\boldsymbol{C}\boldsymbol{x}, \boldsymbol{\zeta}_2 \rangle = \langle \boldsymbol{C}^T\boldsymbol{z}, \boldsymbol{\zeta}_2 \rangle \text{ for all } \boldsymbol{\zeta}_2 \in H_1^{n_x}. \tag{9}$$

Since our discretization is based on a Galerkin discretization of (8)–(9) using wavelets, we briefly introduce their basic properties and indicate their significance in the present context. For a more detailed treatment of wavelets we refer the reader to [11, 13, 14].

2.2 Wavelets

Apart from boundary adaptation wavelet functions $\psi_{j,k}$ are obtained by dilation and translation of a suitable *mother wavelet* ψ, i.e., $\psi_{j,k} = 2^{j/2}\psi(2^j \cdot -k)$, where the normalizing factor $2^{j/2}$ keeps the L_2-norm of the $\psi_{j,k}$ independent of j, k. Thus when ψ is chosen to have bounded support, as it will always be the case in our applications, one has *diam supp* $\psi_{j,k} \sim 2^{-j}$, $j \to \infty$. One of the key features of wavelets is that for suitable ψ the collection $\boldsymbol{\Psi} = \{\psi_{j,k}\}$ forms a *Riesz basis* of L_2; i.e., any $f \in L_2$ has a unique expansion

$$f = \boldsymbol{\Psi}^T \boldsymbol{d}_f = \sum_{j,k} d_{j,k}\psi_{j,k},$$

and the norm equivalence

$$c_1 \parallel \boldsymbol{d}_f \parallel_{\ell_2} \leq \parallel f \parallel_{L_2} \leq c_2 \parallel \boldsymbol{d}_f \parallel_{\ell_2} \tag{10}$$

holds for some positive constants c_1, c_2 independent of f. Clearly (10) means that small perturbations of the expansion coefficients $\boldsymbol{d}_f$ cause only small changes of f in the L_2-norm and vice versa. Depending on the smoothness of the employed wavelets, norm equivalences with respect to other function spaces such as Sobolev and Besov spaces also hold with appropriate scale-dependent weights on the wavelet coefficients. They are, for example, essential for signal analysis and for the preconditioning of discretized differential (see section 2.4) and integral operators.

In general, (10) implies the existence of a *dual basis* $\tilde{\boldsymbol{\Psi}}$ in L_2 that is *biorthogonal* to $\boldsymbol{\Psi}$, i.e., $\langle\langle \boldsymbol{\Psi}, \tilde{\boldsymbol{\Psi}} \rangle\rangle = \boldsymbol{I}$, where we use the shorthand notation $\langle\langle \boldsymbol{\Psi}, \tilde{\boldsymbol{\Psi}} \rangle\rangle = (\langle \psi_{j,k}, \tilde{\psi}_{i,l} \rangle)_{(j,k),(i,l)}$ with $\langle \cdot, \cdot \rangle$ denoting the standard L_2-inner product. Thus the wavelet coefficients $d_{j,k}$ are given by $\langle f, \tilde{\psi}_{j,k} \rangle$.

The starting point for the construction of wavelet bases is usually a so-called *multiresolution sequence* of nested spaces $S_{j_0} \subset \cdots \subset S_j \subset S_{j+1} \subset \cdots$ whose union is dense in L_2. Here j_0 stands for some coarsest discretization level. The spaces S_j are spanned by scaling functions, $S_j = \text{span}\{\varphi_{j,k} : k \in \mathcal{I}_j\}$, which are obtained by a suitable, compactly supported function φ. Due to the compact support of φ we can view 2^{-j} as the uniform mesh size of S_j. Given a suitable dual sequence $\tilde{S}_j$, a successive decomposition of S_j leads to the multiscale splitting $S_J = S_{j_0} \bigoplus_{j=j_0}^{J-1} W_j$, where $W_j = \text{span}\{\psi_{j,k} : k \in \mathcal{I}'_j\} \perp \tilde{S}_j$. Analogous splittings for the spaces $\tilde{S}_j$ yield complement spaces $\tilde{W}_j \perp S_j$. The bases $\boldsymbol{\Psi}$ and $\tilde{\boldsymbol{\Psi}}$ for the spaces W_j and $\tilde{W}_j$, respectively, are biorthogonal.

All the above properties can be realized when choosing φ as the cardinal B-spline of order m. Moreover, one can construct for any $\tilde{m} \in \mathbb{N}, \tilde{m} \geq m, m + \tilde{m}$ even, a compactly supported dual scaling function $\tilde{\varphi}$ such that the dual multiresolution spaces $\tilde{S}_j$ contain all polynomials up to order $\tilde{m}$. As a consequence of biorthogonality, one has the *moment conditions* $\langle (\cdot)^r, \psi \rangle = 0$ for $r = 0, \ldots, \tilde{m} - 1$. Likewise one has $\langle (\cdot)^r, \tilde{\psi} \rangle = 0$ for $r = 0, \ldots, m - 1$. Therefore smooth functions have wavelet expansions with rapidly decaying wavelet coefficients so that scalewise truncation of wavelet expansions provides good approximations. However, as soon as sharp transitions or even singularities occur, different nonuniform selections of basis functions are expected to provide more economical approximations. In fact, combining the above statements with the locality of wavelets, large wavelet coefficients reflect a large local change of the function. In particular, retaining only the first N largest coefficients in modulus provides —up to a uniform constant— the best approximation to f that can be composed from any selection of N basis functions (*best N-term approximation*). This is the basis for nonlinear compression techniques. Accordingly, the objective of *adaptive* approximation is to successively track the most significant coefficients of the function that is to be determined.

2.3 Discretization: Wavelets as Trial and Test Functions

The Galerkin discretization of the optimization problem based on wavelets is given by using a representation of each function with respect to an appropriately chosen wavelet basis in the first order necessary conditions (8)–(9). In our case we use piecewise linear, continuous wavelets $\boldsymbol{\Psi}$ as trial functions for $\boldsymbol{x}$ and $\boldsymbol{\zeta}_2$ (i.e. the scaling function is a B-spline of order $m = 2$ fulfilling the moment conditions with $\tilde{m} = 2$), and the functions $\boldsymbol{z}$ are discretized by the dual wavelets $\tilde{\boldsymbol{\Psi}}$, while we employ piecewise constant wavelets, so-called Haar wavelets $\boldsymbol{\Psi}^H$ ($m = \tilde{m} = 1$) for $\boldsymbol{u}$, the Lagrange multipliers $\boldsymbol{\lambda}$ and the test functions $\boldsymbol{\zeta}_1$. This corresponds to the minimal regularity for a conforming discretization. Consequently, we use the representations $\boldsymbol{x} = \boldsymbol{\Psi}^T_{\Lambda_x}\boldsymbol{d_x}$ and $\boldsymbol{\lambda} = (\boldsymbol{\Psi}^H_{\Lambda_\lambda})^T\boldsymbol{d_\lambda}$, etc. in (8)–(9). We obtain a system of equations for the unknown wavelet coefficients $\boldsymbol{d_x}$ and $\boldsymbol{d_\lambda}$. The discretization matrix can be described in a rather condensed form with the aid of Kronecker products. The Kronecker product of a matrix $\boldsymbol{A} \in I\!R^{n\times m}$ with a matrix $\boldsymbol{B}$ is given by $\boldsymbol{A} \otimes \boldsymbol{B} := (a_{i,j}\boldsymbol{B})_{\substack{i=1,\ldots,n\\ j=1,\ldots,m}}$. Defining

$$\boldsymbol{R}_\Lambda = \boldsymbol{R} \otimes \boldsymbol{I}_\Lambda, \quad \boldsymbol{Q}_\Lambda = \boldsymbol{C}^T\boldsymbol{C} \otimes \langle\langle \boldsymbol{\Psi}_\Lambda, \boldsymbol{\Psi}_\Lambda \rangle\rangle, \quad \boldsymbol{C}_\Lambda = \boldsymbol{C} \otimes \boldsymbol{I}_\Lambda, \tag{11}$$

$$\boldsymbol{D}_\Lambda = \boldsymbol{I}_{n_x} \otimes \langle\langle \boldsymbol{\Psi}^H_{\Lambda'}, \dot{\boldsymbol{\Psi}}_\Lambda \rangle\rangle, \quad \boldsymbol{A}_\Lambda = \boldsymbol{A} \otimes \langle\langle \boldsymbol{\Psi}^H_{\Lambda'}, \boldsymbol{\Psi}_\Lambda \rangle\rangle, \tag{12}$$

yields the following system of equations, which is equivalent to (8)–(9) and still *infinite dimensional* but *discretized*:

$$\begin{pmatrix} -\boldsymbol{R}_\Lambda & \boldsymbol{D}_\Lambda - \boldsymbol{A}_\Lambda \\ & \\ (\boldsymbol{D}_\Lambda - \boldsymbol{A}_\Lambda)^T & \boldsymbol{Q}_\Lambda \end{pmatrix} \begin{pmatrix} \boldsymbol{d_\lambda} \\ \\ \boldsymbol{d_x} \end{pmatrix} = \begin{pmatrix} \boldsymbol{d_u} \\ \\ \boldsymbol{C}^T_\Lambda \boldsymbol{d_z} \end{pmatrix}. \tag{13}$$

While the above formulation still provides the exact solution, for the numerical treatment we have to choose finite dimensional approximations. Hence for each function involved we have to choose a finite set of basis functions, described by their index set $\Lambda_{\boldsymbol{x}}$, etc.. As mentioned before, we will start with some initial, possibly small index sets $\Lambda^0_{\boldsymbol{x}}$, etc. and adaptively enlarge these index sets to $\Lambda^{\ell+1}_{\boldsymbol{x}} = \Lambda^\ell_{\boldsymbol{x}} \cup \Gamma^\ell_{\boldsymbol{x}}$ in each refinement step. Moreover, the Haar basis consists of one less scaling function on the coarsest scale but consists of the same number of wavelets. We have $\Lambda' = \Lambda \setminus \{(1,2)\}$. The involved wavelet mass matrices exhibit so-called finger structure with $O(\#\Lambda log(\#\Lambda))$ entries.

