

Numerical Linear Algebra for Model Reduction in Control and Simulation

Peter Benner

Professur Mathematik in Industrie und Technik

Fakultät für Mathematik

Technische Universität Chemnitz

D-09107 Chemnitz

Germany

`benner@mathematik.tu-chemnitz.de`

Abstract

Model reduction is an ubiquitous tool in analysis and simulation of dynamical systems, control design, circuit simulation, structural dynamics, CFD, etc. In the past decades many approaches have been developed for reducing the order of a given model. Often these methods have been derived in parallel in different disciplines with particular applications in mind. We will discuss some of the most prominent methods used for linear systems, compare their properties and highlight similarities. In particular, we will emphasize the role of recent developments in numerical linear algebra in the different approaches. Efficiently using these new techniques, the range of applicability of some of the methods has considerably widened.

Keywords. Model reduction, modal truncation, Padé approximation, balanced truncation, Lyapunov equation, Krylov subspace method.

MSC (2000). 65F30,93B11,41A20,65F50.

1 Introduction

Consider the linear, time-invariant (LTI) system

$$\begin{aligned}\dot{x}(t) &= Ax(t) + Bu(t), & t > 0, & \quad x(0) = x^0, \\ y(t) &= Cx(t) + Du(t), & t \geq 0,\end{aligned}\tag{1}$$

where $A \in \mathbb{R}^{n \times n}$ is the state matrix, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{p \times n}$, $D \in \mathbb{R}^{p \times m}$, and $x^0 \in \mathbb{R}^n$ is the initial state of the system. Here, n is the order (or state-space dimension) of the system and $x(t) \in \mathbb{R}^n$, $y(t) \in \mathbb{R}^p$, $u(t) \in \mathbb{R}^m$ are the state, output and input of the system, respectively. In some application areas like structural dynamics, only the differential equation in (1) is used to describe the model dynamics while in other areas like control or circuit simulation, the system description provided in (1) almost always contains the (algebraic) output equation. If the output equation is not present in the mathematical model used to describe the investigated physical process, one might simply set $y(t) = x(t)$, i.e., $C = I_n$, $D = 0$, if a method is to be

used that needs the C and D matrices. But often it is also natural in these applications to define specific variables that can serve as outputs as the complete state is seldom measurable in practice.

Applying the Laplace transform to (1) under the assumption that $x(0) = 0$, we obtain a set of algebraic equations from which an input-to-output mapping can be defined as follows:

$$Y(s) = (C(sI - A)^{-1}B + D)U(s),$$

where s is the Laplace variable and Y, U are the Laplace transforms of y and u , respectively. The associated transfer function matrix (TFM) then is

$$G(s) = C(sI - A)^{-1}B + D. \quad (2)$$

Note that any *state-space transformation* $x \mapsto Tx$, $T \in \mathbb{R}^{n \times n}$ nonsingular, yielding a new system description via

$$(A, B, C, D) \mapsto (TAT^{-1}, TB, CT^{-1}, D) \quad (3)$$

leaves the dynamics of the system and its transfer function invariant as can be seen from

$$(CT^{-1})(sI - TAT^{-1})^{-1}(TB) + D = C(sI_n - A)^{-1}B + D = G(s).$$

Therefore, there exist infinitely many matrix quadruples (A, B, C, D) representing the same LTI system. Each element of the associated equivalence class is called a *realization* of the LTI system. It is easy to see that there exist realizations of (1) of arbitrarily high order, but there is a lower limit on the order n of the system. This number is called the *McMillan degree* of the system and will be denoted here by \hat{n} . A realization of (1) of order $n = \hat{n}$ is called a *minimal realization*. In the model reduction methods discussed throughout this paper, we will use several specific realizations of LTI systems.

The model reduction problem now consists of finding a reduced-order LTI system,

$$\begin{aligned} \dot{\hat{x}}(t) &= \hat{A}\hat{x}(t) + \hat{B}u(t), & t > 0, & \quad \hat{x}(0) = \hat{x}^0, \\ \hat{y}(t) &= \hat{C}\hat{x}(t) + \hat{D}u(t), & t \geq 0, & \end{aligned} \quad (4)$$

of order r , $r \ll n$, with the same number of inputs m , the same number of outputs p , and associated TFM $\hat{G}(s) = \hat{C}(sI - \hat{A})^{-1}\hat{B} + \hat{D}$, so that for the same input function $u \in L_2(0, \infty; \mathbb{R}^m)$, we have $y(t) \approx \hat{y}(t)$, or, in frequency domain, $Y(s) \approx \hat{Y}(s)$. Employing the Paley-Wiener theorem (Parseval's identity) and the operator norm induced by the 2-norm in the frequency domain \mathcal{L}_2 , defined for real-rational TFMs by

$$\|G\|_\infty := \sup_{\omega \in \mathbb{R}} \sigma_{\max}(G(j\omega)) \quad (j := \sqrt{-1}, \quad \sigma_{\max} = \text{maximum singular value}), \quad (5)$$

the approximation error can be quantified as

$$\|y - \hat{y}\|_2 = \|Y - \hat{Y}\|_2 \leq \|G - \hat{G}\|_\infty \|U\|_2 = \|G - \hat{G}\|_\infty \|u\|_2. \quad (6)$$

Here, $\|\cdot\|_2$ denotes the 2-norm either in the input space $L_2(0, \infty; \mathbb{R}^m)$ or the frequency domain \mathcal{L}_2 .

Note that model reduction of discrete-time LTI systems (i.e., linear systems where the dynamics is driven by difference equations) can be formulated in an analogous manner using Z- instead of Laplace transformations; see, e.g., [57].

Model (order) reduction is a common task within the simulation, control, and optimization of complex physical processes. Often, large systems arise due to accuracy requirements on the spatial discretization of control problems for fluids or structures described by partial differential equations, in the context of lumped-circuit approximations of distributed circuit elements, such as the interconnect or package of VLSI chips. or in simulations of micro-electro-mechanical systems (MEMS), which have both electrical and mechanical components, and many other areas. Dimension reduction is generally required for purposes of computational feasibility and/or storage reduction.

Various reduction techniques have been devised, but many of these are described in terms that are discipline-oriented or even application-specific even though they share many common features and origins. See the recent monographs and surveys [2, 3, 11, 6, 28, 57, 67]. In case of linear systems, it seems that three approaches play the most prominent role, these are

1. modal truncation and the related techniques of substructuring and static condensation,
2. Padé and Padé-type approximations, and
3. balancing-related truncation techniques.

All three approaches rely on efficient numerical linear algebra techniques to be applicable to very large scale problems with state-space dimensions of order in the thousands or even in the millions. We will discuss the basic principles of all three approaches and the related techniques from numerical linear algebra. It is well-known that the first two approaches listed above can be applied to very large-scale problems, see, e.g., [6, 17, 26, 28]. In contrast to common belief, it is also possible to apply balanced truncation techniques for large-scale problems. As this is based on relatively new developments, we will spend a little more effort to explain the underlying new techniques from numerical linear algebra. Further information on these developments can be found in [5, 9, 10, 12, 42, 48, 49, 60, 61].

The rest of this paper is structured as follows: in Sections 2, 3, and 4, we review the three approaches listed above, i.e., model truncation, Padé approximation, and balanced truncation. A comparison of the three concepts from a theoretical point of view is given in Section 5 while Section 6 provides some numerical examples for the application of the different approaches to some benchmark problems from several application areas. Moreover, a fairly extensive bibliography is provided which can be used as a basis for further reading.

2 Modal Approaches

Modal methods are based on catching the behaviour of the dominant modes of the system in the reduced-order model. In the language of linear algebra, this means that the state-space is projected on the subspace spanned by the eigenvectors (“eigenforms”) corresponding to specified eigenvalues of the state matrix A (or equivalently, poles of $G(s)$). Modal truncation is probably one of the oldest model reduction techniques [21, 24, 51] and still in use, mainly in structural dynamics. It is for instance available in the popular software packages ANSYS [1] or Nastran [54]. The eigenvalues and -vectors can be computed by subspace iteration methods which is often used in structural dynamics packages. Usually, more efficient algorithms for large-scale problems result from utilizing modern methods of sparse numerical linear algebra like Krylov subspace method (Lanczos or Arnoldi iterations) or the Jacobi-Davidson algorithm [7, 17, 47, 68].

