COMPUTATIONAL EXPERIENCE WITH STRUCTURE-PRESERVING HAMILTONIAN SOLVERS IN OPTIMAL CONTROL

Vasile Sima
National Institute for Research & Development in Informatics, 8-10 Bd. Mareşal Averescu, Bucharest, Romania

Keywords: Linear-quadratic optimization, Optimal control, Structure-preserving algorithms.

Abstract: Structure-preserving techniques for solving essential computational problems in optimal control are presented. The techniques use possibly extended skew-Hamiltonian/Hamiltonian matrix pencils, and specialized algorithms to exploit their structure: the symplectic URV decomposition, periodic QZ algorithm, solution of periodic Sylvester-like equations, etc. The structure-preserving approach has the potential to avoid the numerical difficulties which are encountered for a traditional, non-structured solution, returned by the currently available software tools. Preliminary computational results are presented.

1 INTRODUCTION

Several basic computational problems in optimal and robust systems analysis and design involve structured, e.g., Hamiltonian and symplectic, matrix pencils. Two important problems, with many applications, are discussed below. One such basic computation is the evaluation of the L_∞- and H_∞-norms, which are used, e.g., to quantify the trade-off between performance and robust stability. Quadratically convergent algorithms (Boyd et al., 1989; Bruinsma and Steinbuch, 1990) for the computation of the these norms use the purely imaginary eigenvalues of a matrix or matrix pencil at each iteration. This matrix (pencil) is structured, Hamiltonian or symplectic, in the continuous- and discrete-time case, respectively. (Actually, the pencils arising in the continuous-time descriptor case are skew-Hamiltonian/Hamiltonian.) Some details are given in (Sima, 2006) (and the references therein), where the Hamiltonian structure is exploited in the matrix case. The state-of-the-art function `norm` in the MATLAB® Control System Toolbox computes the eigenvalues using the standard eigensolver `eig`, which does not take the structure into account. But the detection of purely imaginary eigenvalues is a delicate numerical problem if a non-structured algorithm is used. Several simple examples are given in Section 3.

Another fundamental computation in control systems design is the solution of continuous-time and discrete-time algebraic Riccati equations (CAREs and DAREs). CAREs and DAREs arise in many applications, such as, stabilization and linear-quadratic regulator problems, Kalman filtering, LQG—linear-quadratic Gaussian (H_2-) optimal control problems, computation of (sub)optimal H_∞ controllers, etc. In applications, usually the *stabilizing solution* is required, which can be used to stabilize the closed-loop system matrix or matrix pencil. A very important class of CARE/DARE solvers makes use of stable invariant or deflating subspaces of some matrices or pencils, assuming certain nonsingularity and eigenvalue dichotomy assumptions (Laub, 1979; Pappas et al., 1980). The associated CARE/DARE solvers used matrix inversions (for instance, of the control weighting matrix, or of the system matrix, for DAREs), but this can sometimes ruin the accuracy of the results. Better results are obtained using stable deflating subspaces of extended matrix pencils, with no inversion involved (Bender and Laub, 1987a; Bender and Laub, 1987b; Lancaster and Rodman, 1995; Mehrmann, 1991; Sima, 1996; Van Dooren, 1981):

– *extended pencil for CARE*:

$$N - \lambda M = \begin{bmatrix} A & 0 & B \\ Q & A^H & L \\ L^H & B^H & R \end{bmatrix} - \lambda \begin{bmatrix} E & 0 & 0 \\ 0 & -E^H & 0 \\ 0 & 0 & 0 \end{bmatrix};$$

– *extended pencil for DARE*:

$$N - \lambda M = \begin{bmatrix} A & 0 & B \\ Q & -E^H & L \\ L^H & 0 & R \end{bmatrix} - \lambda \begin{bmatrix} E & 0 & 0 \\ 0 & -A^H & 0 \\ 0 & -B^H & 0 \end{bmatrix}$$

