SOFTWARE FOR CONTROL SYSTEM ANALYSIS AND DESIGN, SINGULAR VALUE DECOMPOSITION

The singular value decomposition (SVD) goes back to the beginning of this century. In a paper of Beltrami (1) it was shown for the first time that any $n \times n$ matrix A can be diagonalized via orthogonal row and column transformations. For the more general case of an $n \times n$ complex matrix A, the result says that there exist unitary matrices U and V of dimension $n \times n$ and a *real* diagonal matrix $\Sigma = \text{diag}\{\sigma_1, \ldots, \sigma_n\}$ such that

$$A = U\Sigma V^*$$

where $\sigma_1 \geq \cdots \geq \sigma_n \geq 0$. If A is real then U and V are also real. This implies indeed that U and V "diagonalize" A since

$$U^*AV = \Sigma$$

The decomposition is in nature close to an eigenvalue decomposition, which was well known at the time. But this new decomposition is also very different since singular values are always positive real, whereas eigenvalues are in general complex. Also, the transformations are unitary in this decomposition, whereas in the eigenvalue decomposition they are just nonsingular, and hence can be quite "badly conditioned" (see the next section for a more detailed discussion). The use of this new decomposition was not apparent from the very beginning, but nowadays it has become an invaluable tool in several application areas such as statistics, signal processing, and control theory.

The first important property that was observed is the perturbation result for the singular values of a matrix A. If ΔA is a small perturbation of the matrix A, then its singular values σ_i are perturbed by an amount that can be bounded by the norm of the perturbation $\|\Delta A\|$. The fact that the sensitivity of singular values to perturbations is rather low makes them a good candidate for measuring certain variations in an observed phenomenon or in a model for it, and this is also its principal use in engineering applications. More formally, one can show that several matrix norms can actually be expressed in terms of its singular values. The most important ones are the 2-norm and the Frobenius norm:

$$\|A\|_2 \doteq \max_{x \neq 0} \frac{\|Ax\|_2}{\|x\|_2} = \sigma_1(A), \quad \|A\|_F \doteq \sqrt{\sum_{i,j} |a_{i,j}|^2} = \sqrt{\sum_i \sigma_i^2}$$

But the singular values are also used to measure angles. The well-known formula for the angle θ between two *real vectors* a and b of norm 1

$$a^{\mathrm{T}}a = 1, b^{\mathrm{T}}b = 1, \cos\theta \doteq |a^{\mathrm{T}}b|$$

extends to two spaces $\mathscr{A}$ and $\mathscr{B}$ of dimension m spanned by the orthonormal columns of two given matrices A and B:

$$A^*A = I_m, B^*B = I_m$$

The definition uses now the singular value decomposition of the "inner product"

$$A^*B = U\Sigma V^*$$

Indeed, if one takes $\tilde{A} = AU$ and $\tilde{B} = BV$, then the columns of $\tilde{A}$ and $\tilde{B}$ span the same spaces and still are orthonormal:

$$\tilde{A}^*\tilde{A} = I_m, \tilde{B}^*\tilde{B} = I_m$$

Moreover,

$$\tilde{A}^*\tilde{B} = \Sigma = \text{diag}\{\sigma_1, \ldots, \sigma_m\}$$

and from this diagonal form one can define the *canonical angles* between the spaces $\mathscr{A}$ and $\mathscr{B}$ as $\cos\theta_i = \sigma_i$, $i = 1, \ldots, m$.

The second important property is that the singular value decomposition yields a direct construction of "best" lower-rank approximation to a given matrix A. Let us indeed rewrite the SVD in its dyadic form:

$$A = \sum_{i=1}^{n} \sigma_i u_i v_i^*$$

Then the matrices

$$A_r = \sum_{i=1}^{r} \sigma_i u_i v_i^*$$

are rank r matrices and the error

$$\Delta A_r \doteq A - A_r = \sum_{i=r+1}^{n} \sigma_i u_i v_i^*$$

has obviously norm $\|\Delta A_r\|_2 = \sigma_{r+1}$, which is the minimal norm among all possible rank r approximations (see Ref. 2 and the section entitled "Robustness in Systems and Control"). This leads to the important concept of the ϵ_A rank of a matrix, defined in terms of the machine accuracy ϵ of the computer used, and the norm $\|A\|_2$ of the given matrix. For $\epsilon_A \doteq \epsilon\|A\|_2$ one defines the ϵ_A rank of A as the smallest rank r of A_r within ϵ_A distance of A. It turns out that this is the most reliable way to recognize rank deficiency of a given matrix and hence it is an important tool in engineering. In several applications the order of the system identified from observed measurements and the minimality of a constructed model indeed amount to rank determination problems (see the section entitled "Applications in Systems and Control").

A final important feature of the SVD is that it puts the matrix A in a diagonal form under orthogonal (or unitary) transformations of its columns and rows, and these transformations have good numerical properties. Interpreting $y = Ax$ as a mapping from a space $\mathscr{X}$ to a space $\mathscr{Y}$, we have thus transformed the coordinate systems in both spaces (by a well-behaved transformation) such that the mapping becomes diagonal. It is obvious that this coordinate system will reveal

special properties of the original mapping since now it is "decoupled" in a set of scalar equations. We will see that in the context of dynamical systems this coordinate system plays a fundamental role in what are called "balanced" realizations (see the section entitled "Balanced Realization").

THE SINGULAR VALUE DECOMPOSITION

Before we give more details about its basic properties, we first recall the singular value decomposition in its most general form.

Theorem 1. Let A be a $m \times n$ complex matrix of rank r. Then there exist unitary matrices U and V of dimensions $m \times m$ and $n \times n$, respectively, such that

$$A = U\Sigma V^* \tag{1}$$

where

$$\Sigma = \begin{bmatrix} \Sigma_r & 0 \\ 0 & 0 \end{bmatrix} \in \mathbb{R}^{m\times n} \tag{2}$$

and $\Sigma_r = \text{diag}\{\sigma_1, \ldots, \sigma_r\}$ with $\sigma_1 \geq \cdots \geq \sigma_r > 0$. If A is real then U and V are also real.

The numbers $\sigma_1, \ldots, \sigma_r$ together with $\sigma_{r+1} = 0, \ldots, \sigma_{\min(m,n)} = 0$ are called the *singular values* of A. The columns $\{u_i, i = 1, \ldots, m\}$ of U are called the *left singular vectors* of A and the columns $\{v_i, i = 1, \ldots, n\}$ of V are called the *right singular vectors* of A. They also appear in the dyadic decomposition of A:

$$A = \sum_{i=1}^{r} \sigma_i u_i v_i^* \tag{3}$$

which is nothing but an alternative way of writing $A = U\Sigma V^*$. Written in block form, this becomes

$$A = U_1 \Sigma_r V_1^* \tag{4}$$

where U_1 and V_1 are the submatrices of U and V, respectively, containing their first r columns. This decomposition can be viewed as a more compact form of writing Eqs. (1) and (2). It is also called a *rank factorization* of A since the factors have the rank r of A as at least one of their dimensions.

The proof of the preceding theorem is based on the eigendecomposition of the Hermitian matrices AA^* and A^*A. From the SVD one can indeed see that

$$AA^* = U\Sigma\Sigma^{\mathrm{T}}U^*, \quad A^*A = V\Sigma^{\mathrm{T}}\Sigma V^*$$

where $\Sigma\Sigma^{\mathrm{T}}$ and $\Sigma^{\mathrm{T}}\Sigma$ are clearly diagonal. Hence the left singular vectors are the eigenvectors of AA^*, the right singular vectors are the eigenvectors of A^*A, and the nonzero singular values are the square roots of the nonzero eigenvalues of both AA^* and A^*A. Deriving Theorem 1 from these connections is in fact quite simple, but we refer to Ref. 2 for the details.

NUMERICAL BACKGROUND

The importance of the SVD is strongly tied to numerical aspects. For this reason, we first give a very brief discussion of numerical stability and conditioning, which play a very important role in the study of numerical algorithms. For more details we refer to standard textbooks such as Refs. 3 and 4. We also choose the example of the singular value decomposition to introduce the relevant concepts.