An obvious, though certainly not always optimal, choice of dominant modes is to select those eigenvalues of A having nonnegative or small negative real values. Basically, these eigenvalues dominate the long-term dynamics of the solution of the linear ordinary differential equation describing the dynamics of (1)—solution components corresponding to large negative real parts decay rapidly and mostly play a less important (negligible) role in vibration analysis or control design. This viewpoint is rather naive as it neither takes into account the transient behaviour of the dynamical system nor the oscillations caused by large imaginary parts or the sensitivity of the eigenvalues with respect to small perturbations. Nevertheless, this approach is often successful when A comes from an FEM analysis of an elliptic operator as the ones arising in linear elasticity or heat transfer processes. Here, a large range of eigenvalues are insignificant as they result only from discretizing the differential operator and approximate a continuous part of its spectrum. The corresponding modes can safely be neglected as they have no impact on the dynamic response of the system.

An advantage of modal truncation is that the poles of the reduced-order system (eigenvalues of \hat{A}) are poles of the original system (eigenvalues of A). This is important in applications like vibration analysis as the modes correspond to the resonance frequencies of the original system; the most important resonances are then retained in the reduced-order model. On the other hand, the exact location of eigenvalues plays a less important role in the analysis of dynamical systems.

The error of modal truncation can easily be quantified. The case of a diagonal matrix A is treated in [38, Lemma 9.2.1] which is trivially generalized to the following result.

Theorem 2.1 *Suppose A in (1) is stable and unitarily diagonalizable as*

$$A = U \begin{bmatrix} \Lambda_1 & \\ & \Lambda_2 \end{bmatrix} U^H,$$

where $U \in \mathbb{C}^{n \times n}$ is unitary and $\Lambda_1 \in \mathbb{C}^{r \times r}$, $\Lambda_2 \in \mathbb{C}^{(n-r) \times (n-r)}$ are diagonal matrices with Λ_1 containing the desired eigenvalues $\lambda_1, \dots, \lambda_r$. If $U =: [U_1 \ U_2]$ is partitioned accordingly and the reduced-order model (4) is computed as

$$\hat{A} = U_1^H A U_1, \quad \hat{B} = U_1^H B, \quad \hat{C} = C U_1, \quad \hat{D} = D,$$

then the reduced-order model satisfies

$$\|G - \hat{G}\|_\infty \leq \|C_2\|_2 \|B_2\|_2 \max_{\lambda \in \{\lambda_{r+1}, \dots, \lambda_n\}} \frac{1}{|\operatorname{Re}(\lambda)|}, \quad (7)$$

where $B_2 = U_2^H B$, $C_2 = C U_2$.

Proof. With $G(s) = C U (sI - U^H A U)^{-1} U^H B + D$ and the construction of the reduced-order model it follows immediately that

$$G(s) - \hat{G}(s) = C_2 (sI - \Lambda_2)^{-1} B_2.$$

Now, as A is assumed to be stable, we have no purely imaginary poles so that

$$\begin{aligned} \|G - \hat{G}\|_\infty &= \max_{\omega \in \mathbb{R}} \|C_2 (j\omega I - \Lambda_2)^{-1} B_2\|_2 \\ &\leq \|B_2\|_2 \|C_2\|_2 \max_{\omega \in \mathbb{R}} \|(j\omega I - \Lambda_2)^{-1}\|_2 \\ &= \|B_2\|_2 \|C_2\|_2 \max_{\omega \in \mathbb{R}} \frac{1}{\min_{r+1 \leq j \leq n} |j\omega - \lambda_j|}, \end{aligned}$$

where the 2-norms in the right-hand side are the matrix 2-norms (or spectral norms). The error bound now follows by observing that the maximum of the right-hand side is taken for the eigenvalue closest to the imaginary axis. \square

The error bound (7) is computable if the full matrix U (or all eigenvectors of A) is (are) available. Otherwise, the following bound can be used if the eigenvalues in Λ_1 are selected as the eigenvalues of A with smallest real parts and ordered as

$$|\operatorname{Re}(\lambda_1)| \leq |\operatorname{Re}(\lambda_2)| \leq \dots \leq |\operatorname{Re}(\lambda_r)| \leq |\operatorname{Re}(\lambda_{r+\ell})| \quad \forall \ell = 1, \dots, n - r.$$

By noting that $\|B_2\|_2 \leq \|B\|_2$, $\|C_2\|_2 \leq \|C\|_2$, and $|\operatorname{Re}(\lambda_j)| \geq |\operatorname{Re}(\lambda_r)|$, we get

$$\|G - \hat{G}\|_\infty \leq \|C\|_2 \|B\|_2 \frac{1}{|\operatorname{Re}(\lambda_r)|}. \quad (8)$$

In case A is not unitarily diagonalizable, the situation becomes slightly more complicated. If A is diagonalizable by a non-unitary similarity transformation with $T \in \mathbb{C}^{n \times n}$ nonsingular, then the same argument as in the proof of Theorem 2.1 yields

$$\|G - \hat{G}\|_\infty \leq \operatorname{cond}_2(T) \|C_2\|_2 \|B_2\|_2 \max_{\lambda \in \{\lambda_{r+1}, \dots, \lambda_n\}} \frac{1}{|\operatorname{Re}(\lambda)|}, \quad (9)$$

where $T^{-1}AT = \Lambda$ is the spectral decomposition of A and $\operatorname{cond}_2(T)$ is the spectral norm condition number of the eigenvector matrix T . But note that for state matrices A that are not unitarily diagonalizable, the eigenvector basis may be ill-conditioned. Therefore, it is advisable to use an orthogonal basis for the A -invariant subspace spanned by the desired eigenvectors.

This brief error analysis shows that minimization of the error bounds for modal truncation corresponds to keeping the eigenvalues with small imaginary parts in the reduced-order model which is in full accordance with the engineering significance of the corresponding modes. More advanced dominance measures have been suggested which also take into account the influence of the j th mode on the input-output behaviour of the LTI system. For a discussion of several dominance measures and some enhancement strategies see [69]. But note that all suggested dominance measures require the (block-)diagonalization of A —a task which is usually infeasible for really large-scale problems unless the model description directly leads to an A matrix in such a special form.

In order to better approximate the static behaviour of a system, it is often advised to include static solutions in the reduced-order model. The static response of the LTI system (1) corresponds to setting $\dot{x}(t) = 0$. If A is invertible, we therefore get

$$x(t) = -A^{-1}Bu(t).$$

In the model reduction process this can be taken into account by *base enrichment*. One possibility is to simply add $\tilde{U} := -A^{-1}B$ to U_1 so that the reduced-order model becomes

$$\hat{A} = V^H AV, \quad \hat{B} = V^H B, \quad \hat{C} = CV, \quad \hat{D} = D,$$

where $V = [U_1 \quad \tilde{U}]$. Somewhat more sophisticated combinations of static reduction methods (also called *Guyan reduction* [43]) and modal truncation lead to the Craig-Bampton method [21] which is the basis of modern *substructuring* and *component mode synthesis (CMS)* methods, see, e.g., [17, 64] and the references therein. Also note that in the context of substructuring and CMS methods, new dominance measures have been derived in [34, 73] for applications in structural dynamics.

3 Padé Approximation

In this section we consider systems of the form

$$E\dot{x}(t) = Ax(t) + Bu(t), \quad y(t) = Cx(t), \quad (10)$$

with corresponding transfer function $G(s) = C(sE - A)^{-1}B$. This formulation slightly generalizes (1) in allowing a non-identity matrix $E \in \mathbb{R}^{n \times n}$ in front of the derivative \dot{x} which is not necessarily invertible. (Note that modal truncation as well as balanced truncation can easily be adapted to this situation, too.) On the other hand, we assume $D = 0$. The latter is not a restriction as otherwise, the system can be re-written in the form (1) with a slightly enlarged generalized state-space, see, e.g., [23].