Sima V..
COMPUTATIONAL EXPERIENCE WITH STRUCTURE-PRESERVING HAMILTONIAN SOLVERS IN OPTIMAL CONTROL.
DOI: 10.5220/0003534100910096
In *Proceedings of the 8th International Conference on Informatics in Control, Automation and Robotics* (ICINCO-2011), pages 91-96
ISBN: 978-989-8425-74-4

where $A,E,Q \in \mathbb{C}^{n\times n}$, $B,L \in \mathbb{C}^{n\times m}$, $R \in \mathbb{C}^{m\times m}$, $Q = Q^H$, $R = R^H$. If $\left[\begin{array}{ccc} U_1^T & U_2^T & U_3^T \end{array}\right]^T$ spans the stable right deflating subspace of $N-\lambda M$, then the stabilizing solution of the corresponding algebraic Riccati equation is $X_* = U_2(EU_1)^{-1}$ (if E is nonsingular). The solvers currently available, e.g., in MATLAB® Control System Toolbox, and SLICOT (Benner et al., 1999; Benner et al., 2010), are using the standard QZ algorithm for reordering the eigenvalues, to determine the stable deflating subspaces. The special structure of the matrix pencils involved is not exploited. But the use of structure-preserving algorithms might improve the numerical properties of the Riccati solvers.

Recently, structure-exploiting techniques have been investigated for solving skew-Hamiltonian/Hamiltonian eigenproblems, see, e.g., (Benner et al., 2002; Benner et al., 2007). These techniques can be employed for CARE solvers. For solving DAREs, the pencils can be preprocessed by an extended Cayley transformation, which only involves matrix additions and subtractions (Xu, 2006), to obtain equivalent skew-Hamiltonian/Hamiltonian pencils.

The paper presents some preliminary results obtained by the author using new software, developed in cooperation with Technical University Chemnitz, for computing the eigenvalues and stable deflating subspaces (with application in solving CAREs) based on structure-exploiting algorithms for skew-Hamiltonian/Hamiltonian matrix pencils. To the author's knowledge, this is the first attempt to use such algorithms in Riccati solvers.

This section is finished with few definitions. A matrix pencil $N-\lambda M$ is *Hamiltonian* if $N\mathcal{J}M^H = -M\mathcal{J}N^H$, and it is *symplectic* if $N\mathcal{J}N^H = M\mathcal{J}M^H$, where

$$\mathcal{J} := \left[\begin{array}{cc} 0 & I_n \\ -I_n & 0 \end{array}\right], \qquad \mathcal{J}^T = -\mathcal{J} = \mathcal{J}^{-1},$$

the superscripts H and T denote the conjugate-transpose and transpose, respectively, and I_n denotes the identity matrix of order n. If $M = I_{2n}$, definitions for Hamiltonian and symplectic matrices are obtained; for instance, N is *Hamiltonian* if $(N\mathcal{J})^H = N\mathcal{J}$, and it is *skew-Hamiltonian* if $(N\mathcal{J})^H = -N\mathcal{J}$. A matrix pencil $\lambda M - N$ is *skew-Hamiltonian/Hamiltonian* if M is skew-Hamiltonian, and N is Hamiltonian. These pencils have spectra which are symmetric with respect to the imaginary axis. In the sequel, the pencils $\lambda M - N$ will be represented in the numerically better form $\alpha M - \beta N$, with $\lambda = \alpha/\beta$ (possibly ∞).

2 COMPUTATION OF EIGENVALUES AND STABLE DEFLATING SUBSPACES

Let $\alpha\mathcal{S} - \beta\mathcal{H}$ be skew-Hamiltonian/Hamiltonian, i.e., $(\mathcal{S}\mathcal{J})^H = -\mathcal{S}\mathcal{J}$, $(\mathcal{H}\mathcal{J})^H = \mathcal{H}\mathcal{J}$. By definition, these pencils have even size. After eventual extension (to an even size, $2(n+\ell)$), permutation and scaling, the pencils corresponding to CARE have the following form

$$\alpha\mathcal{S} - \beta\mathcal{H} = \alpha \left[\begin{array}{cccc} E & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & E^H & 0 \\ 0 & 0 & 0 & 0 \end{array}\right] - \beta \left[\begin{array}{cccc} A & B_1 & 0 & B_2 \\ L_2^H & R_{12}^H & B_2^H & R_{22} \\ -Q & -L_1 & -A^H & -L_2 \\ -L_1^H & -R_{11} & -B_1^H & -R_{12} \end{array}\right], \quad (1)$$