Let A be an arbitrary $m \times n$ matrix. Then it is well known that there always exist unitary matrices U and V such that

$$U^* \cdot A = \begin{bmatrix} R \\ \overline{0} \end{bmatrix}, \; A \cdot V = \begin{bmatrix} C | 0 \end{bmatrix} \tag{5}$$

where R and C have, respectively, r linearly independent rows and columns. This implies, of course, that r is the rank of the matrix A. We call such transformations a row and column compression of the matrix A, respectively, and R and C are said to be of *full row rank* and *full column rank,* respectively. These decompositions can, for example, be computed with the singular value decomposition Eqs. (1) and (2). It is easy to verify that U^*A and AV yield, respectively, a row and a column compression of the matrix A. In this new coordinate system, the kernel and image of the map U^*AV are also simple to express. Indeed,

$$\text{Im}\, U^*AV = \text{Im} \begin{bmatrix} I_r \\ 0 \end{bmatrix}, \quad \text{Ker}\, U^*AV = \text{Im} \begin{bmatrix} 0 \\ I_{n-r} \end{bmatrix}$$

As a consequence we also have (2)

$$\text{Im}\, A = \text{Im}\, U \begin{bmatrix} I_r \\ 0 \end{bmatrix} = \text{Im}\, U_1, \quad \text{Ker}\, A = \text{Im}\, V \begin{bmatrix} 0 \\ I_{n-r} \end{bmatrix} = \text{Im}\, V_2$$

where U_1 is the submatrix of the first r columns of U and V_2 is the submatrix of the last $n - r$ columns of V. The computation of the preceding decomposition is, of course, subject to rounding errors. Denoting computed quantities by an overbar, we generally have, for some error matrix ΔA,

$$\overline{A} \doteq A + \Delta A = \overline{U}\,\overline{\Sigma}\,\overline{V}^* \tag{6}$$

Hence, the computed decomposition does not correspond exactly to the given matrix A but rather to a perturbed version $A + \Delta A$. When using the SVD algorithm available in the literature (5,6), this perturbation ΔA can be bounded by

$$\|\Delta A\| \leq c_A \epsilon_A = c_A \epsilon \|A\| \tag{7}$$

where ϵ is the machine accuracy and c_A is some known polynomial expression in m and n (3). Very often, this is a rough upper bound and one prefers to replace c_A by some statistical estimate $\hat{c}_A$, usually close to 1. The error ΔA induced by this algorithm—called the *backward error* because it is interpreted as an error on the *data*—thus has roughly the same norm as the input error Δ_{in} generated when reading in the data A in the computer. When such a bound exists for the perturbation ΔA induced by a numerical algorithm, it is called *backward stable.* We can make this definition more rigorous by considering a function $X = f(A)$ with data A and solution X.

If the computed solution $\overline{X} = \bar{f}(A)$ satisfies $\overline{X} = f(A + \Delta A)$ and $\|\Delta A\| \approx \epsilon\|A\|$, then the algorithmic implementation $\bar{f}(\cdot)$ of the function $f(\cdot)$ is said to be *backward stable*.

Notice that backward stability does not warrant any bounds on the errors in the results $\overline{U}$, $\overline{\Sigma}$, and $\overline{V}$. This depends indeed on how perturbations on the data (namely ΔA) affect the resulting decomposition (or the differences $\Delta U \doteq \overline{U} - U$, $\Delta\Sigma \doteq \overline{\Sigma} - \Sigma$, and $\Delta V \doteq \overline{V} - V$). In other words, it depends on the sensitivity of the function $f(\cdot)$ with input A and solution X. This sensitivity is commonly measured by the condition of f at A:

$$\kappa[f(A)] \doteq \lim_{\delta \to 0} \sup_{\|A-\overline{A}\|=\delta} \frac{\|X - \overline{X}\|}{\delta}, \quad X = f(A), \quad \overline{X} = f(\overline{A}) \tag{8}$$

Notice that we have not specified what norms are used in this definition, but in principle one can use different norms in the data and solution spaces (7). From this definition it is clear that the condition number $\kappa[f(A)]$ is some sort of "derivative" of the function $X = f(A)$ that we want to compute. When $\kappa[f(A)]$ is infinite, the problem of determining $X = f(A)$ from A is ill posed (as opposed to well posed). When $\kappa[f(A)]$ is finite and relatively large (or relatively small, the problem is said to be ill conditioned (or well conditioned). Further details can be found in Ref. 7.

It is important to note that backward stability is a property of an algorithm, while conditioning is associated with a problem and the specific data for that problem. The errors in the result depend on both the stability of the algorithm used and the conditioning of the problem solved. A good algorithm should therefore be backward stable since the size of the errors in the result is then mainly due to the condition of the problem and not due to the algorithm. An unstable algorithm, on the other hand, may yield a large error even when the problem is well conditioned.

We point out that if $f(A)$ has a Taylor expansion around A, then we can write

$$f(\overline{A}) = f(A) + \nabla x f(A)(X - \overline{X}) + O(\|X - \overline{X}\|^2) \tag{9}$$

Setting $X = f(A)$ and $\overline{X} = f(\overline{A})$ and taking norms, we then have

$$\|X - \overline{X}\| \le \|\nabla x f(A)\|\|A - \overline{A}\| + O(\|X - \overline{X}\|^2) \approx \kappa[f(A)]\|A - \overline{A}\|$$

This is a very powerful inequality, which indicates that forward errors $\|X - \overline{X}\|$ are bounded in norm by the sensitivity $\kappa[f(A)]$ and the backward error $\|A - \overline{A}\|$. Forward errors depend thus on two factors: the sensitivity of the problem and the backward error induced by the algorithm, and these two factors multiply each other in the preceding bound.

Bounds of the type in Eq. (7) are obtained by an error analysis of the algorithm used; see, for example, Ref. 8. The condition of the problem is obtained by a sensitivity analysis; see, for example, Refs. 4 and 8.

NUMERICAL ROBUSTNESS OF SINGULAR VALUES

One of the most important features of the singular value decomposition is that the singular values can be computed in a numerically reliable manner. There are indeed numerical methods available to compute the SVD in a backward stable manner, that is, such that the backward error ΔA satisfies Eqs. (6) and (7). Moreover, it is known that the singular values of any matrix A are in fact well conditioned:

$$\kappa[f\Sigma(A)] = 1 \tag{10}$$

This is derived from the variational properties of singular values (2,3) and leads to the following theorem.

Theorem 2. Let A be an arbitrary $m \times n$ complex matrix and ΔA an arbitrary perturbation of the same dimensions. Then the corresponding SVD of A and $A + \Delta A$ satisfy the following strict bounds:

$$\|\Sigma - \overline{\Sigma}\|_2 \le \|\Delta A\|_2, \quad \|\Sigma - \overline{\Sigma}\|_F \le \|\Delta A\|_F$$

This result is also proven using variational inequalities and we refer again to Ref. 2 for a proof. Suppose now that the computed singular values $\overline{\sigma}_i$ are such that

$$\overline{\sigma}_1 \ge \cdots \ge \overline{\sigma}_r > \epsilon_A \ge \overline{\sigma}_{r+1} \ge \cdots \ge \overline{\sigma}_{\min(m,n)}$$

