

Emre Mengi*

A subspace framework for $\mathcal{L}_\infty$ model reduction

<https://doi.org/10.1515/jnma-2023-0115>

Received September 15, 2023; accepted October 4, 2024; published online January 15, 2025

Abstract: We propose an approach for the $\mathcal{L}_\infty$ model reduction of descriptor systems based on the minimization of the $\mathcal{L}_\infty$ objective by means of smooth optimization techniques. Direct applications of smooth optimization techniques are not feasible even for systems of modest order, since the optimization techniques converge at best at a linear rate requiring too many evaluations of the costly $\mathcal{L}_\infty$ -norm objective to be practical. We replace the original system with a system of smaller order interpolating the original system at points on the imaginary axis, minimize the $\mathcal{L}_\infty$ objective after this replacement, and refine the smaller system based on the minimization. We also describe how asymptotic stability constraints on the reduced system sought can be incorporated into our approach. The numerical experiments illustrate that the approach leads to locally optimal solutions to the $\mathcal{L}_\infty$ model reduction problem, and its capability to deal with systems of order a few ten thousands.

Keywords: H_∞ model reduction; descriptor system; quasi-Newton methods; Petrov–Galerkin projection; Hermite interpolation

MSC 2010: 65D05; 65F15; 65L80; 90C53; 93A15; 93C05

1 Introduction

Various applications give rise to linear time-invariant (LTI) descriptor systems; see, e.g., [1]–[4], and references therein. A model order reduction problem for such a system typically concerns the approximation of the system with a system of smaller and prescribed order. Here we deal with one of them, namely the $\mathcal{L}_\infty$ model reduction problem formally defined below.

We assume that the LTI descriptor system under consideration is available in a state-space representation

$$Ex'(t) = Ax(t) + Bu(t), \quad y(t) = Cx(t) + Du(t) \quad (1)$$

for given matrices $E, A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{p \times n}$, $D \in \mathbb{R}^{p \times m}$. The transfer function of the descriptor system in (1) (defined over $s \in \mathbb{C}$ such that $sE - A$ is invertible) is

$$H(s) = C(sE - A)^{-1}B + D \quad (2)$$

with the $\mathcal{L}_\infty$ norm

$$\|H\|_{\mathcal{L}_\infty} := \sup_{\omega \in \mathbb{R}} \sigma_{\max}(H(i\omega)) = \sup_{\omega \in \mathbb{R}, \omega \geq 0} \sigma_{\max}(H(i\omega)),$$

where $\sigma_{\max}(\cdot)$ denotes the largest singular value of its matrix argument, and the last equality holds as A, B, C, D, E are real matrices. Note that we customarily set $\|H\|_{\mathcal{L}_\infty} = \sup_{\omega \in \mathbb{R}} \sigma_{\max}(H(i\omega)) = \infty$ if H has a pole on the

*Corresponding author: **Emre Mengi**, Department of Mathematics, Koç University, Rumeli Feneri Yolu 34450, Sarıyer, İstanbul, Türkiye, E-mail: emengi@ku.edu.tr. <https://orcid.org/0000-0003-0788-0066>

imaginary axis, or the norm of its restriction to the imaginary axis is not bounded. Let $\mathbb{C}^+ := \{z \in \mathbb{C} \mid \operatorname{Re} z > 0\}$, and L_2 denote a space of functions $f: \mathbb{R} \rightarrow \mathbb{R}^k$ satisfying $\|f\|_{L_2} := \sqrt{\int_{-\infty}^{\infty} \|f(t)\|_2^2 dt} < \infty$ for some $k \in \mathbb{Z}^+$. To be specific, in our setting, L_2 consists of either input functions $u(t)$ or output functions $y(t)$ (so $k = m$ or $k = p$), and which one L_2 refers to will be clear from the context. If the system in (1) is asymptotically stable with poles in the open left half of the complex plane, then the $\mathcal{L}_\infty$ norm of H is the same as the $\mathcal{H}_\infty$ norm of H defined as

$$\|H\|_{\mathcal{H}_\infty} := \sup_{s \in \mathbb{C}^+} \sigma_{\max}(H(s)),$$

which in turn is equal to the induced norm of the operator $\varphi: L_2 \rightarrow L_2$ associated with (1) in the time domain mapping u to y as $\varphi u = y$ (for an explicit expression for φ see Ref. [3, Theorem 2.29]), that is given by

$$\|\varphi\|_{L_2} := \sup \{ \|\varphi u\|_{L_2} / \|u\|_{L_2} \mid u \in L_2, u \neq 0 \}.$$

Hence, under the asymptotic stability assumption on the descriptor system in (1), we have $\|H\|_{\mathcal{L}_\infty} = \|H\|_{\mathcal{H}_\infty} = \|\varphi\|_{L_2}$.

The $\mathcal{L}_\infty$ model reduction problem considered in this work for a given descriptor system of order n and for a prescribed positive integer $r < n$ concerns finding a reduced descriptor system of order r that is closest to the given system of order n with respect to the $\mathcal{L}_\infty$ norm. Formally, let $S^{\text{red}} = (A^{\text{red}}, E^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}})$ denote a system of order r with the state-space representation

$$E^{\text{red}} x'(t) = A^{\text{red}} x(t) + B^{\text{red}} u(t), \quad y(t) = C^{\text{red}} x(t) + D^{\text{red}} u(t), \quad (3)$$

described by the matrices $E^{\text{red}}, A^{\text{red}} \in \mathbb{R}^{r \times r}$, $B^{\text{red}} \in \mathbb{R}^{r \times m}$, $C^{\text{red}} \in \mathbb{R}^{p \times r}$, $D^{\text{red}} \in \mathbb{R}^{p \times m}$, and with the transfer function

$$H(s; S^{\text{red}}) = C^{\text{red}}(sE^{\text{red}} - A^{\text{red}})^{-1}B^{\text{red}} + D^{\text{red}}. \quad (4)$$

Furthermore, let $S = (A, E, B, C, D)$ be the given system of order n and with the transfer function H as in (2). Setting $\sigma(\omega; S^{\text{red}}) := \sigma_{\max}(H(i\omega) - H(i\omega; S^{\text{red}}))$, the $\mathcal{L}_\infty$ model reduction problem involves finding a descriptor system $S_\star^{\text{red}}$ of order r that minimizes the objective

$$\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty} = \sup_{\omega \in \mathbb{R}} \sigma(\omega; S^{\text{red}}) = \sup_{\omega \in \mathbb{R}, \omega \geq 0} \sigma(\omega; S^{\text{red}}) \quad (5)$$

over all descriptor systems S^{red} of order r . The objective in (5) is non-convex, and here we aim to determine a local minimizer of this objective numerically. The quality of the determined local minimizer also matters, however this issue is largely dependent upon with which reduced system of order r our approach is initialized.

Two important remarks are in order regarding the minimization of the objective in (5). First, in addition to non-convexity, an additional difficulty is the nonsmooth nature of the problem. The objective in (5) as a function of S^{red} is typically not differentiable when $\sigma(\omega; S^{\text{red}})$ has multiple global maximizers over $\omega \geq 0$. Secondly, under asymptotic stability assumptions on the original system and the reduced system, the error $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ gives a uniform upper bound on how much the outputs of the original and reduced systems can differ. To be precise, suppose that the system S of order n , and the reduced system S^{red} of order r are asymptotically stable. Furthermore, let us denote with φ and φ_r the operators in the time domain corresponding to the systems in (1) and (3), respectively. For every $u \in L_2$, we have

$$\|y - y_r\|_{L_2} \leq \|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty} \|u\|_{L_2},$$

where y, y_r are such that $y = \varphi u$ and $y_r = \varphi_r u$. This means that if a small error $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ can be ensured by the minimization of (5), then the output $y_r = \varphi_r u$ of the minimizing reduced system approximates the original output $y = \varphi u$ well uniformly over every input u of prescribed norm.

1.1 Literature and contributions

For an asymptotically stable descriptor system with the transfer function H , the H_∞ model reduction problem – that is, for a given small order r , finding an asymptotically stable system S^{red} of order r such that $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{H}_\infty}$ is small – has been under consideration for a long time. One of the classical approaches for the $\mathcal{H}_\infty$ model reduction problem is balanced truncation, which determines a state transformation so that the observability and controllability Gramians are the same diagonal matrix, and truncates the system matrices after applying this state transformation [5]–[8]. The reduced system by balanced truncation is typically not a local minimizer of $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{H}_\infty}$ over S^{red} , though it usually is a good quality approximation of H with respect to the $\mathcal{H}_\infty$ norm [9]. The major difficulty with balanced truncation that limits its applicability to larger systems is that it requires the solution of two Lyapunov equations involving matrices of size equal to the order of the system. With iterative approaches for solving Lyapunov equations, such as ADI methods [8], [10], [11], balanced truncation is tractable for systems with higher order.

A classical alternative is finding a best approximation with respect to the Hankel norm (HNA) [12] rather than the $\mathcal{H}_\infty$ norm. Approaches to compute a globally optimal solution to HNA in polynomial time are proposed [12]. Generalized HNA approaches are introduced for larger sparse descriptor systems [13]. However, a globally optimal solution to HNA is again usually not a local minimizer of $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{H}_\infty}$. Furthermore, finding a globally optimal solution to HNA is even costlier than balanced truncation. Even with the efficient use of computational linear algebra tools [14], solving HNA for systems with high order is out of reach.

The iterative rational Krylov algorithm (IRKA) [15] was introduced in order to find a reduced system of prescribed order that is locally optimal with respect to the $\mathcal{H}_2$ norm defined as $\|H\|_{\mathcal{H}_2} = \sqrt{\frac{1}{2\pi} \int_{-\infty}^{\infty} \text{trace}(H(i\omega)^* H(i\omega)) d\omega}$ for a system with the transfer function H . Formally, IRKA is an iterative interpolatory approach that finds a local minimizer of $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{H}_2}$ over all systems S^{red} of order r . In Ref. [16], starting from the reduced-order system of order r generated by IRKA, an optimization based approach is proposed to find a locally optimal solution of $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{H}_\infty}$ for single-input-single-output (SISO) systems but, denoting with $\mathbf{e}$ the vector of ones, with respect to particular rank-one modifications $\Delta_A = \varepsilon \mathbf{e} \mathbf{e}^T$, $\Delta_B = -\varepsilon \mathbf{e}$, $\Delta_C = -\varepsilon \mathbf{e}^T$, $\Delta_D = \varepsilon$ of the system matrices $A^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}}$ generated by IRKA over $\varepsilon \in \mathbb{R}$. In the reported results in Ref. [16], this optimization improves the accuracy of the reduced system returned by IRKA by a factor of 2–4 with respect to the $\mathcal{H}_\infty$ norm. But again the eventual system is usually not a local minimizer of the objective $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{H}_\infty}$ over systems S^{red} of order r .

Here, we propose an approach to compute a local minimizer of $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ over all systems S^{red} of order r . Nonsmooth optimization techniques, specifically bundle methods [17], [18], and the gradient sampling algorithm [19], [20], have been employed to locally optimize such nonsmooth objectives. Bundle methods are especially suitable for convex objectives, but, to our knowledge, they could not be applied as effectively to nonconvex objectives. The gradient sampling algorithm is applicable to nonsmooth and nonconvex objectives. Indeed, it is shown to converge to a Clarke stationary point under local Lipschitz continuity assumption and other mild conditions on the objective [20, Section 3]. The difficulty with the gradient sampling algorithm is that it computes the gradient at sample points around the current iterate, which may be too costly for some objectives, e.g., for the $\mathcal{L}_\infty$ error objective here [21]. Consequently, due to efficiency considerations, our approach to compute a local minimizer of $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ uses smooth optimization techniques, which have been employed for solving nonsmooth optimization problems [22], [23] in the last 15 years. For instance, quasi-Newton methods, most often, seem to be capable of locating locally optimal solutions [23], [24], even though some counter examples are known [23]–[25]. When convergence occurs with a smooth optimization technique to a nonsmooth locally optimal solution, it usually occurs slowly at best at a linear rate. As a result, as we shall see below, a direct application of them to minimize $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ is prohibitively expensive even for systems of medium order, as it requires the computation of the objective, that is the $\mathcal{L}_\infty$ norm, too many times. Instead, we replace H with an approximation $\tilde{H}$ of small order greater than r . Rather than $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$, we minimize $\|\tilde{H} - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$, then update $\tilde{H}$ based on the minimizer, and repeat. The approximation $\tilde{H}$ is built using

subspace projections, and an update involves the expansion of the projection subspaces. We show that the proposed subspace framework converges quadratically under simplicity assumptions. We also describe how the asymptotic stability constraints can be imposed on the variable S^{red} when minimizing $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ in case the original system in (1) is asymptotically stable. As the system corresponding to $H - H(\cdot; S^{\text{red}})$ is asymptotically stable when S and S^{red} are asymptotically stable, the incorporation of this constraint into our approach leads to a local minimization of $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{H}_\infty}$ over all asymptotically stable systems S^{red} of order r , i.e., locally optimal solution of the $\mathcal{H}_\infty$ model reduction problem.

Smooth optimization techniques have been recently used for model order reduction of asymptotically stable structured systems [26]. It is also observed over there that a direct minimization of the $\mathcal{H}_\infty$ -norm error objective via smooth optimization techniques is too costly. However, the proposed approach in Ref. [26] is not based on subspace projections. Instead, the $\mathcal{H}_\infty$ -norm error objective is replaced by a smooth objective that involves a prescribed number of largest singular values of the difference between the transfer functions of the original and reduced system at sample points on the imaginary axis. The minimization of this smoothed objective is illustrated to yield models more accurate than those by various structure-preserving model order reduction techniques on some benchmark examples. The reduced system obtained by the method in Ref. [26] is a local minimizer of the smoothed objective, which is related to but different than the $\mathcal{H}_\infty$ objective. In addition to nonsmoothness, the nonconvexity of the $\mathcal{H}_\infty$ error objective is a challenge. In Ref. [27], a convex relaxation of the $\mathcal{H}_\infty$ error objective is proposed for $\mathcal{H}_\infty$ model reduction.

On a related note, our recent work [28] concerns the minimization of the $\mathcal{H}_\infty$ norm of a descriptor system with large order with respect to a modest number of parameters. At the center of that work is a subspace framework to cope with the large order of the system. It may seem plausible to look at the current work from that perspective. However, we have too many optimization parameters here. Attempting to apply the framework of [28] to attain quick convergence in the setting here is not feasible, as doing so yields projection subspaces growing rapidly (i.e., see Algorithm 2 in Ref. [28] to attain superlinear convergence). In the framework here, it suffices to add a small number of new directions independent of r into the subspaces at every iteration. In particular, if $m = p$, then only $4m$ new directions are added into the subspaces at every iteration. Moreover, we observe quick convergence, so the subspaces remain small throughout. Also in the context of parametric descriptor systems, in Refs. [29], [30], an $\mathcal{H}_\infty$ error objective is minimized on a discrete set of parameter values to choose interpolation points for interpolatory model order reduction techniques.

