Operator Theory:
Advances and Applications, Vol. 167, 69–85

A Hierarchical Semi-separable Moore-Penrose Equation Solver

Patrick Dewilde and Shivkumar Chandrasekaran

Abstract. The main result of the present paper is a method to transform a matrix or operator which has a hierarchical semi-separable (HSS) representation into a URV (Moore-Penrose) representation in which the operators U and V represent collections of efficient orthogonal transformations and the block upper matrix R still has the HSS form. The paper starts with an introduction to HSS-forms and a survey of a recently derived multi resolution representation for such systems. It then embarks on the derivation of the main ingredients needed for a Moore-Penrose reduction of the system while keeping the HSS structure. The final result is presented as a sequence of efficient algorithmic steps, the efficiency resulting from the HSS structure that is preserved throughout.

Mathematics Subject Classification (2000). Primary: 65F; Secondary: 15A.

Keywords. Hierarchically semi-separable systems, hierarchically quasi-separable systems, Moore-Penrose inverse, structured matrices.

1. Introduction

Many physical systems are modeled by systems of differential equations, integral equations or combinations of them. Solving these systems requires discretization of the equations and leads to large systems of algebraic equations. In the case of systems governed by linear equations, the resulting system of equations will be linear as well and can be solved either directly or iteratively, involving a preconditioner and a Lanczos-type recursion. In the case of non-linear systems, the discretized system will be non-linear as well and an iterative procedure has to be set up to find the solution. Such a procedure is, e.g., of the 'Newton-Raphson' type and would in turn require the solution of a system of linear equations, now involving a differential such as a Jacobian. In all cases the resulting systems tend to lead to matrices of very large dimensions, even for fairly small problems, so that solvers using standard numerical procedures quickly run out of steam. Luckily, many systems exhibit quite a bit of structure that can be exploited to make the

solver more efficient. In the case of a discretized system for a differential equation, the resulting matrix structure will be very sparse, as only entries corresponding to nearby points will be different from zero. In this case, the matrix-vector multiplication can be efficiently executed. Iterative methods are well suited to exploit this fact, but they are dependent on a low-complexity approximant of the inverse of the original system, the so-called preconditioner. Approximate solutions of the system can be iteratively constructed via low-complexity calculations, provided the pre-conditioner exhibits the necessary structure.

In the case of integral equations, it has been remarked by Gohberg, Kailath and Koltracht [9] and Rokhlin [12] that low rank approximations of the integral kernel lead to large submatrices of low rank in the resulting system of equations. Exploiting this structure which was termed 'Semi-Separable' leads to solution procedures that are linear in the size of the matrix and quadratic in the size of the approximation. A systematic method to obtain such low rank approximations was proposed by Greengard and Rokhlin [10] and is known as the Fast Multipole Method. These original approaches suffered from numerical problems as the use of backward stable orthogonal transformations in this context was not yet well understood. The introduction of time-varying system theory to model the system of equations [13] provided for the necessary structure to allow for more general types of transformations than those used by the original authors cited. A survey of these techniques can be found in the book [6]. Based on these ideas, the Semi-Separable structure was extended to a more generic form called 'Quasi-Separable' and numerically stable system solvers were developed for this structure by a number of authors [8, 3, 7].

Although these developments lead to a satisfactory and useful theory, it was also evident that they did not exploit the structural properties of most physical systems sufficiently. Two examples may suffice to illustrate this point. In the case of a partial differential equation in 3D space, discretization coordinates will have three indices, say $\{i, j, k\}$, and interaction between values in close-by points may be expected. To construct the interaction matrix, each discretization coordinate has to be assigned a single index. In case of a regular grid of dimension N^3, the index assignment would run as $i + Nj + N^2k$ and the resulting matrix would have a hierarchical structure consisting of a diagonal bands of blocks of dimension N^2 each consisting of diagonal bands of blocks of dimension N, which in turn consist of scalar diagonal bands. The Semi-Separable or Quasi-Separable theory is insufficient to handle such types of matrices, it gets a good grip only on the top level of the hierarchy, while the structure of the lower hierarchical levels is greatly disturbed [5]. Also in the case of the multipole method and assuming the distribution of 'objects' handled by the method in 3D space to be fairly general (assuming of course that the multipole assumption holds as well), a similar problem will arise: many submatrices will be of low rank, but they will have a 'hierarchical ordering', restricting the applicability of the Semi-Separable method.

In the present paper we deal with an intermediate structure which has a nice hierarchical (or equivalently multi-resolution) structure and is capable to cope

with the problem mentioned in the previous paragraph. The structure was first presented in [4], and a few solvers for it were presented [11]. In particular, [2] shows how the structure can be reduced to a sparse, directly solvable system, using a state space model of intermediate variables, much as was done for the Semi- or Quasi-separable case. In contrast to the latter the state-space model turns out to have a hierarchical (multi-resolution) structure, which can efficiently be exploited. These straight solvers assume the original system to be square non-singular, allowing for partial recursive elimination of unknowns and recursive back-substitutions as the algorithm proceeds. In the present paper we propose a new, backward stable solver that finds the Moore-Penrose solution for a general system of equations, namely a system that is not assumed to be square, non-singular. Our goal is to obtain the same order of numerical complexity as the straight solvers, but now for the Moore-Penrose case.

2. HSS representations

The Hierarchical Semi-Separable representation of a matrix A is a layered representation of the multi-resolution type, indexed by the hierarchical level. At the top level 1, it is a 2×2 block matrix representation of the form:

$$A = \begin{bmatrix} A_{1;1,1} & A_{1;1,2} \\ A_{2;2,1} & A_{2;2,2} \end{bmatrix} \tag{2.1}$$

in which we implicitly assume that the ranks of the off-diagonal blocks is low so that they can be represented by an 'economical' factorization ('H' indicates Hermitian transposition, for real matrices just transposition), as follows:

$$A = \begin{bmatrix} D_{1;1} & U_{1;1}B_{1;1,2}V_{1;2}^H \\ U_{1;2}B_{1;2,1}V_{1;1}^H & D_{1;2} \end{bmatrix}. \tag{2.2}$$