Then it is reasonable to assume that

$$\sigma_{r+1} = \cdots = \sigma_{\min(m,n)} = 0$$

The level of accuracy of the computed singular values is indeed ϵ_A (which is a very small quantity) and hence only σ_1 to σ_r are guaranteed to be nonzero. While in general there may be little justification for setting the small ϵ_A singular values equal to zero, it is indeed a sound choice in several applications. A typical example is the determination of Im A and Ker A, which, for example, plays an important role in least-squares solutions of the equation $Ax = b$. Notice that orthonormal bases for Im A and Ker A are given by the columns of U_1 and V_2 defined earlier. The condition of Im A and of Ker A is thus connected to the sensitivity of the transformations U and V of the SVD. Consider, for example, the computation of Im A (where we assume $m > n$). As the distance function between two spaces $\mathscr{X}$ and $\mathscr{Y}$ we use the gap $\gamma(\mathscr{X}, \mathscr{Y}) \doteq \|P_{\mathscr{X}} - P_{\mathscr{Y}}\|_2$, where $P_{\mathscr{S}}$ is the orthogonal projector on the space $\mathscr{S}$. If A has full column rank n, then

$$\kappa[\text{Im } A] = \sigma_n^{-1} \tag{11}$$

as illustrated by the following example ($\sigma_1 = 1$, $\sigma_2 = a < 1$):

$$A \doteq \begin{bmatrix} 1 & 0 \\ 0 & a \\ 0 & 0 \\ 0 & 0 \end{bmatrix}; \qquad \Delta A \doteq \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & \delta \\ 0 & 0 \end{bmatrix};$$

$$\overline{A} \doteq A + \Delta A = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & c & -s & 0 \\ 0 & s & c & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 0 & \overline{a} \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$$

with $\overline{a} \doteq \sqrt{a^2 + \delta^2}$, $c \doteq a/\overline{a}$, $s \doteq \delta/\overline{a}$. The second basis vector of Im A is rotated in Im $\overline{A}$ over an angle θ, where $\sin\theta = s$,

and one easily checks that $\gamma(\text{Im } A, \text{Im } \overline{A}) = s$. Therefore

$$\lim_{\delta \to 0} \frac{\gamma(\text{Im } A, \text{Im } \overline{A})}{\delta} = \frac{1}{a} = \sigma_2^{-1}$$

The fact that the singular values have a low sensitivity to perturbations does of course not mean that every algorithm will compute them to high accuracy. The link with the eigenvalue problems AA^* and A^*A indicates that there can be no finite algorithm for computing singular values. An early iterative procedure was actually based on these connected Hermitian eigenvalue problems, but they have been shown to be unreliable because of the intermediate construction of "squared matrices" (3). It was shown in Refs. 5 and 6 that the unitary transformations U and V of the decomposition can be constructed via an iterative procedures that works directly on A to give the SVD. This algorithm first computes unitary matrices U_1 and V_1 such that $B = U_1^*AV_1$ is in bidiagonal form, that is, only the elements on its diagonal and first superdiagonal are non-zero. Then it uses an iterative procedure to compute unitary matrices U_2 and V_2 such that $U_2^*BV_2$ is diagonal and non-negative. The SVD defined in Eqs. (1) and (2) is then given by $\Sigma = U^*BV$, where $U = U_1U_2$ and $V = V_1V_2$. The computed U and V are unitary to approximately the working precision, and the computed singular values can be shown to be the exact σ_i's for $A + \Delta A$, where $\|\Delta A\|/\|A\|$ is a modest multiple of ϵ.

Other alternative methods to compute the singular values of a matrix A were proposed later and are based on Jacobi-like methods (9). They have been shown to have speed and accuracy comparable to the Golub-Kahan algorithm (10,11). As a consequence of the discussion in the preceding section, the singular values are thus computed with small absolute error. More recent results suggest that in particular cases the singular values of matrices can sometimes be computed to high *relative* accuracy as well (12–14). We finally remark that although the singular value decomposition is the most reliable method for determining the numerical rank of a given matrix, it is considerably more expensive than, for example, the QR factorization with column pivoting, which can usually give equivalent information with less computation (15).

APPLICATIONS IN SYSTEMS AND CONTROL

The problems considered in this article arise in the study of dynamical systems that can be modeled as state-space models:

$$\begin{aligned} \dot{x}(t) &= Ax(t) + Bu(t) \\ y(t) &= Cx(t) + Du(t) \end{aligned} \tag{12}$$

Here, $x(t)$ is an n-vector of states, $u(t)$ is an m-vector of controls or inputs, and $y(t)$ is an p-vector of outputs. The standard discrete-time analog of Eq. (12) takes the form

$$\begin{aligned} x_{k+1} &= Ax_k + Bu_k \\ y_k &= Cx_k + Du_k \end{aligned} \tag{13}$$

Typically the matrices A, B, C, and D are real. In case the results for real matrices are different we will explicitly state it.

In the subsections that follow, we survey a number of problems from systems and control theory that rely heavily on the singular value decomposition. We shall only discuss the numerical aspects here; for the system theoretical background, we refer the reader to the systems and control literature.

Impulse Response Realization

Let $H(z)$ be an $p \times m$ transfer function of a discrete-time causal system, and let its impulse response be given by

$$H(z) = \sum_{i=0}^{\infty} H_i z^{-i} \tag{14}$$

The realization problem is to find the transfer function $H(z)$ in state-space description,

$$H(z) = D + C(zI_n - A)^{-1}B \tag{15}$$

when the impulse response sequence $\{H_i\}$ is given. In the scalar case, this problem is related to the Padé approximation problem, for which fast methods exist (see Refs. 16 and 17 for a survey). In Ref. 16, it is shown that the Padé approach is in fact unstable, and it is better to consider a more general technique based on matrix decompositions of the Hankel matrix:

$$\boldsymbol{H}_{k,l} = \begin{bmatrix} H_1 & H_2 & \cdots & H_l \\ H_2 & \iddots & \iddots & \vdots \\ \vdots & \iddots & \iddots & \vdots \\ H_k & \cdots & \cdots & H_{k+l-1} \end{bmatrix} \tag{16}$$

Here k and l are upper bounds for the minimal dimension n of the state-space realization Eq. (15) of $H(z)$. From the expansion of $H(z)$ in powers of z^{-1} one finds that

$$H_0 = D, \quad H_i = CA^{i-1}B, \quad i = 1, \ldots, k + l - 1$$

and therefore $\boldsymbol{H}_{k,l}$ can be factorized as follows:

$$\boldsymbol{H}_{k,l} = \boldsymbol{C}_l \boldsymbol{O}_k$$

where

$$\boldsymbol{C}_l \doteq [B \quad AB \cdots \quad A^{l-1}B], \quad \boldsymbol{O}_k \doteq \begin{bmatrix} C \\ CA \\ \cdots \\ CA^{k-1} \end{bmatrix} \tag{17}$$

This implies that $\boldsymbol{H}_{k,l}$ has at most rank n and a simple argument proves that $\boldsymbol{H}_{k,l}$ will have exactly rank n. Since determining the order of the system requires a rank determination, it is natural to use here the SVD Eq. (4):

$$\boldsymbol{H}_{k,l} = U_1 \Sigma_n V_1^{\mathrm{T}}$$

For the construction of the triple $\{A, B, C\}$, let

$$L = U_1 \Sigma_n^{1/2}, \quad R = \Sigma_n^{1/2} V_1^{\mathrm{T}}, \quad \boldsymbol{H}_{k,l} = L \cdot R \tag{18}$$

and partition these "left" and "right" factors as follows:

$$L = \begin{bmatrix} C_1 \\ L_2 \end{bmatrix} = \begin{bmatrix} L_1 \\ C_2 \end{bmatrix}, \quad R = [B_1 \quad R_2] = [R_1 \quad B_2] \tag{19}$$

where C_1 and C_2 have p rows and B_1 and B_2 have m columns. From the two ways [Eqs. (17) to (19)] of writing the factorization of $\boldsymbol{H}_{k,l}$, one derives then that

$$C = C_1, \quad B = B_1$$

and that A can be solved from the overdetermined systems

$$L_1 A = L_2, \quad AR_1 = R_2$$

The particular choice of factors L and R makes the realization unique and we shall see that it is also linked to so-called *balanced realizations* later on (18). This realization algorithm based on the singular value decomposition of $\boldsymbol{H}_{k,l}$ was first given in Refs. 19 and 20.