1.2 Outline

We first consider the direct minimization of $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ over systems S^{red} of order r by means of smooth optimization techniques in Section 2. In this section, we indicate the optimization variables, and spell out expressions for the first derivatives of the objective with respect to these variables. As we shall see, direct optimization is too costly even for systems with moderate order, since smooth optimization techniques converge very slowly and require the evaluation of the $\mathcal{L}_\infty$ objective too many times. Consequently, in Section 3, we replace the transfer function H with an approximating transfer function $\tilde{H}$ of small order greater than r that Hermite interpolates H at several points on the imaginary axis. Then we minimize $\|\tilde{H} - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ (by smooth optimization techniques), and refine $\tilde{H}$ so that Hermite interpolation with H at another point on the imaginary axis is attained based on the computed minimizer of $\|\tilde{H} - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$. We introduce a refinement step on $\tilde{H}$ so that interpolation properties can also be attained between the full objective $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ and the reduced objective $\|\tilde{H} - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$. Then the procedure is repeated with the refined $\tilde{H}$. In Section 4, we investigate the interpolation properties between full objective and the reduced objective. Based on these interpolation properties, we argue in Section 5 that the algorithm converges at a quadratic rate under smoothness and nondegeneracy assumptions. If the original descriptor system is asymptotically stable, it may be natural to minimize $\|\tilde{H} - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ subject to the asymptotic stability constraints on the reduced system S^{red} . We discuss in Section 6 the incorporation of such asymptotic stability constraints on the reduced system into our approach. Section 7 is devoted to the details that need to be taken into account in a practical implementation of the proposed algorithm such as the initialization of the smooth optimization routines, and termination. A MATLAB

implementation of the algorithm is publicly available. In Section 8, we report numerical results obtained with this implementation. The numerical results indicate quick convergence to a locally optimal solution, and the capability to deal with systems having orders in the range of tens of thousands.

2 Use of first-order and second-order derivative-based methods

First-order methods such as the gradient descent algorithm, and second-order methods such as quasi-Newton algorithms equipped with proper line-searches have been successfully applied to nonsmooth optimization problems in recent years [22], [23], [31], [32]. In these circumstances, if a curvature condition is employed in the line-search, this should take into consideration that the directional derivatives need not converge to zero unlike the situation for smooth optimization problems, e.g., if Wolfe conditions are imposed in the line-search, weak Wolfe conditions should be used rather than strong Wolfe conditions [23, Section 4]. Also, for termination small gradient norms should not be required. Instead, for instance, a failure to take a reasonably long step in the objective along the descent search direction may indicate convergence to a locally optimal solution [31, Section 3].

The objective to be minimized in (5) for the $\mathcal{L}_\infty$ model reduction problem can be expressed as

$$\begin{aligned} \mathcal{F}(S^{\text{red}}) &:= \sup_{\omega \geq 0} \sigma_{\max}(H(i\omega) - H(i\omega; S^{\text{red}})) = \sup_{\omega \geq 0} \sigma_{\max}(\mathcal{E}(i\omega; S^{\text{red}})) = \|\mathcal{E}(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}, \\ \mathcal{E}(s; S^{\text{red}}) &:= [C \quad -C^{\text{red}}] \begin{bmatrix} sE - A & 0 \\ 0 & sE^{\text{red}} - A^{\text{red}} \end{bmatrix}^{-1} \begin{bmatrix} B \\ B^{\text{red}} \end{bmatrix} + (D - D^{\text{red}}), \end{aligned} \quad (6)$$

where $H(\cdot; S^{\text{red}})$ is as in (4). Assuming that the reduced system is at most index one and has semi-simple poles, by the Kronecker canonical form, there exist invertible $r \times r$ real matrices W, V such that $WE^{\text{red}}V$ is diagonal, and $WA^{\text{red}}V$ is block diagonal with 2×2 and 1×1 blocks along the diagonal. (Here, a 2×2 block along the diagonal of $WA^{\text{red}}V$ corresponds to a pair of complex conjugate eigenvalues, and a 1×1 block to a real or an infinite eigenvalue of the pencil $L(s) = A^{\text{red}} - sE^{\text{red}}$.) Consequently, the reduced system is equivalent to a system (with the same transfer function) for which $A^{\text{red}}, E^{\text{red}}$ are converted into tridiagonal and diagonal forms, respectively. Hence, under index one and semi-simple pole assumptions, we can perform the minimization over tridiagonal A^{red} and diagonal E^{red} . Recalling the dimensions of $A^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}}, E^{\text{red}}$, there are precisely $4r - 2 + rm + pr + pm$ optimization variables.

The gradient descent algorithm, as well as quasi-Newton algorithms to minimize $\mathcal{F}$ require the gradients of $\mathcal{F}$. To this end, suppose there is a unique $\omega_* \geq 0$ satisfying

$$\sigma_{\max}(H(i\omega_*) - H(i\omega_*; S^{\text{red}})) = \mathcal{F}(S^{\text{red}}) = \sup_{\omega \geq 0} \sigma_{\max}(H(i\omega) - H(i\omega; S^{\text{red}})),$$

ensuring that $\mathcal{F}$ is differentiable at S^{red} . Additionally, let u, v denote a consistent pair of unit left, right singular vectors corresponding to $\sigma_{\max}(H(i\omega_*) - H(i\omega_*; S^{\text{red}}))$, and let us introduce $\tilde{u} := u^* C^{\text{red}} (i\omega_* E^{\text{red}} - A^{\text{red}})^{-1}$, $\tilde{v} := (i\omega_* E^{\text{red}} - A^{\text{red}})^{-1} B^{\text{red}} v$. Then, by employing the analytical formulas for the derivatives of singular value functions [33], [34, Section 3.3], the gradients of $\mathcal{F}$ are given by

$$\begin{aligned} \nabla_{A^{\text{red}}} \mathcal{F}(S^{\text{red}}) &= -\text{diag}(\text{Re}(\tilde{u}^T \odot \tilde{v})) - \text{diag}_{-1}(\text{Re}(\tilde{u}(2:r)^T \odot \tilde{v}(1:r-1))) - \text{diag}_{+1}(\text{Re}(\tilde{u}(1:r-1)^T \odot \tilde{v}(2:r))), \\ \nabla_{E^{\text{red}}} \mathcal{F}(S^{\text{red}}) &= -\omega_* \text{diag}(\text{Im}(\tilde{u}^T \odot \tilde{v})), \quad \nabla_{B^{\text{red}}} \mathcal{F}(S^{\text{red}}) = -\text{Re}(\tilde{u}^T v^T), \\ \nabla_{C^{\text{red}}} \mathcal{F}(S^{\text{red}}) &= -\text{Re}(\tilde{u} \tilde{v}^T), \quad \nabla_{D^{\text{red}}} \mathcal{F}(S^{\text{red}}) = -\text{Re}(\tilde{u} v^T), \end{aligned} \quad (7)$$

where $\text{Re}(\cdot), \text{Im}(\cdot)$ denote the real part, imaginary part, respectively, of their vector or matrix arguments. Also, above $\odot$ denotes the Hadamard product, $\bar{u}$ is the complex conjugate of u , and the notation $\text{diag}(w)$ represents the square diagonal matrix whose diagonal entries are formed of the entries of the vector w . The notations $\text{diag}_{-1}(w)$ and $\text{diag}_{+1}(w)$ are similar to $\text{diag}(w)$ but with the difference that the subdiagonal and superdiagonal entries of the matrix are filled with the entries of w rather than the diagonal entries.

It is essential that a quasi-Newton method such as BFGS generates approximate Hessians that are positive definite. This is traditionally imposed by the line-searches. For instance, if BFGS is to be used to minimize $\mathcal{F}$, then a line-search ensuring the satisfaction of the weak Wolfe conditions may be adopted so that the approximate Hessians remain positive definite. On the other hand, for the gradient descent algorithm to minimize $\mathcal{F}$, it is sufficient to adopt a simpler line-search that guarantees only sufficient reduction in the objective, e.g., an Armijo backtracking line-search.

One difficulty with using methods such as gradient descent and BFGS to minimize $\mathcal{F}$ is that these algorithms converge rather slowly only at a linear rate at best. This may sound surprising especially for BFGS, which typically converges superlinearly for smooth problems. Slower convergence for BFGS is an artifact of nonsmoothness. As a result of linear convergence rates or worse, the objective $\mathcal{F}$ typically needs to be evaluated many times until reaching a prescribed accuracy. This may be prohibitively expensive, as it is apparent from (6) that evaluation of $\mathcal{F}$ involves the computation of the $\mathcal{L}_\infty$ norm of $\mathcal{E}(\cdot; S^{\text{red}})$, which may be costly assuming the original system in (1) is large-scale.

To illustrate the slow convergence issues described in the previous paragraph, and the computational difficulties that come with it, we apply the gradient descent algorithm to the `iss` example from the SLICOT collection. The system associated with this example has order $n = 270$, and $m = p = 3$. We attempt to solve the $\mathcal{L}_\infty$ model reduction problem for $r = 12$ starting with the initial reduced-order model generated by balanced truncation. The errors ($\mathcal{F}$) and the 2-norms of the gradients of the errors ($\|\nabla \mathcal{F}\|_2$) of the iterates of the gradient descent algorithm are reported in Table 1. It takes 37 iterations until the errors in two consecutive iterations differ by no more than 10^{-6} in a relative sense. The initial $\mathcal{L}_\infty$ -norm error 4.5×10^{-3} approximately (of the system obtained from balanced truncation) is reduced to 2.4×10^{-3} approximately after 37 iterations. The eventual reduced model obtained appears to be a local minimizer of $\mathcal{F}$ up to prescribed tolerances, as evidenced by the plots in Figure 1. Note however that according to the last columns in Table 1 the gradients of $\mathcal{F}$ do not seem to be converging to zero, which indicates that the objective is not differentiable at the local minimizer being approached. Meanwhile, the objective $\mathcal{F}$ is evaluated 624 times, since the line-search at each iteration requires several objective function evaluations (i.e., to be precise 8–28 evaluations per iteration) until the satisfaction of the sufficient decrease condition. This results in a total runtime of about 500 s, costly for a system of relatively small order. To conclude, direct applications of the gradient descent algorithm do not seem viable for systems of even modest order (e.g., a few thousands). On the same example, BFGS requires 239 evaluations of the objective $\mathcal{F}$ with a total runtime of about 240 s, so applying BFGS directly is also costly.

Table 1: This concerns the `iss` example with $r = 12$. The objective $\mathcal{F}^{(k)} := \mathcal{F}(A^{(k)}, B^{(k)}, C^{(k)}, D^{(k)}, E^{(k)})$, and the 2-norm of $\nabla \mathcal{F}^{(k)} := \nabla \mathcal{F}(A^{(k)}, B^{(k)}, C^{(k)}, D^{(k)}, E^{(k)})$ for the iterate $(A^{(k)}, B^{(k)}, C^{(k)}, D^{(k)}, E^{(k)})$ by the gradient descent method at the k th iteration are listed.

k	$\mathcal{F}^{(k)}$	$\ \nabla \mathcal{F}^{(k)}\ _2$
0	0.004470060020	1.000093488
1	0.004346739384	0.833556647
2	0.003609940202	1.000097230
3	0.003175718111	0.769359926
4	0.002975716755	1.000095596
5	0.002946113130	0.999918608
6	0.002697635041	0.844275929
7	0.002656707905	0.999952423
30	0.002415516341	0.803721909
31	0.002415479783	1.000008471
32	0.002415475189	0.803718441
33	0.002415456030	1.000008467
34	0.002415454613	0.803716708
35	0.002415444154	1.000008465
36	0.002415439844	0.803714645
37	0.002415438945	1.000008462




Figure 1: The locally minimal reduced system generated by the gradient descent method for the iss example and $r = 12$ is varied, and the error $\mathcal{F}$ is plotted as a function of the variation. In each one of the plots (a)–(h), only the indicated entry of one of the optimal coefficients $A^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}}, E^{\text{red}}$ is varied by amounts in $[-0.5, 0.5]$. Zero variation corresponds to the optimal reduced system.

3 A subspace framework

The computational difficulty in minimizing the objective $\mathcal{F}$ in (6) is due to the large order of the original system $S = (A, E, B, C, D)$. In this section, we propose to replace this system with a system of smaller order $S_q = (A_q, E_q, B_q, C_q, D_q)$ with the state-space representation

$$E_q x'_q(t) = A_q x_q(t) + B_q u(t), \quad y(t) = C_q x_q(t) + Du(t), \quad (8)$$

and solve the resulting $\mathcal{L}_\infty$ model reduction problem, that is minimize

$$\mathcal{F}_q(S^{\text{red}}) := \sup_{\omega \geq 0} \sigma_{\max}(H_q(i\omega) - H(i\omega; S^{\text{red}})) = \sup_{\omega \geq 0} \sigma_{\max}(\mathcal{E}_q(i\omega; S^{\text{red}})), \quad (9)$$

where

$$H_q(s) := C_q(sE_q - A_q)^{-1}B_q + D, \\ \mathcal{E}_q(s; S^{\text{red}}) := \begin{bmatrix} C_q & -C^{\text{red}} \end{bmatrix} \begin{bmatrix} sE_q - A_q & 0 \\ 0 & sE^{\text{red}} - A^{\text{red}} \end{bmatrix}^{-1} \begin{bmatrix} B_q \\ B^{\text{red}} \end{bmatrix} + (D - D^{\text{red}}).$$

The question that we need to address is how do we form a small system $S_q = (A_q, E_q, B_q, C_q, D)$ that is a good representative of the original system near a local minimizer of the original $\mathcal{L}_\infty$ model reduction problem.

Recall how pure Newton's method operates to minimize a function $f: \mathbb{R}^\ell \rightarrow \mathbb{R}$. It approximates f with a quadratic model, and finds a local minimizer $\tilde{x}$ of the quadratic model. Then, assuming f is twice differentiable at $\tilde{x}$, it refines the quadratic model so that the refined quadratic model Q satisfies $f(\tilde{x}) = Q(\tilde{x})$, $\nabla f(\tilde{x}) = \nabla Q(\tilde{x})$ and $\nabla^2 f(\tilde{x}) = \nabla^2 Q(\tilde{x})$. In the context of $\mathcal{L}_\infty$ model reduction, we view $\mathcal{F}_q$ as the model function for $\mathcal{F}$, even though $\mathcal{F}_q$ is not quadratic. We minimize $\mathcal{F}_q$ locally rather than $\mathcal{F}$, and refine the small system in (8) with the hope that the objective error function $\mathcal{F}_{q+1}$ of the refined system interpolates $\mathcal{F}$ and its first two derivatives at the computed minimizer of $\mathcal{F}_q$.

To obtain the small system in (8), we employ projection-based model reduction. In particular, let $\mathcal{V}_q, \mathcal{W}_q$ be two subspaces of $\mathbb{R}^n$ of equal dimension. Denoting with V_q, W_q matrices whose columns form orthonormal bases for $\mathcal{V}_q, \mathcal{W}_q$, the original system is approximated by

$$W_q^T(EV_q x'_q(t) - AV_q x_q(t) - Bu(t)) = 0, \quad y(t) = CV_q x_q(t) + Du(t),$$

giving rise to a system of the form (8) with

$$E_q = W_q^T E V_q, \quad A_q = W_q^T A V_q, \quad B_q = W_q^T B, \quad C_q = C V_q. \quad (10)$$

For the realization of the ideas in the previous paragraph, we need to be equipped with a tool that gives us the capability to interpolate $H(s)$ and its derivatives at a prescribed point in the complex plane with those of the transfer function for the small system. This tool is introduced in the next result, which follows from [35, Theorem 1].