where the four block rows and columns have orders n, ℓ, n, and ℓ, respectively. For some problems, including linear-quadratic optimization applications, $\mathcal{S}$ can be given in a factored form, the so-called *skew-Hamiltonian Cholesky factorization*, defined by $\mathcal{S} = \mathcal{J}\mathcal{Z}^H\mathcal{J}^T\mathcal{Z}$ (with the blocks of $\mathcal{J}$ of order $n+\ell$). For instance, in (1),

$$\mathcal{Z} = \left[\begin{array}{cccc} I_n & 0 & 0 & 0 \\ 0 & I_\ell & 0 & 0 \\ 0 & 0 & E^H & 0 \\ 0 & 0 & 0 & 0 \end{array}\right].$$

Some properties of skew-Hamiltonian/Hamiltonian pencils are proven, e.g., in (Benner et al., 2002).

For convenience, the real case only is dealt with in the sequel. An algorithm for computing the eigenvalues and a basis for the stable right deflating subspace (corresponding to the eigenvalues with strictly negative real part) of a skew-Hamiltonian/Hamiltonian pencil is summarized below, based on Algorithm 4 in (Benner et al., 2007):

1. Compute the following decompositions, defined by the matrices $\mathcal{Q}_1$ and $\mathcal{Q}_2$,

$$\begin{aligned} \mathcal{Q}_1^T \mathcal{S}\mathcal{J}\mathcal{Q}_1\mathcal{J}^T &= \left[\begin{array}{cc} N_1 & N_2 \\ 0 & N_1^T \end{array}\right], \\ (\mathcal{J}\mathcal{Q}_2\mathcal{J}^T)^T\mathcal{S}\mathcal{Q}_2 &= \left[\begin{array}{cc} M_1 & M_2 \\ 0 & M_1^T \end{array}\right], \\ \mathcal{Q}_1^T\mathcal{H}\mathcal{Q}_2 &= \left[\begin{array}{cc} H_{11} & H_{12} \\ 0 & H_{22} \end{array}\right], \end{aligned}$$

where N_1, M_1, and H_{11} are upper triangular, $N_2 = -N_2^T$, $M_2 = -M_2^T$, and H_{22}^T is upper quasi-triangular.

2. Find orthogonal matrices $\mathcal{Q}_3$ and $\mathcal{Q}_4$, such that

$$\begin{aligned}\mathcal{N}_{11} &= \mathcal{Q}_4^T \begin{bmatrix} N_1 & 0 \\ 0 & M_1 \end{bmatrix} \mathcal{Q}_3\,, \\ \mathcal{H}_{11} &= \mathcal{Q}_4^T \begin{bmatrix} 0 & H_{11} \\ -H_{22}^T & 0 \end{bmatrix} \mathcal{Q}_3\,,\end{aligned}$$

where $\mathcal{N}_{11}$ is upper triangular, and $\mathcal{H}_{11}$ is upper quasi-triangular.

3. Update

$$\begin{aligned}\mathcal{N}_{12} &= \mathcal{Q}_4^T \begin{bmatrix} N_2 & 0 \\ 0 & M_2 \end{bmatrix} \mathcal{Q}_4\,, \\ \mathcal{H}_{12} &= \mathcal{Q}_4^T \begin{bmatrix} 0 & H_{12} \\ H_{12}^T & 0 \end{bmatrix} \mathcal{Q}_4\,,\end{aligned}$$

and form

$$\mathcal{R}_{\mathcal{N}} = \begin{bmatrix} \mathcal{N}_{11} & \mathcal{N}_{12} \\ 0 & \mathcal{N}_{11}^T \end{bmatrix}, \quad \mathcal{R}_{\mathcal{H}} = \begin{bmatrix} \mathcal{H}_{11} & \mathcal{H}_{12} \\ 0 & -\mathcal{H}_{11}^T \end{bmatrix}.$$