The second hierarchical level is based on a further but similar decomposition of the diagonal blocks, respect. $D_{1;1}$ and $D_{1;2}$:

$$\begin{aligned} D_{1;1} &= \begin{bmatrix} D_{2:1} & U_{2;1}B_{2;1,2}V_{2;2}^H \\ U_{2;2}B_{2;2,1}V_{2;1}^H & D_{2;2} \end{bmatrix}, \\ D_{1;2} &= \begin{bmatrix} D_{2;3} & U_{2;3}B_{2;3,4}V_{2;4}^H \\ U_{2;4}B_{2;4,3}V_{2;3}^H & D_{2;4} \end{bmatrix} \end{aligned} \tag{2.3}$$

for which we have the further *level compatibility* assumption

$$\operatorname{span}(U_{1;1}) \subset \operatorname{span}\left(\begin{bmatrix} U_{2;1} \\ 0 \end{bmatrix}\right) \oplus \operatorname{span}\left(\begin{bmatrix} 0 \\ U_{2;2} \end{bmatrix}\right), \tag{2.4}$$

$$\operatorname{span}(V_{1;1}) \subset \operatorname{span}\left(\begin{bmatrix} V_{2;1} \\ 0 \end{bmatrix}\right) \oplus \operatorname{span}\left(\begin{bmatrix} 0 \\ V_{2;2} \end{bmatrix}\right) \text{ etc} \ldots \tag{2.5}$$

This spanning property is characteristic for the HSS structure, it is a kind of hierarchical 'Lanczos' property and allows a substantial improvement on the numerical complexity as a multiplication with higher level structures always can be

done using lower level multiplications, using so called 'translation operators'

$$U_{1;i} = \begin{bmatrix} U_{2;2i-1}R_{2;2i-1} \\ U_{2;2i}R_{2;2i} \end{bmatrix}, \quad i = 1, 2, \tag{2.6}$$

$$V_{1;i} = \begin{bmatrix} V_{2;2i-1}W_{2;2i-1} \\ V_{2;2i}W_{2;2i} \end{bmatrix}, \quad i = 1, 2. \tag{2.7}$$

Notice the use of indices: at a given level i rows respect. columns are subdivided in blocks indexed by $1, \ldots, i$. Hence the ordered index $(i; k, \ell)$ indicates a block at level i in the position (k, ℓ) in the original matrix. The same kind of subdivision can be used for column vectors, row vectors and bases thereof (as are generally represented in the matrices U and V).

In [2] it is shown how this multilevel structure leads to efficient matrix-vector multiplication and a set of equations that can be solved efficiently as well. For the sake of completeness we review this result briefly here. Let us assume that we want to solve the system $Ax = b$ and that A has an HSS representation with deepest hierarchical level K. We begin by accounting for the matrix-vector multiplication Ax. At the leave node $(K; i)$ we can compute

$$g_{K;i} = V^H_{K;i}x_{K;i}.$$

If $(k; i)$ is not a leaf node, we can infer, using the hierarchical relations

$$g_{k;i} = V^H_{k;i}x_{k;i} = W^H_{k+1;2i-1}g_{k+1;2i-1} + W^H_{k+1;2i}g_{k+1;2i}.$$

These operations update a 'hierarchical state' $g_{k;i}$ upwards in the tree. To compute the result of the multiplication, a new collection of state variables $\{f_{k;i}$ is introduced for which it holds that

$$b_{k;i} = A_{k;i,i} + U_{k;i}f_{k;i}$$

and which can also be computed recursively downwards by the equations

$$\begin{bmatrix} f_{k+1;2i-1} \\ f_{k+1;2i} \end{bmatrix} = \begin{bmatrix} B_{k+1;2i-1,2i}g_{k+1;2i} + R_{k+1;2i-1}f_{k,i} \\ B_{k+1;2i,2i-1}g_{k+1;2i-1} + R_{k+1;2i}f_{k;i} \end{bmatrix},$$

the starting point being $f_{0;} = []$, an empty matrix. At the leaf level we can now compute (at least in principle - as we do not know x) the outputs from

$$b_{K;i} = D_{K;i}x_{K;i} + U_{K;i}f_{K;i}.$$

The next step is to represent the multiplication recursions in a compact form using matrix notation and without indices. We fix the maximum order K as before. Next we define diagonal matrices containing the numerical information, in breadth first order:

$$\mathbf{D} = \mathrm{diag}[D_{K;i}]_{i=1,\ldots,K}, \ \mathbf{W} = \mathrm{diag}[(W_{1;i})_{i=1,2}, (W_{2;i})_{i=1\cdots4}, \ldots], \ \text{etc}\ldots$$

Next, we need two shift operators relevant for the present situation, much as the shift operator Z in time-varying system theory [6]. The first one is the shift-down operator $Z_\downarrow$ on a tree. It maps a node in the tree on its children and is a nilpotent operator. The other one is the level exchange operator $Z_\leftrightarrow$. At each level

it exchanges children of the same node and is a permutation operator. Finally, we need the leaf projection operator $\mathbf{P}_{\text{leaf}}$ which on a state vector which assembles in breadth first order all the values $f_{k;i}$ produces the values of the leaf nodes (again in breadth first order). The state equations representing the efficient multiplication can now be written as

$$\left\{ \begin{array}{rcccc} \mathbf{g} & = & \mathbf{P}_{\text{leaf}}^H \mathbf{V}^H \mathbf{x} & + & Z_\downarrow^H \mathbf{W}^H \mathbf{g} \\ \mathbf{f} & = & \mathbf{R} Z_\downarrow \mathbf{f} & + & \mathbf{B} Z_\leftrightarrow \mathbf{g} \end{array} \right. \tag{2.8}$$

while the 'output' equation is given by

$$\mathbf{b} = \mathbf{D}\mathbf{x} + \mathbf{U}\mathbf{P}_{\text{leaf}}\mathbf{f}. \tag{2.9}$$

This system of equations is sparse and can always be solved (even efficiently, that is by visiting the given data once), because $(I - \mathbf{W}Z_\downarrow)$ and $(I - \mathbf{R}Z_\downarrow)$ are invertible operators due to the fact that $Z_\downarrow$ is nilpotent. We obtain

$$A = \mathbf{D} + \mathbf{U}\mathbf{P}_{\text{leaf}}(I - \mathbf{R}Z_\downarrow)^{-1}\mathbf{B}(I - Z_\downarrow^H \mathbf{W}^H)^{-1}\mathbf{P}_{\text{leaf}}\mathbf{V}^H)\mathbf{x} = \mathbf{b}. \tag{2.10}$$

Various strategies can be used to solve this sparse system of equations, we refer to the paper mentioned for more information. One elimination procedure that is aesthetically attractive follows the hierarchical ordering of the data bottom up. In a tree that is two levels down the elimination order would be:

$$\begin{array}{c} (f_{2;1}, g_{2;1}, x_{2;1}), (f_{2;2}, g_{2;2}, x_{2;2}), (f_{2;3}, g_{2;3}, x_{2;3}), (f_{2;4}, g_{2;4}, x_{2;4}), \\ (f_{1;1}, g_{1;1}), (f_{1;2}, g_{1;2}), (f_{0;1}, g_{0;1}). \end{array} \tag{2.11}$$

The computation must start at the leaf nodes, where multiplication with the base vectors takes place, in higher up locations there is only multiplication with transfer operators which relate the higher up bases to the bases at the leaf level. In the paper cited it is shown that this procedure hierarchically eliminates unknowns without producing any fill ins in the original sparse matrix describing the system.