Without loss of generality, we consider below only the case $\boldsymbol{W} = \boldsymbol{I}$, i.e., $\boldsymbol{R} = 1/\varepsilon\boldsymbol{I}$, since the following considerations for the iterative solution are extendible to the symmetric indefinite case. Obviously we can eliminate the Lagrange multipliers $\boldsymbol{d_\lambda}$ in (13) and obtain with

$$\boldsymbol{X}_\Lambda \;:=\; \varepsilon\,(\boldsymbol{D}_\Lambda - \boldsymbol{A}_\Lambda)^T\,(\boldsymbol{D}_\Lambda - \boldsymbol{A}_\Lambda) + \boldsymbol{Q}_\Lambda \tag{14}$$

an equation for the wavelet coefficients $\boldsymbol{d_x}$ of the states as follows:

$$\boldsymbol{X}_\Lambda \boldsymbol{d_x} \;=\; \boldsymbol{C}^T_\Lambda \boldsymbol{d_z} + \varepsilon\,(\boldsymbol{D}_\Lambda - \boldsymbol{A}_\Lambda)^T\,\boldsymbol{d_u}. \tag{15}$$

Since (15) is solved iteratively, the matrix $\boldsymbol{X}_\Lambda$, as well as $\boldsymbol{D}_\Lambda, \boldsymbol{A}_\Lambda$, and $\boldsymbol{Q}_\Lambda$, is not assembled. Instead only matrix-vector multiplications are realized. We recall that (15) is solved in the context of nested iteration, i.e., in a sequence of equations for Λ^ℓ, $\ell = 0, 1, \ldots$ with $\Lambda^\ell \subseteq \Lambda^{\ell+1}$ and consequently with increasing dimension.

2.4 Iterative Solver and Eigenvalue Distribution

Since $\boldsymbol{X}_\Lambda$ is symmetric and positive definite we apply the pcg-method. The approximate solution $\boldsymbol{d}_{\boldsymbol{x}}^{\ell-1}$ on $\Lambda^{\ell-1}$ serves as an initial value in the ℓth step. A change of basis is not necessary due to the choice of wavelet basis. The vector $\boldsymbol{d}_{\boldsymbol{x}}^{\ell-1}$ has only to be extended by zeros for the indices in $\Lambda^\ell/\Lambda^{\ell-1}$. The quality of this initial value depends on the chosen refinement $\Lambda^\ell/\Lambda^{\ell-1}$, which not only determines the change in $\boldsymbol{X}_\Lambda$ but also specifies the additional informations of the signal $\boldsymbol{z}$ and the controls $\boldsymbol{u}$, which have to be processed.

Common knowledge for the pcg-method is an error estimation depending on the condition number $\kappa = \lambda_{max}(\boldsymbol{X})/\lambda_{min}(\boldsymbol{X})$. Results concerning the influence of the distribution of the whole spectrum and, in particular, the influence of large gaps between the largest or smallest eigenvalues, are summarized in [2, 19, 23]. Essentially these results show that after several iterations (depending on the gap) the convergence rate is independent of the extreme eigenvalues. Moreover, it is known that the rate of convergence has the tendency to increase, which may be called superlinear convergence [23, 28]. In [2] the convergence behavior is divided into the following three phases: in the initial phase the convergence depends on the number of iterations; in the middle phase the rate depends on the condition number κ; and in the final phase one observes superconvergence. Usually the middle phase is longer if very small eigenvalues are present [23].

To give an impression of a typical eigenvalue distribution occurring in our context we choose as a first example a small problem with only four states and one output function. (Details of this example can be found in [7, Example 1].) Figure 1 displays the eigenvalues for a uniform mesh size of 2^{-9}; i.e., Λ includes all indices of scaling and wavelet functions up to scale 9. The arising condition number can be seen by considering the logarithmic scale, while gaps in the eigenvalue distribution are displayed in the linear scale. The large eigenvalues increase scale to scale; i.e., 2^i eigenvalues correspond to scale i of the wavelet discretization and have a magnitude of order 2^{2i}. This is, roughly speaking, due to the involved differential operator of order 2 in (15). On lower scales the model matrices are dominant in comparison to the differential operator. The scales cannot be distinguished. No isolated eigenvalues are present and the smallest eigenvalues are very close to zero. This leads, as expected, to a long middle phase in the convergence behavior.

Figure 1: Eigenvalue distribution.

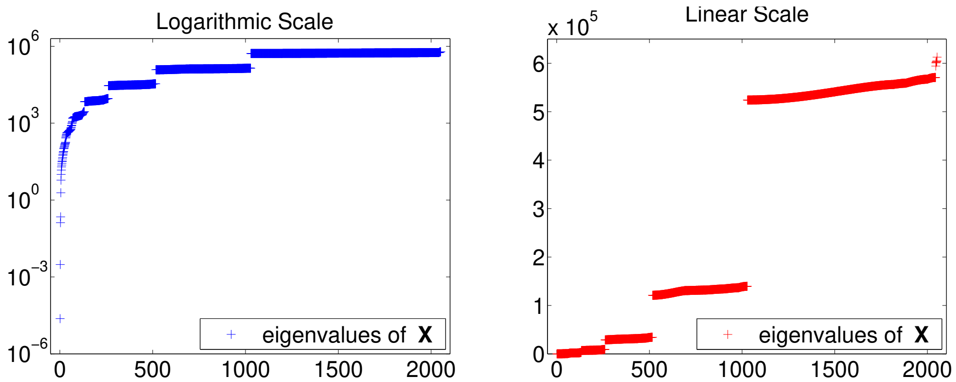

Based on the knowledge of boundedness of the original operator from H^1 to $(H^1)'$

and its inverse, and using specific norm equivalences of wavelet representations, it is shown in [6] that scaling of the wavelet representations guarantees an upper limit for the number of iterations independently of ℓ. This scaling can also be viewed as symmetric, diagonal preconditioning with entries 2^{-j}. Instead of 2^{-j} an approximate Jacobi preconditioner $\boldsymbol{N}_{\Lambda^\ell}$ is used; i.e., the square root of an approximation of the diagonal of $\boldsymbol{X}_\Lambda$ is employed. The entries reflect the necessary scale-dependent scaling. In addition this preconditioner corresponds to normalizing the wavelet basis in the energy norm and usually gives better results. The arising condition numbers fulfill

$$\mathrm{cond}_2(\boldsymbol{N}_{\Lambda^\ell}\boldsymbol{X}_{\Lambda^\ell}\boldsymbol{N}_{\Lambda^\ell}) \leq c \tag{16}$$

independent of ℓ [6]. The scale-to-scale increase of the eigenvalues is eliminated. However, although the boundedness is guaranteed, the size of the constant c can be very large. In numerical experiments we see that the smallest eigenvalues are still close to zero and there is no gap in the lower spectrum (see, for example, Figure 2). Results for other diagonal preconditioners can be found in [7], where also various, more sophisticated preconditioners are discussed.

The smallest eigenvalues are due to inherent features of the estimation problem. The original problem is ill-posed and the measure of observability of the model is typically low. In addition, the stiffness of the involved ordinary differential equation plays a crucial role, since the initial values have to be estimated. To tackle the first problem an appropriate regularization parameter ε is chosen and the states are normalized by a rough guess of their maximum norm (see [7]). Nevertheless, this tuning of the problem formulation is still not sufficient to overcome the second and third problem-inherent features and to obtain a favorable eigenvalue distribution.

3 The Schur Complement Method to Eliminate Eigenvalues

3.1 Elimination of Eigenvalues

The idea of improving the eigenvalue distribution is based on the statements in the theorems given below. They confirm, roughly speaking, that if we choose a submatrix corresponding to the eigenspace of the smallest eigenvalues $0 < \lambda_1 \leq \cdots \leq \lambda_s$, then the Schur complement of this submatrix in the complete matrix has no eigenvalues smaller than λ_s. The small eigenvalues are eliminated. The theorems also include the symmetric indefinite case.