Realization from Input/Output Data

Very often one does not have access to the impulse response of the system but only to a sequence of inputs $\{u_i\}$ and corresponding outputs $\{y_i\}$. In such cases a novel algorithm was derived in Ref. 21 based on the following Hankel matrix:

$$\boldsymbol{H}_{1:k,l} = \begin{bmatrix} z_1 & z_2 & \cdots & z_l \\ z_2 & \iddots & \iddots & \vdots \\ \vdots & \iddots & \iddots & \vdots \\ z_k & \cdots & \cdots & z_{k+l-1} \end{bmatrix}$$

where

$$z_i = \begin{bmatrix} u_i \\ y_i \end{bmatrix} \tag{20}$$

We start by noting that the problem would be much simpler if the sequence of states x_k would be known as well. From Eq. (13), rewritten as

$$\left[\begin{array}{c} x_{k+1} \\ \hline y_k \end{array}\right] = \left[\begin{array}{c|c} A & B \\ \hline C & D \end{array}\right] \left[\begin{array}{c} x_k \\ \hline u_k \end{array}\right] \tag{21}$$

one derives immediately the concatenated form:

$$\begin{bmatrix} x_2 & x_3 & \cdots & x_l \\ y_1 & y_2 & \cdots & y_{l-1} \end{bmatrix} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{bmatrix} x_1 & x_2 & \cdots & x_{l-1} \\ u_1 & u_2 & \cdots & u_{l-1} \end{bmatrix} \tag{22}$$

Under the assumption of persistence of excitation one shows that the right "data matrix" in Eq. (22) has full column rank $n + m$ and has thus a right inverse. Equivalently, Eq. (22) can be solved in a least-squares sense for the evolution matrix

$$\boldsymbol{E} = \begin{bmatrix} A & B \\ C & D \end{bmatrix} \tag{23}$$

So the problem is solved as soon as the states x_i are determined. But those depend on the choice of coordinates chosen for the state-space model. Replace indeed x_i by $\hat{x}_i = Tx_i$; then Eq. (22) becomes the related equation

$$\begin{bmatrix} \hat{x}_2 & \hat{x}_3 & \cdots & \hat{x}_l \\ y_1 & y_2 & \cdots & y_{l-1} \end{bmatrix} = \begin{bmatrix} \hat{A} & \hat{B} \\ \hat{C} & D \end{bmatrix} \begin{bmatrix} \hat{x}_1 & \hat{x}_2 & \cdots & \hat{x}_{l-1} \\ u_1 & u_2 & \cdots & u_{l-1} \end{bmatrix} \tag{24}$$

or

$$\begin{bmatrix} Tx_2 & Tx_3 & \cdots & Tx_l \\ y_1 & y_2 & \cdots & y_{l-1} \end{bmatrix} = \begin{bmatrix} TAT^{-1} & TB \\ CT^{-1} & D \end{bmatrix} \begin{bmatrix} Tx_1 & Tx_2 & \cdots & Tx_{l-1} \\ u_1 & u_2 & \cdots & u_{l-1} \end{bmatrix} \tag{25}$$

So each sequence of states

$$X_{1,l} = [x_1 \quad x_2 \quad \cdots \quad x_l] \tag{26}$$

can only be expected to be known up to an invertible row transformation corresponding to the particular coordinate system of the reconstructed model $\{A, B, C, D\}$. This row transformation T leaves the row space of $X_{1,l}$ unchanged. Also the rank condition for Eq. (22) to be solvable implies that Eq. (23) must be full rank n since this is a submatrix of the right-hand side matrix in Eq. (22). The row space of $X_{1,l}$ is therefore n-dimensional. This row space can now be found with the aid of the singular value decomposition based on the following theorem, proved in Ref. 21.

Theorem 3. Define

$$X_{k+1,l} \doteq [x_{k+1} \quad x_{k+2} \quad \cdots \quad x_{k+l}] \tag{27}$$

and the Hankel matrices $\boldsymbol{H}_{1:k,l}$ and $\boldsymbol{H}_{k+1:2k,l}$ as before, then

$$\operatorname{Im}[\boldsymbol{X}_{k+1,l}^{\mathrm{T}}] = \operatorname{Im}[\boldsymbol{H}_{1:k,l}^{\mathrm{T}}] \cap \operatorname{Im}[\boldsymbol{H}_{k+1:2k,l}^{\mathrm{T}}] \tag{28}$$

provided the input/output (I/O) data are persistently exciting and $k \geq n$, $l \geq (m + p)k$.

In practice, due to perturbations on the data, the row spaces typically do not intersect. An approximate intersection, using the singular value decomposition or some rank revealing QR decomposition, has thus to be constructed. A possible implementation of this idea is the following decomposition:

$$\left[\begin{array}{c} \boldsymbol{H}_{1:k,l} \\ \hline \boldsymbol{H}_{k+1:2k,l} \end{array}\right] = \left[\begin{array}{c|c} I & 0 \\ \hline 0 & Q \end{array}\right] \left[\begin{array}{cc|c} A_{11} & A_{12} & 0_1 \\ A_{21} & 0_3 & 0_2 \\ \hline \times & \times & A_{33} \end{array}\right] V^{\mathrm{T}} \tag{29}$$

where

$[A_{11}A_{12}]$ has full column rank equal to the rank of $\boldsymbol{H}_{1:k,l}$ [which equals $r = mk + n$ under the assumption of persistence of excitation (21)]

A_{33} has full row rank, which must be smaller than r if an intersection is to be detected

A_{21} has full column rank equal to the dimension of the intersection, hence n

$\times$ denotes an arbitrary matrix

The order in which this decomposition is constructed is as follows. First the transformation V^{T} is constructed to compress the columns of $\boldsymbol{H}_{1:k,l}$, yielding the trailing zero matrix 0_1. Then the rows of the trailing bottom matrix are compressed with the transformation Q, yielding 0_2 and a full row rank A_{33}. Then V^{T} is updated to yield the full column rank matrix A_{21} and the trailing zero matrix 0_3. Notice that all three steps involve a rank factorization that can be done with the singular value decomposition (or any other rank revealing factorization). The center matrix in this decomposition has a form that trivially displays the intersection of row spaces of the top and bottom parts, namely,

$$\operatorname{Im}\begin{bmatrix} A_{11}^{\mathrm{T}} \\ A_{12}^{\mathrm{T}} \\ 0 \end{bmatrix} \cap \operatorname{Im}\begin{bmatrix} A_{21}^{\mathrm{T}} & \times \\ 0 & \times \\ 0 & A_{33}^{\mathrm{T}} \end{bmatrix} = \operatorname{Im}\begin{bmatrix} I_n \\ 0 \\ 0 \end{bmatrix}$$

Because of the transformation V^{T} in Eq. (29) one derives that

$$\operatorname{Im}(\boldsymbol{H}_{1:k,l}^{\mathrm{T}}) \cap \operatorname{Im}(\boldsymbol{H}_{k+1;2k,l}^{\mathrm{T}}) = V \operatorname{Im}\begin{bmatrix} I_n \\ 0 \\ 0 \end{bmatrix}$$

that is, the first n rows of V^{T} are a representation of $X_{k+1,l}$. From this we can now construct $\{A, B, C, D\}$ as explained in Eqs. (22) to (25).

Balanced Realization

In the preceding section we pointed out that the realization problem from an impulse response or from input/output data is only defined up to a state-space transformation T, which in principle can be chosen arbitrarily. Is there a particular coordinate system that should be chosen for some reason, and if so, how can we construct it?

We develop here the concept of balanced realization that is based on the singular value decomposition and has several appealing properties. For this we first need to define the controllability Gramian G_{c} and observability Gramian G_{o} of a system. For the continuous-time system equation (12) these are defined as follows:

$$G_{\mathrm{c}}(T) \doteq \int_0^T (e^{At}B)(e^{At}B)^{\mathrm{T}}\,dt, \quad G_{\mathrm{o}}(T) \doteq \int_0^T (Ce^{At})^{\mathrm{T}}(Ce^{At})\,dt \tag{30}$$

while for the discrete-time systems they are defined as follows:

$$G_{\mathrm{c}}(K) \doteq \sum_{k=0}^{K-1} (A^iB)(A^iB)^{\mathrm{T}}, \quad G_{\mathrm{o}}(K) \doteq \sum_{k=0}^{K-1} (CA^i)^{\mathrm{T}}(CA^i) \tag{31}$$

An intuitive interpretation (18) is that the controllability Gramian G_{c} measures the amount of energy needed to control the states of the system using an input sequence of a certain duration, while the observability Gramian G_{o} measures the amount of energy one can observe from the state in an output sequence of a particular duration. Typically K and T are chosen to be infinite if the system is stable, because then these Gramians can be computed efficiently using Lyapunov equations (22). For a given realization $\{A, B, C, D\}$ one easily shows that these matrices are positive semidefinite and that they allow one to detect whether or not the realization is minimal by computing the rank of these matrices (18). This suggests the use of the SVD of both G_{c} and G_{o} in order to find minimal realizations of the system, provided a given system $\{A, B, C, D\}$ is not necessarily minimal (see also the next section).