4. Determine an orthogonal matrix $\widehat{\mathcal{Q}}$, such that $\mathcal{J}\widehat{\mathcal{Q}}^T\mathcal{J}^T\left(\alpha\mathcal{R}_{\mathcal{N}} - \beta\mathcal{R}_{\mathcal{H}}\right)\widehat{\mathcal{Q}}$ is still in structured triangular form and $\Lambda_-(\mathcal{R}_{\mathcal{H}}, \mathcal{R}_{\mathcal{N}})$ is contained in the spectrum of the leading $2p \times 2p$ principal subpencil of $\alpha\mathcal{N}_{11} - \beta\mathcal{H}_{11}$. The notation $\Lambda_-(N,M)$ denotes the stable spectrum of the pencil $\alpha M - \beta N$, and p is the number of eigenvalues in $\Lambda_-(\mathcal{H}, \mathcal{S})$.

5. Set

$$V = \begin{bmatrix} I_{2n} & 0 \end{bmatrix} \left(\mathcal{Y} \begin{bmatrix} \mathcal{J}\mathcal{Q}_1\mathcal{J}^T & 0 \\ 0 & \mathcal{Q}_2 \end{bmatrix} \mathcal{P} \times \begin{bmatrix} \mathcal{Q}_3 & 0 \\ 0 & \mathcal{Q}_4 \end{bmatrix} \widehat{\mathcal{Q}} \right) \begin{bmatrix} I_{2p} \\ 0 \end{bmatrix},$$

where

$$\mathcal{Y} = \frac{\sqrt{2}}{2} \begin{bmatrix} I_{2n} & I_{2n} \\ -I_{2n} & I_{2n} \end{bmatrix}, \quad \mathcal{P} = \begin{bmatrix} I_n & 0 & 0 & 0 \\ 0 & 0 & I_n & 0 \\ 0 & I_n & 0 & 0 \\ 0 & 0 & 0 & I_n \end{bmatrix},$$

and compute an orthogonal basis of the stable deflating subspace.

Step 2 of the algorithm needs to reorder the eigenvalues in the formal matrix product

$$N_1^{-1}H_{11}M_1^{-1}H_{22}^T\,, \qquad (2)$$

where H_{22}^T is upper quasi-triangular, and all the other matrices are upper triangular, so that the triangular form is kept, but the last diagonal blocks correspond to all nonpositive real eigenvalues and the first diagonal blocks correspond to the other eigenvalues. Note that Step 1 also uses the formal matrix product in (2), to reduce the obtained upper Hessenberg matrix H_{22}^T to upper quasi-triangular form, while preserving the other factors upper triangular. The periodic QZ algorithm (Bojanczyk et al., 1992; Sreedhar and Van Dooren, 1994) is used. Techniques for eigenvalue reordering in formal matrix products are discussed in (Sima, 2010) and the references therein. Solutions of certain periodic Sylvester-like equations are used. No factor is actually inverted. If only the eigenvalues are desired, they are returned by the periodic QZ algorithm called in Step 1 of the algorithm.

The structure can be exploited in Step 3 of the algorithm. For instance, $\mathcal{N}_{12} = -\mathcal{N}_{12}^T$ and $\mathcal{H}_{12} = \mathcal{H}_{12}^T$, and so, only their upper triangular parts should be computed. Also, the first block row only of the matrices $\mathcal{R}_{\mathcal{N}}$ and $\mathcal{R}_{\mathcal{H}}$ can be used in Step 4.

The reordering involved in Step 4 does not need the periodic QZ algorithm, but the standard QZ algorithm, for upper block triangular pencils of order 3 or 4. (Actually, the second matrix of the small order pencils is upper triangular.) In addition, reordering of the eigenvalues of special 2×2 or 4×4 skew-Hamiltonian/Hamiltonian pencils is needed. This can be done using relatively simple matrix calculations, as well as the QR factorization, and Givens rotations.