Theorem 3.1 *Suppose the matrix $\boldsymbol{X}$ is symmetric with eigenvalues $\lambda_1 \leq \cdots \leq \lambda_s < 0 < \lambda_{s+1} \leq \cdots \leq \lambda_n$ and eigenvectors $\boldsymbol{v}_i$ $(i = 1, \ldots, n)$. Furthermore, given a non-singular matrix $\boldsymbol{M}$ let us denote by ξ_i the eigenvalues of the symmetrically preconditioned matrix $\boldsymbol{M}\boldsymbol{X}\boldsymbol{M}^T$ with eigenvectors $\boldsymbol{w}_i$, where ξ_i is sorted in a nondecreasing order for $i = 1, \ldots, n$. Then, we have*

$$|\lambda_r| \min_{\boldsymbol{y}\in W_r} \frac{\boldsymbol{y}^T\boldsymbol{M}^T\boldsymbol{M}\boldsymbol{y}}{\boldsymbol{y}^T\boldsymbol{y}} \leq |\xi_r| \leq |\lambda_r| \max_{\boldsymbol{y}\in V_r} \frac{\boldsymbol{y}^T\boldsymbol{M}^T\boldsymbol{M}\boldsymbol{y}}{\boldsymbol{y}^T\boldsymbol{y}} \tag{17}$$

where

$$W_r = \begin{cases} \boldsymbol{M}^{-1} span\,\{\boldsymbol{w}_r, \ldots, \boldsymbol{w}_s\} & if\ r \leq s, \\ \boldsymbol{M}^{-1} span\,\{\boldsymbol{w}_{s+1}, \ldots, \boldsymbol{w}_r\} & if\ r > s, \end{cases} \tag{18}$$

$$V_r = \begin{cases} span\,\{\boldsymbol{v}_r, \ldots, \boldsymbol{v}_s\} & if\ r \leq s, \\ span\,\{\boldsymbol{v}_{s+1}, \ldots, \boldsymbol{v}_r\} & if\ r > s. \end{cases} \tag{19}$$

Proof: For $i = 1, \ldots, n-s$ the ith largest eigenvalue of $\boldsymbol{M}^{-T}\boldsymbol{X}^{-1}\boldsymbol{M}^{-1}$ is $1/\xi_{s+i}$. Hence, using the minimax characterization of the eigenvalues of symmetric matrices [18, 31] we have

$$\begin{aligned}
1/\xi_{s+i} &= \max_{dimU=i} \min_{\boldsymbol{y}\in U} \frac{\boldsymbol{y}^T\boldsymbol{M}^{-T}\boldsymbol{X}^{-1}\boldsymbol{M}^{-1}\boldsymbol{y}}{\boldsymbol{y}^T\boldsymbol{y}} \\
&\geq \min_{\boldsymbol{y}\in \boldsymbol{M}V_{s+i}} \left(\frac{\boldsymbol{y}^T\boldsymbol{M}^{-T}\boldsymbol{X}^{-1}\boldsymbol{M}^{-1}\boldsymbol{y}}{\boldsymbol{y}^T\boldsymbol{M}^{-T}\boldsymbol{M}^{-1}\boldsymbol{y}} \; \frac{\boldsymbol{y}^T\boldsymbol{M}^{-T}\boldsymbol{M}^{-1}\boldsymbol{y}}{\boldsymbol{y}^T\boldsymbol{y}} \right) \\
&\geq 1/\left\{ \lambda_{s+i} \max_{\boldsymbol{y}\in V_{s+i}} \frac{\boldsymbol{y}^T\boldsymbol{M}^T\boldsymbol{M}\boldsymbol{y}}{\boldsymbol{y}^T\boldsymbol{y}} \right\},
\end{aligned}$$

which gives the upper limit in (17). Exploiting the fact that $1/\lambda_{s+i}$ is the ith largest eigenvalue of $\boldsymbol{X}^{-1}$ for $i = 1, \ldots, n-s$ and using the space W_{s+i}, one can follow the above steps to obtain the lower limit. To obtain the estimates for the remaining negative eigenvalues we observe that $-1/\xi_{s+1-i}$ is the ith largest eigenvalue of $-\boldsymbol{M}^{-T}\boldsymbol{X}^{-1}\boldsymbol{M}^{-1}$ and use V_{s+1-i}, while $-1/\lambda_{s+1-i}$ is the ith largest eigenvalue of $-\boldsymbol{X}^{-1}$ and we can use $dimW_{s+1-i} = i$ for $i = 1, \ldots, s$. q.e.d.

A similar result can be obtained by directly considering the eigenvalues (and not their inverses) of $\boldsymbol{X}$ and $\boldsymbol{M}\boldsymbol{X}\boldsymbol{M}^T$. However, for elimination of the in modulo smallest eigenvalues the above result is more useful.

Before we present the next theorem let us define some notation. Suppose we partition the equation $\boldsymbol{X}\boldsymbol{d} = \boldsymbol{r}$ into

$$\begin{pmatrix} \boldsymbol{X}_{0,0} & \boldsymbol{X}_{0,J} \\ \boldsymbol{X}_{J,0} & \boldsymbol{X}_{J,J} \end{pmatrix} \begin{pmatrix} \boldsymbol{d}_0 \\ \boldsymbol{d}_J \end{pmatrix} = \begin{pmatrix} \boldsymbol{r}_0 \\ \boldsymbol{r}_J \end{pmatrix}, \tag{20}$$

where $\boldsymbol{X}_{0,0}$ is nonsingular, and let us denote the Schur complement by

$$\boldsymbol{\mathcal{S}} := \left(\boldsymbol{X}_{J,J} - \boldsymbol{X}_{J,0}\boldsymbol{X}_{0,0}^{-1}\boldsymbol{X}_{0,J} \right). \tag{21}$$

Then the solution of (20) is given by the solution of

$$\boldsymbol{\mathcal{S}}\boldsymbol{d}_J = \boldsymbol{r}_J - \boldsymbol{X}_{J,0}\boldsymbol{X}_{0,0}^{-1}\boldsymbol{r}_0 \tag{22}$$

followed by

$$\boldsymbol{d}_0 = \boldsymbol{X}_{0,0}^{-1}\left(\boldsymbol{r}_0 - \boldsymbol{X}_{0,J}\boldsymbol{d}_J \right). \tag{23}$$

Building the Schur complement $\boldsymbol{\mathcal{S}}$ corresponds to the application of symmetric preconditioning to (20). Namely, we obtain

$$\boldsymbol{M}\boldsymbol{X}\boldsymbol{M}^T = \begin{pmatrix} \boldsymbol{X}_{0,0} & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{\mathcal{S}} \end{pmatrix} \quad \text{with } \boldsymbol{M} := \begin{pmatrix} \boldsymbol{I} & \boldsymbol{0} \\ -\boldsymbol{X}_{J,0}\boldsymbol{X}_{0,0}^{-1} & \boldsymbol{I} \end{pmatrix}. \tag{24}$$

Hence the n eigenvalues of $\boldsymbol{M}\boldsymbol{X}\boldsymbol{M}^T$ are given by the n_0 eigenvalues of $\boldsymbol{X}_{0,0}$ and $n - n_0$ eigenvalues of $\boldsymbol{\mathcal{S}}$. The Sylvester law of inertia immediately yields the following proposition.

Proposition 3.1 *Let $\boldsymbol{X}$ have the inertia (s, z, p) and $\boldsymbol{X}_{0,0}$ the inertia (s_0, z_0, p_0). Then $\boldsymbol{\mathcal{S}}$ has the inertia $(s - s_0, z - z_0, p - p_0)$.*

A related result is given in [27]. There the Schur complement is formed with respect to a coarse finite element space to eliminate the negative eigenvalues for symmetric indefinite linear systems, which arise in elasticity problems. It is illustrated that this, also called the coarse space orthogonalization approach, is a key ingredient in obtaining satisfactory convergence for that class of problems. The Schur complement technique fits into the framework of the deflation technique. Deflation of a subspace $\boldsymbol{Z}$ is obtained by the projection $\boldsymbol{P} = \boldsymbol{I} - \boldsymbol{X}\boldsymbol{Z}(\boldsymbol{Z}^T\boldsymbol{X}\boldsymbol{Z})^{-1}\boldsymbol{Z}^T$ of $\boldsymbol{X}\boldsymbol{d} = \boldsymbol{r}$. For positive definite matrices $\boldsymbol{X}$ it is known that deflation of an invariant subspace cancels the corresponding eigenvalues, leaving the rest of the spectrum untouched (see, for example, [16]). Hence, if $\boldsymbol{Z}$ is the invariant subspace to the ith smallest eigenvalues, then the effective condition number of $\boldsymbol{P}\boldsymbol{X}$ is given by λ_n/λ_{i+1}. However, for the Schur complement $\boldsymbol{Z} = (\boldsymbol{I}_m, \boldsymbol{0})^T$, which is in general not an invariant subspace. In the following theorem we assume the weaker condition that the subspace $\boldsymbol{Z}$ contains the in modulo smallest eigenvalues and quantify for this case the location of the eigenvalues of $\boldsymbol{\mathcal{S}}$.

Theorem 3.2 *Let $\boldsymbol{X}$ have eigenvectors of the form $\boldsymbol{v}_l = (\boldsymbol{v}_{l,0}^T, \boldsymbol{0}^T)^T$ for $l = s-i', \ldots, s+i$. Moreover, let the corresponding eigenvalues λ_l also be the in modulo smallest eigenvalues of $\boldsymbol{X}_{0,0}$. Then $\boldsymbol{\mathcal{S}}$ has eigenvalues ζ_j with*

$$\zeta_1 \leq \cdots \leq \zeta_{s-s_0} \leq \lambda_{s-i'} \leq \lambda_{s+i} \leq \zeta_{s-s_0+1} \leq \cdots \leq \zeta_{n-n_0}. \tag{25}$$

Proof: To prove the statement we show that $\xi_l = \lambda_l$. We use (17) and consider the image of the relevant eigenspaces of $\boldsymbol{X}$ and $\boldsymbol{M}\boldsymbol{X}\boldsymbol{M}^T$ with respect to $\boldsymbol{M}$.