For $s_0 \notin \Lambda(A, E)$, where $\Lambda(A, E)$ denotes the set of generalized eigenvalues of the matrix pencil $A - \lambda E$, we can re-write the transfer function of (10) and expand it into a power (Laurent) series as follows:

$$\begin{aligned} G(s) &= C((s_0E - A) + (s - s_0)E)^{-1}B \\ &= C(I - (s - s_0)(s_0E - A)^{-1}E)^{-1}(s_0E - A)^{-1}B \\ &= M_0 + M_1(s - s_0) + M_2(s - s_0)^2 + \dots \end{aligned}$$

For $s_0 = 0$, the coefficient matrices in the power series are given by $M_j := C(A^{-1}E)^j A^{-1}B$ and are called *moments*. In case $s_0 = \infty$ we have $M_j := C(E^{-1}A)^j E^{-1}B$, called the *Markov parameters* of the system.

As reduced-order model we can now use the r th Padé approximant \hat{G} to G , defined by:

$$G(s) = \hat{G}(s) + \mathcal{O}((s - s_0)^{2r}),$$

i.e., $M_j = \widehat{M}_j$ for $j = 0, \dots, 2r - 1$. The corresponding computational problem is called *moment matching* if $s_0 < \infty$ and *partial realization* if $s_0 = \infty$.

For $m = p = 1$, the solution of the moment matching problem is given by

$$\hat{G}(s) = \frac{\alpha_{r-1}s^{r-1} + \alpha_{r-2}s^{r-2} + \dots + \alpha_1s + \alpha_0}{\beta_r s^r + \beta_{r-1}s^{r-1} + \dots + \beta_1s + 1},$$

where the coefficients of the numerator and denominator polynomials can be computed via the so-called *Asymptotic Waveform Evaluation (AWE)* procedure: solve the linear system of equations

$$M \begin{bmatrix} \beta_r \\ \vdots \\ \beta_1 \end{bmatrix} = \begin{bmatrix} M_r \\ \vdots \\ M_{2r-1} \end{bmatrix},$$

with the *Hankel matrix*

$$M = \begin{bmatrix} M_0 & M_1 & M_2 & \dots & M_{r-1} \\ M_1 & M_2 & & \ddots & M_r \\ M_2 & & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \\ M_{r-1} & M_r & \dots & & M_{2r-2} \end{bmatrix}.$$

Then, with $\beta_0 := 1$ we get $\alpha_j = \sum_{k=0}^j M_k \beta_{j-k}$. Despite the fact that there are efficient methods for solving linear systems for Hankel matrices, see, e.g., [19], AWE is numerically not feasible due to the unavoidable large errors caused by explicitly computing the moments M_j .

A breakthrough for Padé approximation methods was the observation that the moments need not be computed explicitly as moment matching is equivalent to projecting the state-space onto

$$\mathcal{V} = \text{span}\{\tilde{B}, \tilde{A}\tilde{B}, \dots, \tilde{A}^{r-1}\tilde{B}\} =: \mathcal{K}(\tilde{A}, \tilde{B}, r), \quad (11)$$

where $\tilde{A} = (s_0 E - A)^{-1} E$, $\tilde{B} = (s_0 E - A)^{-1} B$, along

$$\mathcal{W} = \text{span}\{C^H, \tilde{A}^H C^H, \dots, (\tilde{A}^H)^{r-1} C^H\} =: \mathcal{K}(\tilde{A}^H, C^H, r). \quad (12)$$

Here, $\mathcal{K}(M, b, \ell)$ denotes the *Krylov subspace*

$$\text{span}\{b, Mb, M^2b, \dots, M^{\ell-1}b\}.$$

Bi-orthogonal bases of these two (block-)Krylov subspaces can be computed via the unsymmetric (block, band) Lanczos method [28, 29, 30]. In addition, a realization of the reduced-order model is computed as a by-product of the unsymmetric Lanczos process. It can be shown that for $n \geq m, p$, the corresponding reduced-order model satisfies

$$G(s) - \hat{G}(s) = \mathcal{O}((s - s_0)^{q(n)}), \quad q(n) \geq \left\lfloor \frac{n}{m} \right\rfloor + \left\lfloor \frac{n}{p} \right\rfloor.$$

The computation of a reduced-order model based on the moment matching idea implemented as unsymmetric Lanczos process is called the *(matrix) Padé-via-Lanczos ((M)PVL) method* [28, 29]. PVL is popular in circuit simulation and micro electronics as it can be applied without changes to descriptor systems which quite commonly arise in these applications if $s_0 \notin \Lambda(A, E)$ and $s_0 < \infty$. (By a descriptor system we denote a system in the form (10) where E is a singular matrix.) There exist several variants of PVL and MPVL adapted to special situations arising in circuit simulation.

Similarly, the ideas described above can be applied for matching the Markov parameters, i.e., $s_0 = \infty$. This problem is also called *partial realization* and the implementation as model reduction method using projection on Krylov subspaces for non-singular E is discussed in [39, 45] while the situation for $s_0 = \infty$ is treated in [16].

Still there are several difficulties with the PVL method.

1. Despite several efforts ([8, 41]), so far there exists in general no computable error estimate or bound for $\|y - \hat{y}\|_2$.
2. The choice of the expansion points s_0 is usually based on heuristic criteria. (In circuit simulation often the operating frequency of the circuit can be used which explains the particular success in this application area.)
3. The reduced-order model provides a good approximation quality only locally.
4. The preservation of physical properties like stability or passivity can only be shown in very special cases; usually some post processing which (partially) destroys the moment matching properties is required. Such a method based on PVL or MPVL for RCL circuits is discussed in [27]; the resulting SyPVL/SyMPVL algorithm computes stable and passive reduced-order models and matches at least half as many moments as the PVL and MPVL methods.

Padé-type methods are also based on the moment matching property, but the approximations need not match the maximum possible number of moments, that is for approximating the transfer function it is only required that

$$G(s) - \hat{G}(s) = \mathcal{O}((s - s_0)^{\tilde{q}(n)}), \quad \tilde{q}(n) \leq q(n).$$

In this sense, the SyPVL and SyMPVL methods are Padé-type methods. Other Padé-type methods result from computing bases of the Krylov subspaces (11) and (12) with the Arnoldi process rather than the Lanczos method. This has the advantage of producing orthogonal bases and therefore of improved numerical stability properties. A band (block) Arnoldi process for model reduction which is based on (11) is devised in [27, 28]. It is also shown there that

$$\tilde{q}(n) \geq \left\lfloor \frac{n}{m} \right\rfloor.$$

A special case of this procedure for RCL circuits is known under the name PRIMA in circuit simulation [58]. Due to the symmetry properties of the realizations describing RCL circuits, this procedure is mathematically equivalent to SyMPVL [27, 28]. Other Krylov subspace methods for model reduction include [40, 44, 45]. A two-sided Arnoldi method that computes bases for both Krylov subspaces in (11)–(12) is suggested in [22]; this approach can be interpreted as a Padé approximation method.

Current research in Padé-type approximation methods include aspects of (multi-point) rational and tangential interpolation [3, 31, 32, 33, 41] which allow to obtain good local approximation properties at several frequencies s_0, \dots, s_k . The projection matrices needed to compute reduced-order models by these approximation techniques can be computed from shift-and-invert, or, more generally, rational Krylov subspace methods [65].

Summarizing this section, the development of efficient Krylov subspace methods based on the Lanczos and Arnoldi algorithms in the past two decades has led to new model reduction techniques based on rational approximation. Without this numerical linear algebra machinery, there would be no way to apply Padé or Padé-like methods to the large-scale problems for which reduced-order models are needed.

