

Parametric model order reduction using pyMOR

Petar Mlinarić*

Stephan Rave†

Jens Saak*

March 13, 2020

Abstract

pyMOR is a free software library for model order reduction that includes both reduced basis and system-theoretic methods. All methods are implemented in terms of abstract vector and operator interfaces, which allows direct integration of pyMOR’s algorithms with a wide array of external PDE solvers. In this contribution, we give a brief overview of the available methods and experimentally compare them for the parametric instationary thermal-block benchmark defined in [13].

1 Introduction

pyMOR is a free software library for building model order reduction applications with the Python programming language [7, 9]. Originally only implementing reduced basis methods, since version 0.5, released in January 2019, it additionally implements system-theoretic methods such as balanced truncation [8] and IRKA [1]. Here, we focus on version 2019.2, released in December 2019, which added support for parametric system-theoretic methods.

We consider model reduction of the thermal-block model defined in [13], which takes the form

$$\begin{aligned} E\dot{x}(t; \mu) &= A(\mu)x(t; \mu) + Bu(t), & x(0; \mu) &= 0, \\ y(t; \mu) &= Cx(t; \mu), \end{aligned}$$

with system matrices $E, A(\mu) \in \mathbb{R}^{n \times n}$, input matrix $B \in \mathbb{R}^{n \times m}$, output matrix $C \in \mathbb{R}^{p \times n}$, state $x(t) \in \mathbb{R}^n$, input $u(t) \in \mathbb{R}^m$, and output $y(t) \in \mathbb{R}^p$, where $\mu \in \mathcal{P} \subset \mathbb{R}^d$ is the parameter. The matrix-valued function A additionally has parameter-affine form $A(\mu) = A_0 + \sum_{i=1}^d \mu_i A_i$, where $\mu = (\mu_1, \mu_2, \dots, \mu_d)$. We also consider a non-parametric version, for which we write A instead of $A(\mu)$.

We begin, in Section 2, with a brief discussion of pyMOR’s software design. In Section 3, we give a brief overview of the methods implemented in pyMOR 2019.2. Next, we give numerical results in Section 4. A conclusion follows in Section 5.

2 Software design

The central goal of pyMOR’s design is to allow an easy integration with external PDE solver libraries. To this end, generic interfaces for vectors and operators have been defined that give pyMOR access to the solver’s internal data structures representing vectors, matrices or nonlinear operators, as well as operations on them, e.g., the computation of inner products or the solution of linear equation system.

*Max Planck Institute for Dynamics of Complex Technical Systems, Sandtorstr. 1, 39106 Magdeburg, Germany, {mlinaric, saak}@mpi-magdeburg.mpg.de

†University of Münster, Orleans-Ring 10, 48149 Münster, Germany, stephan.rave@uni-muenster.de

All high-dimensional model reduction operations in `pyMOR`, for instance POD computation or Petrov-Galerkin projection, are expressed in terms of these interfaces. Compared to a file-based exchange of matrices or solution snapshots, this approach enables the usage of problem adapted solvers implemented in the PDE library or the reduction of very large MPI-distributed problems [7].

3 Overview of model order reduction methods

The majority of MOR methods implemented in `pyMOR` are projection-based methods, i.e., they consist of finding basis matrices V and W and defining the reduced-order model as

$$\begin{aligned}\widehat{E}\widehat{\dot{x}}(t; \mu) &= \widehat{A}(\mu)\widehat{x}(t; \mu) + \widehat{B}u(t), & \widehat{x}(0; \mu) &= 0, \\ \widehat{y}(t; \mu) &= \widehat{C}\widehat{x}(t; \mu),\end{aligned}$$

where $\widehat{E} = W^T E V$, $\widehat{A}(\mu) = W^T A(\mu) V = \widehat{A}_0 + \sum_{i=1}^d \mu_i \widehat{A}_i$, $\widehat{A}_i = W^T A_i V$, $\widehat{B} = W^T B$, and $\widehat{C} = C V$. If $\text{im}(V) = \text{im}(W)$, we call it a Galerkin-projection and otherwise a Petrov-Galerkin projection.

In the following, we give short descriptions of some projection-based methods with remarks on their implementation in `pyMOR`.