A similar algorithm for a factored matrix $\mathcal{S}$ is summarized in (Sima, 2010), based on Algorithm 3 in (Benner et al., 2007), and the called algorithms. In this case, the formal matrix product involves six factors. Moreover, the computations begin with an initial reduction, called *generalized symplectic URV decomposition*, defined as follows (Benner et al., 2007): Given a real $2n \times 2n$ skew-Hamiltonian/Hamiltonian pencil $\alpha\mathcal{S} - \beta\mathcal{H}$, $\mathcal{S} = \mathcal{T}\mathcal{Z}$ ($\mathcal{T} = \mathcal{J}\mathcal{Z}^T\mathcal{J}^T$), orthogonal matrices $\mathcal{Q}_1$, $\mathcal{Q}_2$ and orthogonal symplectic matrices $\mathcal{U}_1$, $\mathcal{U}_2$ are determined, such that

$$\begin{aligned}\mathcal{Q}_1^T\mathcal{T}\,\mathcal{U}_1 &= \begin{bmatrix} T_{11} & T_{12} \\ 0 & T_{22} \end{bmatrix}, \\ \mathcal{U}_2^T\mathcal{Z}\,\mathcal{Q}_2 &= \begin{bmatrix} Z_{11} & Z_{12} \\ 0 & Z_{22} \end{bmatrix}, \\ \mathcal{Q}_1^T\mathcal{H}\,\mathcal{Q}_2 &= \begin{bmatrix} H_{11} & H_{12} \\ 0 & H_{22} \end{bmatrix},\end{aligned}$$

where T_{11}, T_{22}^T, Z_{11}, Z_{22}^T, and H_{11} are upper triangular, and H_{22}^T is upper quasi-triangular. By definition, the matrices $\mathcal{U}_i$, $i = 1,2$, have the following form,

$$\mathcal{U}_i = \begin{bmatrix} U_{i1} & U_{i2} \\ -U_{i2} & U_{i1} \end{bmatrix},$$

so, they can be stored compactly in an implementation (the first n rows only).

Below is a summary about the related software:

- Fortran and MATLAB software for eigenvalues and deflating subspaces have just been developed.

- Both real and complex cases are considered.
- Factored or unfactored versions are covered.
- Optimized kernels for problems of order 2, 3, or 4, called by the general solvers, are available.

3 NUMERICAL RESULTS

This section presents some preliminary numerical results. These results have been obtained on a portable Intel Dual Core computer at 2 GHz, with 2 GB RAM, and relative machine precision $\varepsilon \approx 2.22 \times 10^{-16}$, using Windows XP (Service Pack 2) operating system, Intel Visual Fortran 11.1 compiler, and MATLAB 7.11.0.584 (R2010b).

3.1 Computation of Eigenvalues

Many numerical tests have been performed, to assess the correct behavior of the developed solvers. The matrices

$$\mathcal{S} = \begin{bmatrix} A & D \\ E & A^T \end{bmatrix}, \quad \mathcal{H} = \begin{bmatrix} B & V \\ W & -B^T \end{bmatrix},$$

where A, B, D, E, V, $W \in \mathbb{R}^{m\times m}$, have been generated with MATLAB commands using either uniform (0,1) random generator or the normal random generator, so that D and E be skew-symmetric matrices and V and W be symmetric matrices, resulting skew-Hamiltonian/Hamiltonian pencils.

Few very small skew-Hamiltonian/Hamiltonian examples are used below to illustrate the limitations of the standard, non-structured approach. The generalized eigenvalues computed by a structure-preserving algorithm and the standard QZ algorithm, optimally implemented in the MATLAB function `eig`, have been compared with those delivered by symbolic calculations, using the following MATLAB commands[1]

```
Ss = sym( S );  Hs = sym( H );
evs = double( eig( Ss \ Hs ) );
```