But here we want to focus on another application of the SVD in this context. If one applies a state-space transformation $x \to \hat{x} \doteq Tx$ then the system triple transforms to $\{\hat{A}, \hat{B}, \hat{C}, D\} \doteq \{TAT^{-1}, TB, CT^{-1}, D\}$ and the Gramians to $\hat{G}_{\mathrm{c}} \doteq TG_{\mathrm{c}}T^{\mathrm{T}}$ and $\hat{G}_{\mathrm{o}} \doteq T^{-\mathrm{T}}G_{\mathrm{o}}T^{-1}$. One shows then that T can be chosen such that both new Gramians are equal and diagonal:

$$TG_{\mathrm{c}}T^{\mathrm{T}} = \hat{G}_{\mathrm{c}} = \Sigma = \hat{G}_{\mathrm{o}} = T^{-T}G_{\mathrm{o}}T^{-1} \tag{32}$$

which is exactly the balanced coordinate system. In order to construct T one starts from the Cholesky factorization of G_{c} and G_{o}:

$$G_{\mathrm{c}} = L_1L_1^{\mathrm{T}}, \quad G_{\mathrm{o}} = L_2^{\mathrm{T}}L_2 \tag{33}$$

where L_1 and L_2 are both lower triangular. One then computes the singular value decomposition of the upper triangular matrix $L_1^{\mathrm{T}}L_2^{\mathrm{T}}$:

$$L_2^{\mathrm{T}}L_1^{\mathrm{T}} = U\Sigma V^{\mathrm{T}} \tag{34}$$

using, for example, the efficient algorithm described in Ref. 23. Then, defining $T = \Sigma^{1/2}U^{\mathrm{T}}L_1^{-1}$ and $T^{-1} = L_2^{-1}V\Sigma^{1/2}$, one checks that

$$T^{-1}T = L_1^{-1}(V\Sigma U^{\mathrm{T}})L_2^{-1} = L_1^{-1}(L_1L_2)L_2^{-1} = I \tag{35}$$

and

$$\hat{G}_c = TG_cT^{\mathrm{T}} = \Sigma^{1/2}U^TL_1^{-1}(L_1L_1^{\mathrm{T}})L_1^{-\mathrm{T}}U\Sigma^{1/2} = \Sigma$$
$$\hat{G}_o = T^{-\mathrm{T}}G_oT^{-1} = \Sigma^{1/2}V^{\mathrm{T}}L_2^{-\mathrm{T}}(L_2^{\mathrm{T}}L_2)L_2^{-1}V\Sigma^{1/2} = \Sigma$$

In this new coordinate system one can associate the diagonal element σ_i of Σ with the unit vectors e_i (zero everywhere except a 1 in position i): σ_i gives then the energy needed to control the state e_i as well as the energy observed from this state. Since both these are equal the state-space system is said to be balanced (see Ref. 18 for more details).

It turns out that in addition to an "energy interpretation," this coordinate system has useful properties. If a single-input/single-output system is stable and if one uses the infinite-horizon Gramians, then the evolution matrix of the balanced realization

$$\hat{\boldsymbol{E}} \doteq \left[\begin{array}{c|c} \hat{A} & \hat{b} \\ \hline \hat{c} & d \end{array}\right] \tag{36}$$

is symmetric up to some sign changes, that is, there exists a diagonal matrix S of ± 1 such that $\hat{\boldsymbol{E}}S = S\hat{\boldsymbol{E}}^T$. As a consequence of this, one shows that such systems have low sensitivity to roundoff propagation (24). The sign symmetry is not preserved for multi-input/multi-output systems or when considering finite-horizon Gramians, but nevertheless the robustness properties are (25).

While the singular values can be interpreted as a measure for the energy that is being transferred from the input space to the state space as well as from the state space to the output space, they also play an important role in constructing approximate models that "preserve" this energy as well as possible. Since the singular values are ordered decreasingly one can partition the balanced evolution matrix as follows:

$$\left[\begin{array}{c|c} \hat{A} & \hat{B} \\ \hline \hat{C} & D \end{array}\right] \doteq \left[\begin{array}{cc|c} \hat{A}_{11} & \hat{A}_{12} & \hat{B}_1 \\ \hat{A}_{21} & \hat{A}_{22} & \hat{B}_2 \\ \hline \hat{C}_1 & \hat{C}_2 & D \end{array}\right] \tag{37}$$

and the subsystem $\{\hat{A}_{11}, \hat{B}_1, \hat{C}_1, D\}$ will be a good approximation of the original system in the sense that it keeps the states that were responsible for the largest part of the energy transfer (i.e., the largest singular values). A more formal approximation measure is the so-called $\mathscr{H}_\infty$ norm, which can be bounded as a function of the balanced singular values if the balancing was based on infinite-horizon Gramians. This so-called *balanced truncation* problem is also related to that of optimal approximation in the Hankel norm (26,27) and has led to a renewed interest in the partial realization problem and related topics (28).

Finally, we point out here that the realization algorithm described in the section entitled "Impulse Response Realization" for discrete-time systems in fact constructs immediately a balanced realization. Let the matrices $\boldsymbol{C}_l$ and $\boldsymbol{O}_k$ be defined as in Eq. (17). Then

$$\begin{aligned} \boldsymbol{C}_l \boldsymbol{C}_l^{\mathrm{T}} &= \sum_{i=0}^{l-1} (A^i B)(A^i B)^{\mathrm{T}} \doteq G_{\mathrm{c}}(l), \\ \boldsymbol{O}_k^{\mathrm{T}} \boldsymbol{O}_k &= \sum_{i=0}^{k-1} (CA^i)^{\mathrm{T}} (CA^i) \doteq G_{\mathrm{o}}(k) \end{aligned} \tag{38}$$

and from the choice of factorization, Eq. (18), it follows that

$$\boldsymbol{C}_l = \sum\nolimits_n^{1/2} V_1^{\mathrm{T}}, \quad \boldsymbol{O}_k = U_1 \sum\nolimits_n^{1/2} \tag{39}$$

Since U_1 and V_1 have orthonormal columns, one obtains

$$\begin{aligned} \boldsymbol{C}_l \boldsymbol{C}_l^{\mathrm{T}} &= \sum\nolimits_n^{1/2} V_1^{\mathrm{T}} V_1 \sum\nolimits_n^{1/2} = \sum\nolimits_n, \\ \boldsymbol{O}_k^{\mathrm{T}} \boldsymbol{O}_k &= \sum\nolimits_n^{1/2} U_1^{\mathrm{T}} U_1 \sum\nolimits_n^{1/2} = \sum\nolimits_n \end{aligned} \tag{40}$$

Controllability and Observability

The concepts of controllability and observability play a fundamental role in systems and control theory. A system in the form of Eq. (12) is indeed a minimal representation of the input/output behavior of the system if and only if it is both controllable and observable. The same holds for a discrete-time system, Eq. (13), except that there one talks about *reachability* rather then controllability. The conditions for controllability (reachability) and observability are equivalent to the following rank conditions, respectively (29),

$$\operatorname{rank} \boldsymbol{C}_n = n, \quad \operatorname{rank} \boldsymbol{O}_n = n \tag{41}$$

In principle one could use the SVD to check these rank conditions, but it is not recommended to construct these matrices since they contain powers of the matrix A, which could lead to a considerable buildup of rounding errors (see Ref. 30). It is well known (see e.g., Ref. 29) that the infinite-horizon Gramians, $G_{\mathrm{c}}(\infty)$ and $G_{\mathrm{o}}(\infty)$, have the same rank as $\boldsymbol{C}_n$ and $\boldsymbol{O}_n$, respectively. Since the latter can be computed as the solution of Lyapunov equations, this seems a possible alternative, but it turns out to be a sensitive roundabout as well (31). A third way is to find an appropriate coordinate system for the pair $\{A, B\}$ or $\{A, C\}$ such that the rank of the matrices in Eq. (41) becomes apparent. Since observability and controllability are dual to each other we discuss this for controllability only. The following theorem proposes such a coordinate change, which is orthogonal and based on a succession of singular value decompositions (31).