The present paper aims at presenting a QR-type elimination procedure capable of deriving the Moore-Penrose inverse of a (presumably singular or ill conditioned) problem. The additional difficulty here is that elimination of variables cannot be done on the fly, because the Moore-Penrose solution can only be determined after the whole structure has been visited. Therefore we will aim at constructing the Moore-Penrose inverse rather than at solving the equations recursively as they appear.

3. Preliminaries

We shall use a number of reduction theorems (in a well-specified order).

Proposition 1. *Let* $V = \left[\frac{V_1}{V_2}\right]$ *be a (tall) full rank matrix of size* $(k+m) \times k$ *with* $m \geq k$*, then an orthogonal transformation* Q *that reduces* V *to the form*

$$QV = \left[\frac{0}{R}\right] \tag{3.12}$$

with R square non-singular exists and can be chosen of the form

$$Q = \left[\begin{array}{c|c} w_1 & -L^H w_2^H \\ \hline w_2 & I - w_2 K w_2^H \end{array}\right] \tag{3.13}$$

in which $w_1 = V_1 R^{-1}$, $w_2 = V_2 R^{-1}$, K is a hermitian matrix that satisfies

$$I - w_2 K w_2^H = (I - w_2 w_2^H)^{1/2} \tag{3.14}$$

and L a unitary matrix that satisfies

$$(I - K w_2^H w_2) = w_1 L. \tag{3.15}$$

Proof. The theorem claims the existence of K and L implicitly. If V_1 happens to be invertible, then this is straightforward, the difficulty is when V_1 is not invertible, we use an implicit proof to cover the general case. R is defined as a square matrix that satisfies

$$V_1^H V_1 + V_2^H V_2 = R^H R \tag{3.16}$$

and its non-singularity follows from the non-singularity assumption on V. Next let $w_1 = V_1 R^{-1}$ and $w_2 = V_2 R^{-1}$, then $w = \left[\begin{array}{c} w_1 \\ \hline w_2 \end{array}\right]$ is isometric, $w_1^H w_1$ is contractive and an eigenvalue decomposition

$$w_1^H w_1 = v_1 \sigma^2 v_1^H \tag{3.17}$$

can be chosen such that the positive diagonal matrix σ satisfies $0 \leq \sigma \leq I$. Since w_1 is square, a unitary u_1 will exist such that $w_1 = u_1 \sigma v_1^H$ (the proof goes as in the proof of the SVD). Next, $w_2^H w_2 = I - w_1^H w_1$ and an eigenvalue decomposition for it is

$$v_1 (I - \sigma^2) v_1^H. \tag{3.18}$$

Since w_2 is tall by assumption, there will exist u_2 isometric such that

$$w_2 = u_2 (I - \sigma^2)^{1/2} v_1^H. \tag{3.19}$$

It is now easy to verify directly that

$$Q = \left[\begin{array}{c|c} u_1 \sigma v_1^H & -u_1 (I - \sigma^2)^{1/2} u_2^H \\ \hline u_2 (I - \sigma^2)^{1/2} v_1^H & I - u_2 (I - \sigma) u_2^H \end{array}\right] \tag{3.20}$$

is a unitary matrix. Putting $K = v(I + \sigma)^{-1} v^H$ and $L = v u_1^H$ produces the desired form of Q. The converse check that any Q with the given form is unitary is also immediate by direct verification. □

The theorem shows that a unitary matrix close to identity (where 'close' means 'the difference is a low rank matrix') can be constructed that reduces a tall matrix to a small triangular matrix. In Numerical Analysis one traditionally uses 'Householder transformations' for this purpose, the transformation presented here has the advantage that its determinant can be controlled more easily.

Proposition 2. *Let* $T = \left[\begin{array}{c} UV^H \\ \hline D \end{array}\right]$ *in which* U *is a tall, isometric matrix of rank* δ, T *is of dimension* $(k+m)\times m$, *accordingly partitioned and of full column rank, and* $k \leq m$. *Let* $N^H N = VV^H + D^H D$, *in which* N *is a square matrix. Then* N *is non-singular, and there exists a unitary matrix* Q *such that*

$$T = Q \left[\begin{array}{c} 0 \\ \hline N \end{array}\right]. \tag{3.21}$$

Moreover, Q *can be chosen of the form*

$$Q = \left[\begin{array}{c|c} d_1 & Uv_r^H \\ \hline u_\ell v_\ell^H & d_2 \end{array}\right] \tag{3.22}$$

in which u_ℓ *has at most the same rank as* U *and* d_1 *is a rank* δ *perturbation of the unit matrix.*

Proof. The non-singularity of N follows directly from the full column rank assumption. The proof is then based on 'joint' SVD's of $UV^H N^{-1}$ and DN^{-1} and then completing the form. More precisely, let $N^{-T}VUU^HV^HN^{-1} = v\sigma^2 v^H$ be an eigenvalue decomposition with unitary v and a positive diagonal matrix σ of dimension $m \times m$. Let $\pi_k = \left[\begin{array}{c|c} I_k & 0_{m-k} \end{array}\right]$. Then there will exist unitary matrices u_1 and u_2 so that

$$\left[\begin{array}{c} UV^H N^{-1} \\ \hline DN^{-1} \end{array}\right] = \left[\begin{array}{c} u_1 \pi_k \sigma v^H \\ \hline u_2 (I-\sigma^2)^{1/2} v^H \end{array}\right]. \tag{3.23}$$

(It is not hard to check that the right-hand side form is indeed isometric!) Let moreover $\sigma_k = \pi_k \sigma \pi_k^H$. A Q that satisfies the requirements of the theorem is now easily constructed as

$$Q = \left[\begin{array}{c|c} u_1 (I-\sigma_k^2)^{1/2} u_1^H & u_1 \pi_k \sigma v^H \\ \hline -u_2 \sigma \pi_k^H u_1^H & u_2 (I-\sigma^2)^{1/2} v^H \end{array}\right]. \tag{3.24}$$