4 Balanced Truncation

The basic idea of balanced truncation is to compute a special realization of an LTI system which is called a *balanced realization*. Such a realization is defined by the property that the controllability and the observability Gramians W_c and W_o , respectively, of the LTI system are diagonal and equal, i.e., they satisfy

$$W_c = W_o = \hat{\Sigma} = \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_n \end{bmatrix},$$

where $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n > 0$. The values σ_j are called the *Hankel singular values* of the LTI system. In general they are defined as the square roots of the eigenvalues of the product $W_c W_o$ and are system invariants as they are independent of the particular realization of the LTI system. Note that the Gramians W_c and W_o are given by the solutions of the two dual *Lyapunov equations*

$$AW_c + W_c A^T + BB^T = 0, \quad A^T W_o + W_o A + C^T C = 0, \quad (13)$$

with A, B, C as in (1). Suppose we know a state-space transformation (defined via $T \in \mathbb{R}^{n \times n}$ nonsingular) which transforms the given realization of the LTI system via

$$(TAT^{-1}, TB, CT^{-1}, D) = \left(\begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}, \begin{bmatrix} B_1 \\ B_2 \end{bmatrix}, [C_1 \ C_2], D \right) \quad (14)$$

to a balanced realization, where $A_{11} \in \mathbb{R}^{r \times r}$, $B_1 \in \mathbb{R}^{r \times m}$, $C_1 \in \mathbb{R}^{p \times r}$ with $r < n$. Then the basic idea is to use as the reduced-order model the truncated realization

$$(\hat{A}, \hat{B}, \hat{C}, \hat{D}) = (A_{11}, B_1, C_1, D). \quad (15)$$

This idea goes essentially back to [53, 55]. Collecting results from [53, 35], we can summarize the properties of balanced truncation.

Proposition 4.1 *Let (A, B, C, D) be a minimal realization of a stable LTI system with transfer function $G(s)$ and let $(\hat{A}, \hat{B}, \hat{C}, \hat{D})$ with associated transfer function \hat{G} be computed as in (14)–(15). Then the following holds:*

a) *The reduced-order system \hat{G} is balanced, minimal, and stable. Its Gramians are*

$$\hat{W}_c = \hat{W}_o = \begin{bmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_r \end{bmatrix} =: \hat{\Sigma}.$$

b) *The absolute error bound*

$$\|G - \hat{G}\|_\infty \leq 2 \sum_{k=r+1}^{\hat{n}} \sigma_k =: \delta_{BT}. \quad (16)$$

holds and implies

$$\|y - \hat{y}\|_2 = \|Y - \hat{Y}\|_2 \leq \delta_{BT} \|U\|_2 = \delta_{BT} \|u\|_2,$$

where $\|\cdot\|_2$ denotes the L_2 -norm in either state-space or frequency domain and \hat{Y} is the Laplace transform of the output \hat{y} of the reduced-order model.

Of particular importance is the error bound (16) as it allows an adaptive choice of the order of the reduced-order model based on a prescribed tolerance threshold for the approximation quality. (The error bound (16) can be improved in the presence of Hankel singular values with multiplicity greater than one—they need to appear only once in the sum on the right-hand side.)

It is easy to check that for a controllable and observable (minimal) system, i.e., a system with nonsingular Gramians, a balancing state-space transformation can be defined via

$$T = \Sigma^{\frac{1}{2}} U^T R^{-T}, \quad (17)$$

where $W_c = R^T R$ (e.g., a Cholesky factorization) and $U \Sigma^2 U^T$ is a singular value or spectral decomposition of the positive definite matrix $R W_o R^T$. A nice observation in [66, 46] allows

to compute (15) also for non-minimal systems without the need to compute the full matrix T . The first part of this observation is that for $W_o = S^T S$,

$$S(W_c W_o)S^{-1} = (SR^T)(SR^T)^T = (U\Sigma V^T)(V\Sigma U^T) = U\Sigma^2 U^T$$

so that U, Σ can be computed from an SVD of SR^T ,

$$SR^T = \begin{bmatrix} U_1 & U_2 \end{bmatrix} \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix} \begin{bmatrix} V_1^T \\ V_2^T \end{bmatrix}, \quad \begin{aligned} \Sigma_1 &= \text{diag}(\sigma_1, \dots, \sigma_r), \\ \Sigma_2 &= \text{diag}(\sigma_{r+1}, \dots, \sigma_n), \end{aligned} \quad (18)$$

where $\sigma_r > \sigma_{r+1}$. The second part needed is the fact that computing

$$T_l = \Sigma_1^{-1/2} V_1^T R, \quad T_r = S^T U_1 \Sigma_1^{-1/2}, \quad (19)$$

and

$$\hat{A} := T_l A T_r, \quad \hat{B} := T_l B, \quad \hat{C} := C T_r \quad (20)$$

is equivalent to first computing a minimal realization of (1), then balancing the system as in (14) with T as in (17), and finally truncating the balanced realization as in (15). In particular, the realizations obtained in (15) and (20) are the same, T_l contains the first r rows of T and T_r the first r columns of T^{-1} —those parts of T needed to compute A_{11}, B_1, C_1 in (14). Also note that the product $T_r T_l$ is a projector onto an r -dimensional subspace of the state-space and model reduction via (20) can therefore be seen as projecting the dynamics of the system onto this subspace. For non-minimal systems, choosing $r = \hat{n}$, this approach can also be used to compute a minimal realization of the LTI system.

The algorithm resulting from (20) is often referred to as the *SR method* for balanced truncation. In [46, 66] and all textbooks treating balanced truncation, S and R are assumed to be the (square, triangular) Cholesky factors of the system Gramians. In [13] it is shown that everything derived so far remains true if instead of Cholesky factors, full-rank factors of the system Gramians are used. That is, in the factorizations $W_c = R^T R$, $W_o = S^T S$, the factors satisfy $R \in \mathbb{R}^{\text{rank}(W_c) \times n}$, $S \in \mathbb{R}^{\text{rank}(W_o) \times n}$ and hence may be rectangular if the Gramians are rank-deficient. This yields a much more efficient implementation of balanced truncation whenever $\hat{n} \ll n$ (numerically). Low numerical rank of the Gramians usually results from a rapid decay of their eigenvalues as shown in Figure 1 and implies a rapid decay of the Hankel singular values. Partial analysis of this phenomenon is given in [4, 36, 62]. An algorithm exploiting this fact by directly computing numerical full-rank factors is derived in [13], see also [12, 14].

It is often stated that balanced truncation is not suitable for large-scale problems as it requires the solution of two Lyapunov equations, followed by an SVD and that both steps require $\mathcal{O}(n^2)$ storage and $\mathcal{O}(n^3)$ flops. This is no longer true due to several recent developments in numerical linear algebra, allowing to implement balanced truncation at a cost essentially proportional to the number of nonzeros in A if it is a sparse matrix (see [61, 49, 42]) or in $\mathcal{O}(n \log^2(n))$ (see [9]) if A is approximated by a hierarchical matrix [37].

