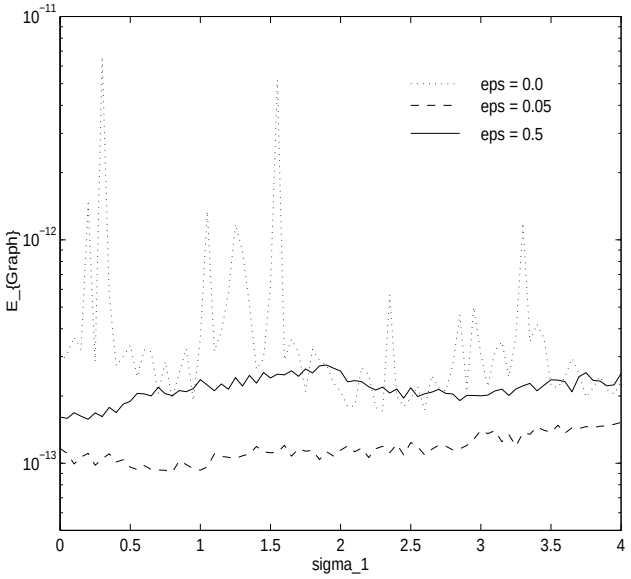

Figure 2: Distance of the given subspace graph (S) and the modelled subspace span $(T\,X_0)$. The distance is measured as the gap metric between the two subspaces.

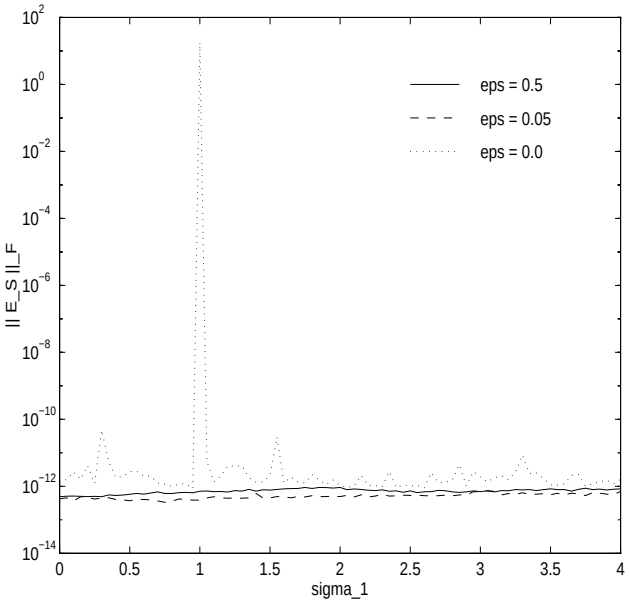

Figure 3: Frobenius norm of the error matrix $S - \hat{S}$ between the originally given matrix S and the modelled matrix $\hat{S}$ as a function of the singular value σ_1 for three values for the parameter ϵ.

5. CONCLUSION

We have demonstrated that the numerical quality of certain matrix factorizations based on noncompact groups can be significantly improved when studying Lie groups operating on graphs of matrices (Grassmannian approch) instead of applying conventional row and column transformations on the matrices. This can be achieved exploiting the inherent degree of freedom to perform *locally* a change of basis during the algorithm if numerical problems are detected. Experimental results on the noncompact Lie group $\mathbf{O}_{n,n}$ so far indicate that based on this approach one can impose an priori upper bound on the norm of elementary transfomation matrices (e.g. plane reflections). The problems discussed here partly are characteristic for lossless or $\mathbf{O}_{n,n}$ transformations due to the mirror image property of the eigenvalues of these "rigid motions" with respect to the unit circle. Nevertheless, there is no doubt that one may achieve improved numerical results even in more general cases. Clearly, the preliminary implementation of the investigated modifications leads to an increase in computational complexity. No algorithmic optimization in any way has been performed so far. We simply wanted to show that (i) there are new perspectives for algorithm design, especially when dealing with noncompact groups, (ii) there is much work to be done in order to fully exploit the potential gains and insights of this approach.

6. REFERENCES

[1] J. Della-Dora. Numerical linear algorithms and group theory. *Lin. Alg. Appl.*, 10:267–283, 1975.