Given the special structure of the eigenvectors of $\boldsymbol{X}$ for $l \in \{s-i', \ldots, s+i\}$, we have $\boldsymbol{X}_{0,0}\boldsymbol{v}_{l,0} = \lambda_l \boldsymbol{v}_{l,0}$ and, in addition, $\boldsymbol{X}_{J,0}\boldsymbol{v}_{l,0} = \boldsymbol{0}$. It follows that $\boldsymbol{M}\boldsymbol{v}_l = \boldsymbol{v}_l$. Hence, $\boldsymbol{M}\boldsymbol{y} = \boldsymbol{y}$ for all $\boldsymbol{y} \in V_l$ and we deduce that $|\xi_l| \leq |\lambda_l|$.

For $\xi_r = \lambda_l$ we have $\boldsymbol{w}_r = \boldsymbol{v}_l$ and consequently, $\boldsymbol{M}\boldsymbol{w}_r = \boldsymbol{w}_r$. The eigenvector corresponding to $\xi_r = \zeta_j$ is given by $\boldsymbol{w}_r = (\boldsymbol{0}^T, \boldsymbol{w}_{r,J}^T)^T$ and we can deduce that $\boldsymbol{M}\boldsymbol{w}_r = \boldsymbol{w}_r$. Furthermore, if we denote by σ_j the remaining eigenvalues of $\boldsymbol{X}_{0,0}$, then $\boldsymbol{w}_r = (\boldsymbol{w}_{r,0}^T, \boldsymbol{0}^T)^T$ and $\boldsymbol{X}_{J,0}\boldsymbol{w}_{r,0} \neq \boldsymbol{0}$ hold for $\xi_r = \sigma_j$. This leads to $\boldsymbol{M}\boldsymbol{w}_r \neq \boldsymbol{w}_r$. However, in accordance to the assumption, we have $\sigma_j \leq \lambda_{s-i'}$ or $\lambda_{s+i} \geq \sigma_j$, which implies $r < s-i'$ or $r > s+i$ for $\xi_r = \sigma_j$. Thus $\boldsymbol{w}_r \notin \{\boldsymbol{w}_{s-i'}, \ldots, \boldsymbol{w}_{s+i}\}$ and we conclude $\boldsymbol{M}\boldsymbol{y} = \boldsymbol{y}$ for all $\boldsymbol{y} \in W_l$, which yields $|\lambda_l| \leq |\xi_l|$. q.e.d.

Using $\lambda_{max} = \max_{\boldsymbol{y}\neq 0}(\boldsymbol{y}^T\boldsymbol{X}\boldsymbol{y})/(\boldsymbol{y}^T\boldsymbol{y})$ and the subspace $\boldsymbol{y}_0 = -\boldsymbol{X}_{0,0}^{-1}\boldsymbol{X}_{0,J}\boldsymbol{y}_J$, we find that the following corollary holds for positive definite systems.

Corollary 3.1 *Let $\boldsymbol{X}$ be symmetric positive definite and in agreement with the assumption in Theorem 3.2. Then*

$$\lambda_i \leq \zeta_1 \leq \zeta_{n-n_0} \leq \lambda_n, \tag{26}$$

giving a reduced condition number cond$(\boldsymbol{\mathcal{S}}) \leq \lambda_n/\lambda_i$.

Obviously we can also apply the above ideas to eliminate the in modulo largest eigenvalues.

3.2 The Choice of the Subspace

Based on Theorem 3.2, efficient preconditioning is built on the following consideration: If we would find a permutation of $\boldsymbol{X}$ such that $\boldsymbol{X}_{0,0}$ inhabits the ill-conditioning of $\boldsymbol{X}$, i.e., has the smallest eigenvalues of $\boldsymbol{X}$, and if the size of $\boldsymbol{X}_{0,0}$ is very small compared with $\boldsymbol{\mathcal{S}}$, then the Schur complement $\boldsymbol{\mathcal{S}}$ is still large in size but is so well conditioned that the application of an iterative solver is very efficient. In general, it is difficult to determine a subspace and with it a suitable permutation of $\boldsymbol{X}$ such that the assumptions of the theorem are fulfilled. It may even happen that there is no suitable reordering apart from the trivial choice $\boldsymbol{X}_{0,0} = \boldsymbol{X}$. Here, the basis in which $\boldsymbol{X}$ represents the corresponding linear map is crucial. In our application the map is represented in a wavelet basis. The arising matrices $\boldsymbol{X}_{\Lambda^\ell}$ are discretization matrices and we can associate with the submatrix $\boldsymbol{X}_{0,0}$ a discretization space V_0 and, respectively, an index set Λ_0. Several questions arise: Is it necessary to fulfill $\boldsymbol{v}_l \in V_0$ exactly? Does a set Λ_0 exist that is coarse? How can we find such a set? Of course there are many more open questions, for example: How many eigenvalues would we like to eliminate and how do we invert the ill-conditioned but small matrix $\boldsymbol{X}_{0,0}$? From the reflections below and the numerical analysis in section 4 we see that employing the Schur complement method yields encouraging answers to the first three questions.

For the choice of a suitable subspace we keep to the spirit of wavelet discretization. Namely, the advantage of using wavelets is that the (infinite dimensional) wavelet discretization is a representation of the original problem. The system of functional equations is transformed into a system of discrete equations. This discrete system is then projected on a hierarchy of subspaces to find approximate solutions. Hence, corresponding to the smallest eigenvalues of the infinite dimensional problem we choose a subspace determined by an index set Λ_0, which shall capture good approximations of the eigenvectors, or, equivalently, an approximation space $\{\psi_{j,k} | \{j,k\} \in \Lambda_0\}$ for the eigenfunctions. Then, this subspace shall also serve as a subspace for the Schur complement method for all principal submatrices $\boldsymbol{X}_{\Lambda^\ell}$, $\ell = 0, 1, \ldots$. Here, we take into account the fact that only the large eigenvalues increase with the scale, while the small eigenvalues stay roughly the same. Exploiting the Riesz basis property and with it the good compression property of wavelets, one expects to find a very coarse space, respectively, a set Λ_0 with few indices, which contains a good approximation of the relevant eigenspace.

Now let us consider the original problem. For $\boldsymbol{W} = \boldsymbol{I}$ the first order conditions (8)–(9) reduce to $\langle \bar{\boldsymbol{\mathcal{S}}} \boldsymbol{x}, \xi \rangle = \langle \boldsymbol{z}, \boldsymbol{C}\xi \rangle + \varepsilon \langle \boldsymbol{u}, \dot{\xi} - \boldsymbol{A}\xi \rangle$ with $\langle \bar{\boldsymbol{\mathcal{S}}} \boldsymbol{x}, \xi \rangle := \varepsilon \langle \dot{\boldsymbol{x}} - \boldsymbol{A}\boldsymbol{x}, \dot{\xi} - \boldsymbol{A}\xi \rangle + \langle \boldsymbol{C}\boldsymbol{x}, \boldsymbol{C}\xi \rangle$. Using the connection with the Sturm–Liouville eigenvalue problem one can show the following (details can be found in [10]).

Proposition 3.2 *If we consider the estimation problem (1)–(3) with one-dimensional constraints only, then the operator $\bar{\boldsymbol{\mathcal{S}}} : H^1 \to (H^1)'$ has infinite, countable many eigenvalues λ_i, where λ_i behaves like i^2. The smallest eigenvalue is $\lambda_0 = C^2$ with eigenfunction* $\exp(At)$.

Hence, Proposition 3.2 motivates us to employ an adapted index set Λ_0, which provides for large A a good wavelet approximation of $exp(At)$. This corresponds to the choice of a nonuniform grid with grid points close to the origin.

3.3 Preconditioning the Schur Complement

Before we present the numerical results, we would like to bring to the reader's attention the fact that symmetric preconditioning with diagonal blocks of sizes as in $\boldsymbol{M}$ and constructing the Schur complement can be interchanged. This follows since

$$\underbrace{\begin{pmatrix} \boldsymbol{N}_0 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{N}_J \end{pmatrix}}_{:=\boldsymbol{N}} \begin{pmatrix} \boldsymbol{I} & \boldsymbol{0} \\ -\boldsymbol{X}_{J,0}\boldsymbol{X}_{0,0}^{-1} & \boldsymbol{I} \end{pmatrix} = \underbrace{\begin{pmatrix} \boldsymbol{I} & \boldsymbol{0} \\ -\bar{\boldsymbol{X}}_{J,0}\bar{\boldsymbol{X}}_{0,0}^{-1} & \boldsymbol{I} \end{pmatrix}}_{:=\overline{\boldsymbol{M}}} \begin{pmatrix} \boldsymbol{N}_0 & \boldsymbol{0} \\ \boldsymbol{0} & \boldsymbol{N}_J \end{pmatrix},$$

where $\bar{\boldsymbol{X}}_{J,0} = \boldsymbol{N}_J \boldsymbol{X}_{J,0} \boldsymbol{N}_0^T$ and $\bar{\boldsymbol{X}}_{0,0} = \boldsymbol{N}_0 \boldsymbol{X}_{0,0} \boldsymbol{N}_0^T$ are the matrices of the preconditioned matrix of $\boldsymbol{N}\boldsymbol{X}\boldsymbol{N}^T$. This ensures that the preconditioning presented in subsection 2.4 also provides an efficient preconditioner for the Schur complement.