We will briefly describe an implementation of balanced truncation for large-scale sparse systems. We follow here the description given in [10]. For this purpose, we note that the first step of any balanced truncation procedure consists in solving the Lyapunov equations (13). Here we want to go for low-rank approximations to factors of W_c, W_o . Hence, we start by considering the solution of Lyapunov equations of the form

$$FQ + QF^T = -WW^T,$$




Figure 1: Decay of the eigenvalues of the controllability Gramian of a random LTI system with $n = 500$, $m = 10$.

where F is stable. The ADI iteration for solving such equations can be written as follows [70]:

$$\begin{aligned} (F + p_j I)Q_{(j-1)/2} &= -WW^T - Q_{j-1}(F^T - p_j I), \\ Q_j(F^T + \bar{p}_j I) &= -WW^T - (F - \bar{p}_j I)Q_{(j-1)/2}. \end{aligned}$$

If the shift parameters p_j are chosen appropriately, then $\lim_{j \rightarrow \infty} Q_j = Q$ with a superlinear convergence rate. Starting this iteration with $Q_0 = 0$, the ADI iteration can be re-written as fixed-point iteration

$$Q_j = (F - \bar{p}_j I)(F + p_j I)^{-1}Q_{j-1}(F^T + \bar{p}_j I)^{-1}(F^T - p_j I) - 2\text{Re}(p_j)(F + p_j I)^{-1}WW^T(F^T + \bar{p}_j I)^{-1}. \quad (21)$$

Observing that for stable F , Q is positive semidefinite, we can assume that $Q_j = Y_j Y_j^T$ for some $Y_j \in \mathbb{R}^{n \times r_j}$, where $r_j \geq \text{rank}(Q_j)$. Inserting this into (21), re-arranging terms and combining two iteration steps, we obtain the following *factored ADI iteration*:

$$\begin{aligned} V_1 &\leftarrow \sqrt{-2\text{Re}(p_1)}(F + p_1 I)^{-1}W, & Y_1 &\leftarrow V_1 \\ \text{FOR } j &= 2, 3, \dots \\ V_j &\leftarrow \frac{\sqrt{\text{Re}(p_j)}}{\sqrt{\text{Re}(p_{j-1})}} (I - (p_j + \bar{p}_{j-1})(F + p_j I)^{-1}) V_{j-1}, \\ Y_j &\leftarrow \begin{bmatrix} Y_{j-1} & V_j \end{bmatrix}. \\ \text{END FOR} \end{aligned}$$

It should be noted that all V_j 's have the same number of columns as $W \in \mathbb{R}^{n \times w}$, i.e., at each iteration j , we have to solve w linear systems of equations with the same coefficient matrix $F + p_j I$. Hence, if convergence with respect to a suitable stopping criterion is achieved after j_{\max} steps, $Y_{j_{\max}} = [V_1 \ \dots \ V_{j_{\max}}] \in \mathbb{R}^{n \times j_{\max} w}$. For large n and small w we can therefore

expect that $r_j = j_{\max}w \ll n$. In that case, we have computed a low-rank approximation $Y_{j_{\max}}$ to a factor Y of the solution, i.e., $Q = YY^T \approx Y_{j_{\max}}Y_{j_{\max}}^T$. For an implementation of this method, we need a strategy to select the shift parameters. We will not discuss this here any further; see [71, 72] for a detailed discussion of this problem and an optimal choice in case of symmetric negative definite F . A numerically inexpensive, heuristic parameter selection that gives good performance in practice can be found in [59]. Usually, a finite number of shifts is computed in advance and applied cyclically if the ADI method needs more iterations than the number of available shifts.

Also note that the above algorithm can be implemented in real arithmetic by combining two steps, even if complex shifts need to be used (which may be the case if F is not symmetric). A complexity analysis of the factored ADI method is hard to give. It depends on the method used to solve the linear systems in each iteration step. In any case, all linear systems of equations to be solved in one iteration step have the same coefficient matrix $F + p_j I_n$. If this is a banded matrix or can be re-ordered to become banded, then a direct solver can be employed. This remains true for some other sparsity patterns that allow a sparse Cholesky or LU factorization, see [18, 52]. If workspace permits, it is desirable to compute all the factorizations of $F + p_j I_n$ for each different shift parameter beforehand (usually, very few parameters are used). These factorizations can then be used in the iteration steps of the ADI iteration so that only forward and backward substitution are needed to compute the new V_j . In particular, if F is symmetric negative definite as will be the case in many applications, and can be re-ordered to a (narrow) band matrix, then each factorization is $\mathcal{O}(n)$ and the total cost $\mathcal{O}(k_{\max}j_{\max}n)$ scales with n as desired. On the other hand, if iterative solvers are employed for the linear systems, it should be noted that due to the shift invariance of Krylov subspaces [25], only one Krylov space needs to be computed from which all linear systems of equations can be solved (see [49] for details). Hence, we obtain again an efficient implementation of the factored ADI iteration.

Stopping criteria for the factored ADI iteration can be based either on the fact that $\|V_j\| \rightarrow 0$ very rapidly or on the residual norm $\|FY_jY_j^T + Y_jY_j^TF^T + WW^T\|$; see [62] for an efficient way to compute the Frobenius norms of the residuals. Alternatively, we can employ the relative change in the iterates:

$$\frac{\|Q_j - Q_{j-1}\|_F}{\|Q_j\|_F} = \frac{\|Z_jZ_j^T - Z_{j-1}Z_{j-1}^T\|_F}{\|Z_jZ_j^T\|_F} = \frac{\|V_jV_j^T\|_F}{\|Z_jZ_j^T\|_F} = \frac{\|V_j^TV_j\|_F}{\|Z_j^TZ_j\|_F},$$

where the latter equality (based on $\|MM^T\|_F = \|M^TM\|_F$ for any matrix M) is cheap to evaluate whenever the number of columns in Z remains small with respect to n .

It should also be noted that for balanced truncation, the two factored ADI iterations needed to compute the solution factors of the Gramians can be coupled. As the iteration parameters are only based on the spectrum of F (with $F \in \{A, A^T\}$ and A from (13)), they will be the same for both iterations and thus need only be computed once. Moreover, solving a linear system of equations with coefficient matrix $A^T + p_j I_n$ as needed in the iteration for the observability Gramian can be cheaply implemented using the factorization of $A + p_j I_n$ computed in the iteration for the controllability Gramian. Suppose we have computed a permuted LU decomposition (with the permutation based on pivoting and fill-in minimization)

$$A + p_j I_n = PLUQ,$$

with P, Q permutation matrices, L lower triangular with unit diagonal, U upper triangular, then

$$(A + p_j I)^{-1} V_{j-1}^B = (PLUQ)^{-1} V_{j-1}^B$$

is computed by permutation according to P , forward substitution with respect to L , backward substitution with respect to U and permutation with respect to Q , while

$$(A^T + p_j I)^{-1} V_{j-1}^C = (Q^T U^T L^T P^T)^{-1} V_{j-1}^C$$

can be implemented in the order: permutation with respect to Q^T , forward substitution with respect to U^T , backward substitution with respect to L^T , and permutation with respect to P^T .

Summarizing, new and efficient algorithms for Lyapunov equations based on the (approximate) low-rank nature of solutions yield balanced truncation algorithms that can be applied to problems of the same complexity as modal and Padé(-type) methods. In Section 6 it will be shown that such a balanced truncation implementation is more than competitive with the other methods.

5 Comparison

A joint feature of most model reduction methods is that they can be interpreted as *Galerkin* or *Petrov-Galerkin* type projection of the state-space onto a low-dimensional subspace \mathcal{V} along the orthogonal complement of another subspace \mathcal{W} : assume

$$x \approx VW^T x =: \tilde{x},$$

where

$$\text{range}(V) = \mathcal{V}, \quad \text{range}(W) = \mathcal{W},$$

and W is chosen such that $W^T V = I_r$. Then, with $\hat{x} = W^T x$, we obtain $x \approx V\hat{x}$ and

$$\|x - \tilde{x}\| = \|x - V\hat{x}\|.$$

For modal truncation as described in Section 2, \mathcal{V} is simply the eigenspace corresponding to the chosen eigenvalues, i.e., $V = U_1$, and $\mathcal{W} = \mathcal{V}$ for unitarily diagonalizable A while

$$\mathcal{W} = \text{span}\{w_1, \dots, w_r\},$$

where

$$\begin{bmatrix} w_1 & \dots & w_n \end{bmatrix} = T^{-T}$$

with T as in Section 2. In any case, modal truncation has to choose \mathcal{V} (for diagonalizable A) from the $\binom{n}{r}$ many eigenspaces of A .

Padé-via-Lanczos methods can be interpreted analogously with \mathcal{V} and \mathcal{W} as in (11) and (12), respectively. That is, once s_0 is determined, PVL has no choice to pick the subspace \mathcal{V} in any other way than in (11).