Theorem 4. There always exists an orthogonal state-space transformation U such that $[U^{\mathrm{T}}B \| U^{\mathrm{T}}AU] \doteq [\hat{B} \| \hat{A}]$ has the form

$$\begin{aligned} [\hat{B} \| \hat{A}] &= \left[\begin{array}{c|c\|c|c} B_{\mathrm{c}} & A_{\mathrm{c}} & \times \\ \hline 0 & 0 & A_{\bar{\mathrm{c}}} \end{array}\right] \\ &\doteq \left[\begin{array}{c\|ccccc|c} X_1 & A_{1,1} & A_{1,2} & \cdots & A_{1,k} & A_{1,k+1} \\ 0 & X_2 & A_{2,2} & & \vdots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \vdots & \vdots \\ \vdots & & \ddots & X_k & A_{k,k} & A_{k,k+1} \\ \hline 0 & \cdots & \cdots & 0 & 0 & A_{k+1,k+1} \end{array}\right] \end{aligned} \tag{42}$$

where $A_{i,i}$, $i = 1, \ldots, k$, are $r_i \times r_i$ matrices, and X_i, $i = 1, \ldots, k$, are $r_i \times r_{i-1}$ matrices of full row rank r_i (with $r_0 \doteq m$).

The matrices X_i are constructed recursively as the result or row compressions, using, for example, the SVD (31). In this new coordinate system, one easily sees that the controllability matrix $\hat{\boldsymbol{C}}_n = U^{\mathrm{T}} \boldsymbol{C}_n$ has the form

$$\begin{aligned} \hat{\boldsymbol{C}}_n &= [\hat{B} \quad \hat{A}\hat{B} \ldots \hat{A}^{n-1}\hat{B}] \\ &= \begin{bmatrix} X_{1:1} & \times & \cdots & \times & \times & \cdots & \times \\ 0 & X_{1:2} & & \vdots & \vdots & & \vdots \\ \vdots & 0 & \ddots & \vdots & \vdots & & \vdots \\ \vdots & \vdots & & X_{1:k} & \times & \cdots & \times \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 0 \end{bmatrix} \end{aligned} \tag{43}$$

where $X_{1:i}$ is the product $X_1 \ldots X_i$. Since these products have full row rank r_i by construction, the factorization $\boldsymbol{C}_n = U\hat{\boldsymbol{C}}_n$ has a second factor that is row compressed, and the result thus follows. The *controllable* subspace is defined as the space spanned by the columns of $\boldsymbol{C}_n$. It follows from Eq. (43) that this space has dimension $n_{\mathrm{c}} = \sum_{i=1}^{k} r_i$ and that

$$\operatorname{Im} \hat{\boldsymbol{C}}_n = \operatorname{Im} \begin{bmatrix} I_{n_{\mathrm{c}}} \\ 0 \end{bmatrix}$$

Since $\boldsymbol{C}_n = U\hat{\boldsymbol{C}}_n$ this space is spanned in the original coordinate system $\{A, B\}$ by the first n_{c} columns of U. The matrix pair $\{A_{\mathrm{c}}, B_{\mathrm{c}}\}$ is shown to be controllable, and the eigenvalues

of $A_{\bar{c}}$ are called the uncontrollable modes of the pair $\{A, B\}$. The indices r_i are the so-called controllability indices of the pair $\{A, B\}$ (31).

A dual form of the staircase algorithm applied to $\{A, C\}$ yields a dual result:

$$\left[\begin{array}{c} V^{\mathrm{T}}AV \\ \hline\hline CV \end{array}\right] = \left[\begin{array}{c|c} A_{\bar{o}} & \times \\ \hline 0 & A_{\mathrm{o}} \\ \hline\hline 0 & C_{\mathrm{o}} \end{array}\right] \tag{44}$$

where the subsystem $\{A_{\mathrm{o}}, C_{\mathrm{o}}\}$ is observable. If $n_{\bar{o}}$ is the number of rows or columns of $A_{\bar{o}}$ then the first $n_{\bar{o}}$ columns of V span Ker $\boldsymbol{O}_n$, which is called the *unobservable subspace* of the pair $\{A, C\}$.

Combining the preceding decompositions, one constructs an orthogonal state-space transformation yielding a transformed system $\{\hat{A}, \hat{B}, \hat{C}, D\}$ which has the form

$$\left[\begin{array}{c|c} \hat{A} & \hat{B} \\ \hline \hat{C} & D \end{array}\right] = \left[\begin{array}{ccc|c} A_{11} & A_{12} & A_{13} & B_1 \\ 0 & A_{22} & A_{23} & B_2 \\ 0 & 0 & A_{33} & 0 \\ \hline 0 & C_2 & C_3 & D \end{array}\right] \tag{45}$$

and where the subsystem $\{A_{22}, B_2, C_2, D\}$ is minimal, that is, both observable and controllable (reachable). Moreover, the transfer functions of $\{\hat{A}, \hat{B}, \hat{C}, D\}$ and $\{A_{22}, B_2, C_2, D\}$ are equal. This form is closely related to the Kalman decomposition and for its construction we refer to Ref. 31.

In addition to the controllable and unobservable subspaces of a system, there are other spaces that play a fundamental role in the control of systems modeled as Eqs. (12) and (13). Two other fundamental objects in the so-called geometric system theory (32) are the supremal (A, B)-invariant and controllability subspaces contained in a given subspace. As shown in Refs. 31 and 33 they can also be computed via a matrix recurrence based on a sequence of SVDs constructing an orthogonal basis for the relevant spaces. The role of the SVD in these "staircase" algorithms is not only the reliable rank determination of the subsequent steps, but at the same time the singular values allow one to assess the sensitivity of the computed bases (31).

Robustness in Systems and Control

In the last decade, there has been a significant growth in the theory and techniques of robust control. These developments mainly center around two concepts: H_∞ (34) and the structured singular value (35). They both provide a framework for synthesizing robust controllers for linear systems, in the sense that they achieve a desired system performance in the presence of a significant amount of uncertainty in the system.

In this section, we first focus on H_∞ techniques. The H_∞ norm of a stable rational transfer matrix $H(s)$ (continuous-time) or $H(z)$ (discrete-time) is defined as

$$\|H(s)\|_\infty \doteq \sup_{\omega\in\mathbb{R}} \sigma_{\max}[H(j\omega)], \quad \|H(z)\|_\infty \doteq \sup_{\omega\in[0,2\pi]} \sigma_{\max}[H(e^{j\omega})] \tag{46}$$

where $\sigma_{\max}[\cdot]$ denotes the largest singular value of a (complex) matrix. We explain how this quantity comes about by starting from a basic robustness problem. Consider the homogeneous systems

$$\dot{x} = Ax, \quad x_{k+1} = Ax_k \tag{47}$$

which are assumed to be stable against C_s, that is, the eigenvalues of the matrix A are in a region C_s of the complex plane, which is the open left-half plane for a continuous-time system and the open unit disk for a discrete-time system, respectively.