Proposition 3.3 *Let $\boldsymbol{N}_\Lambda$ be given as in (16). Then, the condition numbers of the preconditioned Schur complements*

$$\kappa^\ell := cond_2(\boldsymbol{N}_{\Lambda^\ell_{J,J}} \boldsymbol{\mathcal{S}}_{\Lambda^\ell} \boldsymbol{N}_{\Lambda^\ell_{J,J}}) \leq c \tag{27}$$

are bounded independently of the refinement step ℓ.

Proof: Since for symmetric positive definite matrices $\boldsymbol{X}$ the estimations $cond(\boldsymbol{\mathcal{S}}) \leq cond(\boldsymbol{X})$ and $cond(\boldsymbol{X}_{0,0}) \leq cond(\boldsymbol{X})$ hold (see [2]), we have $cond(\boldsymbol{M}\boldsymbol{X}\boldsymbol{M}^T) \leq cond(\boldsymbol{X})$. Hence, using (16) we have

$$\begin{aligned} cond(\boldsymbol{N}_{\Lambda^\ell_{J,J}} \boldsymbol{\mathcal{S}}_{\Lambda^\ell} \boldsymbol{N}_{\Lambda^\ell_{J,J}}) &\leq cond(\boldsymbol{N}_{\Lambda^\ell} \boldsymbol{M} \boldsymbol{X}_{\Lambda^\ell} \boldsymbol{M}^T \boldsymbol{N}_{\Lambda^\ell}) \\ &= cond(\overline{\boldsymbol{M}} \boldsymbol{N}_{\Lambda^\ell} \boldsymbol{X}_{\Lambda^\ell} \boldsymbol{N}_{\Lambda^\ell} \overline{\boldsymbol{M}}^T) \\ &\leq cond(\boldsymbol{N}_{\Lambda^\ell} \boldsymbol{X}_{\Lambda^\ell} \boldsymbol{N}_{\Lambda^\ell}) \leq c \end{aligned}$$

independent of the refinement step ℓ. q.e.d.

In practice we employ instead of $\boldsymbol{N}_{\Lambda^\ell_{J,J}}$ an approximation $\boldsymbol{B}^\ell$ of the square root of the diagonal of the Schur complement.This matrix exhibits the necessary scale-dependent scaling as explained in subsection 2.4 and corresponds to normalizing the wavelet basis in the energy norm with respect to $\boldsymbol{\mathcal{S}}_{\Lambda^\ell}$. Numerical investigations confirmed that $\boldsymbol{B}^\ell$ clearly yields the best results compared with various diagonal preconditioners having entries with the relevant scale dependency, such as diagonal entries 2^{-j}, the diagonal of $\boldsymbol{X}_{J,J}$, and the matrix $\boldsymbol{N}_{\Lambda^\ell_{J,J}}$. A factor of 100 between the condition numbers is not unusual, although though all preconditioners provide asymptotically a scale-independent behavior of the condition numbers.

4 Numerical Investigations

4.1 The Algorithm

In summary, we obtain the following sketched algorithm. Here the upper index ℓ denotes the ℓth nested iteration while the lower indices 0 and J correspond to the splitting of the matrix $\boldsymbol{X}^\ell = \boldsymbol{X}_{\Lambda^\ell}$ for the Schur complement. Due to the real-time restriction the nested iteration is limited by time. However, this stopping criteria can easily be replaced, for example, by residual-based criteria using the residual $\boldsymbol{X}^{\ell+1}\boldsymbol{d}^\ell - \boldsymbol{r}^{\ell+1}$ on the next refinement. We would also like to remind the reader that the matrix $\boldsymbol{X}^\ell$ is a result of a reduction step (see (14)). Hence the matrix $\boldsymbol{X}^\ell$ is not only an extension of $\boldsymbol{X}^{\ell-1}$. However, we omit the details for easier readability.

initialize Λ_0
set $\ell = 0$

- *define* $\Lambda^\ell := \Lambda_0$
- *determine* $\boldsymbol{r}^\ell$ *and* $\boldsymbol{X}^\ell$
- *solve with a direct method* $\boldsymbol{X}^\ell \boldsymbol{d}^\ell = \boldsymbol{r}^\ell$

while time allows set $\ell = \ell + 1$ *and:*

- *determine adaptively* Λ^ℓ
- *set* $\Lambda^\ell_0 := \Lambda_0$ *and* $\Lambda^\ell_J := \Lambda^\ell \setminus \Lambda_0$
- *update* $\boldsymbol{X}^0$ *to* $\boldsymbol{X}^\ell_{0,0}$
- *assemble wavelet matrices* $\langle\langle \boldsymbol{\Psi}_{\Lambda_0}, \boldsymbol{\Psi}_{\Lambda^\ell_J} \rangle\rangle$, *etc. which are involved in* $\boldsymbol{X}^\ell_{J,0}$ *and* $\boldsymbol{X}^\ell_{J,J}$ *to realize their applications to a vector*
- *determine the preconditioner* $\boldsymbol{B}^\ell$ *using only the diagonals of all involved matrices*

- *update* $\boldsymbol{r}^0$ *to* $\boldsymbol{r}_0^\ell$ *and determine* $\boldsymbol{r}_J^\ell$
- *determine* $\boldsymbol{r}_{\mathcal{S}}^\ell = \boldsymbol{r}_J^\ell - \boldsymbol{X}_{J,0}^\ell (\boldsymbol{X}_{0,0}^\ell)^{-1} \boldsymbol{r}_0^\ell$ *using a direct solver for* $(\boldsymbol{X}_{0,0}^\ell)^{-1}\boldsymbol{r}_0^\ell$ *and applying* $\boldsymbol{X}_{J,0}^\ell$
- *extend* $\boldsymbol{d}_J^{\ell-1}$ *with zeros to define the initial guess* $\boldsymbol{d}_J^\ell$
- *apply the pcg method to solve* $\boldsymbol{\mathcal{S}}^\ell \boldsymbol{d}_J^\ell = \boldsymbol{r}_{\mathcal{S}}^\ell$ *using the initial guess* $\boldsymbol{d}_J^\ell$ *and the preconditioner* $\boldsymbol{B}^\ell$ *until a given tolerance is reached or the number of maximal iterations is exceeded (the application of* $\boldsymbol{\mathcal{S}}^\ell$ *to a vector is realized by applying* $\boldsymbol{X}_{J,0}^\ell$, $\boldsymbol{X}_{J,J}^\ell$, *and a direct solver for* $(\boldsymbol{X}_{0,0}^\ell)^{-1}$*)*
- *solve* $\boldsymbol{X}_{0,0}^\ell \boldsymbol{d}_0^\ell = \boldsymbol{r}_0^\ell - \boldsymbol{X}_{0,J}^\ell \boldsymbol{d}_J^\ell$
- *set* $\boldsymbol{d}^\ell = ((\boldsymbol{d}_0^\ell)^T, (\boldsymbol{d}_J^\ell)^T)^T$

end

The complexity of one nested iteration step ℓ is therefore given by the following operations: the setup of $\boldsymbol{X}_{0,0}^\ell$, which results in $O(n_x^3(\#\Lambda_0)^2\#\Lambda_J^\ell) + O(n_x^3(\#\Lambda_0)^3)$ operations; the setup of $(\boldsymbol{X}_{0,0}^\ell)^{-1}$ (for example with Cholesky decomposition) with $O(n_x^3(\#\Lambda_0)^3)$ operations; the determination of $\boldsymbol{B}^\ell$ with $O(n_x\#\Lambda_J^\ell)$ operations; and then for each pcg iteration steps the applications of $\boldsymbol{X}_{J,0}^\ell$, $\boldsymbol{X}_{J,J}^\ell$, and $(\boldsymbol{X}_{0,0}^\ell)^{-1}$, which result in $O(n_x^2\#\Lambda_J^\ell log(\#\Lambda_J^\ell)) + O(n_x^2(\#\Lambda_0)^2)$ operations. Considering $\ell \to \infty$ one can think of applying the algorithm proposed in [12] for the matrix-vector multiplication involving $\boldsymbol{X}_{J,0}^\ell$, $\boldsymbol{X}_{J,J}^\ell$. This may lead to a complexity without $log(\#\Lambda_J^\ell)$. However, a detailed study of the applicability is not yet completed. Moreover, one has to keep in mind that in practice n_x can be very large, and then the real-time limitation restricts $\#\Lambda^\ell$. The number of necessary pcg iterations in each nested iteration step ℓ is bounded independently of ℓ due to (27).

4.2 Eigenvalue Distribution and Condition Numbers

Let us first illustrate the effects of the Schur complement method and of the diagonal preconditioning with our example from subsection 2.4. As Λ_0 we choose all wavelets up to scale four and then choose up to scale eight the two wavelets at the boundary $t = 0$, i.e., $\Lambda_0 = \{(0,0), \ldots, (3,7); (i,0), (i,1) \text{ for } i = 4, \ldots, 7\}$. The condition numbers behave as listed in Table 1. As can be seen, wavelet preconditioning alone would not be sufficient, while the combination with the Schur complement technique reduces the condition number by a factor of 1.e+6.

Table 1: Condition numbers on scale 9 ($h = 2^{-9}$).