As the rank of σ is at most equal to the rank of U and $UV^H N^{-1} = u_1 \pi_k \sigma v^H$ there exist v_r such that $u_1 \pi_k \sigma v^H = U v_r^H$. Finally, $(I-\sigma_k^2)^{1/2} = I - \sigma_\delta$ for some positive diagonal matrix σ_δ of rank at most δ – as σ_k itself has rank at most δ. Hence d_1 is at most a matrix of rank δ different from the unit matrix of dimension k. □

The exact form of the reduction matrices (the 'Q'-matrices in the procedure) is of importance for the complexity of the algorithms that will be derived next, because they have to be propagated to other submatrices than the one they are originally derived from. We shall measure computation complexity by tallying how many times data items are visited. If an $m \times n$ matrix multiplies a vector of dimension n with a direct algorithm, then the complexity by this measure is exactly $m * n$, all data items in the matrix are visited once (the operation tally is higher because each multiplication is followed by an addition except the last). The complexity of the row absorption procedure of the previous theorem is δm^2

where δ is the number of columns in the U or V matrix, because the result can be obtained through a QR factorization on the matrix $\left[\begin{array}{c} V^H \\ \hline D \end{array}\right]$.

4. HSS row absorption procedure

Our next goal is to derive a so-called 'HSS row absorption' matrix – one of the main ingredients in the HSS Moore-Penrose reduction process. Starting point is the form

$$\left[\begin{array}{c|c} w_1 v_\ell^H & w_2 v_r^H \\ \hline d_1 & u_r v_r^H \\ \hline u_\ell v_\ell^H & d_2 \end{array}\right] \tag{4.25}$$

in which we assume the top entries to be 'skinny', i.e., of low rank compared to the dimensions of the matrix (this assumption is not used explicitly but underlies the usefulness of the procedure – as discussed above, the HSS compatibility has to be preserved!). We introduce a somewhat simplified notation, whenever clear from the context we drop the level indication. E.g., in the above matrix, the notation U_r and $U_{1;r}$ or $U_{1;1}$ would be equivalent. The goal of the procedure is to reduce the matrix to a row independent HSS form using orthogonal transformations. Important in the reduction procedure are the properties of the overall orthogonal reducing matrix, namely which block entries in it have an HSS form and which are 'skinny' – i.e., have low rank of the same order as the rank of the top entries to be absorbed. The procedure uses the properties derived in the first section.

Step 1. Find orthogonal q so that

$$\left[\begin{array}{c|c} q_{11} & q_{12} \\ \hline q_{21} & q_{22} \end{array}\right] \left[\begin{array}{c} w_2 \\ \hline u_r \end{array}\right] = \left[\begin{array}{c} r \\ \hline 0 \end{array}\right] \tag{4.26}$$

with r square non-singular and q_{22} close to a unit matrix, and apply the transformation to the original after embedding:

$$\left[\begin{array}{c|c} q & \\ \hline & I \end{array}\right] \left[\begin{array}{c|c} w_1 v_\ell^H & w_2 v_r^H \\ \hline d_1 & u_r v_r^H \\ \hline u_\ell v_\ell^H & d_2 \end{array}\right] = \left[\begin{array}{c|c} q_{11} w_1 v_\ell^H + q_{12} d_1 & r v_r^H \\ \hline q_{21} w_1 v_\ell^H + q_{22} d_1 & 0 \\ \hline u_\ell v_\ell^H & d_2 \end{array}\right]. \tag{4.27}$$

Let

$$d_1' = q_{21} w_1 v_\ell^H + q_{22} d_1 \tag{4.28}$$

the product of a lower level form (d_1) with a 'skinny' perturbation of the unit matrix (q_{22}) followed by a 'skinny' additive perturbation ($q_{21} w_1 v_\ell^H$). The new look of the matrix is, after an exchange of block-rows

$$\left[\begin{array}{c|c} d_1' & 0 \\ \hline q_{11} w_1 v_\ell^H + q_{12} d_1 & r v_r^H \\ \hline u_\ell v_\ell^H & d_2 \end{array}\right] \tag{4.29}$$

in which the product $q_{12}d_1$ is 'skinny', but increases the rank of that term beyond the rank of v_ℓ – as can be expected.

Step 2. We now work on the right part of the matrix. Let p be an orthogonal matrix that reduces

$$\left[\begin{array}{c} r v_r^H \\ \hline d_2 \end{array}\right]. \tag{4.30}$$

Since d_2 is square (this assumption is not really necessary!) the result will have the form

$$\left[\begin{array}{c} 0 \\ \hline d_2' \end{array}\right]. \tag{4.31}$$

This procedure amounts to a lower level HSS row absorption problem (p cannot be claimed to be 'skinnily' away from a unit, it will have whatever structure it inherits from the lower level operation, which is isomorphic to what is happening at this level). Applying p to the left column will produce

$$p \left[\begin{array}{c} q_{11} w_1 v_\ell^H + q_{12} d_1 \\ \hline u_\ell v_\ell^H \end{array}\right] = \left[\begin{array}{c} v'^T \\ \hline u_\ell' v_\ell'^T \end{array}\right]. \tag{4.32}$$

The matrix has now been brought to the form

$$\left[\begin{array}{c|c} d_1' & 0 \\ \hline\hline v'^T & 0 \\ \hline u_\ell' v_\ell'^T & d_2' \end{array}\right]. \tag{4.33}$$

Now, if the original matrix is non-singular (has full column rank), then d_2' will have full row rank (also by construction) and a further, lower level absorption is needed (using a new transformation matrix s) on

$$\left[\begin{array}{c} d_1' \\ \hline v'^T \end{array}\right] \tag{4.34}$$

to yield the end result in the form

$$\left[\begin{array}{c|c} d_1'' & 0 \\ \hline 0 & 0 \\ \hline u_\ell' v_\ell'^T & d_2' \end{array}\right], \tag{4.35}$$

in which both d_1'' and d_2' have full row rank (and in case the original matrix was non-singular, also full column rank and hence will be square). It can be remarked that however one looks at it, the full matrix has to be involved in the procedure, but all operations are either elementary or absorptions at a lower level or involve a very skinny transformation at the higher level (the q matrix).