It was not possible to symbolically solve problems with $m \geq 5$. Based on the symmetry properties of the eigenvalues of the $(\mathcal{H},\mathcal{S})$ pencils, just eigenvalues with real parts larger than or equal to 0, and, for purely imaginary eigenvalues, those with positive imaginary parts, are reported. For instance, with

$$\mathcal{S} = \begin{bmatrix} 47 & 86 & 0 & 17 \\ 31 & 92 & -17 & 0 \\ 0 & -10 & 47 & 31 \\ 10 & 0 & 86 & 92 \end{bmatrix},$$

$$\mathcal{H} = \begin{bmatrix} 2 & 86 & 88 & 15 \\ 10 & 69 & 15 & 2 \\ 15 & 67 & -2 & -10 \\ 67 & 95 & -86 & -69 \end{bmatrix},$$

the structured algorithm found the eigenvalues

$$0.483611677311569,\ 1.310473800979598\imath$$

the MATLAB function `eig` returned

$$0.4836116773115708,$$
$$2.140945364757078\cdot 10^{-15} + 1.310473800979599\imath$$

and the symbolic MATLAB function `eig` computed

$$0.4836116773115688,\ 1.310473800979598\imath$$

where $\imath$ denotes the purely imaginary unit. The relative error norms of the first two solvers, compared to the symbolic solver, have the values $1.19\cdot 10^{-16}$ and $2.33\cdot 10^{-15}$, respectively. The first value is about 20 times smaller than the second one.

Fig. 1 and Fig. 2 show a comparison between the eigenvalues computed by the factored version of the structured algorithm and the standard algorithm `eig` for two examples of order 4 ($m = 2$).

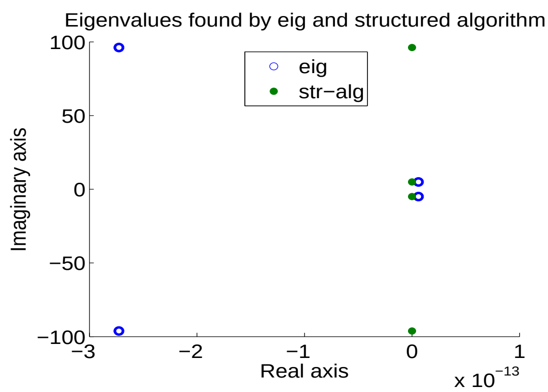

Figure 1: Eigenvalue scatter plot for an example of order 4.

For larger matrices, the differences between the results produced by the structured solver and by `eig` were more pronounced. An example of order 8 had two eigenvalues with real parts of order 10^{-10}, and an example of order 14 had two eigenvalues with real parts of order 10^{-8}, while the structured solver correctly found zero real parts for those eigenvalues.

[1]Unfortunately, there is no MATLAB generalized symbolic eigensolver, so the `mldivide` (or `mrdivide`) operator has been used, but the condition numbers of the tried skew-Hamiltonian matrices were very small, with one exception, for which $\mathcal{S}$ was singular.

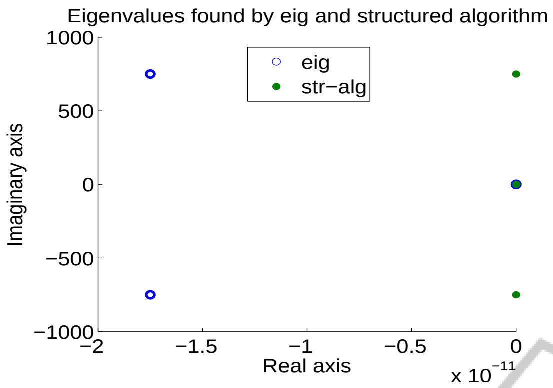

Figure 2: Eigenvalue scatter plot for another example of order 4. Two eigenvalues computed by `eig` are close between them and close to the corresponding eigenvalues computed by the structured solver.

3.2 Computation of Right Deflating Subspaces

Thousands of tests have been performed with random matrices for computing right deflating subspaces of skew-Hamiltonian/Hamiltonian matrix pencils. The results computed by the structured solver have been in good agreement to those obtained by the standard solver. In addition, the solvers have been compared for example problems from the SLICOT CARE benchmark collection (Abels and Benner, 1999). Most of them are difficult numerical examples. Three alternative options have been used for orthogonalizing the subspace basis—QR factorization (QR, for short), QR factorization with column pivoting (QRP), and singular value decomposition (SVD). The results have been compared with those delivered by the MATLAB function `care`.

Table 1 defines the parameters of the CARE examples. The codification of the column "parameter" is as follows: a value of -1 means that the default parameter value(s) are used (see (Abels and Benner, 1999)); a value of 1 means that the other parameter value(s) defined in (Abels and Benner, 1999) are used; a value 0 means that there are no parameters.

