

COMPUTATIONAL EXPERIENCE IN SOLVING LINEAR MATRIX EQUATIONS FOR AUTOMATIC CONTROL

Vasile Sima *

* *National Institute for Research & Development in Informatics
Bd. Mareşal Averescu, Nr. 8–10, 011455 Bucharest 1, Romania*

Abstract: State-of-the-art, uni-processor linear matrix equation solvers for automatic control computations are investigated and compared for various problem sizes. General-purpose SLICOT solvers are the most efficient ones for small-size problems, but they cannot compete for larger problems with specialized solvers designed for certain problem classes. *Copyright ©2005 IFAC*

Keywords: Computer-aided control systems design, Linear control systems, Lyapunov equation, Numerical algorithms, Numerical solutions

1. INTRODUCTION

Human-made systems are heavily dependent on computer technology and automatic control concepts and algorithms. Control systems analysis and design procedures often require the solution of general or special linear or quadratic matrix equations. Examples are: invariant or deflating subspaces of matrices or matrix pairs, block-diagonalization and computation of matrix functions, controllability and observability Gramians, Hankel singular values, model and controller reduction, linear-quadratic optimization, condition estimation for eigenvalue problems and linear or quadratic matrix equations, etc. Theoretical results devoted to matrix equations and related topics abound both in systems and control, as well as in the linear algebra literature. There are also many associated software implementations, either commercial (e.g., in MATLAB¹ (MathWorks, 1998, 1999)), copyrighted freeware (e.g., in the SLICOT Library (Benner *et al.*, 1999; Van Huffel and Sima, 2002; Van Huffel *et al.*, 2004)), or in the public domain (e.g., in Scilab (Gomez, 1999)). The reliability, efficiency, and functionality of various solvers differ significantly from package to package.

Although numerical algorithms for linear matrix equations in control theory were published since 1960, this is still a hot research topic. The challenge for solving larger and larger equations has not yet been fully answered. The proposed techniques are usually not general enough. There are few comparative studies. It is the purpose of this paper to investigate the performances of several powerful solvers for linear matrix equations.

The capabilities and limitations of the general-purpose solvers available in the SLICOT Library and MATLAB are studied, in comparison with some specialized solvers. SLICOT Library (Subroutine Library In Control Theory) provides Fortran 77 implementations of many numerical algorithms in systems and control theory, as well as standardized interfaces to MATLAB and Scilab. Built around a nucleus of basic numerical linear algebra subroutines from the state-of-the-art software packages LAPACK (Anderson *et al.*, 1999), BLAS (Dongarra *et al.*, 1988, 1990; Lawson *et al.*, 1979), this library enables to exploit the potential of modern high-performance computer architectures. The SLICOT solvers for linear matrix equations offer improved efficiency, reliability, and functionality over the corresponding solvers in other computer-aided control system design packages.

¹ MATLAB is a registered trademark of The MathWorks, Inc.

2. SLICOT LINEAR MATRIX EQUATION SOLVERS CAPABILITIES

The following equation classes are considered:

- Continuous-time Sylvester equations:

$$\text{op}(\mathbf{A}) \mathbf{X} \pm \mathbf{X} \text{op}(\mathbf{B}) = \sigma \mathbf{C}; \quad (1)$$

- Generalized Sylvester equations (2) and the “transposed” equations (3):

$$\mathbf{A}\mathbf{X} - \mathbf{Y}\mathbf{B} = \sigma \mathbf{G}, \quad \mathbf{E}\mathbf{X} - \mathbf{Y}\mathbf{F} = \sigma \mathbf{H}; \quad (2)$$

$$\mathbf{A}^T \mathbf{X} + \mathbf{E}^T \mathbf{Y} = \sigma \mathbf{G}, \quad \mathbf{X}\mathbf{B}^T + \mathbf{Y}\mathbf{F}^T = -\sigma \mathbf{H}; \quad (3)$$

- Discrete-time Sylvester equations:

$$\text{op}(\mathbf{A}) \mathbf{X} \text{op}(\mathbf{B}) \pm \mathbf{X} = \sigma \mathbf{C}; \quad (4)$$