3.1 Reduced basis method

We consider a weak POD-Greedy algorithm [6] to build a basis matrix V for which the maximum state-space approximation error

$$\max_{\mu \in \mathcal{S}_{\text{train}}} \sum_{i=1}^N \|x(t_i; \mu) - V\widehat{x}(t_i; \mu)\|^2$$

for constant input $u \equiv 1$ over some training set $\mathcal{S}_{\text{train}}$ of parameters is minimized in the H^1 -norm. To this end, in each iteration of the greedy algorithm the current reduced-order model is solved for all $\mu \in \mathcal{S}_{\text{train}}$ and the parameter μ_{max} is selected for which an (online-efficient) estimate of the MOR error is maximized. For this parameter, the matrix of FOM solution snapshots

$$X = [x(t_1; \mu_{\text{max}}) \quad x(t_2; \mu_{\text{max}}) \quad \cdots \quad x(t_N; \mu_{\text{max}})],$$

is computed, and the first left-singular vectors of its H^1 -orthonormal projection onto the H^1 -orthogonal complement of $\text{im}(V)$ are added to V .

3.2 System-theoretic methods

3.2.1 Balanced truncation

For non-parametric models, balanced truncation (BT) consists of solving two Lyapunov equations

$$\begin{aligned}A P E^T + E P A^T + B B^T &= 0, \\ A^T Q E + E^T Q A + C^T C &= 0.\end{aligned}\tag{1}$$

Based on the solutions P and Q , it computes V and W of the Petrov-Galerkin projection. `pyMOR` provides bindings to dense Lyapunov equation solvers in SciPy [14], Slycot [11] (Python wrappers for SLICOT [10]), and Py-M.E.S.S. [5]. For reduction of large-scale models, there are bindings for low-rank solvers in Py-M.E.S.S.. Since Py-M.E.S.S. does not allow generic vectors, there is also an implementation of the alternating direction implicit iteration in `pyMOR` [2].

It is known that BT preserves asymptotic stability and has a priori \mathcal{H}_∞ and \mathcal{H}_2 error bounds depending on the truncated Hankel singular values (the square roots of the eigenvalues of $E^T Q E P$).

For parametric models, there are several possible extensions of BT [3, 15, 12]. We focus on the simplest global basis approach by concatenating several local basis matrices. Let $\mu^{(1)}, \mu^{(2)}, \dots, \mu^{(\ell)} \in \mathcal{P}$ be parameter samples and $V^{(1)}, V^{(2)}, \dots, V^{(\ell)}$ and $W^{(1)}, W^{(2)}, \dots, W^{(\ell)}$ corresponding local basis matrices. To guarantee asymptotic stability, we use Galerkin projection with

$$[V^{(1)} \ V^{(2)} \ \dots \ V^{(\ell)} \ W^{(1)} \ W^{(2)} \ \dots \ W^{(\ell)}]$$

after orthogonalization and rank truncation.

3.2.2 LQG balanced truncation

LQG balanced truncation (LQGBT) is a variant of BT related to the linear quadratic Gaussian (LQG) optimal control problem. Unlike BT, LQGBT consists of solving Riccati equations

$$\begin{aligned} APE^T + EPA^T - EPC^T CPE^T + BB^T &= 0, \\ A^T QE + E^T QA - E^T QBB^T QE + C^T C &= 0. \end{aligned}$$

Similar to BT, it guarantees preservation of asymptotic stability and has an a priori error bound. As for Lyapunov equations, pyMOR provides bindings for external Riccati equation solvers and an implementation of the low-rank RADI method [4].

Additionally, there is bounded-real BT in pyMOR, but it currently relies on a dense solver which does not respect the vector and operator interfaces, so it is not possible to use it with a PDE solver.