[2] K. Diepold and R. Pauli. Schur parametrization of symmetric indefinite matrices based on a network theoretic model. *Int. J. Electron. Commun. (AEÜ)*, 45:375–385, 1991. (cf. also Proc. ICASSP'91, 3401-3404).

[3] A. J. van der Veen. A Schur method for low-rank matrix approximation. *SIAM J. Matrix Anal. Appl.*, 17:139–160, 1996.

[4] K. Diepold and R. Pauli. A recursive algorithm for lossless embedding of non-passive systems. In M. Verhaegen P. Dewilde, M.A. Kaashoek, editor, *Challenges of a Generalized System Theory*, pages 209–222. Royal Netherlands Academy of Sciences, North-Holland, Amsterdam, 1993.

[5] K. Diepold and R. Pauli. A Schur-type algorithm for the triangular factorization of positive semidefinite matrices . In E. Deprettere and A.-J. Van der Veen, editors, *Algorithms and Parallel VLSI Architectures.*, volume B Proceedings, pages 33–42. Elsevier, 1991.

[6] E. N. Atkinson. Computing $A^T A - B^T B = L^T D L$ using generalized hyperbolic transformations. *Lin. Algebra Appl.*, 194:135–147, 1993.

is made to determine the parameter for a hyperbolic plane reflection.

In comparison to the procedure proposed in [4] the following algorithmic changes have been investigated to achieve the goal stated in the previous section:

1. When at step i of the recursion the parameters for a hyperbolic plane reflection shall be determined, the value of the parameter ρ is checked to satisfy $|\rho| < 1 - \epsilon$ for a given $0 \leq \epsilon < 1$.
2. If the constraint (6) on $|\rho|$ is not met, try to find appropriate orthogonal matrices $T_Q \in \mathbf{O}_n \times \mathbf{O}_n$ and $Q_M \in \mathbf{O}_n$ to perform the transformation

$$\begin{bmatrix} \tilde{A}_i \\ \tilde{B}_i \end{bmatrix} = T_Q \begin{bmatrix} A_i \\ B_i \end{bmatrix} Q_M \tag{7}$$

   such that the constraint (6) is satisfied, and the $n \times n$–matrices $A_i, B_i, \tilde{A}_i, \tilde{B}_i$ are all upper triangular shaped.
3. If appropriate orthogonal matrices cannot be found, a totally isotropic subspace has been detected.
4. The matrices T_Q and Q_M are produced as the product of elementary euclidean plane reflections exclusively.

4. NUMERICAL EXPERIMENTS

For this particular case study an implementation of the generalized Schur algorithm as given in [2] and [4] has been used. The implementation has been enhanced to include the algorithmic changes described in the previous section. Numerical experiments were performed to get a first grip at the conjectured numerical effects of the modifications without paying attention to the computational efficiency of the method.

To perform the numerical experiments, families of matrices $S \in \mathsf{R}^{6\times 6}$ were generated as

$$S = U \cdot \mathrm{diag}(\sigma_i) \cdot V^T \in \mathsf{R}^{6\times 6},$$

with orthogonal matrices U and V. The singular values $\sigma_i, \quad i = 2, 3, \ldots 6$ were randomly drawn from a Rayleigh distribution. One of the singular values was varied as $0.0 \leq \sigma_1 \leq 4.0$ and used a parameter.

Starting out with such matrices and the formulation as in eq. (1) the generalized Schur algorithm including the mentioned modifications was used to determine a solution to eq.(3), i.e. the matrices T, X_0 and M according to eq. (3) were computed.

To study the effects of the algorithmic modification on the numerical performance we reconstruct the originally given matrix S as