- Generalized continuous-time and discrete-time Lyapunov equations

$$\text{op}(\mathbf{A})^T \mathbf{X} \text{op}(\mathbf{E}) + \text{op}(\mathbf{E})^T \mathbf{X} \text{op}(\mathbf{A}) = \sigma \mathbf{C}; \quad (5)$$

$$\text{op}(\mathbf{A})^T \mathbf{X} \text{op}(\mathbf{A}) - \text{op}(\mathbf{E})^T \mathbf{X} \text{op}(\mathbf{E}) = \sigma \mathbf{C}; \quad (6)$$

where the notation $\text{op}(\mathbf{M})$ denotes either the matrix \mathbf{M} , or its transpose, \mathbf{M}^T , \mathbf{A} , \mathbf{B} , \mathbf{E} , and \mathbf{F} , are $n \times n$, $m \times m$, $n \times n$, and $m \times m$ given matrices, respectively, \mathbf{C} , \mathbf{G} , and \mathbf{H} are given matrices of appropriate dimensions, \mathbf{X} and \mathbf{Y} are unknown matrices of appropriate dimensions, and σ is a scaling factor, usually equal to one, but possibly set less than one, in order to prevent overflow in the solution matrix. The equations (1) and (2) are called *one-sided*, since the unknown matrices are either premultiplied or postmultiplied by known matrices, while the equations (4)–(6) are similarly called *two-sided*.

Taking $\mathbf{E} = \mathbf{I}_n$, the identity matrix of order n , in (5) and (6), standard *Lyapunov equations* are obtained. Moreover, when the matrix \mathbf{A} (or the matrix pair (\mathbf{A}, \mathbf{E})) in a (generalized) Lyapunov equation is stable—in the continuous- or discrete-time sense—and the right hand side term $\sigma \mathbf{C}$ is specified in a factored form as $-\sigma^2 \text{op}(\mathbf{D})^T \text{op}(\mathbf{D})$, where $\text{op}(\mathbf{D})$ is $m \times n$, then the corresponding Lyapunov equation has a non-negative definite solution which can be computed in a factored form, $\mathbf{X} = \text{op}(\mathbf{U})^T \text{op}(\mathbf{U})$, with \mathbf{U} upper triangular. Such equations are called (*generalized*) *stable continuous- and discrete-time Lyapunov equations*.

All equation classes and their specializations mentioned above are solvable using the SLICOT solvers, thus illustrating their extended functionality (see also (Sima and Benner, 2003; Slowik *et al.*, 2004)). Let $\mathcal{E}(\mathcal{D}, \mathcal{U}) = \mathcal{R}$ be a shorthand notation for any of the above equations, where \mathcal{E} , \mathcal{D} , \mathcal{U} , and \mathcal{R} denote the corresponding equation formula, data, unknowns, and right hand side term, respectively. For general matrices, the solution is obtained by a *transformation method* (see, e.g., (Sima, 1996, page 144)). Specifically, the data \mathcal{D} are transformed to some simpler forms, $\tilde{\mathcal{D}}$ (usually corresponding to the real Schur form (RSF) of \mathbf{A} , or generalized RSF of a matrix pair),

the right hand side term is transformed accordingly to $\tilde{\mathcal{R}}$, the *reduced equation*, $\mathcal{E}(\tilde{\mathcal{D}}, \tilde{\mathcal{U}}) = \tilde{\mathcal{R}}$, is solved in $\tilde{\mathcal{U}}$, and finally, the solution of the original equation is recovered from $\tilde{\mathcal{U}}$.

The methods implemented in SLICOT are basically the following: the Schur method (also known as Bartels–Stewart method) (Bartels and Stewart, 1972) for Sylvester or Lyapunov equations, with the variant from (Barraud, 1977) for the discrete-time case; the Hessenberg-Schur method in (Golub *et al.*, 1979) for standard Sylvester equations, i.e., with $\text{op}(\mathbf{M}) = \mathbf{M}$; Hammarling’s variant (Hammarling, 1982) of the Bartels–Stewart method for stable Lyapunov equations; and extensions of the above methods for generalized Sylvester (Kågström and Poromaa, 1996) and Lyapunov equations (Penzl, 1998).

The ability to work with the $\text{op}(\cdot)$ operator is important in many control analysis and design problems. For instance, the controllability Gramians can be defined as solutions of stable Lyapunov equations with $\text{op}(\mathbf{A}) = \mathbf{A}^T$, while observability Gramians can be defined as solutions of stable Lyapunov equations with $\text{op}(\mathbf{A}) = \mathbf{A}$. When both controllability and observability Gramians are needed (e.g., in model reduction computations), then the same real Schur form of \mathbf{A} can be used by a solver able to cope with $\text{op}(\cdot)$, and this significantly improves the efficiency.