	$\boldsymbol{X}$	$\boldsymbol{NXN}^T$	$\boldsymbol{\mathcal{S}}$	$\boldsymbol{B\mathcal{S}B}^T$
cond.	2.59e+10	2.42e+07	2.28e+06	7.36e+04

Figure 2 displays the eigenvalue distributions of $\boldsymbol{NXN}^T$ and $\boldsymbol{B\mathcal{S}B}^T$ on scale 9. Since the Schur complement $\boldsymbol{\mathcal{S}}$ has less eigenvalues we associated the indices of the eigenvalues with the employed index set Λ, respectively, with Λ/Λ_0, and listed them along the horizontal axis. We see that the scale-to-scale increase is extinguished. In addition, the Schur complement technique eliminates the smallest and the eleven largest eigenvalues of $\boldsymbol{NXN}^T$. Despite that the preconditioned Schur complement $\boldsymbol{B\mathcal{S}B}^T$ still has two small eigenvalues $(2.6e-5, 0.02)$, we see from the linear scaled plot that a large gap arises to the third smallest eigenvalue (0.51). This has, as mentioned in subsection 2.4, a large impact on the convergence behavior of the cg-method, and the effect can be seen on the iteration numbers listed in the next subsection.

As a second test problem we consider an example which inherits two main features occurring in process monitoring. First, the involved differential equations are stiff and, second, the model is nearly not observable. It is based on a simple model

Figure 2: Eigenvalue distribution of $\boldsymbol{NXN}^T$ and $\boldsymbol{BSB}^T$.

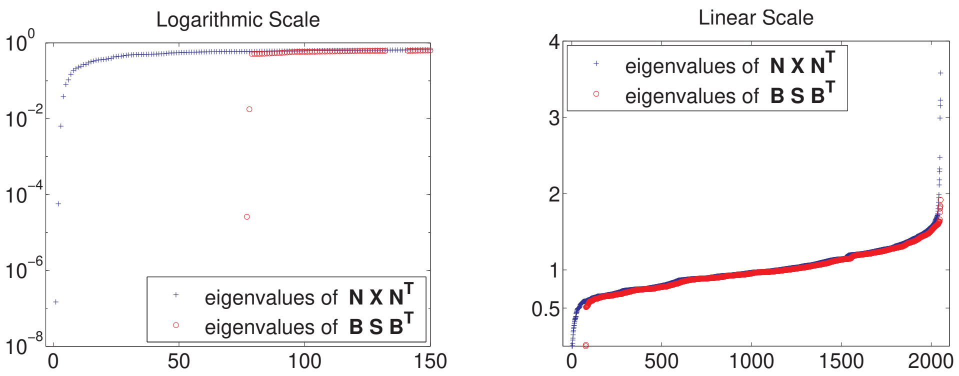

of a chemical process which produces ethylene glycol from ethylene oxide and glycol. The nonlinear model lumps a number of related states and outputs. Further simplification is accomplished by linearization at some stationary operating point and by subsequent elimination of the algebraic state variables to result in a linear model with 12 states and 5 measurements. Details on the model can be found in [7] and references therein. To ensure that the relative errors in the states are equally weighted, the model matrices are scaled such that the components of the nominal solution are roughly of unit size in L_∞. For brevity we do not give the numerical values of the system matrices. However, the eigenvalues of the model matrix $\boldsymbol{A}$ provide insight into system behavior and the values are given by $\{-189.90, -155.14, -3.78, -2.34, -1.98, -1.94, -1.15, -1.08 \pm i\,0.12, -0.89, -6.17e-6, 0\}$. The state x_2 corresponds to the eigenvalue 0 and depends completely on the initial value, which has to be estimated. Moreover, this state has a very low observability measure.

Given the original problem formulation, we were faced with condition numbers around **1e+20** already on scale 7. Scaling of the states to roughly unit size in L_∞ and regularization (with $\varepsilon = 2e-4$) lead to condition numbers around 1e+13 on scale 9. However, only the application of the Schur complement technique reduced the condition number to a reasonable size, namely 280, as can be seen in Table 2 for scale 9. We employed an index set with only 13 indices for each state, to be more precise, $\Lambda_0 = \{(0,0); \ldots; (2,3); (3,0); (4,0); (5,0); (6,0)\}$.

Table 2: Condition numbers on scale 9 ($h = 2^{-9}$).

	$\boldsymbol{X}$	$\boldsymbol{NXN}^T$	$\boldsymbol{\mathcal{S}}$	$\boldsymbol{B\mathcal{S}B}^T$
cond.	**4.52e+13**	1.79e+09	9.51e+04	**2.80e+02**

Figure 3 displays on the left the ten smallest and largest eigenvalues in logarithmic scale while on the right displays only the ten smallest eigenvalues in linear scale. The largest eigenvalues of $\boldsymbol{X}$ and $\boldsymbol{\mathcal{S}}$ as well as of $\boldsymbol{NXN}^T$ and $\boldsymbol{B\mathcal{S}B}^T$ are (nearly) the same. However, the factor between the smallest eigenvalue of $\boldsymbol{X}$ and $\boldsymbol{B\mathcal{S}B}^T$ is 1e+10. Moreover, as can be seen in the plot in linear scale, the Schur complement technique gives rise to large gaps between the smallest eigenvalues, much in favor of the cg-method.

4.3 Iteration Numbers and L_2-Errors

The iteration numbers are reduced more drastically than the condition number due to the arising gaps in the eigenvalue distribution. For instance, for our first example they drop from 973 to 22 for the applied pcg-method given the tolerance $tol = 1.e-6$ for the relative residual (see Table 3). Moreover, a comparable error in the case of

Figure 3: The ten smallest and largest eigenvalues.

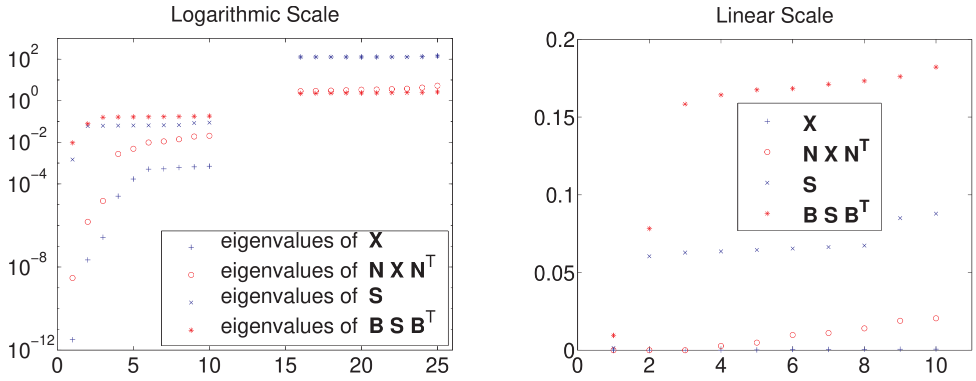

employing $\boldsymbol{X}$ is only reached with $tol = 1.e-10$ and requires 1650 iterations. Obviously, the presented wavelet-Schur-complement preconditioning yields not only lower iteration numbers but also a considerably decreased error. This effect is observed in general.

Table 3: pcg-method applied to $\boldsymbol{X}$ and $\boldsymbol{B}\boldsymbol{\mathcal{S}}\boldsymbol{B}^T$ on scale 9.

	tol.	iter.	$\|\|e_1\|\|_{L_2}$	$\|\|e_2\|\|_{L_2}$	$\|\|e_3\|\|_{L_2}$	$\|\|e_4\|\|_{L_2}$
$\boldsymbol{X}$	1.e-6	973	1.59e-1	6.56e-2	2.63e-2	3.61e-3
$\boldsymbol{X}$	1.e-10	1650	4.36e-5	1.22e-5	6.33e-6	3.45e-6
$\boldsymbol{B}\boldsymbol{\mathcal{S}}\boldsymbol{B}^T$	1.e-6	22	4.77e-5	1.53e-5	5.30e-6	2.91e-6

For our second example (where we started with a condition number of about 1e+20), we obtain iteration numbers around 23, independent of the chosen degree of refinement, which can be seen in Table 4. As tolerance for the relative residuals we use $tol = 1e-04$.

Table 4: Various scales for $\boldsymbol{B}\boldsymbol{\mathcal{S}}\boldsymbol{B}^T$.

Scale	7	8	9	10
Number of unknowns	3729	7441	14865	29713
Iterations	22	25	23	22

The resulting L_2-norms of the errors between the obtained approximations and the exact state functions are listed for the first five states in Table 5, while increasing the employed index set Λ^ℓ scalewise. With the presented preconditioning technique the wavelet-Galerkin discretization yields reasonable results even for the hardly observable state x_2. The rate of convergence behaves as expected, i.e., for step-sizes $h_i = 2^i$ we observe the quotient $4 = (h_i/h_{i+1})^2$ between the errors of succeeding scales. A tighter choice for the tolerance tol does not lead to higher accuracy. This confirms that $tol = 1e-04$ is sufficient. In contrast the SQOPT algorithm [17], designed in particular for quadratic optimization problems, stops after 500 iterations without result on scales larger than 6.