For balanced truncation, the subspaces \mathcal{V} and \mathcal{W} are determined during the computation via

$$\mathcal{V} = \text{range}(T_l), \quad \mathcal{W} = \text{range}(T_r).$$

As the subspaces are not fixed prior to the computation, balanced truncation can—in principle—choose \mathcal{V} from the *Grassman manifold*

$$\mathcal{G}(n, r) := \{\mathcal{V} \subset \mathbb{R}^n : \dim \mathcal{V} = r\}.$$

If we now consider the minimization problem

$$\min_{V \in \mathcal{G}(n, r)} \|x - V\hat{x}\|,$$

it becomes clear that the larger variety in the selection of the projection subspace makes it likely that balanced truncation produces a model with smaller approximation error.

Note that for any of the discussed methods, the state of the original LTI system can be approximated by $V\hat{x}$, where \hat{x} is obtained from simulating the reduced-order model, as V is readily available in all the discussed approaches.

Of course, this is quite a simple discussion (see also [20] for similar observations). It does not take into account any modifications of modal truncation like base enrichment or the fact that an informed choice on the frequency s_0 in PVL methods will necessarily lead to good approximation qualities in a desired region. Thus, the comparison is not always favourable for balanced truncation, but examples in the next section demonstrate that often, a better approximation is achieved if the same reduced-order is used for the methods under discussion or, for achieving the same error tolerance, a model of smaller order is often sufficient when using balanced truncation.

6 Numerical Examples

In the numerical experiments reported in this section, we have used the balanced truncation implementation provided in Lyapack [63] which closely follows the description in Section 4, a simple implementation of modal truncation keeping those modes closest to the imaginary axis based on MATLAB’s `eigs` which is based on ARPACK [47], and a PVL implementation¹ for single-input, single-output (SISO) systems ($m = p = 1$) corresponding to [6]. All experiments were performed using MATLAB Release 14 (SP2). It will become obvious from the presented results that the naive implementation of modal truncation employed here is not competitive to PVL and balanced truncation in frequency response analysis. Usually, the results for modal truncation look better in a time domain analysis. Moreover, implementing more sophisticated dominance measures will certainly lead to much better results.

Example 1. Our first example is a MEMS device, called a microthruster. This device co-integrates solid fuel with a silicon micromachined system and is used for steering “nanosatellites”, but also for gas generation or as high-energy actuator. To produce an impulse, fuel is ignited by heating a resistor at the top of a particular microthruster.

The design problem is to reach the ignition temperature within the fuel, but at the same time staying below the critical temperature at neighbouring microthrusters. Moreover, the resistor temperature during the heating pulse should not become too high as this leads to the destruction of the membrane. Figure 2 shows a microthruster array, further information is available in [50, 56].

Here, we use two different 2D thermo-electrical models of a single microthruster as provided at the Oberwolfach Benchmark site [56]. A finite-element discretization using linear

¹Available at <http://www.cs.ucdavis.edu/~bai/ROMmatlab/>




Figure 2: A microthruster array at ignition (left) and schematic view of a cross-section (right). The array is about the size of a thumbnail. (By courtesy of C. Rossi, LAAS-CNRS.)

elements leads to $n = 4,257$ while using quadratic elements results in $n = 11,445$. The model has one input and seven outputs. As the PVL implementation used here can only treat SISO systems, we only use the first output in the experiments. That is, all frequency responses shown correspond to the transfer function from the input to the first output.

Figure 3 shows the frequency response (Bode magnitude plot) and the pointwise relative errors in the transfer function, i.e.,

$$\frac{|G(j\omega) - \hat{G}(j\omega)|}{|G(j\omega)|},$$

for the full $n = 4257$ system as well as reduced-order models of order $r = 21$ computed by balanced truncation (BT), PVL with $s_0 = 1$ and $s_0 = 10^4$, as well as modal truncation (MT). Obviously, the model computed by balanced truncation gives a satisfactory approximation over the whole frequency range while PVL shows good approximation errors in a frequency range around the selected expansion point s_0 . Close to s_0 , the relative approximation error of PVL is even better than that of balanced truncation, but grows to more than 100% in frequency ranges far from s_0 . The balanced truncation approximation error stays below 0.001% over the whole frequency range! The $r = 21$ model computed by modal truncation is quite useless except for very tiny frequencies. It should be noted that for this example, base enrichment using static solutions does not improve the accuracy of modal truncation—in contrast, the approximation error even grows and is therefore not displayed. Note that the ADI based balanced truncation implementation used six different parameters p_j in the ADI iteration for solving the Lyapunov equations, i.e., the essential cost of balanced truncation comes here from six sparse matrix factorizations as compared to only one sparse matrix factorization in the PVL and modal truncation methods. Thus the computation time needed for balanced truncation is roughly six times that of PVL and modal truncation.

This example demonstrates the claim from Section 5 that usually, balanced truncation yields better approximation for the same order r of the reduced-order models than the other approaches discussed in this paper. This is also supported by the very similar situation in the $n = 11,445$ case, displayed in Figure 4, where a reduced-order model of order $r = 28$ is used.




Figure 3: Example 1 ($n = 4,257$), frequency response (left) and pointwise relative errors (right) of the reduced-order models.




Figure 4: Example 1 ($n = 11,445$), frequency response (left) and pointwise relative errors (right) of the reduced-order models.

Figure 5 demonstrates that in order to get a reduced-order model matching the approximation quality of balanced truncation, we need to more than double the order for the PVL-computed system and that even if we go up to $r = 1,000$ for modal truncation, there is still a significant error in the high frequency range. (The full order model, the BT ($r = 21$) and the PVL ($r = 50$) frequency response are on top of each other in Figure 5 (left).)

Example 2. Here we consider an LTI system of the form (10) coming from an optimal cooling problem for steel profiles. This problem arises in a rolling mill when different steps in the production process require different temperatures of the raw material. To achieve a high production rate, economical interests suggest to reduce the temperature as fast as possible to the required level before entering the next production phase. At the same time, the cooling process, which is realized by spraying cooling fluids on the surface, has to be controlled so that material properties do not degrade. Therefore, a primary design goal is to obtain a preferably even temperature distribution. For a picture of such a cooling plant see Figure 6.

We assume an infinitely long steel profile so that we may restrict ourselves to a 2D model.




Figure 5: Example 1 ($n = 4, 257$), frequency response of the full order model and reduced-order models of several orders computed by PVL (left) and modal truncation (right).




Figure 6: Initial mesh, partitioning of the boundary, and a picture of a cooling plant.

Exploiting the symmetry of the workpiece, the computational domain $\Omega \subset \mathbb{R}^2$ is chosen as the half of a cross section of the rail profile. The heat distribution is modelled by the instationary linear heat equation on Ω :

$$\begin{aligned}
 c\rho\partial_t x(t, \xi) - \lambda\Delta x(t, \xi) &= 0 && \text{in } \mathbb{R}_{>0} \times \Omega, \\
 x(0, \xi) &= x_0(\xi) && \text{in } \Omega, \\
 \lambda\partial_\nu x(t, \xi) &= g_i && \text{on } \mathbb{R}_{>0} \times \Gamma_i, \partial\Omega = \bigcup_i \Gamma_i,
 \end{aligned} \tag{22}$$

where x is the temperature distribution ($x \in H^1([0, \infty], X)$ with $X := H^1(\Omega)$ being the state space), c the specific heat capacity, λ the heat conductivity and ρ the density of the rail profile. We split the boundary into several parts Γ_i on which we have different boundary functions g_i , allowing us to vary the controls on different parts of the surface. By ν we denote the outer normal of the boundary.

Applying a standard Galerkin approach for discretizing the heat transfer model in space using a linear finite element basis results in a first-order ordinary differential equation. By




Figure 7: Example 2, frequency response of the full order model and several reduced-order models (left); pointwise relative error of the reduced-order models (right).

refining the initial Delaunay triangulation shown in Figure 6 by bisection thrice, we obtain an LTI model with $n = 20,209$, corresponding to a maximum meshsize $h \approx 1.38 \cdot 10^{-2}$. For more information, see [15] and [56].