The solvers for stable Lyapunov equations directly compute the Cholesky factor \mathbf{U} of the solution \mathbf{X} . Whenever feasible, the use of the stable solvers instead of the general ones is to be preferred, for several reasons, including the following: • the matrix product $\text{op}(\mathbf{D})^T \text{op}(\mathbf{D})$ need not be computed; • definiteness of \mathbf{X} is guaranteed. Moreover, often the Cholesky factors themselves are actually needed, e.g., for model reduction or for computing the Hankel singular values.

When solving any matrix equation, it is useful to have estimates of the *problem conditioning* and of the solution accuracy, e.g., *error bounds*. Such measures are returned by several routines of the SLICOT Library (Sima *et al.*, 2000), allowing to assess the quality of the computed solution, and the problem sensitivity to small perturbations in its data. This illustrates the increased reliability and functionality of the SLICOT software, in comparison with many other control packages. The condition estimator for a certain equation operator uses a fast iterative procedure, which solves (at each step) the equation and its dual (with transposed data matrices), for suitable right hand sides, to estimate the 1-norm of the inverse operator.

3. SPECIALIZED LINEAR MATRIX EQUATION SOLVERS

The results in (Sima and Benner, 2003; Slowik *et al.*, 2004) and other papers, as well as those included below, show that the high-level MATLAB interfaces

to the SLICOT codes offer improved efficiency (at comparable accuracy) over the existing standard software tools. However, the SLICOT solvers do have some limitations, mainly coming from their generality. These solvers cannot compete in terms of efficiency with specialized solvers designed for specific classes of large-scale problems. Two types of specialized solvers are considered in this investigation: iterative algorithms for stable Lyapunov equations with low rank solutions, and recursive blocked algorithms for quasi-triangular linear matrix equations.

3.1 Iterative algorithms for stable, low rank Lyapunov equations

The approach for solving large-scale Lyapunov equations implemented in the MATLAB package LYAPACK (LYApunov PACKage) (Penzl, 2000) can be applied to structured or sparse stable continuous-time equations of the form

$$\mathbf{F}^T \mathbf{X} + \mathbf{X} \mathbf{F} = -\mathbf{D}^T \mathbf{D}, \quad (7)$$

where $\mathbf{F} \in \mathbb{R}^{n \times n}$ and $\mathbf{D} \in \mathbb{R}^{m \times n}$. In many applications, for instance, model reduction or algebraic Riccati equations, it is sufficient to obtain a factorization of the solution matrix \mathbf{X} , $\mathbf{X} = \mathbf{Z}^T \mathbf{Z}$. For solving continuous-time Riccati equations iteratively, one may use at each iteration linear equations (7), where the matrix \mathbf{F} has the form $\mathbf{F} = \mathbf{A} - \mathbf{B} \mathbf{K}^T$, \mathbf{A} and \mathbf{B} being the matrices of the system state equation, and \mathbf{K}^T the regulator gain matrix.

Besides the limitations imposed by the form of the equation (7) and stability hypothesis, it is also assumed that the number of rows m is small in comparison with n , $m \ll n$, and that the matrix \mathbf{F} is structured so that efficient solution of linear systems with coefficient matrices $\mathbf{F} - p \mathbf{I}_n$, where $p \in \mathbb{C}$, as well as efficient computation of matrix-vector products are possible. Moreover, the order n should be large enough, for instance, $n > 500$, and the equations be sufficiently well-conditioned.

The LYAPACK approach, implemented in the function `lp_lradi`, uses the *low rank Cholesky factor* technique, in combination with *alternating directions* method, abbreviated as LRCF-ADI (Low Rank Cholesky Factor Alternating Directions Implicit) iterations. The efficiency of LRCF-ADI depends on certain ADI *shift parameters*, p_i , computed by an heuristic algorithm. The LRCF technique is based on the observation that in many problems (7) with $m \ll n$, the eigenvalues of the solution matrix \mathbf{X} decay very fast, which suggests the possible existence of very accurate approximations of rank much smaller than n . The ADI iteration for Lyapunov equation (7) is given by