Collecting the transformations we find for the overall Q (the fat entries are full, lower level matrices, q_{22} is skinnily away from unitary and the non-fat entries are skinny):

$$Q = \left[\begin{array}{cc|c} \mathbf{s}_{11} & s_{12} & \\ s_{21} & s_{22} & \\ \hline & & I \end{array}\right] \left[\begin{array}{c|cc} I & & \\ \hline & p_{11} & p_{12} \\ & p_{21} & \mathbf{p}_{22} \end{array}\right] \left[\begin{array}{cc|c} q_{21} & \mathbf{q}_{22} & \\ q_{11} & q_{12} & \\ \hline & & I \end{array}\right]. \tag{4.36}$$

Working out (should not be done in practice) produces:

$$Q = \left[\begin{array}{c|c|c} \mathbf{s}_{11}q_{21} + s_{12}p_{11}q_{11} & \mathbf{s}_{11}\mathbf{q}_{22} + s_{12}p_{11}q_{12} & s_{12}p_{12} \\ \hline s_{21}q_{21} + s_{22}p_{11}q_{12} & s_{21}\mathbf{q}_{22} + s_{22}p_{11}q_{22} & s_{22}p_{12} \\ \hline p_{21}q_{11} & p_{21}q_{12} & \mathbf{p}_{22} \end{array}\right]. \tag{4.37}$$

The most critical term is the 1,2 where the product of two full matrices occurs: $\mathbf{s}_{11}\mathbf{q}_{22}$. But $\mathbf{q}_{22}$ is only 'skinnily' away from a unit matrix, hence this product also has a 'skinny' algorithm. The final form for the matrix Q produces:

$$\left[\begin{array}{c|c|c} Q_{11} & \mathbf{Q}_{12} & Q_{13} \\ \hline Q_{21} & Q_{22} & Q_{23} \\ \hline Q_{31} & Q_{32} & \mathbf{Q}_{33} \end{array}\right] \left[\begin{array}{c|c} w_1 v_\ell^H & w_2 v_r^H \\ \hline d_1 & u_r v_r^H \\ \hline u_\ell v_\ell^H & d_2 \end{array}\right] = \left[\begin{array}{c|c} d_1'' & 0 \\ \hline 0 & 0 \\ \hline u_\ell' v_\ell'^T & d_2' \end{array}\right] \tag{4.38}$$

in which the not-boldface entries of Q have low rank (skinny products), the boldface ones are of HSS form, and both d_1'' and d_2' have full column rank, and d_1'', d_2' have full row rank.

Complexity calculation

As the absorption procedure turns out to be the main workhorse in the reduction procedure, we proceed to its complexity tally. Let us assume that the vectors to be absorbed are of dimension δ_1, while the off-blocks in the lower HSS representation are of (effective) rank $\delta_2 \geq \delta_1$. Let $\mathbf{M}(n, \delta)$ indicate the computational complexity of multiplying an HSS matrix of dimension n with a block of δ vectors, and $\mathbf{C}(n, \delta)$ the complexity of the absorption procedure of a vector block of dimension δ by a dimension n HSS matrix. The complexity of the HSS absorption procedure can now be tallied. An important observation is that the computation of $q_{11}w_1v_\ell^H$ in this step can be postponed until the bottom level of the procedure is obtained. Let us assume that this bottom level is characterized by matrices of dimensions $N_{\text{bot}} \times N_{\text{bot}}$, then the complexity count would be as follows:

Step 1: $\delta_1 N_{\text{bot}} + 2\mathbf{M}(\frac{n}{2}, \delta_1)$;
Step 2: $\mathbf{C}(\frac{n}{2}, \delta_1) + \mathbf{M}(\frac{n}{2}, 2\delta_1) = \mathbf{C}(\frac{n}{2}, \delta_1) + 2\mathbf{M}(\frac{n}{2}, \delta_1)$;
Step 3: $\mathbf{C}(\frac{n}{2}, 2\delta_1) = 2\mathbf{C}(\frac{n}{2}, \delta_1)$.

Hence the total tally is

$$\delta_1 N_{\text{bot}} + 4\mathbf{M}(\frac{n}{2}, \delta_1) + 3\mathbf{C}(\frac{n}{2}, \delta_1). \tag{4.39}$$

We see that the complexity is not directly dependent on the top level dimension n, and just linearly on the lower level dimensions, where presumably similar computations will take place, to be shown in the next section. It is of course dependent on the local rank (δ_1) but also in a linear fashion. Another important point is in what state the matrix is left behind after the reduction procedure, more precisely whether the HSS relations with the other submatrices in the original are still valid. This point will be taken up in the next section where we consider the overall procedure.

The dual of the absorption will be needed as well in the sequel. At this point there is no substantial difference between the two procedures, one works on the

rows, the other on the columns of the HSS matrix, producing a low complexity calculation that preserves the HSS structure.

5. An HSS Moore-Penrose reduction method

For ease of discussion, we assume that the system is a 'flat' system of full row rank, and furthermore given in the traditional HSS form. In case the assumption does not hold then the later steps in the algorithm will have to be modified, but this would not entail major difficulties. We start out with a matrix in HSS form, i.e., it is of the form

$$\left[\begin{array}{c|c} D_1 & U_u V_u^H \\ \hline U_\ell V_\ell^H & D_2 \end{array}\right] \tag{5.40}$$

in which the low rank matrices U_u, V_u, U_ℓ, V_ℓ are *HSS compatible*, i.e., can be generated from the lower HSS levels, as explained in the introduction. To keep low complexity calculations it will be imperative to preserve the HSS compatibility structure whenever appropriate. The first step is the replacement of the original HSS problem (MPHSS) by an equivalent modified set of equations as follows:

$$\left[\begin{array}{c|c} D_1 & U_u V_u^H \\ \hline U_\ell V_\ell^H & D_2 \end{array}\right] \Rightarrow \left[\begin{array}{c|cc} D_1 - D_1 V_\ell V_\ell^H & D_1 V_\ell & U_u V_u^H \\ \hline 0 & U_\ell & D_2 \end{array}\right]. \tag{5.41}$$