Theorem 1. *Let $\mu \in \mathbb{C}$ be such that $A - \mu E$ is invertible. Suppose*

$$\bigoplus_{j=0}^{\kappa} \text{Re} \left[\{(A - \mu E)^{-1} E\}^j (A - \mu E)^{-1} B \right] \subseteq \mathcal{V}_q, \\ \bigoplus_{j=0}^{\kappa} \text{Im} \left[\{(A - \mu E)^{-1} E\}^j (A - \mu E)^{-1} B \right] \subseteq \mathcal{V}_q, \\ \bigoplus_{j=0}^{\kappa} \text{Re} \left[C(A - \mu E)^{-1} \{E(A - \mu E)^{-1}\}^j \right]^* \subseteq \mathcal{W}_q,$$

$$\bigoplus_{j=0}^{\kappa} \operatorname{Im} [C(A - \mu E)^{-1} \{E(A - \mu E)^{-1}\}^j]^* \subseteq \mathcal{W}_q.$$

Then, with A_q, E_q, B_q, C_q defined as in (10), if $A_q - \mu E_q$ is invertible, we have

- (i) $H(\mu) = H_q(\mu)$ and $H(\bar{\mu}) = H_q(\bar{\mu})$;
- (ii) $H^{(j)}(\mu) = H_q^{(j)}(\mu)$ and $H^{(j)}(\bar{\mu}) = H_q^{(j)}(\bar{\mu})$ for $j = 1, \dots, 2\kappa + 1$.

Our proposed subspace framework at iteration r first finds a minimizer of $\mathcal{F}_q(S_q^{\text{red}})$, say $S_q^{\text{red}} = (A_q^{\text{red}}, B_q^{\text{red}}, C_q^{\text{red}}, D_q^{\text{red}}, E_q^{\text{red}})$. This is followed by the computation of an $\omega_q \in \mathbb{R}$, $\omega_q \geq 0$ such that

$$\mathcal{F}(S_q^{\text{red}}) = \sup_{\omega \geq 0} \sigma_{\max}(H(i\omega) - H(i\omega; S_q^{\text{red}})) = \sigma_{\max}(H(i\omega_q) - H(i\omega_q; S_q^{\text{red}})).$$

Computing such an ω_q requires the large-scale $\mathcal{L}_\infty$ -norm computation in (6) but by replacing $S^{\text{red}} = (A^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}}, E^{\text{red}})$ with $S_q^{\text{red}} = (A_q^{\text{red}}, B_q^{\text{red}}, C_q^{\text{red}}, D_q^{\text{red}}, E_q^{\text{red}})$. Then subspaces are expanded so that H and its first three derivatives are interpolated at $i\omega_q$ by those of the transfer function for the small system. A formal description of the framework is given in Algorithm 1 below. As the subspaces $\mathcal{V}_q$ and $\mathcal{W}_q$ are required to be of equal dimension, the description assumes that the number of inputs and the outputs are equal, i.e., $m = p$. Even if it is omitted here for simplicity, it is straightforward to modify the directions $\tilde{V}_{q+1}, \tilde{W}_{q+1}$ in lines 11–12 to be added to the subspaces $\mathcal{V}_q, \mathcal{W}_q$ in order to deal with the systems for which $m \neq p$; see, e.g., [36, Lemma 3.1]. The final refinement step in line 15 aims at the satisfaction of the interpolation condition $\mathcal{F}(S_q^{\text{red}}) = \mathcal{F}_{q+1}(S_q^{\text{red}})$, as well as the interpolation conditions on the derivatives of $\mathcal{F}(S^{\text{red}})$ and $\mathcal{F}_{q+1}(S^{\text{red}})$ at S_q^{red} . This step is elaborated on in the next subsection.

Algorithm 1 (subspace framework for $\mathcal{L}_\infty$ model reduction).

Input: System $S = (A, E, B, C, D)$ as in (1), the order $r \in \mathbb{Z}^+$ of the reduced system sought, and an initial estimate

$S_0^{\text{red}} = (A_0^{\text{red}}, E_0^{\text{red}}, B_0^{\text{red}}, C_0^{\text{red}}, D_0^{\text{red}})$ of order r for a minimizer of $\mathcal{F}$ as in (6).

Output: Estimate $S_\star^{\text{red}} = (A_\star^{\text{red}}, E_\star^{\text{red}}, B_\star^{\text{red}}, C_\star^{\text{red}}, D_\star^{\text{red}})$ for a minimizer of $\mathcal{F}$ as in (6).

1: Choose the initial subspaces $\mathcal{V}_0, \mathcal{W}_0$ and orthonormal bases V_0, W_0 for them.

2: Form A_0, B_0, C_0, E_0 using (10), and let $S_0 = (A_0, E_0, B_0, C_0, D)$.

% main loop

3: **for** $q = 0, 1, \dots$ **do**

4: **if** $q \geq 1$ **then**

5: $S_q^{\text{red}} \leftarrow$ a minimizer of $\mathcal{F}_q(S^{\text{red}})$ (for $\mathcal{F}_q$ defined as in (9)).

6: **end if**

7: $\omega_q \leftarrow$ a maximizer of $\sigma(\omega; S_q^{\text{red}}) = \sigma_{\max}(H(i\omega) - H(i\omega; S_q^{\text{red}}))$ over $\omega \geq 0$.

8: **if** $q \geq 1$ **then**

9: **Return** if convergence has occurred with $S_\star^{\text{red}} \leftarrow S_q^{\text{red}}$.

10: **end if**

% expand the subspaces to interpolate at $i\omega_q$

11: $\tilde{V}_{q+1} \leftarrow [\operatorname{Re}[(i\omega_q E - A)^{-1}B] \quad \operatorname{Re}[(i\omega_q E - A)^{-1}E(i\omega_q E - A)^{-1}B] \quad \operatorname{Im}[(i\omega_q E - A)^{-1}B] \quad \operatorname{Im}[(i\omega_q E - A)^{-1}E(i\omega_q E - A)^{-1}B]]$.

12: $\tilde{W}_{q+1} \leftarrow [\operatorname{Re}[(i\omega_q E - A)^{-*}C^*] \quad \operatorname{Re}[(i\omega_q E - A)^{-*}E(i\omega_q E - A)^{-*}C^*] \quad \operatorname{Im}[(i\omega_q E - A)^{-*}C^*] \quad \operatorname{Im}[(i\omega_q E - A)^{-*}E(i\omega_q E - A)^{-*}C^*]]$.

13: $V_{q+1} \leftarrow \operatorname{orth}(\begin{bmatrix} V_q & \tilde{V}_{q+1} \end{bmatrix})$ and $W_{q+1} \leftarrow \operatorname{orth}(\begin{bmatrix} W_q & \tilde{W}_{q+1} \end{bmatrix})$.

% update the small system

14: Form $A_{q+1}, B_{q+1}, C_{q+1}, E_{q+1}$ using (10), and let $S_{q+1} = (A_{q+1}, E_{q+1}, B_{q+1}, C_{q+1}, D)$.

% refine the small system

15: Refine V_{q+1}, W_{q+1} and S_{q+1} if necessary (using Algorithm 2 below in Section 3.1).

16: **end for**

3.1 Refinement step

First we make a few observations regarding the relation between $\mathcal{F}(S_q^{\text{red}})$ and $\mathcal{F}_{q+1}(S_q^{\text{red}})$ at the q th subspace iteration in Algorithm 1 right before the refinement step.

At the q th iteration of Algorithm 1 right after line 14, by Theorem 1, we have

$$H(i\omega_q) = H_{q+1}(i\omega_q), \quad H^{(j)}(i\omega_q) = H_{q+1}^{(j)}(i\omega_q) \quad (11)$$

for $j = 1, 2, 3$ under the assumptions that $A - i\omega_q E$ and $A_{q+1} - i\omega_q E_{q+1}$ are invertible. Consequently, $H(i\omega_q) - H(i\omega_q; S_q^{\text{red}})$ and $H_{q+1}(i\omega_q) - H(i\omega_q; S_q^{\text{red}})$ are equal, and share the same set of left and right singular vectors. It immediately follows that setting

$$\sigma_{q+1}(\omega; S^{\text{red}}) := \sigma_{\max}(H_{q+1}(i\omega) - H(i\omega; S^{\text{red}})), \quad (12)$$

and recalling $\sigma(\omega; S^{\text{red}}) = \sigma_{\max}(H(i\omega) - H(i\omega; S^{\text{red}}))$, we have

$$\sigma(\omega_q; S_q^{\text{red}}) = \sigma_{q+1}(\omega_q; S_q^{\text{red}}). \quad (13)$$

Indeed, as the singular values and vectors of $H(i\omega_q) - H(i\omega_q; S_q^{\text{red}})$ and $H_{q+1}(i\omega_q) - H(i\omega_q; S_q^{\text{red}})$ are the same, and the first two derivatives of $H(i\omega) - H(i\omega; S_q^{\text{red}})$ and $H_{q+1}(i\omega) - H(i\omega; S_q^{\text{red}})$ at $\omega = \omega_q$ are equal due to (11), we also have

$$\frac{d^j \sigma}{d\omega^j}(\omega_q; S_q^{\text{red}}) = \frac{d^j \sigma_{q+1}}{d\omega^j}(\omega_q; S_q^{\text{red}}) \quad (14)$$

for $j = 1, 2$. Now ω_q is a global maximizer of $\sigma(\omega; S_q^{\text{red}})$ over ω implying

$$\frac{d\sigma}{d\omega}(\omega_q; S_q^{\text{red}}) = 0, \quad \frac{d^2 \sigma}{d\omega^2}(\omega_q; S_q^{\text{red}}) \leq 0.$$

Assuming that the last inequality on the second derivative above holds strictly, (14) implies ω_q is also a local maximizer of $\sigma_{q+1}(\omega; S_q^{\text{red}})$.

Regarding $\mathcal{F}(S_q^{\text{red}})$ and $\mathcal{F}_{q+1}(S_q^{\text{red}})$, the following relation always holds:

$$\mathcal{F}(S_q^{\text{red}}) = \sup_{\omega \geq 0} \sigma(\omega; S_q^{\text{red}}) = \sigma(\omega_q; S_q^{\text{red}}) = \sigma_{q+1}(\omega_q; S_q^{\text{red}}) \leq \sup_{\omega \geq 0} \sigma_{q+1}(\omega; S_q^{\text{red}}) = \mathcal{F}_{q+1}(S_q^{\text{red}}), \quad (15)$$

where the third equality is due to the interpolatory property in (13). As argued in the previous paragraph, the point ω_q is not only a global maximizer of $\sigma(\omega; S_q^{\text{red}})$, but also generically a local maximizer of $\sigma_{q+1}(\omega; S_q^{\text{red}})$. If it happens that ω_q is also a global maximizer of $\sigma_{q+1}(\omega; S_q^{\text{red}})$ beyond being a local maximizer, then the inequality in equation (15) above becomes an equality, and we have the interpolation property

$$\mathcal{F}(S_q^{\text{red}}) = \mathcal{F}_{q+1}(S_q^{\text{red}}). \quad (16)$$

In the refinement step, if it happens that ω_q is merely a local maximizer of $\sigma_{q+1}(\omega; S_q^{\text{red}})$, but not a global maximizer, then we find a global maximizer $\omega_q^{(0)}$ of $\sigma_{q+1}(\omega; S_q^{\text{red}})$ over $\omega \geq 0$ (equivalently compute the $\mathcal{L}_\infty$ norm of $H_{q+1}(\cdot) - H(\cdot; S_q^{\text{red}})$). Observe that finding such a global maximizer has a small computational cost, as the orders of S_{q+1} and S_q^{red} are small. Then, by employing Theorem 1, we expand the subspaces $\mathcal{V}_{q+1}$, $\mathcal{W}_{q+1}$ further so that the interpolatory properties are attained between $H_{q+1}(i\omega)$ after this refinement and $H(i\omega)$ at $\omega = \omega_q^{(0)}$, which in turn implies interpolatory properties between $\sigma_{q+1}(\omega; S_q^{\text{red}})$ and $\sigma(\omega; S_q^{\text{red}})$ at $\omega = \omega_q^{(0)}$. If ω_q after this refinement of S_{q+1} is still only a local maximizer of $\sigma_{q+1}(\omega; S_q^{\text{red}})$, but not a global maximizer, then we repeat

this refinement procedure of S_{q+1} up until ω_q becomes a global maximizer of $\sigma_{q+1}(\omega; S_q^{\text{red}})$ (in practice up to prescribed tolerances). A formal description of the refinement step is given below in Algorithm 2. For simplicity, in line 3 of Algorithm 2 it is assumed that ω_q is the unique global maximizer of $\sigma(\omega; S_q^{\text{red}})$. More generally, all of the global maximizers of $\sigma(\omega; S_q^{\text{red}})$ can be returned in line 7 of Algorithm 1 (e.g., by employing the level-set methods [37], [38] to compute the $\mathcal{L}_\infty$ norm), and whether $\omega_q^{(j)}$ is equal to any of these global maximizers can be checked in line 3 of Algorithm 2.

Assuming $\sigma(\omega; S_q^{\text{red}})$ is Lipschitz continuous, $\sigma_{q+1}(\omega; S_q^{\text{red}})$ is Lipschitz continuous with a uniform Lipschitz constant over the iterations of Algorithm 2, and the maximizer $\omega_q^{(j)}$ of $\sigma_{q+1}(\omega; S_q^{\text{red}})$ over $\omega \geq 0$ in line 2 of Algorithm 2 at every j is required to be in a prescribed bounded interval, the gap $|\omega_q^{(j)} - \omega_q|$ can be made less than any prescribed amount after finitely many iterations of Algorithm 2. At this point, the interpolation condition (16) is also met up to a multiple of the prescribed amount.

Algorithm 2 (refinement step).

```

1: for  $j = 0, 1, \dots$  do
2:    $\omega_q^{(j)} \leftarrow$  a maximizer of  $\sigma_{q+1}(\omega; S_q^{\text{red}}) = \sigma_{\max}(H_{q+1}(i\omega) - H(i\omega; S_q^{\text{red}}))$  over  $\omega \geq 0$ .
3:   if  $\omega_q^{(j)} = \omega_q$  (up to prescribed tolerances) then
4:     Terminate with  $V_{q+1}$ ,  $W_{q+1}$  and  $S_{q+1}$ .
5:   end if
6:   % expand the subspaces to interpolate at  $i\omega_q^{(j)}$ 
7:    $\tilde{V}_{q+1} \leftarrow \begin{bmatrix} \text{Re}[(i\omega_q^{(j)}E - A)^{-1}B] & \text{Re}[(i\omega_q^{(j)}E - A)^{-1}E(i\omega_q^{(j)}E - A)^{-1}B] & \text{Im}[(i\omega_q^{(j)}E - A)^{-1}B] & \text{Im}[(i\omega_q^{(j)}E - A)^{-1}E(i\omega_q^{(j)}E - A)^{-1}B] \end{bmatrix}$ .
8:    $\tilde{W}_{q+1} \leftarrow \begin{bmatrix} \text{Re}[(i\omega_q^{(j)}E - A)^{-*}C^*] & \text{Re}[(i\omega_q^{(j)}E - A)^{-*}E(i\omega_q^{(j)}E - A)^{-*}C^*] & \text{Im}[(i\omega_q^{(j)}E - A)^{-*}C^*] & \text{Im}[(i\omega_q^{(j)}E - A)^{-*}E(i\omega_q^{(j)}E - A)^{-*}C^*] \end{bmatrix}$ .
9:    $V_{q+1} \leftarrow \text{orth}(\begin{bmatrix} V_{q+1} & \tilde{V}_{q+1} \end{bmatrix})$  and  $W_{q+1} \leftarrow \text{orth}(\begin{bmatrix} W_{q+1} & \tilde{W}_{q+1} \end{bmatrix})$ .
10:  % update the small system
11:  Form  $A_{q+1}$ ,  $B_{q+1}$ ,  $C_{q+1}$ ,  $E_{q+1}$  using (10),
12:  and let  $S_{q+1} = (A_{q+1}, E_{q+1}, B_{q+1}, C_{q+1}, D)$ .
13: end for