$$\begin{aligned} (\mathbf{F}^T + p_i \mathbf{I}_n) \mathbf{X}_{i-1/2} &= -\mathbf{D}^T \mathbf{D} - \mathbf{X}_{i-1} (\mathbf{F} - p_i \mathbf{I}_n), \\ (\mathbf{F}^T + \bar{p}_i \mathbf{I}_n) \mathbf{X}_i^T &= -\mathbf{D}^T \mathbf{D} - \mathbf{X}_{i-1/2}^T (\mathbf{F} - \bar{p}_i \mathbf{I}_n), \end{aligned}$$

for $i = 1, 2, \dots$, where $\mathbf{X}_0 = 0$. Each iteration involves matrix-vector products and solutions of structured or sparse linear systems. The convergence is accelerated using the parameters p_i . This method generates a sequence of matrices \mathbf{X}_i which converges often very fast to the solution, provided that p_i are suitably chosen. The efficient implementation of the ADI method replaces the iterates \mathbf{X}_i by their low rank Cholesky factors, $\mathbf{X}_i = \mathbf{Z}_i^T \mathbf{Z}_i$. Let \mathcal{P}_j be a real negative number, or a pair of complex conjugated numbers with negative real part. If the matrix $\mathbf{X}_i = \mathbf{Z}_i^T \mathbf{Z}_i$ is generated by a *proper* set of parameters $\{p_1, p_2, \dots, p_i\} = \{\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_i\}$, then \mathbf{X}_i is a real matrix. The problem of finding (sub)optimal ADI parameters, $\mathcal{P} = \{\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_\ell\}$, is strongly connected to the rational minimax problem applied to the function

$$s_{\mathcal{P}}(t) = \frac{|(t - p_1) \cdot \dots \cdot (t - p_\ell)|}{|(t + p_1) \cdot \dots \cdot (t + p_\ell)|}.$$

This problem is stated as $\min_{\mathcal{P}} \max_{t \in \sigma(\mathbf{F})} s_{\mathcal{P}}(t)$, where $\sigma(\mathbf{F})$ denotes the spectrum of the matrix \mathbf{F} . The implementation of the heuristic technique first generates a discrete set, which approximates the spectrum, using a pair of Arnoldi processes. The first process, acting on the matrix \mathbf{F} , produces k_+ Ritz values which tend to approximate the eigenvalues farthest from the origin. The second process, acting on the matrix \mathbf{F}^{-1} , produces k_- Ritz values, approximations of the eigenvalues close to the origin. The set of shift parameters is then chosen as a subset of the Ritz values, as an heuristic, suboptimal solution of the resulting discrete optimization problem.

The use of the LYAPACK package implies that the user writes specific routines performing the following operations with a structured or sparse matrix \mathbf{M}

$$\begin{aligned} \mathbf{Y} &\leftarrow \mathbf{M} \mathbf{Y} \quad \text{or} \quad \mathbf{Y} \leftarrow \mathbf{M}^T \mathbf{Y}, \\ \mathbf{Y} &\leftarrow \mathbf{M}^{-1} \mathbf{Y} \quad \text{or} \quad \mathbf{Y} \leftarrow \mathbf{M}^{-T} \mathbf{Y}, \\ \mathbf{Y} &\leftarrow (\mathbf{M} + p_i \mathbf{I}_n)^{-1} \mathbf{Y} \quad \text{or} \quad \mathbf{Y} \leftarrow (\mathbf{M}^T + p_i \mathbf{I}_n)^{-1} \mathbf{Y}, \end{aligned}$$

where \mathbf{M} is \mathbf{F} or \mathbf{A} , and $\mathbf{Y} \in \mathbb{C}^{n \times t}$, $t \ll n$.

3.2 Recursive blocked algorithms for quasi-triangular linear matrix equations

An approach (Jonsson and Kågström, 2002a, b) which can be applied to all classes of linear matrix equations with quasi-triangular matrices is based on the use of recursive blocked algorithms. The basic ideas are to recursively decompose the quasi-triangular matrices in blocks until the obtained equations are small enough for being solved in the very fast cache memory, and to use some “superscalar” computational kernels for equations of small dimensions. The sizes of the blocks are variable, and this fact enables their automatic adaptation to the computational platform used, and the efficient exploitation of the existing memory hierarchies on modern computing machines.