$$\hat{X} := \begin{bmatrix} 1_6 \\ \hat{S} \end{bmatrix} = T X_0 M,$$

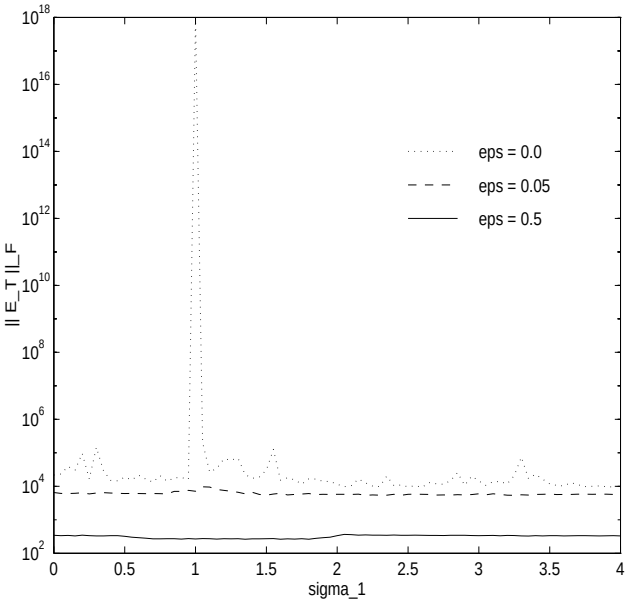

Figure 1: Deviation of T from orthogonality as a function of the singular value σ_1 for three values for the parameter ϵ.

and measure the subspace error as

$$E_{Graph} := \mathrm{gap}(\mathrm{span}\,(X),\ \mathrm{span}\,(\hat{X})).$$

The error on the individual matrix elements is measured by means of

$$||E_S||_F = ||S - \hat{S}||_F.$$

The impact of the modifications on the condition number and the operator norm of the matrix T are analysed by calculating

$$||E_T||_F = ||1_{2n} - T^T T||_F.$$

One set of typical results is displayed in Figs. 1–3. All results are computed as the average over 100 idependent trials for 6×6 matices S. In Fig. 1 it can be seen that deviation of T from orthogonality, and hence the 2–norm and the condition number of T can be reduced by several orders of magnitude, when numerical problems with hyperbolic transformations occur. Furthermore, a numerical improvement can be observed even for more generic cases by means of a more stringent bound on $|\rho|$.

Fig. 2 depicts the corresponding results for the numerical quality of the subspace. Here also, the positive effect of bounding ρ can be seen, while the subspace measure does not exhibit such disastrous numerical behaviour after all. The results are confirmed by the numbers depicted in Fig. 3, where the Frobenius norm has been used to measure the error between S and $\hat{S}$.

2. NUMERICAL PROBLEMS

We base our discussion on the generalized Schur algorithm, which is described e.g. in [2] – [3]. The recursive computation of solutions for (3) corresponds to decompositions of $T = \prod T_i$ into a product of transformations T_i that are "elementary" in some sense. For brevity, we focus on decompositions of $\mathbf{O}_{n,n}$ into the standard euclidean (Givens) and hyperbolic plane transformations. Numerical problems arise from applications of hyperbolic transformations like

$$T_h = \frac{1}{\sqrt{1-\rho^2}} \begin{bmatrix} 1 & -\rho \\ \rho & -1 \end{bmatrix}, \quad \rho^2 < 1\,, \tag{4}$$

and we shall briefly review the reason for these problems.

The eigenvalues of $T_h^T T_h$ can easily be computed as

$$\lambda_1 = \frac{1-\rho}{1+\rho}, \quad \lambda_2 = \frac{1}{\lambda_1} = \frac{1+\rho}{1-\rho},$$

which show that theoretically $\det T_h^T T_h = 1$ even though $\lambda_1 \to 0, \quad \lambda_2 \to \infty$ when $|\rho| \to 1$. Dramatic cancellation errors ($\rho^2 \approx 1$) or breakdown of algorithms ($\rho^2 = 1$) have frequently been reported in these situations.

The case $\rho^2 = 1$ results in a breakdown of the algorithm and is due to isotropic vectors in the basis X. It occurs in exactly two cases [4], [5]:

1. An isotropic vector in X corresponds to a totally isotropic subspace of span (X), (i.e. $n_0 \neq 0$). If this situation is not detected correctly, the algorithm tries to assign span (X_0) to another $\mathbf{O}_{n,n}$-orbit than span (X). Since the J-signature (n_+, n_-, n_0) is an invariant, the algorithm is misled to switch to another J-signature by proceeding along geodesics on the group manifold towards ∞.
2. An isotropic vector in X corresponds to a vanishing principle minor of the J-gramian that is due to $n_0 \neq 0$ or due to ill-conditioned data – a problem that may be easily circumvented by a local change of basis (see below). One should note that not a few numerical problems of this type arise from the fact that transformations acting along preassigned coordinate (hyper)planes may run inevitably into ill-conditioned data constellations.