```

4 Interpolation properties of the subspace framework

Suppose that ω_q is a global maximizer of $\sigma_{q+1}(\omega; S_q^{\text{red}})$ by the termination of the refinement step, in which case the interpolation condition (16) holds due to (15). It can be shown that, assuming $\mathcal{F}$ and $\mathcal{F}_{q+1}$ are twice differentiable at S_q^{red} , indeed all of the first two derivatives of $\mathcal{F}$ and $\mathcal{F}_{q+1}$ are equal at S_q^{red} as well. To this end, let x_1, x_2 be any two entries of the matrix variables $A^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}}, E^{\text{red}}$ of $\mathcal{F}$ and $\mathcal{F}_{q+1}$. Recalling

$$\mathcal{E}(i\omega; S^{\text{red}}) = H(i\omega) - H(i\omega; S^{\text{red}}), \quad \mathcal{E}_{q+1}(i\omega; S^{\text{red}}) = H_{q+1}(i\omega) - H(i\omega; S^{\text{red}}),$$

and by employing (11), it is apparent that

$$\mathcal{E}(i\omega_q; S_q^{\text{red}}) = \mathcal{E}_{q+1}(i\omega_q; S_q^{\text{red}}), \quad (17)$$

$$\frac{\partial \mathcal{E}}{\partial y}(i\omega_q; S_q^{\text{red}}) = \frac{\partial \mathcal{E}_{q+1}}{\partial y}(i\omega_q; S_q^{\text{red}}), \quad (18)$$

$$\frac{\partial^2 \mathcal{E}}{\partial y \partial z}(i\omega_q; S_q^{\text{red}}) = \frac{\partial^2 \mathcal{E}_{q+1}}{\partial y \partial z}(i\omega_q; S_q^{\text{red}}), \quad (19)$$

for all $y, z \in \{\omega, x_1, x_2\}$. By exploiting

$$\mathcal{F}(S^{\text{red}}) = \sup_{\omega \geq 0} \sigma_{\max}(\mathcal{E}(i\omega; S^{\text{red}})), \quad \mathcal{F}_{q+1}(S^{\text{red}}) = \sup_{\omega \geq 0} \sigma_{\max}(\mathcal{E}_{q+1}(i\omega; S^{\text{red}})),$$

and using implicit differentiation

$$\frac{\partial \mathcal{F}}{\partial y}(S_q^{\text{red}}) = \frac{\partial(\sigma_{\max} \circ \mathcal{E})}{\partial y}(\mathrm{i}\omega_q; S_q^{\text{red}}) = \frac{\partial(\sigma_{\max} \circ \mathcal{E}_{q+1})}{\partial y}(\mathrm{i}\omega_q; S_q^{\text{red}}) = \frac{\partial \mathcal{F}_{q+1}}{\partial y}(S_q^{\text{red}}) \quad (20)$$

for $y \in \{x_1, x_2\}$, where the second equality is due to (18), as well as (17) implying the fact that $\mathcal{E}(\mathrm{i}\omega_q; S_q^{\text{red}})$ and $\mathcal{E}_{q+1}(\mathrm{i}\omega_q; S_q^{\text{red}})$ have the same left and right singular vectors. We remark that, for the first and third equalities in (20), we use the fact ω_q is a global maximizer of $\sigma(\omega; S_q^{\text{red}}) = \sigma_{\max}(\mathcal{E}(\mathrm{i}\omega; S_q^{\text{red}}))$ and $\sigma_{q+1}(\omega; S_q^{\text{red}}) = \sigma_{\max}(\mathcal{E}_{q+1}(\mathrm{i}\omega; S_q^{\text{red}}))$, respectively.

Moreover, for any $y, z \in \{\omega, x_1, x_2\}$, we have

$$\frac{\partial^2(\sigma_{\max} \circ \mathcal{E})}{\partial y \partial z}(\mathrm{i}\omega_q; S_q^{\text{red}}) = \frac{\partial^2(\sigma_{\max} \circ \mathcal{E}_{q+1})}{\partial y \partial z}(\mathrm{i}\omega_q; S_q^{\text{red}})$$

due to (18) and (19) combined with the fact that $\mathcal{E}(\mathrm{i}\omega_q; S_q^{\text{red}})$, $\mathcal{E}_{q+1}(\mathrm{i}\omega_q; S_q^{\text{red}})$ have the same singular values and vectors due to (17). Consequently,

$$\begin{aligned} \frac{\partial^2 \mathcal{F}}{\partial y \partial z}(S_q^{\text{red}}) &= \frac{\partial^2(\sigma_{\max} \circ \mathcal{E})}{\partial y \partial z}(\mathrm{i}\omega_q; S_q^{\text{red}}) + \frac{\partial^2(\sigma_{\max} \circ \mathcal{E})}{\partial y \partial \omega}(\mathrm{i}\omega_q; S_q^{\text{red}}) \\ &\quad \times \left\{ -\frac{\partial^2(\sigma_{\max} \circ \mathcal{E})}{\partial \omega \partial z}(\mathrm{i}\omega_q; S_q^{\text{red}}) / \frac{\partial^2(\sigma_{\max} \circ \mathcal{E})}{\partial \omega^2}(\mathrm{i}\omega_q; S_q^{\text{red}}) \right\} \\ &= \frac{\partial^2(\sigma_{\max} \circ \mathcal{E}_{q+1})}{\partial y \partial z}(\mathrm{i}\omega_q; S_q^{\text{red}}) + \frac{\partial^2(\sigma_{\max} \circ \mathcal{E}_{q+1})}{\partial y \partial \omega}(\mathrm{i}\omega_q; S_q^{\text{red}}) \\ &\quad \times \left\{ -\frac{\partial^2(\sigma_{\max} \circ \mathcal{E}_{q+1})}{\partial \omega \partial z}(\mathrm{i}\omega_q; S_q^{\text{red}}) / \frac{\partial^2(\sigma_{\max} \circ \mathcal{E}_{q+1})}{\partial \omega^2}(\mathrm{i}\omega_q; S_q^{\text{red}}) \right\} = \frac{\partial^2 \mathcal{F}_{q+1}}{\partial y \partial z}(S_q^{\text{red}}) \quad (21) \end{aligned}$$

for any $y, z \in \{x_1, x_2\}$.

5 A quadratic convergence result regarding Algorithm 1

In this section, we establish quadratic convergence of the iterates of Algorithm 1 under additional smoothness assumptions. Here and in the next section, $\mathcal{D}^{r,m,p}$ denotes the set consisting of every descriptor system S^{red} of order r and index at most one with semi-simple poles, m inputs, p outputs. Throughout this section, we make use of the vectorization $\mathcal{V}(S^{\text{red}})$ of the system $S^{\text{red}} = (A^{\text{red}}, E^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}})$ defined as

$$\mathcal{V}(S^{\text{red}}) := [\text{vec}(A^{\text{red}})^T \text{ vec}(E^{\text{red}})^T \text{ vec}(B^{\text{red}})^T \text{ vec}(C^{\text{red}})^T \text{ vec}(D^{\text{red}})^T]^T, \quad (22)$$

where $\text{vec}(M)$ denotes the vector obtained by stacking up the columns of matrix M . The gradients $\nabla \mathcal{F}(S^{\text{red}})$ and $\nabla \mathcal{F}_{q+1}(S^{\text{red}})$ are vectors formed of the first partial derivatives of $\mathcal{F}(S^{\text{red}})$ and $\mathcal{F}_{q+1}(S^{\text{red}})$ based on the ordering of the variables, i.e., the entries of $A^{\text{red}}, E^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}}$, in the vectorization in (22). Similarly, $\nabla^2 \mathcal{F}(S^{\text{red}})$ and $\nabla^2 \mathcal{F}_{q+1}(S^{\text{red}})$ denote the Hessians of $\mathcal{F}(S^{\text{red}})$ and $\mathcal{F}_{q+1}(S^{\text{red}})$ based on the ordering of the variables according to (22).

We assume that there are two consecutive iterates S_q^{red} and S_{q+1}^{red} of Algorithm 1 that are sufficiently close to a local minimizer S_*^{red} of $\mathcal{F}(S^{\text{red}})$. Moreover, we silently assume throughout that the interpolation properties in (20) and (21) hold at S_q^{red} . We also keep the assumption stated below that guarantees that $\mathcal{F}(S^{\text{red}})$ is real analytic at S_*^{red} .

Assumption 2. The maximum of $\sigma(\omega; S_*^{\text{red}})$ over all $\omega \geq 0$ is attained at a unique ω_* . Furthermore, $\sigma(\omega_*; S_*^{\text{red}}) = \sigma_{\max}(\mathcal{E}(\mathrm{i}\omega_*; S_*^{\text{red}}))$ is a simple singular value of $\mathcal{E}(\mathrm{i}\omega_*; S_*^{\text{red}})$.

An assumption regarding the smoothness of $\mathcal{F}_{q+1}(S^{\text{red}})$ that we rely on is given next. Recalling $\|v\|_2$ for a vector v denotes the 2-norm of v , we make use of the distance $\|\tilde{S}^{\text{red}} - \hat{S}^{\text{red}}\| := \|\mathcal{V}(\tilde{S}^{\text{red}}) - \mathcal{V}(\hat{S}^{\text{red}})\|_2$ for systems $\tilde{S}^{\text{red}}, \hat{S}^{\text{red}} \in \mathcal{D}^{r,m,p}$, and the ball $B(\hat{S}^{\text{red}}, \delta) := \{\tilde{S}^{\text{red}} \in \mathcal{D}^{r,m,p} \mid \|\tilde{S}^{\text{red}} - \hat{S}^{\text{red}}\| < \delta\}$ for a system $\hat{S}^{\text{red}} \in \mathcal{D}^{r,m,p}$ and positive real number δ . We remark that part (i) of Assumption 3 guarantees that $\mathcal{F}_{q+1}(S^{\text{red}})$ is real-analytic in the ball $B(S_q^{\text{red}}, \delta_q)$, and so three times differentiable in this ball, which we depend on in part (ii) of Assumption 3.

Assumption 3.

- (i) For every $S^{\text{red}} \in B(S_q^{\text{red}}, \delta_q)$ with $\delta_q := \|S_{q+1}^{\text{red}} - S_q^{\text{red}}\|$ the following conditions hold:
 - the maximum of $\sigma_{q+1}(\omega; S^{\text{red}})$ defined in (12) over all $\omega \geq 0$ is attained at a unique point, say at $\tilde{\omega}$;
 - the singular value $\sigma_{q+1}(\tilde{\omega}; S^{\text{red}}) = \sigma_{\max}(\mathcal{E}_{q+1}(i\tilde{\omega}; S^{\text{red}}))$ of $\mathcal{E}_{q+1}(i\tilde{\omega}; S^{\text{red}})$ is simple.
- (ii) Moreover, all of the third derivatives of $\mathcal{F}_{q+1}(S^{\text{red}})$ can be bounded by quantities independent of q at all $S^{\text{red}} \in B(S_q^{\text{red}}, \delta_q)$.

We state and prove the main result that relates $\|S_{q+1}^{\text{red}} - S_*^{\text{red}}\|$ and $\|S_q^{\text{red}} - S_*^{\text{red}}\|$ below.

Theorem 4. Suppose that two consecutive iterates S_q^{red} and S_{q+1}^{red} of Algorithm 1 are sufficiently close to a local minimizer S_*^{red} of $\mathcal{F}(S^{\text{red}})$. Moreover, suppose Assumptions 2, 3 hold, and $\nabla^2 \mathcal{F}(S_*^{\text{red}})$ is invertible. Then there is a constant C independent of q such that

$$\|S_{q+1}^{\text{red}} - S_*^{\text{red}}\| \leq C \|S_q^{\text{red}} - S_*^{\text{red}}\|^2.$$

Proof. By continuity $\sigma(\omega; S^{\text{red}}) = \sigma_{\max}(\mathcal{E}(i\omega; S^{\text{red}}))$ remains simple at all ω and $S^{\text{red}} \in \mathcal{D}^{r,m,p}$ such that ω is sufficiently close to ω_* and S^{red} is sufficiently close to S_*^{red} , where ω_* is as in Assumption 2. Thus, by the analytic implicit function theorem, there is $\tilde{\delta} > 0$ such that $\mathcal{F}(S^{\text{red}})$ is real analytic at all $S^{\text{red}} \in B(S_*^{\text{red}}, \tilde{\delta})$ (see, e.g., [39, Lemma 16] for the details in the analogous context of the distance instability). By the assumption that $\nabla^2 \mathcal{F}(S_*^{\text{red}})$ is invertible, and continuity of the second partial derivatives of $\mathcal{F}(S^{\text{red}})$ in $B(S_*^{\text{red}}, \tilde{\delta})$, the Hessian $\nabla^2 \mathcal{F}(S^{\text{red}})$ remains invertible in a ball $B(S_*^{\text{red}}, \delta)$ for some $\delta < \tilde{\delta}$. Furthermore, without loss of generality, we assume $S_q^{\text{red}}, S_{q+1}^{\text{red}}$ are close enough to S_*^{red} so that $S_q^{\text{red}}, S_{q+1}^{\text{red}} \in B(S_*^{\text{red}}, \delta)$, and the ball $B(S_q^{\text{red}}, \delta_q)$ in Assumption 3 is contained in $B(S_*^{\text{red}}, \delta)$. We let $\beta := \min_{S^{\text{red}} \in B(S_*^{\text{red}}, \delta)} \sigma_{\min}(\nabla^2 \mathcal{F}(S^{\text{red}})) > 0$, and note that $\nabla^2 \mathcal{F}(S^{\text{red}})$ is Lipschitz continuous in $B(S_*^{\text{red}}, \delta)$, say with the Lipschitz constant γ .

By an application of Taylor's theorem with integral remainder, we have

$$\begin{aligned} 0 &= \nabla \mathcal{F}(S_*^{\text{red}}) = \nabla \mathcal{F}(S_q^{\text{red}}) + \int_0^1 \nabla^2 \mathcal{F}(S_q^{\text{red}} + t(S_*^{\text{red}} - S_q^{\text{red}})) (\mathcal{V}(S_*^{\text{red}}) - \mathcal{V}(S_q^{\text{red}})) \, dt \\ &= \nabla \mathcal{F}(S_q^{\text{red}}) + \nabla^2 \mathcal{F}(S_q^{\text{red}}) (\mathcal{V}(S_*^{\text{red}}) - \mathcal{V}(S_q^{\text{red}})) + \mathcal{O}(\|S_*^{\text{red}} - S_q^{\text{red}}\|^2), \end{aligned} \quad (23)$$

where, for the third equality, we have used the Lipschitz continuity of $\nabla^2 \mathcal{F}(S^{\text{red}})$ in $B(S_*^{\text{red}}, \delta)$. Additionally, by Taylor's theorem with second order Lagrange remainder,

$$\begin{aligned} 0 &= \nabla \mathcal{F}_{q+1}(S_{q+1}^{\text{red}}) = \nabla \mathcal{F}_{q+1}(S_q^{\text{red}}) + \nabla^2 \mathcal{F}_{q+1}(S_q^{\text{red}}) (\mathcal{V}(S_{q+1}^{\text{red}}) - \mathcal{V}(S_q^{\text{red}})) + \mathcal{O}(\|S_{q+1}^{\text{red}} - S_q^{\text{red}}\|^2) \\ &= \nabla \mathcal{F}(S_q^{\text{red}}) + \nabla^2 \mathcal{F}(S_q^{\text{red}}) (\mathcal{V}(S_{q+1}^{\text{red}}) - \mathcal{V}(S_q^{\text{red}})) + \mathcal{O}(\|S_{q+1}^{\text{red}} - S_q^{\text{red}}\|^2), \end{aligned} \quad (24)$$

where the third equality is due to $\nabla \mathcal{F}_{q+1}(S_q^{\text{red}}) = \nabla \mathcal{F}(S_q^{\text{red}})$ and $\nabla^2 \mathcal{F}_{q+1}(S_q^{\text{red}}) = \nabla^2 \mathcal{F}(S_q^{\text{red}})$, that are consequences of (20) and (21).