To illustrate, consider the case of a standard continuous-time Sylvester equation,

$$\mathbf{A}\mathbf{X} + \mathbf{X}\mathbf{B} = \mathbf{C}, \quad (8)$$

where \mathbf{A} and \mathbf{B} are either upper triangular or in real Schur form. Depending on the values of m and n , three alternative *recursive block decompositions* can be considered. One such alternative is illustrated below. If $1 \leq m \leq n/2$, \mathbf{A} is decomposed by rows and columns, and \mathbf{C} is decomposed by rows:

$$\begin{bmatrix} \mathbf{A}_{11} & \mathbf{A}_{12} \\ 0 & \mathbf{A}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{bmatrix} + \begin{bmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{bmatrix} \mathbf{B} = \begin{bmatrix} \mathbf{C}_1 \\ \mathbf{C}_2 \end{bmatrix}, \quad (9)$$

or, equivalently,

$$\begin{aligned} \mathbf{A}_{11}\mathbf{X}_1 + \mathbf{X}_1\mathbf{B} &= \mathbf{C}_1 - \mathbf{A}_{12}\mathbf{X}_2, \\ \mathbf{A}_{22}\mathbf{X}_2 + \mathbf{X}_2\mathbf{B} &= \mathbf{C}_2. \end{aligned} \quad (10)$$

Two quasi-triangular Sylvester equations have been obtained. The second equation is solved in \mathbf{X}_2 , and after a GEMM-type update, $\mathbf{C}_1 \leftarrow \mathbf{C}_1 - \mathbf{A}_{12}\mathbf{X}_2$, the first Sylvester equation is solved. For solving each Sylvester equation, one proceeds similarly.

There are three levels of solvers for linear matrix equations. The recursive block solvers are destined to the user. Each of these solvers calls a *sub-system* block solver, when the dimensions m and n of the current subproblem in the recursive decomposition are smaller than a certain block size, `blk_s`. Finally, each sub-system solver calls a *superscalar kernel* for solving equations with $m, n \leq 4$. Besides the advantageous use of the memory hierarchies, the recursive approach allows to consider various forms of parallelism. The major disadvantage of the recursive block solvers is that they merely solve “reduced” equations. The initial reduction to the (generalized) RSF is not covered. The codes are implemented in the RECSY library, and wrappers to the SLICOT solvers are provided, so that general equations can be solved, and condition estimates can be computed.

4. NUMERICAL RESULTS

Some typical results are graphically illustrated in the figures below. The calculations have been done on a PC computer with a 500 MHz Intel processor, 128 Mb memory and the relative machine precision $\epsilon = 2.22 \times 10^{-16}$, using Compaq Visual Fortran V6.5, optimized BLAS provided by MATLAB, and MATLAB 6.5.1.

Application 1 has a band matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ with 5 nonzero diagonals, obtained by discretization of a partial differential equation, using finite differences on an equidistant grid. The right-hand side matrix has the form $\mathbf{D}^T \mathbf{D}$, where $\mathbf{D}^T \in \mathbb{R}^n$. The data matrices have been generated by the LYAPACK example codes `fdm_2d_matrix` and `fdm_2d_vector`.

Figures 1, 2 and 3 show the execution times and speed-up factors for problem orders in the ranges $n \leq 225$,

$196 \leq n \leq 400$, and $400 \leq n \leq 1024$, for the solvers `slylap`, `lyap` and `lp_lradi`, in SLICOT, MATLAB and LYAPACK, respectively. The equations with orders in these ranges could be considered as “small”, “medium”, and “large”, respectively, for the computer used for their solution. The results show

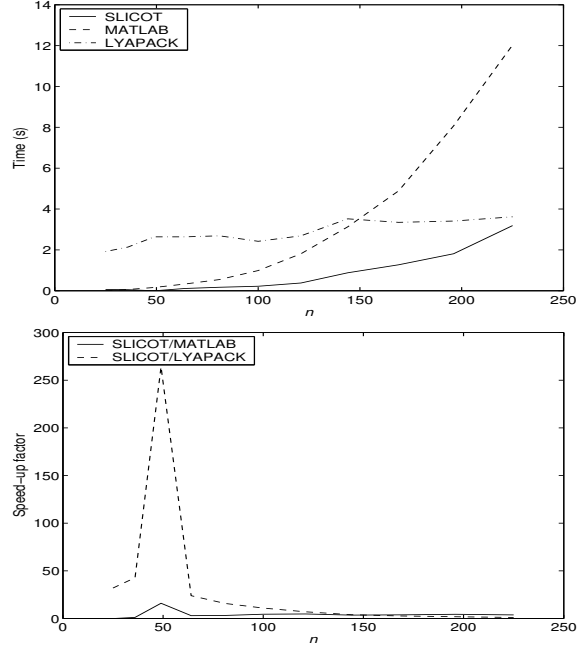


Fig. 1. Application 1, $n \leq 225$. Top: The execution times. Bottom: The speed-up factors.