Fig. 3 presents the relative errors of the structured CARE solver for the three orthogonalizing options: QR, QRP, and SVD. The errors are relative to the exact solution, when known, or to the solution returned by the MATLAB function `care`, otherwise. Fig. 4 presents the relative residuals of the structured CARE solver and `care`. The function `care` uses scaling and permutations of the matrix or pencil, before reducing it. The same scaling, but no permutation, was used by the structured solver.

No orthogonalizing option is the best for all problems. Most examples are solved very well, but the results for some problems are not good enough. A possible explanation might be the fact that the structured algorithm for computing the stable deflating subspace doubles the eigenvalue multiplicities. Further investigation is needed.

Table 1: CARE benchmark examples.

Test	example	n	m	parameter
1	1.1	2	1	0
2	1.2	2	1	0
3	1.3	4	2	0
4	1.4	8	2	0
5	1.5	9	3	0
6	1.6	30	3	0
7	2.1	2	1	1
8	2.1	2	1	-1
9	2.2	2	2	1
10	2.2	2	2	-1
11	2.3	2	1	1
12	2.3	2	1	-1
13	2.3	2	1	10^{-6}
14	2.4	2	2	1
15	2.4	2	2	-1
16	2.5	2	1	1
17	2.5	2	1	-1
18	2.6	3	3	1
19	2.6	3	3	-1
20	2.7	4	1	1
21	2.7	4	1	-1
22	2.8	4	1	1
23	2.8	4	1	-1
24	2.9	55	2	-1
25	3.1	9	5	1
26	3.1	39	20	-1
27	3.2	8	8	1
28	3.2	64	64	-1
29	4.1	21	1	-1
30	4.1	21	1	1
31	4.2	20	1	1
32	4.2	100	1	-1
33	4.3	60	2	-1

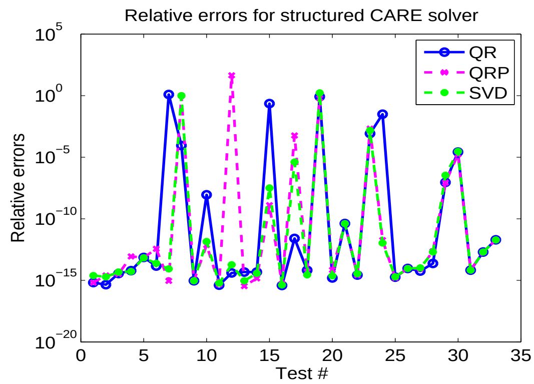

Figure 3: Relative errors of the structured CARE solver for CARE benchmark examples.

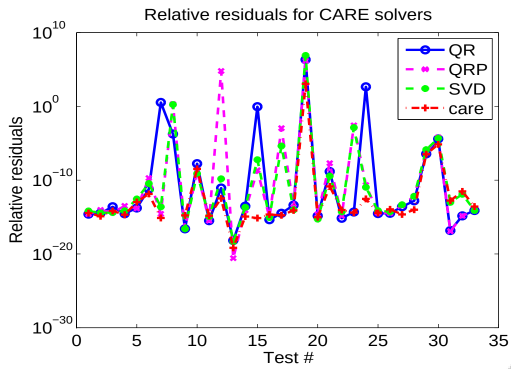

Figure 4: Relative residuals of CARE solvers for CARE benchmark examples.

4 CONCLUSIONS

Main issues related to the structure-preserving algorithms for solving some essential control problems in optimal and robust systems analysis and design are summarized. Eigenvalues and stable right deflating subspaces are computed based on skew-Hamiltonian/Hamiltonian pencils. The results for eigenvalue computations, with applications, e.g., in evaluating L_∞- and H_∞-norms, are very good. The computation of stable deflating subspaces, with applications in CARE/DARE solvers, deserves further investigation for difficult numerical problems.

ACKNOWLEDGEMENTS

The work was partially supported by the German Research Foundation (DFG) and The MathWorks, Inc. The collaboration with Peter Benner and Matthias Voigt from TU Chemnitz is highly acknowledged.