3.2.3 Iterative rational Krylov algorithm

Iterative rational Krylov algorithm (IRKA) is a locally \mathcal{H}_2 optimal MOR method. In each step, it computes (tangential) rational Krylov subspaces

$$\begin{aligned} \text{im}(V) &= \text{span} \left\{ (\sigma_1 E - A)^{-1} B b_1, (\sigma_2 E - A)^{-1} B b_2, \dots, (\sigma_r E - A)^{-1} B b_r \right\}, \\ \text{im}(W) &= \text{span} \left\{ (\sigma_1 E - A)^{-T} C^T c_1, (\sigma_2 E - A)^{-T} C^T c_2, \dots, (\sigma_r E - A)^{-T} C^T c_r \right\}. \end{aligned} \quad (2)$$

The interpolation points $\sigma_1, \sigma_2, \dots, \sigma_r$ for the next step are chosen as reflected poles $-\lambda_1, -\lambda_2, \dots, -\lambda_r$ of the projected matrix pair $\lambda W^T E V - W^T A V$ (vectors b_1, b_2, \dots, b_r and c_1, c_2, \dots, c_r are computed based on the eigenvectors). Even if the original model has real poles, the projected poles can be complex. Since the complex number support is limited in PDE solvers, solving complex shifted linear systems $(\sigma E - A)x = b$ needs to be done using an iterative method. Implementing efficient preconditions for such systems is a future research topic for pyMOR. For this reason, we demonstrate IRKA only on the non-parametric example in Section 4.1. In the parametric case, we only use one-sided IRKA (OS-IRKA), where W in (2) is replaced by V , which guarantees real interpolation points for the heat equation example we consider. To generate the global basis matrix, we concatenate the local basis matrices $V^{(i)}$ and do a rank truncation.

3.2.4 Generating reduced models

All system-theoretic methods in pyMOR can be called similarly. For instance, BT can be run with

```
bt = BTReductor(fom, mu=mu)
rom = bt.reduce(10)
```

where `fom` is the (parametric) full-order model (an instance of `LTIModel`), `mu` is the parameter sample. The `reduce` method of `bt` accepts the reduced order as a parameter (among others) and returns the non-parametric reduced-order model `rom` (again an instance of `LTIModel`). The basis matrices are then available as `VectorArrays` in `bt.V` and `bt.W`.

4 Numerical results

Here, we present results of applying system-theoretic methods to parametric models, in particular the thermal block example. We use the Hardy \mathcal{H}_2 norm to quantify the results, which is defined for non-parametric, asymptotically stable systems

$$\begin{aligned} E\dot{x}(t) &= Ax(t) + Bu(t), & x(0) &= 0, \\ y(t) &= Cx(t), \end{aligned}$$

as the \mathcal{L}_2 norm of the impulse response $\|Ce^{tE^{-1}A}E^{-1}B\|_{\mathcal{L}_2}$. This can be computed using

$$\left\|Ce^{tE^{-1}A}E^{-1}B\right\|_{\mathcal{L}_2}^2 = \text{tr}(CPC^T) = \text{tr}(B^TQB),$$

where P and Q are as in (1). Note that for a reduced-order model

$$\begin{aligned} \widehat{E}\widehat{\dot{x}}(t) &= \widehat{A}\widehat{x}(t) + \widehat{B}u(t), & \widehat{x}(0) &= 0, \\ \widehat{y}(t) &= \widehat{C}\widehat{x}(t), \end{aligned}$$

the error system is of the same form

$$\begin{aligned} \begin{bmatrix} E & 0 \\ 0 & \widehat{E} \end{bmatrix} \begin{bmatrix} \dot{x}(t) \\ \widehat{\dot{x}}(t) \end{bmatrix} &= \begin{bmatrix} A & 0 \\ 0 & \widehat{A} \end{bmatrix} \begin{bmatrix} x(t) \\ \widehat{x}(t) \end{bmatrix} x(t) + \begin{bmatrix} B \\ \widehat{B} \end{bmatrix} u(t), \\ y(t) - \widehat{y}(t) &= \begin{bmatrix} C & -\widehat{C} \end{bmatrix} x(t), \end{aligned}$$

which allows us to compute \mathcal{H}_2 errors.

We begin with the non-parametric version in Section 4.1, comparing system-theoretic methods with POD. Then, in Sections 4.2 and 4.3 we compare methods for parametric versions.

The source code of the implementations used to compute the presented results can be obtained from

<https://doi.org/10.5281/zenodo.3700178>

and is authored by Petar Mlinarić and Stephan Rave.

4.1 Non-parametric version

Figure 1 compares BT, LQGBT, IRKA, OS-IRKA, and POD in terms of relative \mathcal{H}_2 error. The POD model was trained using the step response ($u(t) = 1$ for $t \geq 0$). We see that BT, LQGBT, and IRKA give similar results, while OS-IRKA and POD give worse errors. Interestingly, POD is mostly better than OS-IRKA in this example.