Fig. 2. Application 1, $196 \leq n \leq 400$. Top: The execution times. Bottom: The speed-up factors.

that SLICOT routines always outperform MATLAB calculations (for any order), and also `lp_lradi`, for problems of small size. It should be mentioned that the accuracy is comparable for all these solvers and all equations solved. SLICOT `slylap` is 2–3




Fig. 3. Application 1, $400 \leq n \leq 1024$. Top: The execution times. Bottom: The speed-up factors.

(or much more) times faster than MATLAB `lyap`.² Also, `slyap` is faster (possibly much faster) than `lp_lradi` for Lyapunov equations of order smaller than 275, but slower (possibly much slower) for larger orders. It should, however, be mentioned that, in contrast with `(s)lyap`, `lp_lradi` is not a general solver. Its high efficiency is due to the use of the sparse structure of the matrix \mathbf{A} in operations like $\mathbf{A}\mathbf{b}$ or $\mathbf{A}^{-1}\mathbf{b}$, where \mathbf{b} is a vector. Note also that, besides the stability and sparsity requirements, `lp_lradi` also assumes some additional conditions, such as: the solution matrix has a small rank; the equation order is large enough; the equation is quite well-conditioned.

Figures 4 and 5 present the execution times and the speed-up factors when calling the recursive blocked algorithms for another application, Application 2. In this case, the matrix \mathbf{A} has been obtained starting from a block Jordan matrix \mathbf{A}_0 , and applying a similarity transformation, which filled-up the matrix with nonzero elements and altered its condition number. But `lp_lradi` cannot efficiently solve an equation with a dense matrix \mathbf{A} , so the bidiagonal matrix \mathbf{A}_0 was used; `lp_lradi` becomes more efficient than `slyap` for $n \geq 150$. Figure 6 shows the resulting speed-up factors if \mathbf{A}_0 is also used by `(s)lyap`.

5. CONCLUSIONS

Various state-of-the-art, uni-processor linear matrix equation solvers for automatic control computations have been investigated and compared for various

² The latest `lyap` versions included in MATLAB 7.0.0/1 (2004), is not considered, since it is based on the corresponding SLICOT routines; this version could also be about 20 % slower than `slyap`.




Fig. 4. Application 2, $n \leq 400$, recursive blocked algorithm. Top: The execution times. Bottom: The speed-up factors. `lp_lradi` uses \mathbf{A}_0 .




Fig. 5. Application 2, $400 \leq n \leq 1000$, recursive blocked algorithm. Top: The execution times. Bottom: Speed-up factors. `lp_lradi` uses \mathbf{A}_0 .

problem sizes. The results confirm the natural expectation that general-purpose solvers, such as those currently implemented in the SLICOT Library (and, consequently, in MATLAB 7) cannot compete in efficiency, for large-scale problems, with specialized solvers designed for certain problem classes. However, the SLICOT solvers are the most efficient ones for small-size problems. Moreover, they are general solvers and offer extended functionality and broad computational abilities.




Fig. 6. The speed-up factors for Application 2, using the bidiagonal matrix \mathbf{A}_0 . Top: Non-recursive algorithm. Bottom: Recursive blocked algorithm.