By employing (24) in (23), we deduce

$$\nabla^2 \mathcal{F}(S_q^{\text{red}})(\mathcal{V}(S_*^{\text{red}}) - \mathcal{V}(S_{q+1}^{\text{red}})) = \mathcal{O}(\|S_{q+1}^{\text{red}} - S_q^{\text{red}}\|^2) + \mathcal{O}(\|S_*^{\text{red}} - S_q^{\text{red}}\|^2).$$

Finally, noting $\|\nabla^2 \mathcal{F}(S_q^{\text{red}})(\mathcal{V}(S_*^{\text{red}}) - \mathcal{V}(S_{q+1}^{\text{red}}))\| \geq \beta \|S_*^{\text{red}} - S_{q+1}^{\text{red}}\|$, from the last equality we obtain

$$\|S_{q+1}^{\text{red}} - S_*^{\text{red}}\| \leq \mathcal{O}(\|S_q^{\text{red}} - S_*^{\text{red}}\|^2)$$

as desired. $\square$

6 Dealing with asymptotic stability constraints

In many applications, in addition to the objective that the reduced-order system sought $S^{\text{red}} = (A^{\text{red}}, E^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}})$ of order r is close to the original system with respect to the $\mathcal{L}_\infty$ norm, the system S^{red} sought may also be required to be asymptotically stable; see, e.g., [40]–[42]. As we search through reduced-order systems of index at most one, it follows from the Weierstrass canonical form [43] of $(A^{\text{red}}, E^{\text{red}})$ that the asymptotic stability requirement is equivalent to having all of the poles of the system in the open left half of the complex plane, i.e., to the condition $\alpha(A^{\text{red}}, E^{\text{red}}) < 0$, where $\alpha(A^{\text{red}}, E^{\text{red}})$ is the spectral abscissa of the pencil $L(s) = A^{\text{red}} - sE^{\text{red}}$ defined by

$$\alpha(A^{\text{red}}, E^{\text{red}}) := \max\{\operatorname{Re}(z) \mid z \in \mathbb{C} \text{ s.t. } \det(A - zE) = 0\}.$$

In this setting, with $\mathcal{F}(S^{\text{red}})$ defined as in (6), rather than the unconstrained minimization of $\mathcal{F}(S^{\text{red}})$ over all descriptor systems $S^{\text{red}} \in \mathcal{D}^{r,m,p}$, it may be desirable to solve the constrained minimization problem

$$\min\{\mathcal{F}(S^{\text{red}}) : S^{\text{red}} \in \mathcal{D}^{r,m,p} \text{ s.t. } \alpha(A^{\text{red}}, E^{\text{red}}) \leq \beta\} \quad (25)$$

for a prescribed negative real number β , where $\mathcal{D}^{r,m,p}$ denotes the set of all descriptor systems of order r and index at most one with semi-simple poles, m inputs, p outputs.

Extension of the proposed subspace framework, that is Algorithm 1, to deal with the constrained minimization problem in (25) rather than the unconstrained minimization of $\mathcal{F}(S^{\text{red}})$ is straightforward. The only difference in Algorithm 1 is in line 5, where S_q^{red} is no longer a minimizer of $\mathcal{F}_q(S^{\text{red}})$, but instead a minimizer of the constrained problem

$$\min\{\mathcal{F}_q(S^{\text{red}}) : S^{\text{red}} \in \mathcal{D}^{r,m,p} \text{ s.t. } \alpha(A^{\text{red}}, E^{\text{red}}) \leq \beta\} \quad (26)$$

for the reduced function $\mathcal{F}_q(S^{\text{red}})$ as in (9). The problem in (26) involves only the small systems S_q as well as S^{red} , and is solvable by means of Newton-method based approaches (e.g., using [31], Procedure 2). Such a Newton-method based approach makes use of the gradient of the constraint function $C(S^{\text{red}}) := \alpha(A^{\text{red}}, E^{\text{red}}) - \beta$, in addition to the gradient of the objective $\mathcal{F}_q(S^{\text{red}})$. Let λ be the rightmost eigenvalue of the pencil $L(s) = A^{\text{red}} - sE^{\text{red}}$ with u and v denoting a pair of corresponding left and right eigenvectors normalized so that $u^* E^{\text{red}} v = 1$, and assume λ is a simple eigenvalue and the unique rightmost eigenvalue of $L(s)$, which ensure the differentiability of $C(S^{\text{red}})$. Then, by differentiating the equation $A^{\text{red}} v = \lambda E^{\text{red}} v$ with respect to the entries of A^{red} and E^{red} and multiplying with u^* from left, the partial derivatives of $C(S^{\text{red}})$ are given by

$$\frac{\partial C}{\partial A_{ij}^{\text{red}}}(S^{\text{red}}) = \operatorname{Re}(\bar{u}_i v_j), \quad \frac{\partial C}{\partial E_{jj}^{\text{red}}}(S^{\text{red}}) = -\operatorname{Re}(\lambda \bar{u}_j v_j),$$

where $\operatorname{Re}(z)$ denotes the real part of a complex scalar $z \in \mathbb{C}$, while A_{ij}^{red} is the subdiagonal, superdiagonal or diagonal entry of the matrix variable A^{red} at position (i, j) , and E_{jj}^{red} is the diagonal entry of E^{red} at position (j, j) .

We remark that, assuming ω_q is again a global maximizer of $\sigma_{q+1}(\omega; S_q^{\text{red}})$ after the refinement step, the interpolation properties

$$\mathcal{F}(S_q^{\text{red}}) = \mathcal{F}_{q+1}(S_q^{\text{red}}), \quad \nabla \mathcal{F}(S_q^{\text{red}}) = \nabla \mathcal{F}_{q+1}(S_q^{\text{red}}), \quad \nabla^2 \mathcal{F}(S_q^{\text{red}}) = \nabla^2 \mathcal{F}_{q+1}(S_q^{\text{red}})$$

still hold. Moreover, if a logarithmic barrier approach is adopted for the solution of the constrained problems, then, in essence, constrained problems are turned into unconstrained problems by lifting the constraints to the objective via the logarithmic barrier functions

$$L^\mu(S^{\text{red}}) = \mathcal{F}(S^{\text{red}}) - \mu \log(\beta - \alpha(A^{\text{red}}, E^{\text{red}})), \quad L_q^\mu(S^{\text{red}}) = \mathcal{F}_q(S^{\text{red}}) - \mu \log(\beta - \alpha(A^{\text{red}}, E^{\text{red}}))$$

associated with problems (25), (26), respectively, where $\log(\cdot)$ denotes the natural logarithm of its parameter, μ is a positive real number and represent the barrier parameter. In this case, the interpolation properties extend to the logarithmic barrier functions as well. In particular, we have

$$L^\mu(S_q^{\text{red}}) = L_{q+1}^\mu(S_q^{\text{red}}), \quad \nabla L^\mu(S_q^{\text{red}}) = \nabla L_{q+1}^\mu(S_q^{\text{red}}), \quad \nabla^2 L^\mu(S_q^{\text{red}}) = \nabla^2 L_{q+1}^\mu(S_q^{\text{red}})$$

for every positive real number μ .

There are alternatives to our approach above to enforce asymptotic stability on the reduced system. One possibility is to define the $\mathcal{H}_\infty$ error objective as ∞ for reduced systems that are not asymptotically stable, and deal with the unconstrained minimization of the resulting objective. It is then essential to start with a reduced system that is asymptotically stable, and the sufficient decrease condition in the line-search keeps the iterates away from reduced systems that are not asymptotically stable. Such an approach is adopted in Ref. [21] in the design of a controller with a minimal $\mathcal{H}_\infty$ objective for the closed-loop system. Another alternative is to ensure asymptotic stability by requiring that the reduced system has a certain structure, e.g., port-Hamiltonian structure following the practice in Ref. [26]. Once again this results in the unconstrained minimization of the $\mathcal{H}_\infty$ error objective but over structured systems depending on additional optimization parameters. We opt to depend on the constrained optimization formulation in (25) due to the availability of software such as GRANSO [31] to deal with constraints.

7 Practical issues

Here we spell out a few practical issues regarding Algorithm 1 such as how we form the initial systems S_0, S_0^{red} , when we terminate, the details of bases for projection subspaces, solutions of reduced $\mathcal{L}_\infty$ -norm minimization problems, and $\mathcal{L}_\infty$ -norm computations.

7.1 Initialization

The initial subspaces $\mathcal{V}_0, \mathcal{W}_0$ (in line 1 of Algorithm 1) are chosen so that H_0 , the transfer function of S_0 , interpolates H at the imaginary parts of a prescribed number of dominant poles of H . Formally, suppose H does not have any pole on the imaginary axis, and, for a prescribed ℓ , let $s_1, \dots, s_\ell \in \mathbb{C}$ be the most dominant ℓ poles of H with nonnegative imaginary parts (as the poles of H appear in complex conjugate pairs such that any two complex conjugate poles have the same dominance metric). We set

$$\mathcal{V}_0 = \bigoplus_{k=1}^{\ell} \bigoplus_{j=0}^1 \left\{ \text{Re} \left[\{ (A - i \text{Im}(s_k)E)^{-1} E \}^j (A - i \text{Im}(s_k)E)^{-1} B \right] \oplus \text{Im} \left[\{ (A - i \text{Im}(s_k)E)^{-1} E \}^j (A - i \text{Im}(s_k)E)^{-1} B \right] \right\},$$

$$\mathcal{W}_0 = \bigoplus_{k=1}^{\ell} \bigoplus_{j=0}^1 \left\{ \text{Re} \left[C(A - i \text{Im}(s_k)E)^{-1} \{ E(A - i \text{Im}(s_k)E)^{-1} \}^j \right]^* \oplus \text{Im} \left[C(A - i \text{Im}(s_k)E)^{-1} \{ E(A - i \text{Im}(s_k)E)^{-1} \}^j \right]^* \right\},$$

where $\text{Im}(s_k)$ denotes the imaginary part of $s_k \in \mathbb{C}$. Theorem 1 ensures that

$$\begin{aligned} H(i \text{Im}(s_k)) &= H_0(i \text{Im}(s_k)), & H^{(j)}(i \text{Im}(s_k)) &= H_0^{(j)}(i \text{Im}(s_k)) \\ H(-i \text{Im}(s_k)) &= H_0(-i \text{Im}(s_k)), & H^{(j)}(-i \text{Im}(s_k)) &= H_0^{(j)}(-i \text{Im}(s_k)) \end{aligned}$$

for $j = 1, 2, 3$ and $k = 1, \dots, \ell$.

Additionally, at every subspace iteration with $q > 0$, an initial point is needed for the solution of the minimization problem in line 5 of Algorithm 1 regardless of how it is solved, e.g., via gradient descent or BFGS. This initialization carries significance, as it affects which local minimizer of $\mathcal{F}_q$ is to be converged. At a subspace iteration with $q > 1$, the minimizer is initialized with the optimal reduced system from the previous iteration, that is with S_{q-1}^{red} . For iteration $q = 1$, initial S_0^{red} of order r must be supplied to Algorithm 1. We set S_0^{red} as either

- the model of order r obtained from an application of balanced truncation, or
- the model of order r whose transfer function interpolates H at the imaginary parts of a prescribed number of dominant poles of H .

For the latter choice, we remark that the number of dominant poles used to form S_0 is strictly greater than the number of dominant poles used to form this initial model S_0^{red} for the minimizer. For either choice, we make sure $\dim \mathcal{V}_0 = \dim \mathcal{W}_0 > r$ by using sufficiently many dominant poles of H when forming $S_0 = (A_0, E_0, B_0, C_0, D)$.

An issue that requires attention is that $S_0^{\text{red}} = (A_0^{\text{red}}, E_0^{\text{red}}, B_0^{\text{red}}, C_0^{\text{red}}, D_0^{\text{red}})$ must be such that A_0^{red} is tridiagonal and E_0^{red} is diagonal, whereas balanced truncation or the interpolatory approach yields the system $(\hat{A}, \hat{E}, \hat{B}, \hat{C}, \hat{D})$ of order r such that $\hat{A}$ and $\hat{E}$ do not necessarily have these structures. Let us suppose that the system $(\hat{A}, \hat{E}, \hat{B}, \hat{C}, \hat{D})$ has semi-simple poles, and $\hat{E}$ is invertible. Then we can first compute the eigenvalues of the $r \times r$ pencil $\hat{L}(s) = \hat{A} - s\hat{E}$, and form a block diagonal real matrix T_2 with 2×2 and 1×1 blocks along its diagonal that have the same eigenvalues as $\hat{L}$. The 2×2 and 1×1 blocks of T_2 on its diagonal correspond to a conjugate pair of complex eigenvalues and real eigenvalues of $\hat{L}$, respectively. Here we remark that T_2 is an $r \times r$ matrix. Hence, we can compute its eigenvalue decomposition

$$T_2 = U \Lambda U^{-1} \quad (27)$$

for a diagonal matrix Λ and invertible U at ease. We also have the eigenvalue decomposition

$$\hat{E}^{-1}\hat{A} = V \Lambda V^{-1}$$

at hand. Note that the middle factors in eigenvalue decompositions above are the same, as T_2 has the same eigenvalues as the pencil $\hat{L}$, which in turn has the same eigenvalues as $\hat{E}^{-1}\hat{A}$. But then

$$\begin{aligned} \hat{H}(s) &:= \hat{C}(s\hat{E} - \hat{A})^{-1}\hat{B} + \hat{D} = \hat{C}(sI - \hat{E}^{-1}\hat{A})^{-1}\hat{E}^{-1}\hat{B} + \hat{D} = \hat{C}(sI - V \Lambda V^{-1})^{-1}\hat{E}^{-1}\hat{B} + \hat{D} \\ &= (\hat{C}V)(sI - \Lambda)^{-1}(V^{-1}\hat{E}^{-1}\hat{B}) + \hat{D} = (\hat{C}V)(sI - U^{-1}T_2U)^{-1}(V^{-1}\hat{E}^{-1}\hat{B}) + \hat{D} \\ &= (\hat{C}VU^{-1})(sI - T_2)^{-1}(UV^{-1}\hat{E}^{-1}\hat{B}) + \hat{D}. \end{aligned}$$

Hence, we can use

$$A_0^{\text{red}} := T_2, \quad E_0^{\text{red}} := I, \quad B_0^{\text{red}} := UV^{-1}\hat{E}^{-1}\hat{B}, \quad C_0^{\text{red}} := \hat{C}VU^{-1}, \quad D_0^{\text{red}} := \hat{D}$$

as the matrices of the initial system S_0^{red} .

Remark 1. The invertibility assumption above on $\hat{E}$ amounts to initializing the optimization in line 5 of Algorithm 1 at the first subspace iteration (for $q = 1$) with a descriptor system S_0^{red} with all finite poles, a property stronger than our general assumption that the reduced system sought is at most index one. However, this invertibility assumption initially is not restrictive. In particular, optimization is still over at most index one systems, and is permitted to yield an index one system $S_q^{\text{red}} = (A_q^{\text{red}}, E_q^{\text{red}}, B_q^{\text{red}}, C_q^{\text{red}}, D_q^{\text{red}})$ for $q \geq 1$ with singular E_q^{red} .