Here we assume that V_ℓ is a 'skinny' orthonormal set of columns and that D_1 and D_2 possibly again have HSS forms of lower hierarchical level. Before discussing the merits of this step, and then the further steps to be executed, we verify the algebraic correctness of this first step. Let V_ℓ' be an orthogonal complementary set of columns of V_ℓ, then postmultiplying the original system with an appropriately dimensioned orthogonal matrix embedding $[V_\ell'\ V_\ell]$ produces the equivalent system (the second member has to be adapted in an obvious way, we skip this step):

$$\left[\begin{array}{c|c} D_1 & U_u V_u^H \\ \hline U_\ell V_\ell^H & D_2 \end{array}\right] \left[\begin{array}{c|c} V_\ell' & V_\ell \\ \hline & I \end{array}\right] = \left[\begin{array}{c|cc} D_1 V_\ell' & D_1 V_\ell & U_u V_u^H \\ \hline 0 & U_\ell & D_2 \end{array}\right] \tag{5.42}$$

$$(D_1 - D_1 V_\ell V_\ell^H) V_\ell' = D_1 V_\ell'. \tag{5.43}$$

The 1,1 block entry $D_1 V_\ell'$ may now be replaced by $D_1(I - V_\ell V_\ell^H)$ (increasing the size of the matrix) because the MP solutions of the two systems are closely related: if $\left[\begin{array}{c} y_1 \\ y_2 \end{array}\right]$ solves the latter, then $\left[\begin{array}{c} V_\ell' y_1 \\ y_2 \end{array}\right]$ will solve the former, due to the fact that $V_\ell'^T V_\ell' = I$. The increase in size will be removed in later 'absorption' steps. In this step, D_1 gets modified by a matrix that is only 'skinnily' away from the unit matrix.

Before proceeding, we study the effects $(I - V_\ell V_\ell^H)$ has on D_1 and whether HSS compatibility is preserved in this step. Hence, we assume that D_1 has in turn a lower level HSS form. To avoid a surfeit of indices, we replace V_ℓ by V and

express the lower level HSS form again with the same notation as before:

$$D_1 \leftarrow \left[\begin{array}{c|c} D_1 & U_r V_r^H \\ \hline U_\ell V_\ell^H & D_2 \end{array}\right]. \tag{5.44}$$

Furthermore, because of the HSS relations, we have

$$V = \left[\begin{array}{c} V_1 \\ \hline V_2 \end{array}\right] = \left[\begin{array}{c} V_\ell W_\ell \\ \hline V_r W_r \end{array}\right] \tag{5.45}$$

(which could also be expressed as

$$V = \left[\begin{array}{c} V_\ell \\ \hline V_r \end{array}\right] \odot \left[\begin{array}{c} W_\ell \\ \hline W_r \end{array}\right] \tag{5.46}$$

in which the W's are assumed tall matrices, and the $\odot$ indicates pointwise multiplication of block entries – assuming matching dimensions of course). We find as result for this operation

$$\left[\begin{array}{c|c} D_1 - (D_1 V_1 + U_r V_r^H V_2) V_1^H & U_r V_r^H - (D_1 V_1 + U_r V_r^H V_2) V_2^H \\ \hline U_\ell V_\ell^H - (U_\ell V_\ell^H V_1 + D_2 V_2) V_1^H & D_2 - (U_\ell V_\ell^H V_1 + D_2 V_2) V_2^H \end{array}\right]. \tag{5.47}$$

We see that the ranks of the off-diagonal terms have not increased, but the column basis has changed, and it may be useful to keep the original column vectors although they may not amount to a basis (they are able to generate the column vectors though). Taking this conservative approach we may write the result so far as

$$\left[\begin{array}{c|c} D_1' & [U_{rn}\ U_r] W_{rn}^H V_r^H \\ \hline [U_{\ell n}\ U_\ell] W_{\ell n}^H V_\ell^H & D_2' \end{array}\right] \tag{5.48}$$

in which

$$\left\{\begin{array}{lcl} U_{rn} & = & D_1 V_1 \\ D_1' & = & D_1 - (U_{rn} + U_r V_r^H V_2) V_1^H \\ W_{rn}^H & = & \left[\begin{array}{c} -W_r^H \\ \hline I - V_r^H V_2 W_r^H \end{array}\right] \\ D_2' & = & D_2 - (U_{\ell n} + U_\ell V_\ell^H V_1) V_2^H \\ U_{\ell n} & = & D_2 V_2 \\ W_{\ell n}^H & = & \left[\begin{array}{c} -W_\ell^H \\ \hline I - V_\ell^H V_1 W_\ell^H \end{array}\right] \end{array}\right. \tag{5.49}$$

all quantities that can now be computed at the lowest level, and are either 'skinny' or 'skinny' updates.

Going back to the original procedure, the same strategy can now be applied to the rightmost block in eq. 5.42, this produces:

$$\begin{aligned} &\left[\begin{array}{c|c} U_{1;r} V_{1;r,3}^H & U_{1,r} V_{1;r,4}^H \\ \hline D_{2;3} & U_{2;r} V_{2;r}^H \\ \hline U_{2;\ell} V_{2;\ell}^H & D_{2;4} \end{array}\right] \left[\begin{array}{c|c|c} V_{2;\ell}' & V_{2;\ell} & 0 \\ \hline 0 & 0 & I \end{array}\right] \\ &= \left[\begin{array}{c|c|c} U_{1;r} V_{1;r,3}^H V_{2;\ell}' & U_{1;r} V_{1;r,3}^H V_{2;\ell} & U_{1;r} V_{1;r,4}^H \\ \hline D_{2;3}(I - V_{2;\ell} V_{2;\ell}^H) V_{2;\ell}' & D_{2;3} V_{2;\ell} & U_{2;r} V_{2;r}^H \\ \hline 0 & U_{2;\ell} & D_{2;4} \end{array}\right]. \end{aligned} \tag{5.50}$$

Again, $V'_{2,\ell}$ can be taken out to the second member. Notice also that the size of the system has increased slightly – this is a convenience which keeps the entries in the main block diagonal square, an alternative strategy will be briefly discussed later. We have now reached the 'bottom' as far as this latter part of the matrix is concerned. We can now start eliminating the newly created spurious entries which are all 'skinny' and share column bases. At the bottom we find

$$\left[\begin{array}{c|c\|\|c}\ U_{\ell,4}\ & U_{2;\ell} & D_{2;4}\ \end{array}\right]. \tag{5.51}$$