REFERENCES

Abels, J. and Benner, P. (1999). CAREX—A collection of benchmark examples for continuous-time algebraic Riccati equations (Version 2.0). SLICOT Working Note 1999-14. http://www.slicot.org/.

Bender, D. J. and Laub, A. J. (1987a). The linear-quadratic optimal regulator for descriptor systems. *IEEE Trans. Automat. Contr.*, AC-32(8):672–688.

Bender, D. J. and Laub, A. J. (1987b). The linear-quadratic optimal regulator for descriptor systems: Discrete-time case. *Automatica*, 23(1):71–85.

Benner, P., Byers, R., Losse, P., Mehrmann, V., and Xu, H. (2007). Numerical solution of real skew-Hamiltonian/Hamiltonian eigenproblems. Technical report, Technische Universität Chemnitz, Chemnitz.

Benner, P., Byers, R., Mehrmann, V., and Xu, H. (2002). Numerical computation of deflating subspaces of skew Hamiltonian/Hamiltonian pencils. *SIAM J. Matrix Anal. Appl.*, 24(1):165–190.

Benner, P., Kressner, D., Sima, V., and Varga, A. (2010). Die SLICOT-Toolboxen für Matlab. *at—Automatisierungstechnik*, 58(1):15–25.

Benner, P., Mehrmann, V., Sima, V., Van Huffel, S., and Varga, A. (1999). SLICOT — A subroutine library in systems and control theory. In *Applied and Computational Control, Signals, and Circuits*, vol. 1, ch. 10, 499–539. Birkhäuser, Boston.

Bojanczyk, A. W., Golub, G., and Van Dooren, P. (1992). The periodic Schur decomposition: algorithms and applications. In *Proc. of the SPIE Conference*, vol. 1770, 31–42.

Boyd, S., Balakrishnan, V., and Kabamba, P. (1989). A bisection method for computing the H_∞ norm of a transfer matrix and related problems. *Mathematics of Control, Signals, and Systems*, 2(3):207–219.

Bruinsma, N. A. and Steinbuch, M. (1990). A fast algorithm to compute the H_∞-norm of a transfer function. *Systems & Control Letters*, 14:287–293.

Lancaster, P. and Rodman, L. (1995). *The Algebraic Riccati Equation*. Oxford University Press, Oxford.

Laub, A. J. (1979). A Schur method for solving algebraic Riccati equations. *IEEE Trans. Automat. Contr.*, AC–24(6):913–921.

MATLAB (2010). The MATLAB Control System Toolbox. Version 9.0.

Mehrmann, V. (1991). *The Autonomous Linear Quadratic Control Problem. Theory and Numerical Solution*, volume 163 of *Lect. Notes in Control and Information Sciences*. Springer-Verlag, Berlin.

Pappas, T., Laub, A. J., and Sandell, N. R. (1980). On the numerical solution of the discrete-time algebraic Riccati equation. *IEEE Trans. Automat. Contr.*, AC–25(4):631–641.

Sima, V. (1996). *Algorithms for Linear-Quadratic Optimization*. Marcel Dekker, Inc., New York.

Sima, V. (2006). Efficient algorithm for L_∞-norm calculations. In *5th IFAC Symposium on Robust Control Design, July 5–7, 2006, Toulouse, France*.

Sima, V. (2010). Structure-preserving computation of stable deflating subspaces. In *Proceedings of the 10th IFAC Workshop "Adaptation and Learning in Control and Signal Processing", Antalya, Turkey*, 26–28 August 2010 — Kudret Press & Digital Printing Co.

Sreedhar, J. and Van Dooren, P. (1994). Periodic Schur form and some matrix equations. In *Proc. of the Symposium on the Mathematical Theory of Networks and Systems*, vol. 1, 339–362, Regensburg, Germany.

Van Dooren, P. (1981). A generalized eigenvalue approach for solving Riccati equations. *SIAM J. Sci. Stat. Comput.*, 2(2):121–135.

Xu, H. (2006). On equivalence of pencils from discrete-time and continuous-time control. *Lin. Alg. Appl.*, 414:97–124.