4.2 Single parameter version

In this setting, as the training set we chose 10 logarithmically equi-spaced parameter values from 10^{-6} to 10^2 . For testing, we added additional 9 in-between points. We used BT and OS-IRKA to get reduced models of order 10 for each parameter value and combined BT's local bases as explained in Section 3.2.1. After truncation, BT's global basis was of order 175 and one-sided IRKA's 67. To have a fairer comparison, we further truncated BT's global basis to the same order as one-sided IRKA. Figure 2 shows the \mathcal{H}_2 norm of the full-order model for different parameters and Figure 3 the absolute and relative \mathcal{H}_2 error for BT and one-sided IRKA. Possibly related to BT being a Petrov-Galerkin projection method, its global basis produces worse results than the local bases. On the other hand, OS-IRKA improves with using the global basis.

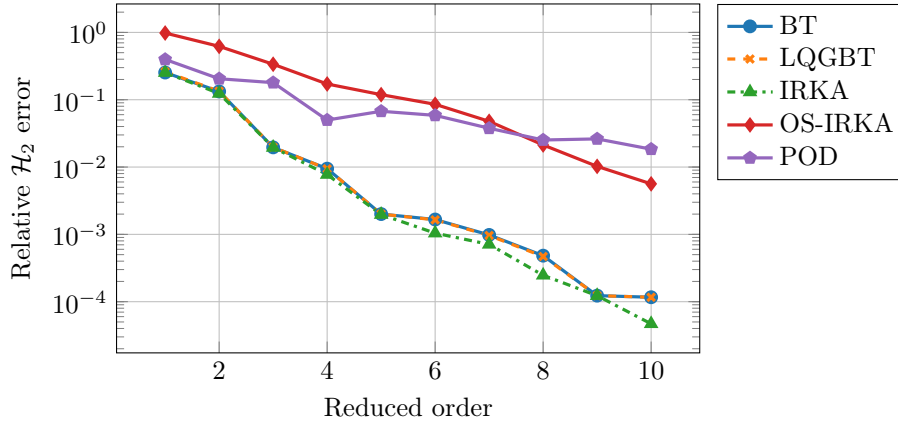


Figure 1: Comparison of methods for the non-parametric model

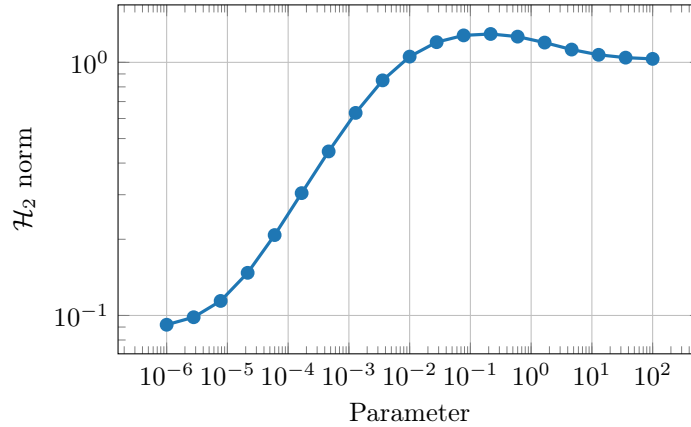


Figure 2: The \mathcal{H}_2 norms of the one-parameter model for different parameter values

Finally, Figure 4 compares BT and one-sided IRKA with RB. For RB, we used the same training set to generate a model of order 67. In this example, one-sided IRKA performed best near the boundaries of the parameter set and comparable to other methods in the middle. On the other hand, BT gave worst results near the boundaries. RB produced an almost flat absolute \mathcal{H}_2 error curve, which is not surprising since it tries to minimize the worst error.

4.3 Four parameter version

Here, we chose 20 uniformly random points in $[-6, 2]^4$ as exponents (with base 10) to generate the training set and 20 additional such points for testing. As before, we used BT and one-sided IRKA to find reduced models of order 10 at each training parameter point. Here, BT's global basis was of order 347 and one-sided IRKA's was 128. Figure 5 compares them, where the first 20 parameter values are from the training set and the other for testing. As we had in the previous example, OS-IRKA gives better results with the global basis.