This can now be reduced by a direct column absorption procedure, dual of the procedure presented in theorem 2. Since we have assumed row-independence, the theorem is directly applicable. This step will reduce the matrix to the form

$$\left[\begin{array}{c\|\|c}\ 0\ & D'_{2;4}\ \end{array}\right] \tag{5.52}$$

in which $D'_{2;4}$ has become square non-singular, and the row basis of the submatrix on top have been modified (without modification of the column basis), and this by 'skinny' calculations. Notice that only submatrices belonging to low-rank off-diagonal blocks are affected by this step, in hierarchical order, as no central matrix resides on top. The adjustment computation can also be restricted to the lowest level in the hierarchy as all the higher levels will follow suit. Since this is an important ingredient in preserving low computational complexity, we make this step more explicit. The full submatrix to be reduced has the following form:

$$\left[\begin{array}{c\|\|c|c}
D_1V_\ell & U_{1;r}V^H_{1;r,3} & U_{1;r}V^H_{1;r,3}V_{2;\ell} & U_{1;r}V^H_{1;r,4} \\
\hline\hline
U_{\ell,3} & D_{2;3}(I - V_{2;\ell}V^H_{2;\ell}) & D_{2;3}V_{2;\ell} & U_{2;r}V^H_{2;r} \\
\hline
U_{\ell,4} & 0 & U_{2;\ell} & D_{2;4}
\end{array}\right]. \tag{5.53}$$

The bottom block rows in this expression have the necessary HSS compatibility relationship to allow for HSS column absorption. D_1V_ℓ involves new data resulting from eliminating the bottom block. The computation of this term is unavoidable as this is the only place in the matrix where the data survives, but it has already been executed as part of the procedure to computer $D_1 - D_1V_\ell V_\ell^H$ described earlier in this section. The entries in the first block-column all have the same reduced row vector $[I_{\delta_1}]$. The same is true one level down (where a similar procedure has been performed), $D_{2;3}V_{2;\ell}$ is what remains from that elimination procedure. The overall row basis matrix that will be affected by the present HSS column absorption procedure has now the form

$$\left[\begin{array}{c\|\|c|c}
I_{\delta_1} & & \\
\hline
 & I_{\delta_2} & \\
\hline
 & & V^H_{2;r}
\end{array}\right] \tag{5.54}$$

where it may be remarked that the columns on top of the zero entries will be unaffected by the absorption procedure. The procedure scrambles this vector to

produce a new row basis, and after reordering of columns produces the right submatrix

$$\left[\begin{array}{c|c|c} U_{1;r}V_{1;r,3}^H & [D_1V_\ell \;\; U_{1;r}V_{1;r,3}^HV_{2;\ell}]V_{2;r.4}^{\prime,H} & U_{1;r}V_{1;r,5}^{\prime H} \\ \hline D_{2;3}(I-V_{2;\ell}V_{2;\ell}^H) & [U_{\ell.3} \;\; D_{2;3}V_{2;\ell}]V_{2;r,4}^{\prime H} & U_{2;r}V_{2;r,5}^{\prime H} \\ \hline 0 & 0 & D'_{2;4} \end{array}\right]. \tag{5.55}$$

It should be clear that the HSS relations still hold (although the rank has necessarily doubled) both for the row bases and for the column bases. This procedure can now be repeated one step further (concerning the rows of hierarchical index 2;3), involving an absorption of $[U_{\ell 3} \;\; D_{2;3}V_{2;\ell}]V_{2;r,4}^{\prime H}$ into $D_{2;3}(I-V_{2;\ell}V_{2;\ell}^H)$, a procedure that is again of the HSS absorption type and can be executed at an even lower level (level 3). The result so far yields the following general form (redefining entries and dropping primes for better visualization)

$$\left[\begin{array}{c|c|c|c|c} D_{2;1} & U_{2;r,1}V_{2;r,2}^H & U_{1;r,1}V_{2;r,3}^H & U_{1;r,1}V_{1;r,4}^H & U_{1;r,1}V_{1;r,5}^H \\ \hline U_{2;\ell,2}V_{2;\ell,1}^H & D_{2;2} & U_{1;r,2}V_{2;r,3}^H & U_{1;r,2}V_{1;r,4}^H & U_{1;r,2}V_{1;r,5}^H \\ \hline 0 & 0 & 0 & D_{2;4} & U_{2;r,3}V_{2;r,5}^H \\ \hline 0 & 0 & 0 & 0 & D_{2;5} \end{array}\right] \tag{5.56}$$

in which all HSS relations hold for as far as applicable. To proceed to the levels with indices 2;1 and 2;2, we must revert to the 1;1 matrix computed earlier (which was the original $D_1(I-V_\ell V_\ell^H)$, and which, as we have shown, still has the HSS form, as shown schematically in the previous equation). This matrix must now undergo the same second level treatment as was done on the original D_2 block, with similar result. The procedure entails the HSS absorption of the third block column in the previous equation into the first two blocks, As this procedure does not affect the fourth and fifth block column, the end result will have the form (again dropping primes)

$$\left[\begin{array}{c|c|c|c|c} D_{2;1} & U_{2;r,1}V_{2;r,2}^H & 0 & U_{1;r,1}V_{1;r,4}^H & U_{1;r,1}V_{1;r,5}^H \\ \hline 0 & D_{2;2} & 0 & U_{1;r,2}V_{1;r,4}^H & U_{1;r,2}V_{1;r,5}^H \\ \hline 0 & 0 & 0 & D_{2;4} & U_{2;r,3}V_{2;r,5}^H \\ \hline 0 & 0 & 0 & 0 & D_{2;5} \end{array}\right] \tag{5.57}$$

where, again, HSS relations have been preserved wherever applicable. If the original system was indeed row independent, then the diagonals blocks in this expression will be square invertible, and a QR factorization of the original system of equations has been achieved, which can simply be solved by an HSS-type back substitution that can be efficiently executed (i.e., by visiting all algebraically relevant data only once, using the HSS relationships). In case the original matrix is not row independent, then a further row reduction can be done on the matrix obtained so far, using a simplified form of the algorithm derived so far.

The dual form of the method presented so far merits some attention. Suppose that our original system is not row- but column independent (i.e., it is a tall system). Is a reduction procedure as used in the previous algorithm that preserves the Moore-Penrose property possible? We show that this is indeed the case.

Proposition 3. *Let*

$$\left[\begin{array}{c|c} D_1 & U_u V_u^H \\ \hline U_\ell V_\ell^H & D_2 \end{array}\right] \tag{5.58}$$

be given as the top-level of a column independent HSS representation in which the matrices $D_i, i = 1, 2$ are HSS compatible, and in which U_u is isometric. Then an equivalent Moore-Penrose system is given by

$$\left[\begin{array}{c|c} (I - U_u U_u^H) D_1 & 0 \\ \hline U_u^H D_1 & V_u^H \\ \hline U_\ell V_\ell^H & D_2 \end{array}\right]. \tag{5.59}$$

Proof. The proof follows from the following two observations:

1. Let U_u' be an isometric matrix whose columns span the orthogonal complement of the (column) range of U_u. Then
$$U_u'^H D_1 = U_u'^H (I - U_u U_u^H) D_1; \tag{5.60}$$
2. if V^H is an isometric matrix then the system $V^H A x = b$ is equivalent to the system $Ay = Vb$ in the sense that if y solves the Moore-Penrose problem of the latter, then x=y is the Moore-Penrose solution of the former.