Remark 2. When the pencil $\hat{L}(s) = \hat{A} - s\hat{E}$ has complex eigenvalues, the matrices V and U from the eigenvalue decompositions above are complex. Yet, T_2 and U satisfying (27) can be chosen so that the product VU^{-1} is real, ensuring in turn that B_0^{red} and C_0^{red} are real matrices.

To see VU^{-1} is real for particular choices of T_2 , U , suppose $a \pm ib$ consist of a pair of complex conjugate eigenvalues of $\hat{L}$, and suppose they appear as the ℓ th, $(\ell + 1)$ th entries along the diagonal of Λ . Moreover, suppose $u \pm iv$ with $u, v \in \mathbb{R}^r$ are eigenvectors of $\hat{L}$ corresponding to $a \pm ib$ appearing on the ℓ th, $(\ell + 1)$ th columns of V . Then we have

$$\begin{bmatrix} a & b \\ -b & a \end{bmatrix} = \underbrace{\begin{bmatrix} 1/\sqrt{2} & 1/\sqrt{2} \\ i/\sqrt{2} & -i/\sqrt{2} \end{bmatrix}}_Z \begin{bmatrix} a+ib & 0 \\ 0 & a-ib \end{bmatrix} \underbrace{\begin{bmatrix} 1/\sqrt{2} & -i/\sqrt{2} \\ 1/\sqrt{2} & i/\sqrt{2} \end{bmatrix}}_{Z^{-1}=Z^*}, \quad (28)$$

so the real matrix on the left has the eigenvalues $a \pm ib$, and we use it on the 2×2 diagonal block of T_2 with row and column indices $\ell, \ell + 1$. Using the decomposition in (28), in particular Z^{-1} , we can set up U so that $(1/\sqrt{2})(e_\ell + e_{\ell+1})$, $(i/\sqrt{2})(-e_\ell + e_{\ell+1})$ are on the ℓ th, $(\ell + 1)$ th columns of U^{-1} , where e_ℓ is the ℓ th column of the $r \times r$ identity matrix. Recalling $u \pm iv$ are the eigenvectors on the ℓ th, $(\ell + 1)$ th columns of V , the ℓ th, $(\ell + 1)$ th columns of VU^{-1} are real and equal to $\sqrt{2}u$, $\sqrt{2}v$, respectively.

7.2 Termination

The termination in line 9 of Algorithm 1 is determined based on the values of $\|H(\cdot) - H(\cdot; S_q^{\text{red}})\|_{\mathcal{L}_\infty}$ at two consecutive subspace iterations. The error $\|H(\cdot) - H(\cdot; S_q^{\text{red}})\|_{\mathcal{L}_\infty}$ is readily available at the q th subspace iteration after line 7, as

$$\|H(\cdot) - H(\cdot; S_q^{\text{red}})\|_{\mathcal{L}_\infty} = \sigma_{\max}(H(i\omega_q) - H(i\omega_q; S_q^{\text{red}})).$$

To be precise, we terminate at the q th subspace iteration in line 9 if $q \geq 1$ and

$$\left| \|H(\cdot) - H(\cdot; S_q^{\text{red}})\|_{\mathcal{L}_\infty} - \|H(\cdot) - H(\cdot; S_{q-1}^{\text{red}})\|_{\mathcal{L}_\infty} \right| \leq \text{tol} \|H(\cdot) - H(\cdot; S_q^{\text{red}})\|_{\mathcal{L}_\infty} \quad (29)$$

for a prescribed tolerance tol .

The termination condition for the minimizer to solve the minimization problem in line 5 of Algorithm 1 also requires some care. Recall that the objective $\mathcal{F}_q$ here is nonsmooth, and, as a result, the norms of the gradients of $\mathcal{F}_q$ at the iterates generated by the minimizer do not have to converge to zero. Instead, the minimizer is terminated if the line-search fails (to return a point that causes sufficient decrease), or the decrease in $\mathcal{F}_q$ at two consecutive iterates is less than $\varepsilon \cdot \text{tol}$ in a relative sense, where we use the same tolerance tol as in (29) and ε is a real number in $(0, 0.5)$.

As for the termination condition of the refinement step (i.e., the condition in line 3 of Algorithm 2) employed in practice, we rely on

$$|\omega_q^{(j)} - \omega_q| \leq \text{tol} |\omega_q|,$$

where we again use the same tolerance tol as in (29).

7.3 Orthonormalization of the bases for the subspaces

Keeping the bases for the subspaces $\mathcal{V}_q$, $\mathcal{W}_q$ (i.e., the columns of V_q , W_q) orthonormal improves the robustness of the algorithm against the rounding errors. For instance, then the system matrices A_q, B_q, C_q, E_q can be formed more accurately in the presence of rounding errors.

This orthonormality property of the bases is attained in line 13 of Algorithm 1, as well as line 8 of Algorithm 2. In line 13 of Algorithm 1, V_q and W_q are already orthonormal bases for $\mathcal{V}_q$ and $\mathcal{W}_q$. The expansion directions $\tilde{V}_{q+1}$,

$\tilde{W}_{q+1}$ to be included in the subspaces to obtain the expanded subspaces $\mathcal{V}_{q+1}$, $\mathcal{W}_{q+1}$ have to be orthonormalized with respect to the existing orthonormal bases V_q , W_q . This is achieved in practice by executing

$$\tilde{V}_{q+1} \leftarrow \tilde{V}_{q+1} - V_q \left(V_q^T \tilde{V}_{q+1} \right), \quad \tilde{W}_{q+1} \leftarrow \tilde{W}_{q+1} - W_q \left(W_q^T \tilde{W}_{q+1} \right). \quad (30)$$

Near convergence the interpolation points $i\omega_q$ of Algorithm 1 start not changing by much in consecutive iterations. This results in the new expansion directions $\tilde{V}_{q+1}$, $\tilde{W}_{q+1}$ that are nearly contained in the existing subspaces $\mathcal{V}_q$, $\mathcal{W}_q$. In this case, the orthonormalization in (30) of $\tilde{V}_{q+1}$, $\tilde{W}_{q+1}$ with respect to existing V_q , W_q suffers from cancellation type rounding errors. Applying the orthonormalization in (30) several times improves the accuracy, and result in directions $\tilde{V}_{q+1}$, $\tilde{W}_{q+1}$ that are better orthonormalized against V_q , W_q . In practice, we apply (30) a few times, e.g., 3–4 times, then orthonormalize the resulting $\tilde{V}_{q+1}$, $\tilde{W}_{q+1}$ via the Gram–Schmidt procedure, and take $V_{q+1} = [V_q \ \tilde{V}_{q+1}]$, $W_{q+1} = [W_q \ \tilde{W}_{q+1}]$ as the matrices whose columns form orthonormal bases for $\mathcal{V}_{q+1}$, $\mathcal{W}_{q+1}$. In line 8 of Algorithm 2, the columns of V_{q+1} and W_{q+1} are similarly orthonormalized. We ultimately use V_{q+1} , W_{q+1} when forming the system S_{q+1} in line 14 of Algorithm 1, and in line 9 of Algorithm 2.

7.4 Solution of the reduced $\mathcal{L}_\infty$ -norm minimization problem

We use BFGS to minimize the reduced $\mathcal{L}_\infty$ objective $\mathcal{F}_q(S^{\text{red}})$ in line 5 of Algorithm 1. To be precise, we have explored two alternatives here; a small variation of a MATLAB implementation of a line-search BFGS due to Michael L. Overton making use of weak Wolfe conditions, and GRANSO [31]. The former is only meant for unconstrained problems when we do not impose the asymptotic stability constraints described in Section 7, whereas the asymptotic stability constraints in Section 6 are also incorporated into this optimization when we use GRANSO.

7.5 Computation of the $\mathcal{L}_\infty$ norm

Algorithm 1 in line 7 requires the computation of the $\mathcal{L}_\infty$ norm of a system whose order is the sum of the order of the original system S and r . If the original system does not have large order, we use the built-in norm command in MATLAB for these $\mathcal{L}_\infty$ -norm computations. Otherwise, if the original system has large order, we use the subspace framework introduced in Ref. [44] for the large-scale $\mathcal{L}_\infty$ -norm computations. Additionally, the minimization of the reduced $\mathcal{L}_\infty$ objective in line 5 via BFGS requires small-scale $\mathcal{L}_\infty$ -norm computations, which we carry out using the getPeakGain command in MATLAB. The use of the routines in the ROSTAPACK package [45], [46], especially the routine normTfMaxPeak, instead of the built-in MATLAB routines norm, getPeakGain may possibly improve the efficiency.

8 Numerical results

In this section, we report the results of numerical experiments performed with a MATLAB implementation of Algorithm 1 taking also the practical issues indicated in the previous section into account. Sections 8.1 and 8.2 concern experiments on rather smaller order systems, Section 8.3 concerns experiments on a system of medium order, while the results of experiments on several large-order systems are reported in Section 8.4. All of the experiments are conducted in MATLAB 2020b on an iMac with Mac OS 12.1 operating system, Intel[®] Core[™] i5-9600K CPU and 32GB RAM.

The numerical experiments are performed using the variation of the MATLAB implementation of BFGS due to Michael L. Overton for the solution of the reduced $\mathcal{L}_\infty$ -norm minimization problems. Hence, the asymptotic stability constraints are not imposed. The original systems in all of the experiments in Sections 8.1–8.3 concerning small- to medium-order systems are asymptotically stable, and the computed optimal reduced systems in these examples also turn out to be asymptotically stable. The tolerance tol for termination (discussed in

Section 7.2) is set equal to 10^{-8} in Sections 8.1–8.3, and 10^{-6} in Section 8.4. As noted in Section 7.1, a prescribed number of most dominant poles with nonnegative imaginary parts are used to form S_0 , and in Section 8.4 to form S_0^{red} as well. Throughout this section, the shorthand NIPs stands for “nonnegative imaginary parts”.

For comparison or initialization purposes, some of the numerical experiments involve the application of balanced truncation for which we use the MATLAB toolbox MORLAB-5.0 [47], in particular the routine `ml_ct_dss_bt` or `ml_ct_ss_bt` depending on whether the system at hand is a descriptor system or more specifically a linear time-invariant system. Moreover, the Hankel singular values computed for smaller systems for comparison purposes are retrieved by calling the built-in routine `hankelsv` in MATLAB. As the first three subsections concern the model reduction of relatively smaller systems, the built-in routine `norm` is employed in line 7 of Algorithm 1 for $\mathcal{L}_\infty$ -norm computations, while the subspace framework in Ref. [44] is employed for this purpose in Section 8.4 that concerns the model reduction of descriptor systems with large order.

8.1 ISS example

We start with the `iss` example of order $n = 270$ that is also considered when optimizing the objective $\mathcal{F}$ directly in Section 2. As before, we seek the nearest reduced descriptor system of order $r = 12$ with respect to the $\mathcal{L}_\infty$ norm. An application of Algorithm 1 with the initial estimate S_0^{red} produced by balanced truncation terminates when $q = 6$. The error $\mathcal{F}(S_\star^{\text{red}}) = 0.0022516$ of the estimate $S_\star^{\text{red}}$ returned is nearly half of the error $\mathcal{F}(S_0^{\text{red}}) = 0.0044701$ of the initial estimate S_0^{red} . The optimal reduced system $S_\star^{\text{red}}$ is indeed slightly better than the estimate returned by the direct optimization at which the objective $\mathcal{F}$ takes the value 0.0024154. Yet the total runtime is about 66 s, much shorter than 500 s, roughly the time required by the direct optimization. The local optimality of $S_\star^{\text{red}} = (A_\star^{\text{red}}, B_\star^{\text{red}}, C_\star^{\text{red}}, D_\star^{\text{red}}, E_\star^{\text{red}})$ is supported by Figure 2, which indicates an increase in the objective $\mathcal{F}$ if one of the entries of one of $A_\star^{\text{red}}, B_\star^{\text{red}}, C_\star^{\text{red}}, D_\star^{\text{red}}$, and $E_\star^{\text{red}}$ is modified. Moreover, the Hankel singular value σ_{r+1} for this example, a lower bound for the minimal error possible for any system of order r , is 0.0022353 smaller than $\mathcal{F}(S_\star^{\text{red}}) = 0.0022516$ only by a slim margin, so $S_\star^{\text{red}}$ must be nearly optimal globally as well.

The largest singular values of the errors $\sigma_{\max}(H(i\omega) - H(i\omega; S_0^{\text{red}}))$ and $\sigma_{\max}(H(i\omega) - H(i\omega; S_\star^{\text{red}}))$ of the initial estimate S_0^{red} and the optimal estimate $S_\star^{\text{red}}$ are plotted as functions of ω in Figure 3. The singular value error function $\sigma_{\max}(H(i\omega) - H(i\omega; S_\star^{\text{red}}))$ for the optimal $S_\star^{\text{red}}$ is extremely flat, as indeed $\sigma_{\max}(H(i\omega) - H(i\omega; S_\star^{\text{red}})) \in [2.02 \cdot 10^{-3}, 2.26 \cdot 10^{-3}]$ at all ω . Furthermore, the error $\sigma_{\max}(H(i\omega) - H(i\omega; S_\star^{\text{red}}))$ is maximized at four distinct points marked by the circles on the right-hand plot. This indicates that the objective $\mathcal{F}$ is not differentiable at the computed optimizer $S_\star^{\text{red}}$.

Information about the progress of Algorithm 1 is given in Table 2. We start with the reduced system S_0 of order 36 that interpolates the original system S of order 270 at three points on the imaginary axis, namely the imaginary parts of the most dominant three poles of S with NIPs. At every iteration, if no refinement step is performed, the order of the reduced system S_q increases by $4m = 12$. Additionally, each refinement step results in an increase of $4m = 12$ in the order of S_q . We observe in the second column that the error $\mathcal{F}(S_q^{\text{red}})$ at the minimizer S_q^{red} of the reduced objective $\mathcal{F}_q$ decays rapidly with respect to q . The total number of objective function evaluations is 492 (i.e., the sum of the function evaluations in the fifth column), however the $\mathcal{L}_\infty$ objective to be minimized involves the reduced system S_q rather than the full system S . For instance, the number of $\mathcal{L}_\infty$ -norm computations performed are 105, 155, 123 at iterations $q = 1, 2, 3$. Yet, these $\mathcal{L}_\infty$ -norm computations involve the reduced system S_q of order 72, 84, 120 for $q = 1, 2, 3$. Observe that the number of BFGS iterations eventually decrease at the later iterations, as the computed optimal S_q^{red} used as the initial estimate when minimizing $\mathcal{F}_{q+1}$ becomes stationary, i.e., as the computed minimizer S_q^{red} of $\mathcal{F}_q$ is also close to a minimizer of $\mathcal{F}_{q+1}$. Refinement steps are needed only at the initial iteration when $q = 0$ and when $q = 2$. No refinement step turns out to be necessary at the later iterations. This is a generic pattern that we observe in vast majority of examples we have experimented on.




Figure 2: The figure is similar to Figure 1 and concerns the iss example with $r = 12$. Only now the minimization of $\mathcal{F}$ is performed using the subspace framework outlined in Algorithm 1. Specifically, each plot depicts $\mathcal{F}$ as a function of the variation of one of the entries of one of $A^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}}, E^{\text{red}}$. Zero variation corresponds to the optimal reduced system by Algorithm 1.