Figure 6 compares the two methods with RB. We see that they give comparable results, with BT giving better errors the most.

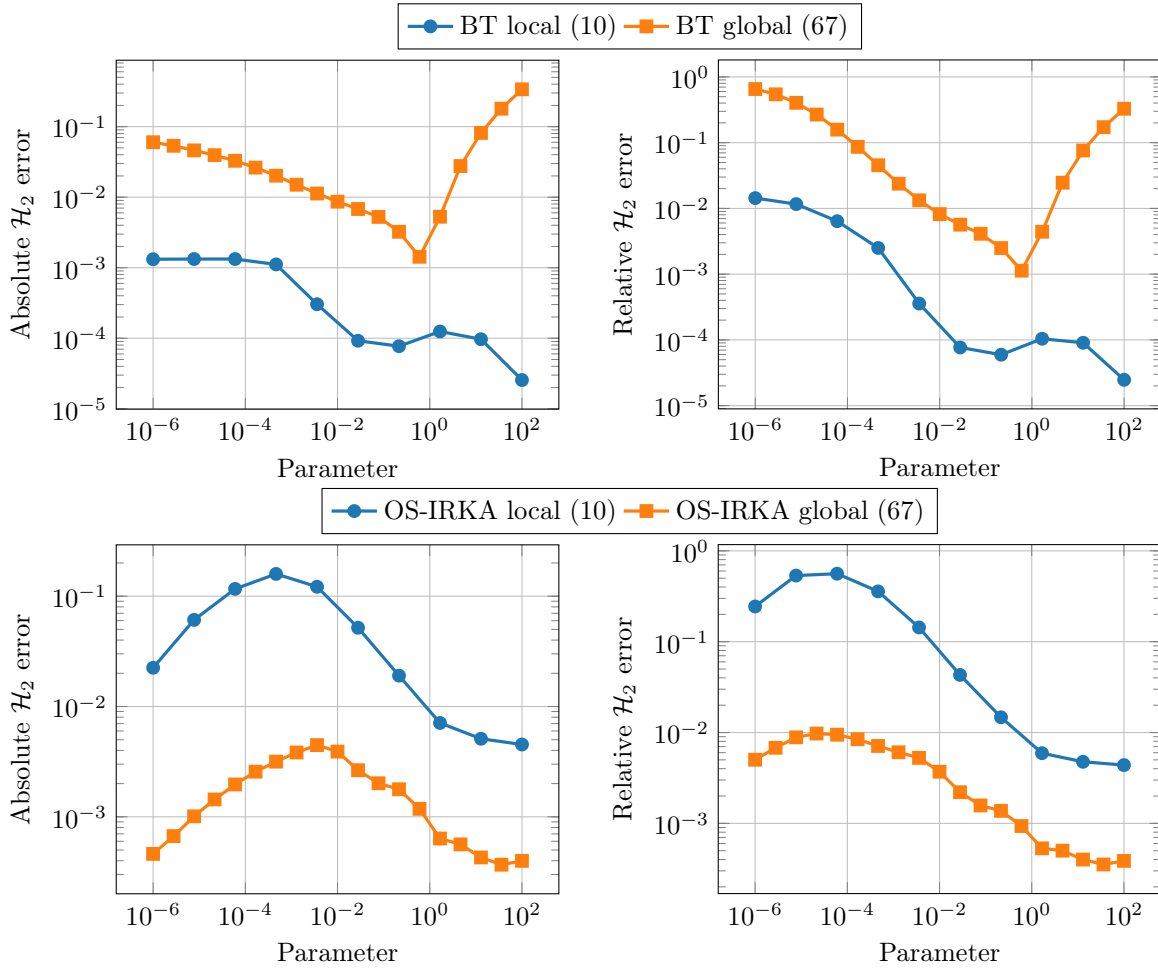


Figure 3: Balanced truncation (BT) and one-sided iterative rational Krylov algorithm (OS-IRKA) for one-parameter model

5 Conclusions

We briefly presented pyMOR, a freely available Python package for MOR, built on generic interfaces for easy integration with external PDE solvers. We then described some of the MOR methods implemented in pyMOR, which includes both system-theoretic and reduced basis methods. Lastly, we compared methods on a thermal block benchmark discretized with FEniCS.