For the discussion of the adaptive refinement of Λ^ℓ we refer to other literature [5, 8]. We would like to mention at this point only that an appropriate choice of Λ^ℓ with much less indices would also guarantee similar L_2-errors. However, we face the same problems concerning the eigenvalue distribution: small eigenvalues are close to zero and large eigenvalues increase with the used scales. By employing the above preconditioning techniques, both features can be controlled. The required iteration numbers remain around 24, independent of the refinement level, and we obtain the maximal convergence order given by the discretization spaces. The SQOPT algorithm failed for the tested adaptive choices of Λ^ℓ with $\#\Lambda^\ell > 100$.

Table 5: L_2-norms of the errors in the states using $\boldsymbol{B\mathcal{S}B}^T$.

Scale	7	8	9	10
x_1	7.7178e-04	1.7699e-04	3.7967e-05	6.5373e-06
x_2	6.1436e+00	1.1187e+00	3.3672e-01	7.2894e-02
x_3	7.7386e-03	1.7650e-03	3.6426e-04	5.8558e-05
x_4	9.4941e-03	2.3140e-03	5.6531e-04	1.3459e-04
x_5	3.2000e-01	8.1444e-02	2.0382e-02	5.3118e-03

In summary we can say that for all examples we tested, the Schur complement technique combined with wavelet preconditioning provides small, scale-independent iteration numbers and in addition an improved estimation quality with lower necessary tolerances in the residuals.

4.4 Influence of the choice of Λ_0

To gain insight into the influence of the choice of Λ_0, we compared several choices of Λ_0 for our second example with respect to resulting eigenvalue distributions, condition numbers, required tolerances, iteration numbers, as well as final cpu times according to $\boldsymbol{B\mathcal{S}B}^T$.

First we calculated the 10 smallest eigenvalues and eigenvectors of the preconditioned matrix $\boldsymbol{NXN}^T$ on scale 9. Then, we approximated the eigenvectors of the 3 or, respectively, 10 smallest eigenvalues based on the norm equivalences for wavelets. We either used the N in modulo largest wavelet coefficients and, hence, employed as a measure for the distance to the relevant eigenspace the sum over the approximations to all considered eigenvectors, or restricted the maximal approximation error δ for each eigenvector. In the first case we considered various numbers N of largest coefficients; in the other case we varied the allowed approximation errors δ. We also considered the choice of Λ_0, as motivated in subsection 3.2, by varying the number of the wavelet functions employed at the beginning of the interval. In particular, we used the choice $\Lambda_0^* := \{(0,0);\ldots;(2,3);(3,0);(4,0);(5,0);(6,0)\}$, which was presented in subsection 4.2. On top of these variations, which have no requirement on the arising structure of index distributions, we either considered the additional requirement of a so-called tree structure of Λ_0, or assumed that the three coarsest levels are contained in the index set, i.e., $\{(0,0);\ldots;(2,3)\} \subseteq \Lambda_0$. The number of employed wavelets varied between 9 and 231 and the approximation error was, in the best case, roughly $9.2e-3$. In the figures below the variation of Λ_0 is sorted along the x-axis, and for each requirement on the index structure a new line or an additional figure is plotted. Numbers 1–19 correspond to the variation of N and δ, while 20–22 belong to the variation at the beginning of the interval. Number 22 corresponds to Λ_0^*.

4.4.1 Eigenvalue Distribution

The variance between the largest eigenvalues arising with the different choices of Λ_0 is, as expected, negligible. The largest eigenvalues stay between 3.585 and 1.463, while the largest eigenvalue of $\boldsymbol{NXN}^T$ is 5.321. Our application of the Schur complement technique aims to influence the smallest eigenvalues. Here large differences arise. While the smallest eigenvalue of $\boldsymbol{NXN}^T$ is $2.966e-9$, for $\boldsymbol{B\mathcal{S}B}^T$ they vary between $2.488e-4$ and $4.224e-1$. Clearly, for all choices $\boldsymbol{B\mathcal{S}B}^T$ has condition numbers which are reduced by at least a factor of $1e-5$. In the best case the condition number reduced to 3.46, although here we employed an index set with $\#\Lambda_0 = 231$ per state (number 19). (The complete index set on scale 9 inhabits 513 indices.) The condition number 29.08 is reached with $\#\Lambda_0 = 42$ and 279.47 with Λ_0^*, which has $\#\Lambda_0^* = 13$.

Using the various approximations of the eigenvectors the minimal eigenvalues are roughly the same as long as we do not incorporate drastically more indices in the set Λ_0 (as for numbers 18 and 19 in the plots). This indicates that the way of measuring the distance to the relevant eigenspace has only a small, if any, influence. Moreover, it is surprising that three eigenvalues were extinguished for any considered Λ_0, while the partly desired elimination of ten eigenvalues was only attained after the additional requirement of a tree structure for the indices. Hence, it seems to be difficult to eliminate the eigenvalues based only on approximating the corresponding eigenvectors. A significant difference can be found for the choices of wavelets at the beginning of the interval. Such behavior is observed in all numerical experiments.

As mentioned, also gaps in the eigenvalue distribution strongly influence the efficiency of the cg-method. Figure 4 displays the distribution of the eight smallest eigenvalues of $\boldsymbol{B}\boldsymbol{\mathcal{S}}\boldsymbol{B}^T$. Obviously without structure Λ_0^* (number 22) has the most favorable distribution, namely large gaps between the smallest, second smallest, and third smallest eigenvalue. Only with a required tree structure are the gaps between the eigenvalues improved to the extent that 1–19 beat 22. However, their index sets include Λ_0 of 22, which is Λ_0^* extended to tree structure. Of course not only the eigenvalue distribution but also the number of employed wavelets play a significant role in the efficiency, as can be seen later on.

Figure 4: The eight smallest eigenvalues of $\boldsymbol{B}\boldsymbol{\mathcal{S}}\boldsymbol{B}^T$.

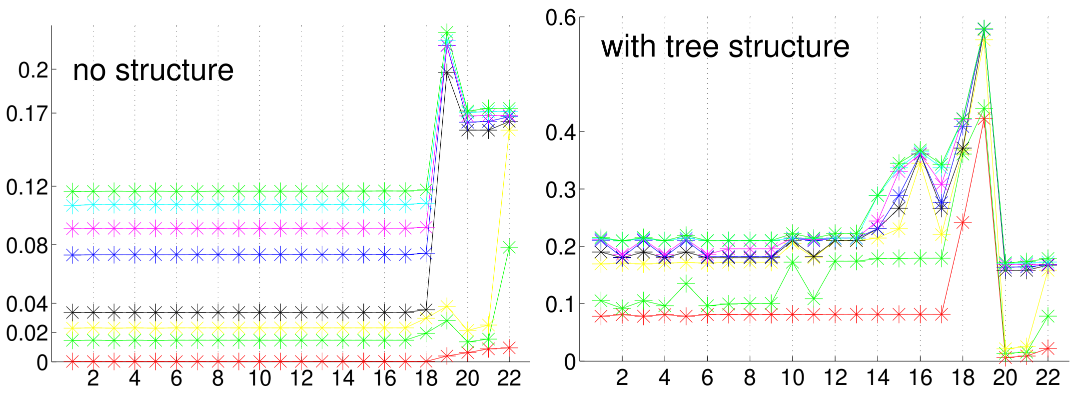

4.4.2 Iteration numbers

One important stopping criteria of an iterative method is the tolerance for the (relative) residual. It turned out that, in general, $tol = 1.e-4$ is not sufficient to guarantee the best attainable approximation error given by the discretization error of the optimization problem. In all cases a tolerance for the relative residual of $tol = 1.e-6$ led to the same approximation errors fulfilling the desired requirements. Hence the results in the next subsections are based on the use of $tol = 1.e-6$.

As the eigenvalue distributions indicated already, Λ_0^* provides the lowest iteration numbers over the scales, neither assuming index structure nor incorporating the coarsest three scales (34, 31, and 30 iterations on scale 8, 9, and 10, respectively). Requiring tree structure, 22 needs 31, 27, and 26 iterations, while 15 with $\#\Lambda_0 = 96$ needs 16, 14, and 16 iterations and 18 with $\#\Lambda_0 = 168$, which can only start on scale 9, needs 10 and 11 iterations. Figure 5 presents graphically the average of the necessary iteration numbers over the scales.

4.4.3 Cpu Time

Figure 6 displays the required cpu times. For scaling reasons only values below 500s are plotted. Here, the computation time is measured, which is required for the

Figure 5: Averaged iteration numbers of the pcg-method applied to $\boldsymbol{B}\boldsymbol{\mathcal{S}}\boldsymbol{B}^T$.

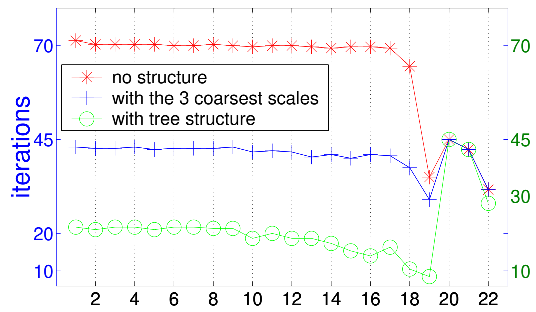

Figure 6: Cpu time of the nested iteration applied to $\boldsymbol{B}\boldsymbol{\mathcal{S}}\boldsymbol{B}^T$.