The complex stability radius measures the robustness of system stability for complex perturbations. This radius is defined as the norm of the smallest complex perturbation Δ such that the perturbed matrix $A + B\Delta C$ becomes unstable (where $B \in \mathbb{C}^{n\times m}$, $C \in \mathbb{C}^{p\times n}$, and hence $\Delta \in \mathbb{C}^{m\times p}$). For $A + B\Delta C$ to be unstable, it must have at least one eigenvalue in the complement of C_s. It is important to note that although C_s is convex, the set of all C_s-stable matrices $\mathrm{St} = \{M\colon \Lambda(M) \subset C_s\}$, where $\Lambda(M)$ denotes the spectrum of M, is nonconvex, as well as its complement Unst of C_s-unstable matrices. The stability radius $r_{\mathbb{C}}$ therefore measures the distance of a stable matrix A to the nonconvex set Unst.

By continuity of the spectrum of a matrix versus perturbations on its entries, the stability radius is clearly equal to the distance from a stable matrix A to an optimal matrix $A + B\Delta C$ lying on ∂Unst. Indeed, when a matrix $A + B\Delta C$ passes the boundary ∂Unst, at least one of its eigenvalue must also cross ∂C_s. The boundary ∂Unst in the matrix space describes matrices with at least one eigenvalue in ∂C_s. Therefore, the stability radius can be written as

$$r_{\mathbb{C}}(A,B,C) = \inf\{\|\Delta\|_2 \colon \Lambda(A + B\Delta C) \cap \partial C_s \neq \emptyset\}$$

Consider a parametrization of the boundary ∂C_s by a real variable ω, such as $\partial C_s = \{j\omega, \omega \in \mathbb{R}\}$ or $\partial C_s = \{e^{j\omega}, \omega \in [0, 2\pi]\}$. The stability radius can then be rewritten as

$$\begin{aligned} r_{\mathbb{C}}(A,B,C) &= \inf_{\lambda\in\partial C_s} [\inf\{\|\Delta\|_2 \colon \det(\lambda I - A - B\Delta C) = 0\}] \\ &= \inf_{\lambda\in\partial C_s} [\inf\{\|\Delta\|_2 \colon \det(I - \Delta C(\lambda I - A)^{-1}B) = 0\}] \end{aligned} \tag{48}$$

the second equality resulting from the stability of the initial matrix A ($\lambda I - A$ is invertible for $\lambda \in \partial C_s$) and from the fact that $\det(I + XY) = 0 \Leftrightarrow \det(I + YX) = 0$. The following classical result allows us to simplify this expression considerably and is based on the singular value decomposition $M = U\Sigma V^*$.

Theorem 5. One has the inequality

$$\inf_{\Delta\in C^{n\times n}} \{\|\Delta\|_2 \colon \det(I - \Delta M) = 0\} \geq \sigma_{\max}^{-1}[M]$$

If there are no constraints on Δ, the bound is attained for $\hat{\Delta} \doteq v_{\max}\sigma_{\max}^{-1}u_{\max}^*$.

Combining this with Eq. (48), it follows that

$$\begin{aligned} r_{\mathbb{C}}(A,B,C) &= \inf_{\lambda\in\partial C_s} \sigma_{\max}^{-1}[C(\lambda I - A)^{-1}B] \\ &= \Big(\sup_{\lambda\in\partial C_s} \sigma_{\max}[C(\lambda I - A)^{-1}B]\Big)^{-1} \end{aligned}$$

which is the H_∞ norm inverse of the rational transfer function $H(\lambda) \doteq C(\lambda I - A)^{-1}B$. This is therefore reduced to a nonconvex optimization problem on ∂C_s, parametrized by a real parameter ω.

Efficient iterative methods are available for computing this norm (36,37), and are based on the relationship between the singular values of $H(j\omega)$ and the imaginary eigenvalues of a Hamiltonian matrix obtained from a state-space realization of $H(\lambda)$ (38). This result is then used to develop a quadratically convergent algorithm for computing the H_∞ norm of a transfer function.

Structured Singular Values

In Theorem 5, it was stated that the lower bound can actually be met when there are no constraints on the perturbation Δ. But Δ will be in general complex, since the matrix M is in general complex [even for transfer functions $H(\lambda)$ with real coefficient matrices A, B and C]. The problem becomes more involved when one imposes Δ to be real. It was shown in Ref. 39 that

$$r_{\mathbb{R}}(A,B,C) = \left(\sup_{\lambda \in \partial C_s} \mu_{\mathbb{R}}[H(\lambda)] \right)^{-1} \tag{49}$$

where, for any $M \in \mathbb{C}^{p \times m}$,

$$\mu_{\mathbb{R}}(M) = \inf_{\gamma \in (0,1]} \sigma_2 \left(\begin{bmatrix} \operatorname{Re} M & -\gamma \operatorname{Im} M \\ \gamma^{-1} \operatorname{Im} M & \operatorname{Re} M \end{bmatrix} \right) \tag{50}$$

This becomes now an optimization problem in two real parameters γ and ω, but in Ref. 39 it is shown that the function Eq. (50) is unimodal in γ. In Ref. 40 an efficient algorithm is then derived to compute this *real* stability radius based on ideas of Ref. 36.

The computation of the real stability radius can be considered a special case of the more general structured stability radius. Structured singular values (35) have been introduced to provide a perturbation bound for structured uncertainties in control system analysis and design. Therefore, the structured singular value approach can be viewed as a complement for the H_∞ approach. In a linear system with multiple independent norm-bounded perturbations, it is always possible by rearranging the system to isolate the perturbations as a single large block-diagonal perturbation Δ. Then, denoting the transfer function from the collective outputs of the perturbation to their inputs by $M(s)$, the stability problem reduces to ensuring that $\det(I - \Delta M) \neq 0$ at all frequencies and for all allowable Δ. Notice that Δ is again not arbitrary anymore and therefore Theorem 5 does not apply. The largest singular value of that theorem has thus to be replaced by the so-called largest structured singular value, which is defined as follows. Let the set of allowable perturbations be denoted by $\mathscr{D} \subseteq \mathbb{C}^{n \times n}$ and be defined as

$$\begin{aligned} \mathscr{D} = \{ \Delta = \text{ block diag}(&\delta_1^r I_{k_1}, \ldots, \delta_p^r I_{k_p}, \delta_1^c I_{k_{p+1}}, \ldots, \delta_q^c I_{k_{p+q}}, \\ &\Delta_1^C, \ldots, \Delta_r^C) : \delta_i^r \in \mathbb{R},\ \delta_i^c \in \mathbb{C},\ \Delta_i^C \in \mathbb{C}^{c_i \times c_i} \} \end{aligned} \tag{51}$$

The largest structured singular value of an $n \times n$ complex matrix M is then defined as

$$\mu_{\mathscr{D}}(M) = \begin{cases} 0 & \text{if } \det(I - \Delta M) \neq 0 \quad \text{for all } \Delta \in \mathscr{D} \\ \left(\min_{\{\Delta \in \mathscr{D} : \det(I - \Delta M) = 0\}} \sigma_{\max}(\Delta) \right)^{-1} & \text{otherwise} \end{cases} \tag{52}$$

Computing $\mu_{\mathscr{D}}(M)$ is a difficult numerical problem. It is a nonconvex optimization problem and its complexity can be nonpolynomial for certain norms (see Ref. 41). One approach, which is computationally rather demanding, is to formulate the problem as a nondifferentiable convex optimization problem involving the maximum singular value of a matrix obtained from M. A more efficient scheme is given in Ref. 42 and uses several smooth optimization problems that do not involve any eigenvalue or singular value computations. The computational complexity of the problem of computing $\mu_{\mathscr{D}}(M)$ has prompted several researchers to look for bounds that are easier to compute (43,44).

CONCLUDING REMARKS

In this paper we have given several uses of the singular value decomposition in analysis and design problems of systems and control. We have considered computational issues and useful properties of this decomposition as well, such as diagonalization, norms, and sensitivity. The list given here is far from complete. Closeness problems (45,46) and canonical forms (31,47) are just a few examples. We expect the number of applications to grow also in the future because of the serious interdisciplinary effort that is under way between the communities of the numerical linear algebra field on the one hand and of the systems and control field on the other hand.