The case $|\rho| \approx 1$ produces serious numerical difficulties due to the noncompactness of $\mathbf{O}_{n,n}$, since in this case $\|T\| \to \infty$. The normalization constant $\sqrt{1-\rho^2}$ is usually considered responsible for this undesirable numerical behaviour due to the unbounded growth of matrix entries. This fact variously led to conjectures about the superiority of alternative algorithmic schemes which minimize the number of used hyperbolic plane transformations [2] (4). The failure of such an approach can be easily seen by looking at the elementary hyperbolic plane reflection when dropping the normalization constant, i.e. considering the transformation

$$\tilde{T}_h = \begin{bmatrix} 1 & -\rho \\ \rho & -1 \end{bmatrix}, \quad |\rho|^2 < 1\,. \tag{5}$$

In this case the, eigenvalues of $\tilde{T}_h^T \tilde{T}_h$ are

$$\lambda_1 = 1 + 3\rho^2, \quad \lambda_2 = 1 - \rho^2.$$

Therefore, once $|\rho| \to 1$ the transformation matrix tends to become singular, even though the matrix elements are not blown up due to the small normalization constant, and the cancellation argument ceases to hold. In general, one singular, or nearly singular transformation in a product decomposition is enough to produce a singular or nearly singular result.

3. MODIFICATION OF ALGORITHM

It is the goal of this investigation to assess the feasibility of bounding $|\rho|$ away from unity, that is, either to avoid the extreme values $|\rho| \approx 1$, or to somehow achieve that ρ satisfies an a priori bound, given in terms of the parameter ϵ, i.e.

$$|\rho| < 1 - \epsilon, \quad 0 \leq \epsilon < 1, \tag{6}$$

while still computing recursively solutions to (3). Such a result may render possible the development of modified algorithms which will produce solutions with a bounded norm $\|T\|$ and improved numerical accurarcy.

The basic idea is as follows: The Grassmannian approach offers the possibility to perform an *a posteriori* improvement of ill–conditioned data in $X \in \mathbf{St}_{n,2n}$ at any stage "i" of the algorithm with actual data X_i. This can be done by use of group actions on $\mathbf{St}_{n,2n}$ that lie in the kernel of the projection $\pi : \mathbf{St}_{n,2n} \to \mathbf{Gr}_{n,2n},\ X_i \mapsto \pi(X_i) = \text{span}\,(X_i)$. Here, we will limit ourselves solely to a change of basis for span (X_i) by postmultiplication of X_i by some $M \in \mathbf{GL}_n$. This comprises known pivoting strategies when M is a permutation matrix [6], [3] as well as a (algorithmically prohibitive) complete orthogonalization of the basis X_i.

In [4] it has been discussed how to detect singular situations, and how these singularirites can be circumvented by means of a change of basis. In the course of a recursive, Schur–type elimination scheme, the occurence of such singular situations is indicated by $|\rho| = 1$, whenever an attempt

[2] Clearly, this number must not exceed the Witt-index of the metric space (the dimension of the maximally totally isotropic subspace). However, it can definitely not be shown that minimization of this number generally leads to better numerical behaviour. Questions of this type amount to the investigation of the sensitivity of $\mathbf{O}_{n,n}$-orbits with respect to different parametrizations of $T = \prod T_i$ under finite precision arithmetic. This problem highly depends on the concrete data X in (3) and answers can be seriously expected only when additional information on the data is given.

ACTIONS OF NONCOMPACT GROUPS AND ALGORITHM DESIGN: A CASE STUDY

Klaus Diepold [1] *and Rainer Pauli* [2]

[1] International Digital Technologies, Munich, Germany.
[2] Technical University of Munich, Germany.

ABSTRACT

Numerical matrix computations involving actions of noncompact transformation groups are known to produce numerical problems since the elements of the pertaining matrix representations are inherently unbounded. In this case study we analyse numerical problems occuring in a class of algorithms that is based on actions of the pseudo-orthogonal group $\mathbf{O}_{n,m}$ – a group that is noncompact (hyperbolic geometry) and well established in signal processing (Schur methods). As a major result, it is shown how to exploit the additional degrees of freedom in defining coordinate frames in a Grassmannian setting in order to impose an a priori bound on the norm of the transformation matrices. This way, numerically disastrous situations can be circumvented systematically. Hence, it becomes possible to develop modified algorithms which exhibit superior numerical performance for a large class of problems based on e.g. hyperbolic transformations.