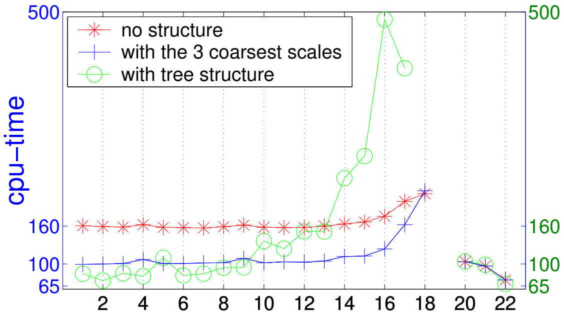

algorithm presented in subsection 4.1, with $\Lambda^1, \Lambda^2, \Lambda^3$ chosen as the complete scales 8, 9, 10, respectively. Only for 18 and 19 did we start at the minimal scale 9, such that the computation times are not really comparable. Nevertheless, it shows that 18 and 19 already require for two scales more computation time than the others for three scales. The computations are executed on an IntelXeon processor with 4GB and 700MHz. Moreover, we would like to emphasize that the code used is not yet time optimal and is used for academic purposes only. However, since the experiments are all run with the same code, the comparison of the times between the various Λ_0 is still valid.

We can clearly observe the interaction between the improved eigenvalue distribution and the number of employed indices with respect to the efficiency measured by the cpu time. Less numbers of indices can reduce the influence of larger condition numbers and vice versa. Due to the best eigenvalue distributions and the very low number of indices, Λ_0^* is the fastest in the case of either no requirement on the index structure or of incorporating the three coarsest levels. This choice leads to 74.69s, while 1–16 need approximately 160s in the first case and approximately 100s in the second case. In case of the tree structure it depends very much on the number of employed indices. Choice 22 is still the most efficient with 68.26s despite its condition number, which is at least 20 times as large as for choice 18 and 19. The originally chosen set Λ_0^* is only 6s slower.

In summary, we have to say that a general guideline for the choice of Λ_0 in order to obtain the optimal efficiency cannot be provided so far. However, for our particular example, the choice Λ_0^* and its extension to a tree structure provides the best results. In the author's opinion this is due to the fact that one main difficulty of the given optimization problem is, as also indicated in Proposition 3.2, to determine backwards in time the initial data from some given data on the interval. The initial data can only be calculated via a differential equation with negative, large in modulo eigenvalues of $\boldsymbol{A}$, and the choice of Λ_0^* reflects this difficulty.

5 Conclusions

The analytical considerations in this paper suggest the use of a Schur complement as preconditioning. Extreme eigenvalues of symmetric matrices can be eliminated with this technique. The way to choose a good subspace is, although theoretically motivated, still heuristic. Here, the Riesz basis property of wavelets provide the connection with the original problem in function spaces. Knowledge about the origin of the ill-conditioning of this problem may lead to an appropriate coarse choice of Λ_0 based on wavelet compression. Nevertheless, further analysis and investigations are necessary.

Despite that, all numerical experiments, whether listed here or not, indicate successful applications. In our second tested example, a simple choice of Λ_0 reduced the condition number by a factor of $1e+7$ and in combination with the wavelet-based preconditioning by a factor of $1e+9$. Obviously, this influences not only the convergence behavior of an iterative method but also the necessary tolerance as stopping criteria for obtaining a sufficient approximation quality. In addition to the extremely small condition numbers the eigenvalue distribution improves too. In most cases large gaps between the smallest eigenvalues arise, which highly favors the efficiency of the pcg-method.

For our application in state estimation, where we are typically faced with very large condition numbers, only the application of the Schur complement technique as a preconditioner enabled us to apply an iterative method in a successful way, providing high estimation quality with low computational cost.

References

[1] F. Allgöwer, T.A. Badgwell, J.S. Qin, J.B. Rawlings, and S.J. Wright. Nonlinear predictive control and moving horizon estimation – an introductory overview. In P.M. Frank, editor, *Advances in Control: Highlights of ECC 1999*, pages 391–449. Springer–Verlag, 1999.

[2] O. Axelsson. *Iterative Solution Methods*. Cambridge University Press, Berlin, New York, Cambridge, 1994.

[3] A. Battermann and M. Heinkenschloss. Preconditioners for Karush-Kuhn-Tucker matrices arising in the optimal control of distributed systems. In W. Desch, F. Kappel, and K. Kunisch, editors, *Control and Estimation of Distributed Parameter Systems*, ISNM 126, pages 15–32. Birkhäuser, Basel, 1998.

[4] T. Binder, L. Blank, W. Dahmen, and W. Marquardt. Towards multiscale dynamic and data reconciliation. In R. Berber and C. Kravaris, editors, *Nonlinear Model Based Process Control, NATO ASI*, pages 623–665. Kluwer Academic Publishers, Dordrecht, the Netherlands, 1998.

[5] T. Binder, L. Blank, W. Dahmen, and W. Marquardt. Grid refinement in multiscale dynamic optimization. In S. Pierucci, editor, *Proceedings of the European Symposium on Computer Aided Process Engineering, (ESCAPE-10)*, pages 31–36. Elsevier, New York, 2000.

[6] T. Binder, L. Blank, W. Dahmen, and W. Marquardt. Iterative algorithms for multiscale state estimation, part 1: Concept. *J. Optim. Theory Appl.*, 111(3):501–527, 2001.

[7] T. Binder, L. Blank, W. Dahmen, and W. Marquardt. Iterative algorithms for multiscale state estimation, part 2: Numerical investigation. *J. Optim. Theory Appl.*, 111(3):529–551, 2001.

[8] T. Binder, L. Blank, W. Dahmen, and W. Marquardt. Multiscale concepts for moving horizon optimization. In M. Grötschel, S.O. Krumke, and J. Rambau, editors, *Online Optimization of Large Scale Systems*, pages 341–361. Springer–Verlag, Berlin, Heidelberg, 2001.

[9] T. Binder, L. Blank, W. Dahmen, and W. Marquardt. On the regularization of dynamic data reconciliation problems. *Journal of Process Control*, 12:557–567, 2002.

[10] L. Blank. State estimation without regularizing the initial data. Preprint, 2003.

[11] C.K. Chui. *An Introduction to Wavelets*. Academic Press, Boston, Massachusetts, 1992.

[12] A. Cohen, W. Dahmen, and R.A. DeVore. Adaptive wavelet methods II. Beyond the elliptic case. *Found. Comput. Math.*, 2(3):203–245, 2002.

[13] W. Dahmen. Wavelet and multiscale methods for operator equations. *Acta Numerica*, 7:55–228, 1997.

[14] I. Daubechies. *Ten Lectures on Wavelets.* SIAM, Philadelphia, Pennsylvania, 1992.

[15] H.C. Elman, D.J. Silvester, and A.J. Wathen. Performance and analysis of saddle point preconditioners for the discrete steady–state Navier–Stokes equations. *Numer. Math.*, 90(4):665–688, 2002.

[16] J. Frank and C. Vuik. On the construction of deflation–based preconditioners. *SIAM J. Sci. Comput.*, 23(2):442–462, 2001.

[17] P.E. Gill, M. Murray, and M.A. Saunders. User's guide for SNOPT 5.3: a Fortran package for large–scale nonlinear programming. Report na 97-5, Department of Mathematics, University of California, San Diego, 1997.

[18] G.H. Golub and C.F. Van Loan. *Matrix Computations.* The John Hopkins University Press, Baltimore, MD, 3rd edition, 1996.

[19] A. Greenbaum. *Iterative Methods for Solving Linear Systems.* SIAM, Philadelphia, 1997.

[20] M.R. Hestenes. *Calculus of Variations and Optimal Control Theory.* John Wiley & Sons, Inc., New York, 1966.

[21] T. Kailath. *Linear Systems.* Prentice Hall, Englewood Cliffs, New Jersey, 1980.

[22] L. Mansfield. Damped Jacobi preconditioning and coarse grid deflation for conjugate gradient iteration on parallel computers. *SIAM J. Sci. Statist. Comput.*, 12:1314–1323, 1991.

[23] G. Meurant. *Computer Solution of Large Linear Systems.* North–Holland, Amsterdam, the Netherlands, 1999.

[24] R.B. Morgan. A restarted GMRES method augmented with eigenvectors. *SIAM J. Matrix Anal. Appl.*, 16(4):1154–1171, 1995.

[25] R.A. Nicolaides. Deflation of conjugate gradients with applications to boundary value problems. *SIAM J. Numer. Anal.*, 24:355–365, 1987.

[26] Y. Saad. Analysis of augmented Krylov subspace methods. *SIAM J. Matrix Anal. Appl.*, 18(2):435–449, 1997.

[27] J.G. Schmidt and G. Starke. Coarse space orthogonalization for indefinite linear systems of equations arising in geometrically nonlinear elasticity. *Appl. Numer. Math.*, 30:267–280, 1999.

[28] A. van der Sluis and H.A. van der Vorst. The rate of convergence of conjugate gradients. *Numer. Math.*, 48:543–560, 1986.

[29] C. Vuik, A. Segal, and J.A. Meijerink. An efficient preconditioned CG method for the solution of a class of layered problems with extreme contrasts in the coefficients. *J. Comput. Phys.*, 152:1–19, 1999.

[30] W. Waldraff, D. Dochain, S. Bourrel, and A. Magnus. On the use of observability measures for sensor location in tubular reactors. *Journal of Process Control*, 8:497–505, 1998.

[31] J.H. Wilkinson. *The Algebraic Eigenvalue Problem.* Oxford University Press, London, 1965.