REFERENCES

- Anderson, E., Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney and D. Sorensen (1999). *LAPACK Users' Guide: Third Edition*. SIAM, Philadelphia.
- Barraud, A.Y. (1977). A numerical algorithm to solve $A^T X A - X = Q$. *IEEE Trans. Automat. Contr.* **AC-22**, 883–885.
- Bartels, R.H. and G.W. Stewart (1972). Algorithm 432: Solution of the matrix equation $AX + XB = C$. *Comm. ACM* **15**, 820–826.
- Benner, P., V. Mehrmann, V. Sima, S. Van Huffel, and A. Varga (1999). SLICOT — A subroutine library in systems and control theory. In: *Applied and Computational Control, Signals, and Circuits* (B.N. Datta, Ed.). Vol. 1, Chap. 10, pp. 499–539. Birkhäuser, Boston.
- Dongarra, J.J., J. Du Croz, I.S. Duff and S. Hammarling (1990). Algorithm 679: A set of Level 3 Basic Linear Algebra Subprograms. *ACM Trans. Math. Softw.* **16**, 1–17, 18–28.
- Dongarra, J.J., J. Du Croz, S. Hammarling and R.J. Hanson (1988). Algorithm 656: An extended set of Fortran Basic Linear Algebra Subprograms. *ACM Trans. Math. Softw.* **14**, 1–17, 18–32.
- Golub, G.H., S. Nash and C.F. Van Loan (1979). A Hessenberg-Schur method for the problem $AX + XB = C$. *IEEE Trans. Automat. Contr.* **AC-24**, 909–913.
- Gomez, C. (Ed.) (1999). *Engineering and Scientific Computing with Scilab*. Birkhäuser, Boston.
- Hammarling, S.J. (1982). Numerical solution of the stable, non-negative definite Lyapunov equation. *IMA J. Numer. Anal.* **2**, 303–323.
- Jonsson, I. and B. Kågström (2002a). Recursive blocked algorithms for solving triangular systems—Part I: One-sided and coupled Sylvester-type matrix equations. *ACM Trans. Math. Softw.* **28**, 392–415.
- Jonsson, I. and B. Kågström (2002b). Recursive blocked algorithms for solving triangular systems—Part II: Two-sided and generalized Sylvester and Lyapunov matrix equations. *ACM Trans. Math. Softw.* **28**, 416–435.
- Kågström, B. and P. Poromaa (1996). LAPACK-style algorithms and software for solving the generalized Sylvester equation and estimating the separation between regular matrix pairs. *ACM Trans. Math. Softw.* **22**, 78–103.
- Lawson, C.L., R.J. Hanson, D.R. Kincaid and F.T. Krogh (1979). Basic Linear Algebra Subprograms for Fortran usage. *ACM Trans. Math. Softw.* **5**, 308–323.
- MathWorks (1998). *Control System Toolbox User's Guide*. The MathWorks, Inc., Natick, MA.
- MathWorks (1999). *Using MATLAB. Version 5*. The MathWorks, Inc., Natick, MA.
- Penzl, T. (1998). Numerical solution of generalized Lyapunov equations. *Advances in Comp. Math.* **8**, 33–48.
- Penzl, T. (2000). LYAPACK Users Guide. Techn. Rep. SFB393/00–33. Technische Universität Chemnitz, Sonderforschungsbereich 393. Chemnitz.
- Sima, V. (1996). *Algorithms for Linear-quadratic Optimization*. Marcel Dekker, Inc., New York.
- Sima, V. and P. Benner (2003). Solving linear matrix equations with SLICOT. In: *Proceedings of the European Control Conference ECC'03*, 1–4 September, 2003. Cambridge, UK. 6 pages.
- Sima, V., P. Petkov and S. Van Huffel (2000). Efficient and reliable algorithms for condition estimation of Lyapunov and Riccati equations. In: *Proceedings CD of the Fourteenth International Symposium of Mathematical Theory of Networks and Systems MTNS-2000*, Perpignan, France, June 19–23, 2000. Session CS 2C, 10 pages.
- Slowik, M., P. Benner and V. Sima (2004). Evaluation of the linear matrix equation solvers in SLICOT. Tech. Rep. SLWN2004-1, Katholieke Universiteit Leuven (ESAT/SISTA), Leuven, Belgium, August, 32 pages.
- Van Huffel, S. and V. Sima (2002). SLICOT and control systems numerical software packages. In: *Proceedings of the 2002 IEEE International Conference on Control Applications and IEEE International Symposium on Computer Aided Control System Design, CCA/CACSD 2002*, Sept. 18–20, 2002, Glasgow, U.K., pp. 39–44. Omnipress.
- Van Huffel, S., V. Sima, A. Varga, S. Hammarling and F. Delebecque (2004). High-performance numerical software for control. *IEEE Control Systems Magazine* **24**, 60–76.