Figure 3: The plots of $\sigma_{\max}(H(i\omega) - H(i\omega; S_0^{\text{red}}))$ (left) and $\sigma_{\max}(H(i\omega) - H(i\omega; S_*^{\text{red}}))$ (right) as functions of ω for the iss example with $r = 12$, where S_0^{red} is the initial estimate, and S_*^{red} is the optimal estimate computed by Algorithm 1. In each plot, the circles mark the points where the largest singular value function attains the largest value.

Table 2: The iterates and information about the progress of Algorithm 1 on the iss example with $r = 12$. The columns of red order, # BFGS iter, # fun evals, and # refine list the order of the system S_q , number of BFGS iterations, number of objective function evaluations performed by BFGS, and number of refinement steps performed at the q th iteration.

q	$F(S_q^{\text{red}})$	Red order	# BFGS iter	# fun evals	# refine
0	0.004470060020	36	—	—	2
1	0.003517059977	72	38	105	0
2	0.002259657400	84	45	155	2
3	0.002252138011	120	35	123	0
4	0.002251613679	132	11	48	0
5	0.002251609387	144	2	29	0
6	0.002251607779	156	2	32	—

8.2 CD player model

Our next example is the CD player model which is available in the SLICOT library. The model is a linear-time invariant system of order $n = 120$ and with $m = 2$ inputs and $p = 2$ outputs. The details of the model can be found in Ref. [48], and the references therein. Our primary purpose here is to compare on this example Algorithm 1 with the approach in Ref. [16] for $\mathcal{H}_\infty$ model reduction based on rank-one modifications of the system matrices. As the approach in Ref. [16] is for SISO systems, the results are reported over there for this example but with only the second input and the first output. We follow the same practice here when applying our approach. The initial estimate S_0^{red} for a minimizer for Algorithm 1 is constructed using balanced truncation. Moreover, the initial reduced system S_0 is of order 12, and is constructed so that it interpolates the full system S at the imaginary parts of its most dominant three poles with NIPs.

The reduced systems S_*^{red} of order $r = 2, 4, 6, 8, 10$ are computed using Algorithm 1. Table 3 lists the relative errors $\|H - H(\cdot; S_*^{\text{red}})\|_{\mathcal{L}_\infty} / \|H\|_{\mathcal{L}_\infty}$ for the reduced system S_*^{red} computed by various approaches. In particular, the columns of IHA, MBT, HNA are the reported results in Ref. [16, Table 4] by using the approach introduced in Ref. [16] initialized with the model returned by, respectively, IRKA, balanced truncation, the best Hankel norm approximation. Moreover, the columns of BT and Lower Bnd correspond to the relative error of the reduced model by balanced truncation, and the theoretical lower bound $\sigma_{r+1} / \|H\|_{\mathcal{L}_\infty}$ for any reduced system of order r for the relative error, where σ_{r+1} is the $(r + 1)$ st largest Hankel singular value of the system. As can be seen in Table 3, our approach produces reduced systems with smaller errors compared to those produced by

Table 3: This table concerns the “CD player model”. Relative errors $\|H - H(\cdot; S_\star^{\text{red}})\|_{\mathcal{L}_\infty} / \|H\|_{\mathcal{L}_\infty}$ are listed for the optimal estimate $S_\star^{\text{red}}$ computed by various methods for finding reduced systems of order $r = 2, 4, 6, 8, 10$, as well as the lower bound $\sigma_{r+1} / \|H\|_{\mathcal{L}_\infty}$.

r	Alg. 1	IHA	MBT	HNA	BT	Lower Bnd
2	3.12×10^{-1}	3.66×10^{-1}	3.68×10^{-1}	3.35×10^{-1}	3.69×10^{-1}	1.95×10^{-1}
4	1.82×10^{-2}	2.14×10^{-2}	2.25×10^{-2}	2.00×10^{-2}	2.25×10^{-2}	1.13×10^{-2}
6	9.44×10^{-3}	1.04×10^{-2}	1.19×10^{-2}	1.23×10^{-2}	1.23×10^{-2}	6.79×10^{-3}
8	4.18×10^{-3}	4.85×10^{-3}	6.40×10^{-3}	5.99×10^{-3}	6.41×10^{-3}	3.20×10^{-3}
10	7.45×10^{-4}	8.99×10^{-4}	1.24×10^{-3}	1.08×10^{-3}	1.32×10^{-3}	5.86×10^{-4}

other approaches in all cases. The reduced systems produced by Algorithm 1 does not seem far away from global optimality either, as their errors are slightly greater than the theoretical lower bounds in terms of the Hankel singular values in the last column.

We give some details of Algorithm 1 applied to find a reduced system of order $r = 8$ in Figures 4 and 5, as well as in Table 4. In particular, Figure 4 supports that the reduced system $S_\star^{\text{red}}$ by Algorithm 1 is locally optimal, i.e., small variations in the entries of the system matrices cause increase in the $\mathcal{L}_\infty$ error objective. Figure 5 displays the error $\sigma_{\max}(H(i\omega) - H(i\omega; S_0^{\text{red}}))$ of the initial model, and the error $\sigma_{\max}(H(i\omega) - H(i\omega; S_\star^{\text{red}}))$ of the model by Algorithm 1 as functions of ω . Once again the error function for the optimal model $S_\star^{\text{red}}$ is flatter, even if it is not as pronounced as for the iss example, compared to that for the initial model S_0^{red} . The error function $\sigma_{\max}(H(i\omega) - H(i\omega; S_\star^{\text{red}}))$ for the optimal model attains its maximum at five different ω values, which implies that the objective $\mathcal{F}$ is not smooth at $S_\star^{\text{red}}$. As displayed in Table 4, the convergence occurs again quite rapidly; indeed four iterations are sufficient to reach prescribed accuracy and terminate. At each iteration, the order of the reduced system increases by $4m = 4$. Additionally, the refinement step performed in the initial iteration causes also an increase of $4m = 4$ in the order of the reduced system. Larger number of BFGS iterations are needed at iterations with $q = 1, 2$, when the objective involves reduced systems of order 20, 24, respectively. The total runtime is about 15 s, and the relative error at termination is $\mathcal{F}(S_\star^{\text{red}}) / \|H\|_{\mathcal{L}_\infty} = \|H - H(\cdot; S_\star^{\text{red}})\|_{\mathcal{L}_\infty} / \|H\|_{\mathcal{L}_\infty} = (2.87 \times 10^{-1}) / (6.87 \times 10^1) = 4.18 \times 10^{-3}$.

8.3 FOM model

We next report numerical results on the FOM example available in the SLICOT library. The FOM example is a linear time-invariant system of order $n = 1,006$, and with $m = p = 1$. The details are given in Ref. [49, Example 3]. Here, we are mainly interested in investigating the quality of the estimates for optimal reduced systems produced by Algorithm 1. To this end, we compare the errors of the reduced systems by Algorithm 1 with those of balanced truncation, as well as the theoretical lower bounds for the errors in terms of Hankel singular values for varying choices of prescribed order r of the reduced system sought. As in Sections 8.1 and 8.2, we set the initial estimate S_0^{red} for a minimizer as the system produced by balanced truncation, and the initial reduced system S_0 is always of order 12 and interpolates the full system S at the imaginary parts of its most dominant three poles with NIPs.

In Figure 6, the $\mathcal{L}_\infty$ error $\|H - H(\cdot; S_\star^{\text{red}})\|_{\mathcal{L}_\infty}$ of the optimal reduced system $S_\star^{\text{red}}$ by Algorithm 1 and balanced truncation are plotted as functions of the prescribed order r of the reduced system sought. Included in the figure is also the plot of the Hankel singular value σ_{r+1} , a theoretical lower bound for the $\mathcal{L}_\infty$ error $\|H - H(\cdot; S^{\text{red}})\|_{\mathcal{L}_\infty}$ of any system S^{red} of order r . Especially when $r \in [2, 6]$, the errors of the reduced systems by Algorithm 1 are quite close to the theoretical lower bound. Indeed, the errors of the reduced systems by Algorithm 1 usually differ by the theoretical lower bound by a factor of two at most. Moreover, in most of cases the errors of the reduced systems by Algorithm 1 is significantly less than the errors of the reduced systems by balanced truncation.




Figure 4: The figure is analogous to Figure 1, but concerns the “CD player model” with $r = 8$. Each plot depicts F as a function of the variation of one of the entries of one of $A^{\text{red}}, B^{\text{red}}, C^{\text{red}}, D^{\text{red}}, E^{\text{red}}$. Zero variation corresponds to the optimal reduced system by Algorithm 1.




Figure 5: The plots illustrate the errors of the initial, optimal models by Algorithm 1 for the “CD player model” with $r = 8$, and are analogous to those in Figure 3.

Table 4: The iterates and information about the progress of Algorithm 1 on the “CD player model” for finding a reduced system or order $r = 8$. The columns represent quantities as in Table 2.

q	$\mathcal{F}(S_q^{\text{red}})$	Red order	# BFGS iter	# fun evals	# refine
0	0.439972058849	12	—	—	1
1	0.291281639337	20	565	1,479	0
2	0.287107598817	24	134	387	0
3	0.287107598817	28	1	32	—




Figure 6: Errors of the reduced systems of order $r \in [2, 11]$ produced by Algorithm 1 and balanced truncation (BT), as well as the $(r + 1)$ st largest Hankel singular value σ_{r+1} for the FOM example.

8.4 Systems with large order

Finally, we report results on systems with large order arising from modeling of power plants due to Rommes and his colleagues. All of these large-scale examples are available on the website of Rommes.¹

Unlike the previous three subsections, we form the initial estimate for the minimizer S_0^{red} using the dominant poles of the system. For each system, we first compute the ten most dominant poles of the system with NIPs

¹ <http://sites.google.com/site/rommes/software>.

using the approach in Ref. [36], in particular its implementation publicly available at <https://zenodo.org/record/5103430>. Then S_0^{red} of order r is constructed so as to interpolate the full system S at the imaginary parts of its $r/(4m)$ most dominant poles with NIPs. Similarly, the initial reduced system S_0 is constructed such that it interpolates S at the imaginary parts of its ℓ most dominant poles with NIPs, where $\ell = 7$ if the system is single-input-single-output (with $m = 1$), and $\ell = 3$ if the system is multiple-input-multiple-output (with $m > 1$). The order of the resulting reduced system S_0 is $4m\ell$. In all of the examples, the prescribed order r is such that $r < 4m\ell$, that is the order of S_0 is greater than the prescribed order r .

Even Algorithm 1 requires the computation of the $\mathcal{L}_\infty$ norm of systems of order $n + r$ a few times (usually not more than 5–6 times in our experiments) in line 7, where n is the large order of the system. The classical level-set approaches [37], [38] for $\mathcal{L}_\infty$ -norm computation and their implementations in MATLAB are usually no more applicable, or when they are applicable, they take an excessive amount of time. Instead, we employ the interpolatory subspace framework in Ref. [44] for these large-scale $\mathcal{L}_\infty$ -norm computations, that is for maximizing $\sigma_{\max}\left(H(i\omega) - H(i\omega; S_q^{\text{red}})\right)$ over ω at the q th iteration. As the approach in Ref. [44] is locally convergent, whether the initial interpolation points are sufficiently close to global maximizers of $\sigma_{\max}\left(H(i\omega) - H(i\omega; S_q^{\text{red}})\right)$ plays a large role in converging to a global maximizer. We choose the initial interpolation points as the union of the imaginary parts of the ten most dominant poles with NIPs, and 15 equally-spaced points on the interval $[-0.1, 2\mathcal{M}]$ with $\mathcal{M}$ denoting the largest of the imaginary parts of the ten most dominant poles with NIPs.

The absolute and relative errors of the computed reduced systems of order r along with the total runtimes are reported in Table 5. For systems S20PI_n, S40PI_n, M40PI_n of order $n = 1, 182$ or $n = 2, 182$, we have also computed reduced systems of order r by balanced truncation. In these examples, the errors of the reduced systems by Algorithm 1 are smaller than those of the reduced systems by balanced truncation. The implementation of balanced truncation in MORLAB-5.0 that we rely on is based on dense linear algebra routines (unlike our approach which benefits from sparsity), so we do not report the runtimes for balanced truncation. As evident from Table 5, Algorithm 1 is able to deal with systems of order ten thousands in a couple of

Table 5: The absolute errors $\|H - H(\cdot; S_\star^{\text{red}})\|_{\mathcal{L}_\infty}$ (error) and relative errors $\|H - H(\cdot; S_\star^{\text{red}})\|_{\mathcal{L}_\infty} / \|H\|_{\mathcal{L}_\infty}$ (rel error) for systems of large order, where $S_\star^{\text{red}}$ is the optimal reduced system by either Algorithm 1 or balanced truncation (BT). Total runtimes for Algorithm 1 in seconds are also listed in the last column.

Example	$n, m = p$	r	Approach	Error	Rel error	Time
S20PI_n	1,182, 1	12	Alg. 1	7.67×10^{-1}	2.23×10^{-1}	19.8
S20PI_n	1,182, 1	16	Alg. 1	7.66×10^{-1}	2.22×10^{-1}	36.2
S20PI_n	1,182, 1	12	BT	1.76×10^0	5.11×10^{-1}	—
S20PI_n	1,182, 1	16	BT	1.32×10^0	3.84×10^{-1}	—
S40PI_n	2,182, 1	12	Alg. 1	9.30×10^{-1}	2.78×10^{-1}	48.1
S40PI_n	2,182, 1	16	Alg. 1	6.71×10^{-1}	2.00×10^{-1}	38.1
S40PI_n	2,182, 1	32	BT	1.75×10^0	5.23×10^{-1}	—
M40PI_n	2,182, 3	12	Alg. 1	1.99×10^0	5.22×10^{-1}	52.2
M40PI_n	2,182, 3	24	Alg. 1	1.70×10^0	4.45×10^{-1}	117.1
M40PI_n	2,182, 3	36	BT	3.07×10^0	8.03×10^{-1}	—
ww_vref_6405	13,251, 1	12	Alg. 1	5.80×10^{-4}	2.04×10^{-1}	9.2
ww_vref_6405	13,251, 1	16	Alg. 1	4.19×10^{-4}	1.48×10^{-1}	15.1
xingo_afonso	13,250, 1	12	Alg. 1	3.55×10^{-2}	8.74×10^{-3}	14.4
xingo_afonso	13,250, 1	16	Alg. 1	3.56×10^{-2}	8.77×10^{-3}	14.0
xingo_afonso	13,250, 1	20	Alg. 1	1.13×10^{-2}	2.79×10^{-3}	26.2
bips07_1998	15,066, 4	16	Alg. 1	1.24×10^1	6.30×10^{-2}	127.6
bips07_1998	15,066, 4	32	Alg. 1	9.67×10^0	4.91×10^{-2}	219.8
bips07_3078	21,228, 4	16	Alg. 1	1.27×10^1	6.06×10^{-2}	200.8
bips07_3078	21,228, 4	32	Alg. 1	1.00×10^1	4.78×10^{-2}	274.1

minutes in the worst case. Most of the runtime of Algorithm 1 is usually taken by BFGS for solving reduced $\mathcal{L}_\infty$ -norm minimization problems in line 5 involving small systems. In the end, rather than performing quite a few large-scale $\mathcal{L}_\infty$ -norm computations, we end up performing quite a few small-scale $\mathcal{L}_\infty$ -norm computations, and only a few large-scale $\mathcal{L}_\infty$ -norm computations. This results in an approach not only computationally feasible but also more reliable, as small-scale $\mathcal{L}_\infty$ -norm computations can be fulfilled accurately, efficiently and reliably using the level-set methods [37], [38] without worrying about local convergence.