BIBLIOGRAPHY

1. E. Beltrami, Sulle funzioni bilineari, *Giornali di Mathematiche,* **11**: 98–106, 1873.
2. P. Lancaster and M. Tismenetsky, *The Theory of Matrices,* 2nd ed., New York: Academic Press, 1985.
3. G. H. Golub and C. F. Van Loan, *Matrix Computations,* 2nd ed., Baltimore, MD: Johns Hopkins Univ. Press, 1989.
4. N. Higham, *Accuracy and Stability of Numerical Algorithms,* Philadelphia, PA: SIAM, 1995.
5. G. H. Golub and W. Kahan, Calculating the singular values and pseudo-inverse of a matrix, *SIAM J. Numer. Anal.,* **2**: 205–224, 1965.
6. G. H. Golub and C. Reinsch, Singular value decomposition and least squares solutions, *Numer. Math.,* **14**: 403–420, 1970.
7. J. R. Rice, A theory of condition, *SIAM J. Numer. Anal.,* **3**: 287–310, 1966.
8. J. H. Wilkinson, *The Algebraic Eigenvalue Problem,* Oxford, England: Oxford University Press, 1965.
9. E. G. Kogbetliantz, Solution of linear equations by diagonalization of coefficient matrix, *Quart. Appl. Math.,* **13**: 123–132, 1955.
10. C. Paige and P. Van Dooren, On the quadratic convergence of Kogbetliatz's algorithm for computing the singular value decomposition, *Lin. Alg. Appl.,* **77**: 301–313, 1986.
11. J.-P. Charlier and P. Van Dooren, On Kogbetliantz's SVD algorithm in the presence of clusters, *Lin. Alg. Appl.,* **95**: 135–160, 1987.
12. J. Demmel and W. Kahan, Accurate singular values of bidiagonal matrices, *SIAM J. Sci. Stat. Comput.,* **11**: 873–912, 1990.
13. V. Fernando and B. Parlett, Accurate singular values and differential *qd* algorithms, *Numer. Math.,* **67**: 191–229, 1994.
14. G. Golub, K. Sølna, and P. Van Dooren, Computing the Singular Values of Products and Quotients of Matrices, in M. Mooren and B. DeMoor (eds.), *SVD in Signal Processing III, Algorithms, Architectures and Applications,* Amsterdam: Elsevier, 1995.

15. T. F. Chan, Rank revealing QR factorizations, *Lin. Alg. Appl.,* **88/89**: 67–82, 1987.
16. L. S. de Jong, Towards a formal definition of numerical stability, *Numer. Math.,* **28**: 211–220, 1977.
17. A. Bultheel, Recursive algorithms for the matrix Padé problem, *Math. Comput.,* **35**: 875–892, 1980.
18. B. C. Moore, Principal component analysis in linear systems: Controllability, observability, and model reduction, *IEEE Trans. Autom. Control.,* **AC-26**: 17–31, 1981.
19. H. Zeiger and A. McEwen, Approximate linear realizations of given dimension via Ho's algorithm, *IEEE Trans. Autom. Control,* **AC-19**: 153, 1974.
20. S. Kung, A New Identification and Model Reduction Algorithm via Singular Value Decompositions, *Proc. 12th Asilomar Conf. Circuits Syst. Comp.,* 1978, pp. 705–714.
21. M. Moonen et al., On- and off-line identification of linear state-space models, *Int. J. Control,* **49**: 219–232, 1989.
22. R. Patel, A. Laub, and P. Van Dooren, *Numerical Linear Algebra Techniques for Systems and Control,* Piscataway, NJ: IEEE, 1993.
23. M. T. Heath et al., Computing the singular value decomposition of a product of two matrices, *SIAM J. Sci. Stat. Comput.,* **7**: 1147–1159, 1986.
24. C. Mullis and R. Roberts, Synthesis of minimum roundoff noise fixed point digital filters, *IEEE Trans. Circuit Syst.,* **CAS-23**: 551–562, 1976.
25. M. Gevers and G. Li, *Parametrizations in Control, Estimation and Filtering Problems,* London: Springer, 1993.
26. V. Adamjan, D. Arov, and M. Krein, Analytic properties of Schmidt pairs for a Hankel operator and the generalized Schur-Takagi problem, *Mat. USSR Sbornik,* **15**: 31–73, 1971.
27. K. Glover, All optimal Hankel norm approximations of linear multivariable systems, and their L_∞ error bounds, *Int. J. Control,* **39**: 1115–1193, 1984.
28. A. Antoulas, New results on the algebraic theory of linear systems: The solution of the cover problems, *Lin. Alg. Appl.,* **50**: 1–45, 1983.
29. T. Kailath, *Linear Systems,* Englewood Cliffs, NJ: Prentice-Hall, 1980.
30. C. Paige, Properties of numerical algorithms related to computing controllability, *IEEE Trans. Autom. Control,* **AC-26**: 130–138, 1981.
31. P. Van Dooren, The generalized eigenstructure problem in linear system theory, *IEEE Trans. Autom. Control,* **AC-26**: 111–129, 1981.
32. W. M. Wonham, *Linear Multivariable Control: A Geometric Approach,* 2nd ed., New York: Springer, 1979.
33. V. C. Klema and A. J. Laub, The singular value decomposition: Its computation and some applications, *IEEE Trans. Autom. Control,* **AC-25**: 164–176, 1980.
34. G. Zames, Feedback and optimal sensitivity: Model reference transformations, multiplicative semi-norms, and approximate inverses, *IEEE Trans. Autom. Control,* **AC-26**: 301–320, 1981.
35. J. C. Doyle, Analysis of feedback systems with structured uncertainties, *IEE Proc., Pt. D,* **129**: 242–250, 1982.
36. S. Boyd, V. Balakrishnan, and P. Kabamba, A bisection method for computing the H_∞ norm of a transfer matrix and related problems, *Math. Control Signals Syst.,* **2**: 207–219, 1989.
37. N. A. Bruinsma and M. Steinbuch, A fast algorithm to compute the H_∞ norm of a transfer matrix, *Syst. Control Lett.,* **14**: 287–293, 1990.
38. R. Byers, A bisection method for measuring the distance of a stable matrix to the unstable matrices, *SIAM J. Sci. Stat. Comput.,* **9**: 875–881, 1988.
39. L. Qiu et al., A formula for computation of the real stability radius, *Automatica,* **31**: 879–890, 1995.
40. J. Sreedhar, P. Van Dooren, and A. L. Tits, A Level-Set Idea to Compute the Real Hurwitz-Stability Radius, *Proc. 34th IEEE Conf. Dec. Control,* New Orleans, 1995, WA05.
41. O. Toker and H. Ozbay, On the NP-Hardness of the Purely Complex μ Computation, Analysis/Synthesis, and Some Related Problems in Multi-Dimensional Systems, *Proc. 1995 Amer. Control Conf.,* 1995, pp. 447–451.
42. M. K. H. Fan and A. L. Tits, Characterization and efficient computation of the structured singular value, *IEEE Trans. Autom. Control,* **AC-31**: 734–743, 1986.
43. P. M. Young, M. P. Newlin, and J. C. Doyle, Practical Computation of the Mixed μ Problem, *Proc. 1992 Amer. Control Conf.,* Chicago, 1992, pp. 2190–2194.
44. C. T. Lawrence, A. L. Tits, and P. Van Dooren, A Fast Algorithm for the Computation of an Upper Bound on the μ-norm, *Proc. 13th IFAC World Congr.,* San Francisco, 1996, vol. H, pp. 59–64.
45. R. Eising, Between controllable and uncontrollable, *Syst. Control Lett.,* **4**: 263–264, 1984.
46. D. Boley, Computing rank-deficiency of rectangular matrix pencils, *Syst. Control Lett.,* **9**: 207–214, 1987.
47. M. M. Konstantinov, P. Hr. Petkov, and N. D. Christov, Invariants and canonical forms for linear multivariable systems under the action of orthogonal transformation groups, *Kybernetika,* **17**: 413–424, 1981.

PAUL VAN DOOREN
Université Catholique de Louvain,
Centre for Systems Engineering
and Applied Mechanics