For comparing the model reduction algorithms, we choose the transfer function from the sixth input to the second output (the original model has 7 inputs and 6 outputs). This is motivated by the geometric vicinity of the corresponding input (spraying nozzle) and output (measured temperature in the area close to Γ_6). Figure 7 demonstrates that the situation here is similar to the one in Example 1. Here, prescribing a tolerance of 10^{-4} for the error bound (16), balanced truncation computes a reduced-order model of order $r = 8$. Again, among the models of order $r = 8$ we get the best global approximation by balanced truncation, while the PVL models exhibit the expected good local approximation quality and this time are only slightly off in the other regions. This is due to the fact that the given transfer function is very smooth and therefore easy to approximate. The model computed by modal truncation for $r = 8$ is completely useless; even the $r = 100$ model (including static modes) does only resemble the original transfer function for very low frequencies.

It should be noted that again, the balanced truncation model is computed using six shift parameters; hence the main computational cost in the ADI-based balanced truncation implementation is six sparse matrix factorizations and therefore about six times the cost of PVL and modal truncation.

Summarizing this section, the presented examples suggest that balanced truncation usually produces the best reduced-order model for a given order in the sense that the approximation error is globally the best. It is also competitive from the point of view of computational cost as it has basically the same complexity as the other methods, though the time to compute a reduced-order model is higher by some linear factor determined mostly by the number of shift parameters in the underlying ADI iteration. This is well compensated if the reduced-order has to be used in control or optimization or for many simulation runs.

7 Conclusions

Advances of numerical linear algebra techniques for sparse matrix computations during the past decades have lead to new and powerful tools for model reduction of linear systems. The Padé and Padé-type approximation techniques have set a standard for computing reduced-order models in circuit simulation and beyond. Balanced truncation has become an alternative for large and sparse problems as the $\mathcal{O}(n^3)$ barrier has been broken by several implementations of balanced truncation that make use of these new sparse matrix techniques.

Research directions in model reduction now shift more and more towards algorithms for parametric and nonlinear models in which parameters and nonlinearities, respectively, are to be preserved. Again, the efficient use of numerical linear algebra techniques will greatly improve the efficiency of these methods.

References

- [1] ANSYS, Inc., <http://www.ansys.com>. *ANSYS*.
- [2] A. Antoulas. *Approximation of Large-Scale Dynamical Systems*. SIAM Publications, Philadelphia, PA, 2005.
- [3] A. Antoulas, D. Sorensen, and S. Guercin. A survey of model reduction methods for large-scale systems. *Contemp. Math.*, 280:193–219, 2001.
- [4] A. Antoulas, D. Sorensen, and Y. Zhou. On the decay rate of Hankel singular values and related issues. *Sys. Control Lett.*, 46(5):323–342, 2002.
- [5] R. Badía, P. Benner, R. Mayo, and E. Quintana-Ortí. Parallel algorithms for balanced truncation model reduction of sparse systems. In *Proc. of PARA'04 – Workshop on State-of-the-Art in Scientific Computing*, Lecture Notes in Computer Science. Springer-Verlag, Berlin, Heidelberg, New York, to appear.
- [6] Z. Bai. Krylov subspace techniques for reduced-order modeling of large-scale dynamical systems. *Appl. Numer. Math.*, 43(1–2):9–44, 2002.
- [7] Z. Bai, J. Demmel, J. Dongarra, A. Ruhe, and H. van der Vorst, editors. *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, volume 11 of *Software - Environments - Tools*. SIAM Publications, Philadelphia, PA, 2000.
- [8] Z. Bai and Q. Ye. Error estimation on the Padé approximation of transfer functions via the Lanczos process. *Electr. Trans. Num. Anal.*, 7:1–17, 1998.
- [9] U. Baur and P. Benner. Factorized solution of the Lyapunov equation by using the hierarchical matrix arithmetic. *Proc. Appl. Math. Mech.*, 4(1):658–659, 2004.
- [10] P. Benner. Solving large-scale control problems. *IEEE Control Systems Magazine*, 14(1):44–59, 2004.
- [11] P. Benner, V. Mehrmann, and D. Sorensen, editors. *Dimension Reduction of Large-Scale Systems*, volume 45 of *Lecture Notes in Computational Science and Engineering*. Springer-Verlag, Berlin/Heidelberg, Germany, 2005.

- [12] P. Benner and E. Quintana-Ortí. Model reduction based on spectral projection methods. Chapter 1 (pages 5–48) of [11].
- [13] P. Benner, E. Quintana-Ortí, and G. Quintana-Ortí. Balanced truncation model reduction of large-scale dense systems on parallel computers. *Math. Comput. Model. Dyn. Syst.*, 6(4):383–405, 2000.
- [14] P. Benner, E. Quintana-Ortí, and G. Quintana-Ortí. State-space truncation methods for parallel model reduction of large-scale systems. *Parallel Comput.*, 29:1701–1722, 2003.
- [15] P. Benner and J. Saak. A semi-discretized heat transfer model for optimal cooling of steel profiles. Chapter 19 (pages 353–356) of [11].
- [16] P. Benner and V. Sokolov. Partial realization of descriptor systems. Preprint 2005-05, Fakultät für Mathematik, TU Chemnitz, 09107 Chemnitz, FRG, 2005. Available from http://www.mathematik.tu-chemnitz.de/preprint/2005/PREPRINT_05.html.
- [17] J. Bennighof and R. Lehoucq. An automated multilevel substructuring method for eigenspace computation in linear elastodynamics. *SIAM J. Sci. Comput.*, 25(6):2084–2106, 2004.
- [18] M. Bollhöfer and O. Schenk. Sparse matrix factorizations. *GAMM Mitteilungen*, to appear.
- [19] A. Böttcher and K. Rost. Topics in the numerical linear algebra of Toeplitz and Hankel matrices. *GAMM Mitteilungen*, 27(2):174–188, 2004.
- [20] C. Brezinski. Model reduction by state space projection. *Comm. Appl. Anal.*, 8(3):323–333, 2004.
- [21] R. Craig and M. Bampton. Coupling of substructures for dynamic analysis. *AIAA J.*, 6:1313–1319, 1968.
- [22] J. Cullum and T. Zhang. Two-sided Arnoldi and nonsymmetric Lanczos algorithms. *SIAM J. Matrix Anal. Appl.*, 24(2):303–319, 2002.
- [23] L. Dai. *Singular Control Systems*. Number 118 in Lecture Notes in Control and Information Sciences. Springer-Verlag, Berlin, 1989.
- [24] E. Davison. A method for simplifying linear dynamic systems. *IEEE Trans. Automat. Control*, AC-11:93–101, 1966.
- [25] R. Freund. Solution of shifted linear systems by quasi-minimal residual iterations. In L. Reichel et al., editor, *Numerical Linear Algebra*, pages 101–121. Walter de Gruyter, Berlin, 1993.
- [26] R. Freund. Reduced-order modeling techniques based on Krylov subspaces and their use in circuit simulation. In B. Datta, editor, *Applied and Computational Control, Signals, and Circuits*, volume 1, chapter 9, pages 435–498. Birkhäuser, Boston, MA, 1999.
- [27] R. Freund. Krylov-subspace methods for reduced-order modeling in circuit simulation. *J. Comput. Appl. Math.*, 123(1–2):395–421, 2000.