9 Software

A MATLAB implementation of Algorithm 1 is publicly available at <https://zenodo.org/record/8344591>. The numerical results reported in the previous section are obtained with this implementation. Scripts are included to reproduce the results for the CD player model in Section 8.2, and the xingo_afonso, bips07_1998 examples in Section 8.4. The results for other benchmark examples can be obtained similarly.

10 Conclusions

We have proposed an approach to find a locally optimal solution of the $\mathcal{L}_\infty$ -norm model reduction problem. Our approach is based on the usage of smooth optimization techniques such as the gradient descent method and BFGS. A direct application of such smooth optimization techniques for the $\mathcal{L}_\infty$ -norm model reduction problem does not seem suitable even for systems with modest order, as smooth optimization techniques converge very slowly and require the evaluation of the costly $\mathcal{L}_\infty$ -norm objective too many times. Hence, our approach replaces the original system of modest or large order with a system of small order, and solves the resulting reduced $\mathcal{L}_\infty$ -norm minimization problem by means of the smooth optimization techniques. Then it refines and increases slightly the order of the reduced system based on the minimizer of this reduced minimization problem. This refinement is performed with an eye to interpolation between the full and reduced $\mathcal{L}_\infty$ objectives. Under smoothness assumptions, admittedly strong in this context, we have given formal arguments for the quick convergence of the approach. We have also described how asymptotic stability constraints on the small system of prescribed order sought can be incorporated into the approach. The numerical experiments on a variety of real benchmark examples indicate that our approach retrieves indeed a locally optimal solution of the $\mathcal{L}_\infty$ -norm model reduction problem in practice. Moreover, on some small benchmark examples, we have obtained reduced systems not far away from being optimal globally according to the theoretical lower bounds in terms of Hankel singular values. Experiments on large benchmark examples illustrate that the approach is usually suitable for systems of order a few ten thousands.

The quality of the converged locally optimal solution depends on the initial guess for the optimal reduced system. To generate the initial guess, we have employed two different strategies based on balanced truncation and dominant poles. The first of these strategies may not be applicable unless the original system is asymptotically stable. On the other hand, there is no asymptotic stability requirement for the second strategy. However, a strategy generating a good initial guess is certainly worth further research. The proposed approach typically requires a few large-scale $\mathcal{L}_\infty$ -norm computations. Performing these $\mathcal{L}_\infty$ -norm computations accurately, especially without getting stagnated at a local maximizer that is not optimal globally, is crucial for the reliability of the proposed approach. We have employed the interpolatory subspace framework in Ref. [44] with the initial interpolation points chosen based on the dominant poles for these large-scale $\mathcal{L}_\infty$ -norm computations. This approach usually seems to work well in practice for large-scale $\mathcal{L}_\infty$ -norm computations. Still, we hope to explore further a good initial interpolation selection strategy for [44] so that it converges globally, leading to the correct $\mathcal{L}_\infty$ norm with very high probability. Other efficient and accurate candidates for large-scale $\mathcal{L}_\infty$ -norm computation are worth studying. In Ref. [16], the original system is replaced by a smaller order system obtained from the Loewner framework [50] to reduce the burden of large-scale $\mathcal{L}_\infty$ -norm computations. We have not attempted here to incorporate the Loewner framework into our approach. As a future work, our approach can possibly

benefit from the Loewner framework; for instance, the initial reduced system replacing the full system can perhaps be obtained using the Loewner framework. Our quick convergence result for the proposed approach is under strong smoothness assumptions. Investigating the order of convergence of the approach in the likely nonsmooth setting (i.e., when the $\mathcal{L}_\infty$ objective at the converged minimizer is nonsmooth) is a possible direction for future research. Last but not the least, the convergence of smooth optimization techniques such as BFGS on a variety of nonsmooth optimization problems is observed empirically, yet these empirical observations could not be supported by a general convergence theory so far. Analyzing and understanding rigorously the convergence of BFGS and other smooth optimization techniques when the objective is nonsmooth at the optimizers sought are important open problems.

Acknowledgment: The author is grateful to two anonymous referees for their invaluable feedback on the initial version of this manuscript.

Research ethics: Not applicable.

Informed consent: Not applicable.

Author contributions: The author has accepted responsibility for the entire content of this manuscript and approved its submission.

Use of Large Language Models, AI and Machine Learning Tools: None declared.

Conflict of interest: The authors state no conflict of interest.

Research funding: None declared.

Data availability: The raw data and codes associated with the current study are publicly available at <https://zenodo.org/record/8344591> and <https://sites.google.com/site/rommes/software>.

References

- [1] K. E. Brenan, S. L. Campbell, and L. R. Petzold, *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*, 2nd ed. Philadelphia, PA, SIAM, 1996.
- [2] Y. Feng and M. Yagoubi, *Robust Control of Linear Descriptor Systems*, Singapore, Springer, 2017.
- [3] P. Kunkel and V. Mehrmann, *Differential-Algebraic Equations: Analysis and Numerical Solution*, Zurich, Switzerland, EMS Press, 2006.
- [4] R. Riaza, *Differential Algebraic Systems. Analytical Aspects and Circuit Applications*. Notes on Mathematics and its Applications, Singapore, World Scientific, 2008.
- [5] A. C. Antoulas, *Approximation of Large-Scale Dynamical Systems, Volume 6 of Adv. Des. Control*, Philadelphia, PA, SIAM Publications, 2005.
- [6] V. Mehrmann and T. Stykel, “Balanced truncation model reduction for large-scale systems in descriptor form,” in *Dimension Reduction of Large-Scale Systems*, P. Benner, D. C. Sorensen, and V. Mehrmann, Eds., Berlin, Springer, vol. 45, 2005, pp. 83–115.
- [7] B. Moore, “Principal component analysis in linear systems: controllability, observability, and model reduction,” *IEEE Trans. Autom. Control*, vol. 26, no. 1, pp. 17–32, 1981.
- [8] T. Stykel, “Gramian-based model reduction for descriptor systems,” *Math. Control Signals Syst.*, vol. 16, no. 4, pp. 297–319, 2004.
- [9] S. Gugercin and A. C. Antoulas, “A survey of model reduction by balanced truncation and some new results,” *Int. J. Control*, vol. 77, no. 8, pp. 748–766, 2004.
- [10] S. Gugercin, D. C. Sorensen, and A. C. Antoulas, “A modified low-rank Smith method for large-scale Lyapunov equations,” *Numer. Algorithms*, vol. 32, no. 1, pp. 27–55, 2003.
- [11] T. Penzl, “A cyclic low-rank Smith method for large sparse Lyapunov equations,” *SIAM J. Sci. Comput.*, vol. 21, no. 4, pp. 1401–1418, 1999.
- [12] K. Glover, “All optimal hankel-norm approximations of linear multivariable systems and their $\mathcal{L}_\infty$ -error bounds,” *Int. J. Control*, vol. 39, no. 6, pp. 1115–1193, 1984.
- [13] P. Benner and S. W. R. Werner, “Hankel-norm approximation of large-scale descriptor systems,” *Adv. Comput. Math.*, vol. 46, Art.no. 40, 2020.
- [14] P. Benner, E. S. Quintana-Orti, and G. Quintana-Orti, “Computing optimal hankel norm approximations of large-scale systems,” in *Proceedings of 43rd IEEE Conference on Decision and Control*, vol. 3, 2004, pp. 3078–3083.
- [15] S. Gugercin, A. C. Antoulas, and C. Beattie, “ $\mathcal{H}_2$ model reduction for large-scale linear dynamical systems,” *SIAM J. Matrix Anal. Appl.*, vol. 30, no. 2, pp. 609–638, 2008.
- [16] G. Flagg, C. A. Beattie, and S. Gugercin, “Interpolatory H_∞ model reduction,” *Syst. Control Lett.*, vol. 62, no. 7, pp. 567–574, 2013.
- [17] C. K. Krzysztof, *Methods of Descent for Nondifferentiable Optimization*, Berlin, New York, Springer-Verlag, 1985.

- [18] M. M. Mäkelä and P. Neittaanmäki, *Nonsmooth Optimization: Analysis and Algorithms with Applications to Optimal Control*, Singapore, World Scientific, 1992.
- [19] J. V. Burke, A. S. Lewis, and M. L. Overton, “Two numerical methods for optimizing matrix stability,” *Linear Algebra Appl.*, vols. 351–352, pp. 117–145, 2002.
- [20] J. V. Burke, A. S. Lewis, and M. L. Overton, “A robust gradient sampling algorithm for nonsmooth, nonconvex optimization,” *SIAM J. Opt.*, vol. 15, no. 3, pp. 751–779, 2005.
- [21] S. W. R. Werner, M. L. Overton, and B. Peherstorfer, “Multifidelity robust controller design with gradient sampling,” *SIAM J. Sci. Comput.*, vol. 45, no. 2, pp. A933–A957, 2023.
- [22] A. Asl and M. L. Overton, “Analysis of the gradient method with an Armijo–Wolfe line search on a class of non-smooth convex functions,” *Optim. Method Softw.*, vol. 35, no. 2, pp. 223–242, 2020.
- [23] A. S. Lewis and M. L. Overton, “Nonsmooth optimization via quasi-Newton methods,” *Math. Program.*, vol. 141, nos. 1–2, pp. 135–163, 2013.
- [24] C. Lemaréchal, *Numerical Experiments in Nonsmooth Optimization*, Laxenburg, Austria, International Institute for Applied System Analysis (IIASA), 1982, pp. 61–84.
- [25] J. V. Burke, F. E. Curtis, A. S. Lewis, M. L. Overton, and L. E. A. Simões, “Gradient sampling methods for nonsmooth optimization,” in *Numerical Nonsmooth Optimization: State of the Art Algorithms*, A. M. Bagirov, M. Gaudioso, N. Karimtsa, M. M. Mäkelä, and S. Taheri, Eds., Cham, Switzerland, Springer International Publishing, 2020, pp. 201–225.
- [26] P. Schwerdtner and M. Voigt, “SOBMOR: structured optimization-based model order reduction,” *SIAM J. Sci. Comput.*, vol. 45, no. 2, pp. A502–A529, 2023.
- [27] A. Megretski, “H-infinity model reduction with guaranteed suboptimality bound,” in *2006 American Control Conference*, 2006, p. 6.
- [28] A. Aliyev, P. Benner, E. Mengi, and M. Voigt, “A subspace framework for H-infinity norm minimization,” *SIAM J. Matrix Anal. Appl.*, vol. 41, no. 2, pp. 928–956, 2020.
- [29] S. Chellappa, L. Feng, V. de la Rubia, and P. Benner, “Adaptive interpolatory mor by learning the error estimator in the parameter domain,” in *Model Reduction of Complex Dynamical Systems*, P. Benner, T. Breiten, H. Faßbender, M. Hinze, T. Stykel, and R. Zimmermann, Eds., Cham, Switzerland, Birkhäuser, 2021, pp. 97–117.
- [30] L. Feng and P. Benner, “A new error estimator for reduced-order modeling of linear parametric systems,” *IEEE Trans. Microw. Theory*, vol. 67, no. 12, pp. 4848–4859, 2019.
- [31] F. E. Curtis, T. Mitchell, and M. L. Overton, “A BFGS-SQP method for nonsmooth, nonconvex, constrained optimization and its evaluation using relative minimization profiles,” *Optim. Method Softw.*, vol. 32, no. 1, pp. 148–181, 2017.
- [32] F. E. Curtis and M. L. Overton, “A sequential quadratic programming algorithm for nonconvex, nonsmooth constrained optimization,” *SIAM J. Opt.*, vol. 22, no. 2, pp. 474–500, 2012.
- [33] P. Lancaster, “On eigenvalues of matrices dependent on a parameter,” *Numer. Math.*, vol. 6, no. 1, pp. 377–387, 1964.
- [34] E. Mengi, E. A. Yildirim, and M. Kiliç, “Numerical optimization of eigenvalues of Hermitian matrix functions,” *SIAM J. Matrix Anal. Appl.*, vol. 35, no. 2, pp. 699–724, 2014.
- [35] C. Beattie and S. Gugercin, “Interpolatory projection methods for structure-preserving model reduction,” *Syst. Control Lett.*, vol. 58, no. 3, pp. 225–232, 2009.
- [36] E. Mengi, “Large-scale estimation of dominant poles of a transfer function by an interpolatory framework,” *SIAM J. Sci. Comput.*, vol. 44, no. 4, pp. A2412–A2438, 2022.
- [37] S. Boyd and V. Balakrishnan, “A regularity result for the singular values of a transfer matrix and a quadratically convergent algorithm for computing its $\mathcal{L}_\infty$ -norm,” *Syst. Control Lett.*, vol. 15, no. 1, pp. 1–7, 1990.
- [38] N. A. Bruinsma and M. Steinbuch, “A fast algorithm to compute the H_∞ -norm of a transfer function matrix,” *Syst. Control Lett.*, vol. 14, no. 4, pp. 287–293, 1990.
- [39] E. Mengi, “Large-scale and global maximization of the distance to instability,” *SIAM J. Matrix Anal. Appl.*, vol. 39, no. 4, pp. 1776–1809, 2018.
- [40] K. Carlberg, R. Tuminaro, and P. Boggs, “Preserving Lagrangian structure in nonlinear model reduction with application to structural dynamics,” *SIAM J. Sci. Comput.*, vol. 37, no. 2, pp. B153–B184, 2015.
- [41] I. Kalashnikova, B. van Bloemen Waanders, S. Arunajatesan, and M. Barone, “Stabilization of projection-based reduced order models for linear time-invariant systems via optimization-based eigenvalue reassignment,” *Comput. Methods Appl. Mech. Eng.*, vol. 272, pp. 251–270, 2014.
- [42] G. Serre, P. Lafon, X. Gloerfelt, and C. Bailly, “Reliable reduced-order models for time-dependent linearized Euler equations,” *J. Comput. Phys.*, vol. 231, no. 15, pp. 5176–5194, 2012.
- [43] F. R. Gantmacher, *The Theory of Matrices*, vol. II, New York, NY, USA, Chelsea Publishing Company, 1959.
- [44] A. Aliyev, P. Benner, E. Mengi, P. Schwerdtner, and M. Voigt, “Large-scale computation of $\mathcal{L}_\infty$ -norms by a greedy subspace method,” *SIAM J. Matrix Anal. Appl.*, vol. 38, no. 4, pp. 1496–1516, 2017.
- [45] P. Benner and T. Mitchell, “Faster and more accurate computation of the $\mathcal{H}_\infty$ norm via optimization,” *SIAM J. Sci. Comput.*, vol. 40, no. 5, pp. A3609–A3635, 2018.
- [46] T. Mitchell, “ROSTAPACK: ROBust STAbility PACKage v3.0,” 2024. <http://www.timtmitchell.com/software/ROSTAPACK/>.

- [47] P. Benner and S. W. R. Werner, “MORLAB — model order reduction LABoratory (version 5.0),” 2019. Available at: <http://www.mpi-magdeburg.mpg.de/projects/morlab>.
- [48] Y. Chahlaoui and P. Van Dooren, “Benchmark examples for model reduction of linear time-invariant dynamical systems,” in *Dimension Reduction of Large-Scale Systems*, P. Benner, D. C. Sorrenson and V. Mehrmann, Eds., Berlin, Springer, vol. 45, 2005, pp. 379–392.
- [49] T. Penzl, “Algorithms for model reduction of large dynamical systems,” *Linear Algebra Appl.*, vol. 415, nos. 2–3, pp. 322–343, 2006.
- [50] A. J. Mayo and A. C. Antoulas, “A framework for the solution of the generalized realization problem,” *Linear Algebra Appl.*, vol. 425, nos. 2–3, pp. 634–662, 2007.