1. MOTIVATION AND BACKGROUND

In contrast to the prevailing use of orthogonal transformations in matrix computations, signal processing algorithms related to identification, approximation or decomposition of physical systems involve actions of noncompact groups. The (real) pseudo-orthogonal group $\mathbf{O}_{n,m}$ preserves energy flow of n-input m-output systems, where

$$\mathbf{O}_{n,m} = \{T \in \mathbf{GL}_{n+m} \,:\, T^T J T = J\}$$

$$J = \begin{bmatrix} 1_n & 0_{n\times m} \\ 0_{m\times n} & -1_m \end{bmatrix}.$$

$\mathbf{O}_{n,m}$ is noncompact (hyperbolic geometry) and has $\mathbf{O}_n \times \mathbf{O}_m$ as its maximal compact subgroup, where $\mathbf{O}_n \times \mathbf{O}_m$ denotes the set of block-diagonal matrices of the form $U \oplus V$, where $U \in \mathbf{O}_n\,, V \in \mathbf{O}_m$.

Conventional matrix factorization algorithms split a given matrix $A \in \mathsf{R}^{n\times n}$ into a product $A = GR$, with G and R belonging to two different subgroups of the general linear group $\mathbf{GL}_n$. Such factorizations with A belonging e.g. to $\mathbf{O}_{n,m}$ have been investigated by Della-Dora [1]. A more comprehensive type of factorizations may be motivated by a network theoretic approach to the design of matrix algorithms (cf.[2]). Instead of a group operating directly on a matrix A, one studies natural group actions on graph subspaces of linear input-output maps, i.e. a matrix $A \in \mathsf{R}^{n\times n}$ is identified with the linear space

$$\begin{aligned} \text{graph}\,(A) &= \{(x, Ax) \,:\, x \in \mathsf{R}^n\} \\ &= \text{span}\,(\begin{bmatrix} 1_n \\ A \end{bmatrix}) \in \mathbf{Gr}_{n,2n}\,, \end{aligned} \tag{1}$$

and hence with a point on a Grassmannian manifold $\mathbf{Gr}_{n,2n}$ of n-dimensional subspaces of $2n$-space[1].

In order to perform computations it is useful to represent these subspaces by equivalence classes of $2n \times n$ basis matrices $X \in \mathbf{St}_{n,2n}$ (the Stiefel manifold of full rank matrices from $\mathsf{R}^{2n\times n}$), where any two bases X and Y are equivalent whenever

$$\begin{aligned} \text{span}\,(X) = \text{span}\,(Y) \iff & \quad Y = XM, \\ X, Y \in \mathbf{St}_{n,2n}\,, & \quad M \in \mathbf{GL}_n\,. \end{aligned} \tag{2}$$

Note that a change of basis induced by $M \in \mathbf{GL}_n$ provides for free parameters in the representation of the matrix $A \Leftrightarrow \text{graph}\,(A)$. The Grassmannian approach leads to a very natural theoretical framework for factorizations such as

$$\begin{aligned} X = \begin{bmatrix} 1_n \\ A \end{bmatrix} &= T X_0 M, \\ X, X_0 \in \mathbf{St}_{n,2n}\,, \quad T \in \mathbf{O}_{n,n}, & \quad M \in \mathbf{GL}_n \end{aligned} \tag{3}$$

Here, depending on the application at hand, M may achieve upper (block-)triangular form and X_0 may be some 'nice' canonical form of X (preferably two stacked $n \times n$ diagonal matrices).

$\mathbf{O}_{n,n}$ does not act transitively on $\mathbf{Gr}_{n,2n}$. More precisely, it exhibits $\binom{n+2}{2}$ orbits that may be labeled by the J-signature (n_+, n_-, n_0) of the space span (X), where n_+, n_- and n_0 $(n_+ + n_- + n_0 = n)$ are the inertia of the J-gramian $X^{\mathrm{T}} J X$. Since span (X) and span (X_0) according to (3) belong to the same $\mathbf{O}_{n,n}$-orbit, the J-gramians of X and X_0 must belong to the same congruence class.

[1] One should note that the "extension" of this framework to non-square matrices is trivial due to the Grassmannian approach.