- [28] R. Freund. Model reduction methods based on Krylov subspaces. *Acta Numerica*, 12:267–319, 2003.
- [29] R. W. Freund and P. Feldmann. Small-signal circuit analysis and sensitivity computations with the PVL algorithm. *IEEE Transactions on Circuits and Systems—II: Analog and Digital Signal Processing*, 43:577–585, 1996.
- [30] K. Gallivan, E. Grimme, and P. Van Dooren. Asymptotic waveform evaluation via a lanczos method. *Appl. Math. Lett.*, 7(5):75–80, 1994.
- [31] K. Gallivan, E. Grimme, and P. Van Dooren. A rational Lanczos algorithm for model reduction. *Numer. Algorithms*, 12:33–63, 1996.
- [32] K. Gallivan, A. Vandendorpe, and P. Van Dooren. Model reduction of MIMO systems via tangential interpolation. *SIAM J. Matrix Anal. Appl.*, 26(2):328–349, 2004.
- [33] K. Gallivan, A. Vandendorpe, and P. Van Dooren. Sylvester equations and projection-based model reduction. *J. Comput. Appl. Math.*, 162(1):213–229, 2004.
- [34] D. Givoli, P. Barbone, and I. Patlashenko. Which are the important modes of a subsystem? *Int. J. Numer. Methods Eng.*, 59(12):1657–1678, 2004.
- [35] K. Glover. All optimal Hankel-norm approximations of linear multivariable systems and their L^∞ norms. *Internat. J. Control*, 39:1115–1193, 1984.
- [36] L. Grasedyck. Existence of a low rank or H -matrix approximant to the solution of a Sylvester equation. *Numer. Lin. Alg. Appl.*, 11:371–389, 2004.
- [37] L. Grasedyck and W. Hackbusch. Construction and arithmetics of \mathcal{H} -matrices. *Computing*, 70(4):295–334, 2003.
- [38] M. Green and D. Limebeer. *Linear Robust Control*. Prentice-Hall, Englewood Cliffs, NJ, 1995.
- [39] E. Grimme. *Krylov projection methods for model reduction*. PhD thesis, Univ. of Illinois at Urbana-Champaign, USA, 1997.
- [40] E. Grimme, D. Sorensen, and P. Van Dooren. Model reduction of state space systems via an implicitly restarted Lanczos method. *Numer. Algorithms*, 12:1–31, 1996.
- [41] S. Gugercin and A. Antoulas. Model reduction of large-scale systems by least squares. *Linear Algebra Appl.*, to appear.
- [42] S. Gugercin and J.-R. Li. Smith-type methods for balanced truncation of large systems. Chapter 2 (pages 49–82) of [11].
- [43] R. Guyan. Reduction of stiffness and mass matrices. *AIAA J.*, 3:380, 1968.
- [44] I. Jaimoukha and E. Kasenally. Krylov subspace methods for solving large Lyapunov equations. *SIAM J. Numer. Anal.*, 31:227–251, 1994.
- [45] I. Jaimoukha and E. Kasenally. Implicitly restarted Krylov subspace methods for stable partial realizations. *SIAM J. Matrix Anal. Appl.*, 18(3):633–652, 1997.

- [46] A. Laub, M. Heath, C. Paige, and R. Ward. Computation of system balancing transformations and other application of simultaneous diagonalization algorithms. *IEEE Trans. Automat. Control*, 34:115–122, 1987.
- [47] R. Lehoucq, D. Sorensen, and C. Yang. *ARPACK User's Guide*, volume 6 of *Software - Environments - Tools*. SIAM Publications, Philadelphia, PA, 1998.
- [48] J.-R. Li and J. White. Reduction of large circuit models via low rank approximate gramians. *Int. J. Appl. Math. Comp. Sci.*, 11(5):1151–1171, 2001.
- [49] J.-R. Li and J. White. Low rank solution of Lyapunov equations. *SIAM J. Matrix Anal. Appl.*, 24(1):260–280, 2002.
- [50] J. Lienemann, E. Rudnyi, and J. Korvink. MST MEMS compact modeling meets model order reduction: Requirements and benchmarks. *Linear Algebra Appl.*, to appear.
- [51] S. Marschall. An approximate method for reducing the order of a linear system. *Contr. Eng.*, 10:642–648, 1966.
- [52] G. Meurant. Gaussian elimination for the solution of linear systems of equations. In P. Ciarlet and J. Lions, editors, *Handbook of Numerical Analysis*, volume 7, pages 3–170. North-Holland/ Elsevier, Amsterdam, 2000.
- [53] B. Moore. Principal component analysis in linear systems: Controllability, observability, and model reduction. *IEEE Trans. Automat. Control*, AC-26:17–32, 1981.
- [54] MSC.Software Corporation, <http://www.mscsoftware.com>. *MSC.Nastran*.
- [55] C. Mullis and R. Roberts. Synthesis of minimum roundoff noise fixed point digital filters. *IEEE Trans. Circuits and Systems*, CAS-23(9):551–562, 1976.
- [56] Oberwolfach model reduction benchmark collection. Available at <http://www.imtek.de/simulation/benchmark/>.
- [57] G. Obinata and B. Anderson. *Model Reduction for Control System Design*. Communications and Control Engineering Series. Springer-Verlag, London, UK, 2001.
- [58] A. Odabasioglu, M. Celik, and L. Pileggi. PRIMA: passive reduced-order interconnect macromodeling algorithm. In *Tech. Dig. 1997 IEEE/ACM Intl. Conf. CAD*, pages 58–65. IEEE Computer Society Press, 1997.
- [59] T. Penzl. *Numerische Lösung großer Lyapunov-Gleichungen*. Logos-Verlag, Berlin, Germany, 1998. Dissertation, Fakultät für Mathematik, TU Chemnitz, 1998.
- [60] T. Penzl. Algorithms for model reduction of large dynamical systems. Technical Report SFB393/99-40, Sonderforschungsbereich 393 *Numerische Simulation auf massiv parallelen Rechnern*, TU Chemnitz, 09107 Chemnitz, FRG, 1999. Available from <http://www.tu-chemnitz.de/sfb393/sfb99pr.html>.
- [61] T. Penzl. A cyclic low rank Smith method for large sparse Lyapunov equations. *SIAM J. Sci. Comput.*, 21(4):1401–1418, 2000.

- [62] T. Penzl. Eigenvalue decay bounds for solutions of Lyapunov equations: the symmetric case. *Sys. Control Lett.*, 40:139–144, 2000.
- [63] T. Penzl. LYAPACK Users Guide. Technical Report SFB393/00-33, Sonderforschungsbereich 393 *Numerische Simulation auf massiv parallelen Rechnern*, TU Chemnitz, 09107 Chemnitz, FRG, 2000. Available from <http://www.tu-chemnitz.de/sfb393/sfb00pr.html>.
- [64] D. Rixen. A dual Craig–Bampton method for dynamic substructuring. *J. Comput. Appl. Math.*, 168(1–2):383–391, 2004.
- [65] A. Ruhe. Rational Krylov: A practical algorithm for large sparse nonsymmetric matrix pencils. *SIAM J. Sci. Comput.*, 19(5):1535–1551, 1998.
- [66] M. Tombs and I. Postlethwaite. Truncated balanced realization of a stable non-minimal state-space system. *Internat. J. Control*, 46(4):1319–1330, 1987.
- [67] I. Troch, P. Müller, and K.-H. Fasol. Modellreduktion für Simulation und Reglerentwurf. *Automatisierungstechnik*, 40(2/3/4):45–53/93–99/132–141, 1992.
- [68] H. van der Vorst. Computational methods for large eigenvalue problems. In P. Ciarlet and J. Lions, editors, *Handbook of Numerical Analysis*, volume 8, pages 3–179. North-Holland/ Elsevier, Amsterdam, 2002.
- [69] A. Varga. Enhanced modal approach for model reduction. *Math. Model. Syst.*, 1(2):91–105, 1995.
- [70] E. Wachspress. Iterative solution of the Lyapunov matrix equation. *Appl. Math. Letters*, 107:87–90, 1988.
- [71] E. Wachspress. The ADI model problem, 1995. Available from the author.
- [72] E. Wachspress. ADI iteration parameters for the Sylvester equation, 2000. Available from the author.
- [73] C. Yang, W. Gao, Z. Bai, X. Li, L. Lee, P. Husbands, and E. Ng. An algebraic substructuring algorithm for large-scale eigenvalue calculation. *SIAM J. Sci. Comput.*, to appear.