The required transformation now follows from the identity

$$\left[\begin{array}{c|c} U_u'^H & \\ \hline U_u^H & \\ \hline & I \end{array}\right] \left[\begin{array}{c|c} D_1 & U_u V_u^H \\ \hline U_\ell V_\ell^H & D_2 \end{array}\right] = \left[\begin{array}{c|c} U_u'^H D_1 & 0 \\ \hline U_u^H D_1 & V_u^H \\ \hline U_\ell V_\ell^H & D_2 \end{array}\right]. \tag{5.61}$$

An application of the properties mentioned above allows one to eliminate $U_u'^H$ so that the modification of the entries only involves the product $U_u^H D_1$, which can efficiently be relegated to the lower hierarchical level. □

Notice that just as before the HSS relations remain in place. We can summarize the procedures in the following conceptual theorem. In the formulation 'URV-type' means: using orthogonal transformation on the rows (for the U-factor) and the columns (for the V factor).

Theorem 1. *A matrix represented in HSS form can be constructively reduced to a non-singular block upper matrix in HSS form using efficient, HSS compatible transformations of the URV type.*

6. Discussion and conclusions

A number of observations seem relevant:

1. the HSS form of an invertible matrix is stable under inversion, it is not too hard to prove that the inverse of such a matrix has an HSS representation of the same complexity as the original. However, an efficient algorithm to compute it has not been presented yet to our knowledge. The present paper goes some way towards this end, it shows at least that the URV form can be

computed efficiently. The result even extends to non-invertible matrices and their Moore-Penrose inverse;
2. the theoretical results presented in this paper have not been tested numerically. This will be done in the near future;
3. as mentioned in the introduction, the HSS form is not as generally available as one would wish. 1D and 2D scattering problems can be brought into the form through the use of multipole theory. For 3D problems, the reduction is far from evident. Here also we are still lacking a good reduction theory. The same can be said about finite element systems: for 1D and 2D cases there is a forthcoming reduction possible as is indicated in the thesis of T. Pals [11];
4. another interesting but as yet unsolved problem is the determination of (close to optimal) preconditioners in HSS form; rank reductions in the HSS form would amount to multi-resolution approximations;
5. the representation for HSS forms discussed in the introduction amounts to a 'system theory on a tree' much as is the case for the multi-resolution theory of Alpay and Volok [1]. However, the representations are fundamentally different: in our case the system theory represents *computational states*, i.e., intermediate data as they are stored in a hierarchical computation, while in the Alpay-Volok case, the states parallel the multi-resolution, a higher up state consists of a summary or average of lower lying states. Although the states have therefore very different semantics, from a system theoretical point of view they are states indexed by a tree and hence operators acting on these states will have similar effects. In particular, the shift operator and its adjoint as well as the level exchange operator presented in the introductory section have the same meaning in both theories (the definition of these operators differ somewhat due to different normalizations). A major difference, however, is that in our case the state space structure is not uniform, while in the Alpay-Volok it is, so that the system in their case can be advantageously reduced to a sequentially semi-separable form.

Acknowledgements

The authors wish to thank a number of colleagues for discussions and interactions with them on the general theme of HSS systems, to wit Alle-Jan van der Veen, Ming Gu, Timoghy Pals, W. Lyons and Zhifeng Sheng.

References

[1] D. Alpay and D. Volok. Point evaluation and hardy space on a homogeneous tree. *Integral Equations and Operator Theory*, vol 53 (2005), pp. 1–22.

[2] S. Chandrasekaran, P. Dewilde, M. Gu, W. Lyons, and T. Pals. A fast solver for hss representations via sparse matrices. In *Technical Report*. Delft University of Technology, August 2005.

[3] S. Chandrasekaran, M. Gu, and T. Pals. A fast and stable solver for smooth recursively semi-separable systems. In *SIAM Annual Conference, San Diego and SIAM Conference of Linear Algebra in Controls, Signals and Systems, Boston*, 2001.

[4] S. Chandrasekaran, M. Gu, and T. Pals. Fast and stable algorithms for hierarchically semi-separable representations. In *Technical Report*. University of California at Santa Barbara, April 2004.

[5] P. Dewilde, K. Diepold, and W. Bamberger. A semi-separable approach to a tridiagonal hierarchy of matrices with application to image flow analysis. In *Proceedings MTNS*, 2004.

[6] P. Dewilde and A.-J. van der Veen. *Time-varying Systems and Computations*. Kluwer, 1998.

[7] P. Dewilde and A.-J. van der Veen. Inner-outer factorization and the inversion of locally finite systems of equations. *Linear Algebra and its Applications*, page *to appear*, 2000.

[8] Y. Eidelman and I. Gohberg. On a new class of structured matrices. *Notes distributed at the 1999 AMS-IMS-SIAM Summer Research Conference*, Structured Matrices in Operator Theory, Numerical Analysis, Control, Signal and Image Processing, 1999.

[9] I. Gohberg, T. Kailath, and I. Koltracht. Linear complexity algorithms for semiseparable matrices. *Integral Equations and Operator Theory*, 8:780–804, 1985.

[10] L. Greengard and V. Rokhlin. A fast algorithm for particle simulations. *J. Comp. Phys.*, 73:325–348, 1987.

[11] T. Pals. *Multipole for Scattering Computations: Spectral Discretization, Stabilization, Fast Solvers*. PhD thesis, Department of Electrical and Computer Engineering, University of California, Santa Barbara, 2004.

[12] V. Rokhlin. Applications of volume integrals to the solution of pde's. *J. Comp. Phys.*, 86:414–439, 1990.

[13] A.J. van der Veen. Time-varying lossless systems and the inversion of large structured matrices. *Archiv f. Elektronik u. Übertragungstechnik*, 49(5/6):372–382, September 1995.

Patrick Dewilde
Department of Electrical Engineering
Delft University of Technology
Delft, The Netherlands
e-mail: `p.dewilde@ewi.tudelft.nl`

Shivkumar Chandrasekaran
Department of Electrical and Computer Engineering
University of California
Santa Barbara, CA 93106-9560, USA
e-mail: `shiv@ece.ucsb.edu`