References

- [1] A. C. Antoulas, C. A. Beattie, and S. Gugercin. Interpolatory model reduction of large-scale dynamical systems. In Javad Mohammadpour and Karolos M. Grigoriadis, editors, *Efficient Modeling and Control of Large-Scale Systems*, pages 3–58. Springer US, 2010.
- [2] L. Balicki. Low-rank alternating direction implicit iteration in pyMOR. *GAMM Archive for Students*, 2(1):1–13, February 2020.

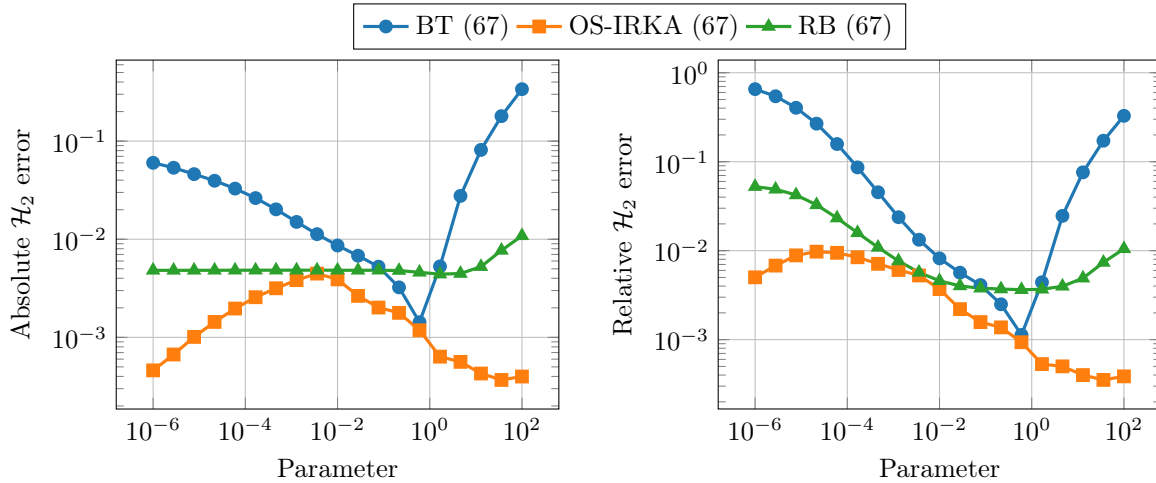


Figure 4: Comparison of methods for the one-parameter model

- [3] U. Baur and P. Benner. Modellreduktion für parametrisierte Systeme durch balanciertes Abschneiden und Interpolation (Model Reduction for Parametric Systems Using Balanced Truncation and Interpolation). *at-Automatisierungstechnik*, 57(8):411–420, 2009.
- [4] P. Benner, Z. Bujanović, P. Kürschner, and J. Saak. RADI: A low-rank ADI-type algorithm for large scale algebraic Riccati equations. *Numer. Math.*, 138(2):301–330, February 2018.
- [5] Peter Benner, Martin Köhler, and Jens Saak. M.E.S.S. – the matrix equations sparse solvers library. <https://www.mpi-magdeburg.mpg.de/projects/mess>.
- [6] B. Haasdonk. Convergence rates of the POD-Greedy method. *ESAIM: Math. Model. Numer. Anal.*, 47(3):859–873, 2013.
- [7] R. Milk, S. Rave, and F. Schindler. pyMOR – generic algorithms and interfaces for model order reduction. *SIAM J. Sci. Comput.*, 38(5):S194–S216, 2016.
- [8] B. C. Moore. Principal component analysis in linear systems: controllability, observability, and model reduction. *IEEE Trans. Autom. Control*, AC–26(1):17–32, 1981.
- [9] pyMOR developers and contributors. pyMOR - model order reduction with Python. <https://pymor.org>
- [10] SLICOT. <http://www.slicot.org>.
- [11] Slycot developers and contributors. Slycot. <https://github.com/python-control/Slycot>.
- [12] N. T. Son and T. Stykel. Solving parameter-dependent Lyapunov equations using the reduced basis method with application to parametric model order reduction. *SIAM J. Matrix Anal. Appl.*, 38(2):478–504, 2017.
- [13] The MORwiki Community. Thermal block. hosted at MORwiki – Model Order Reduction Wiki, 2020.
- [14] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero,

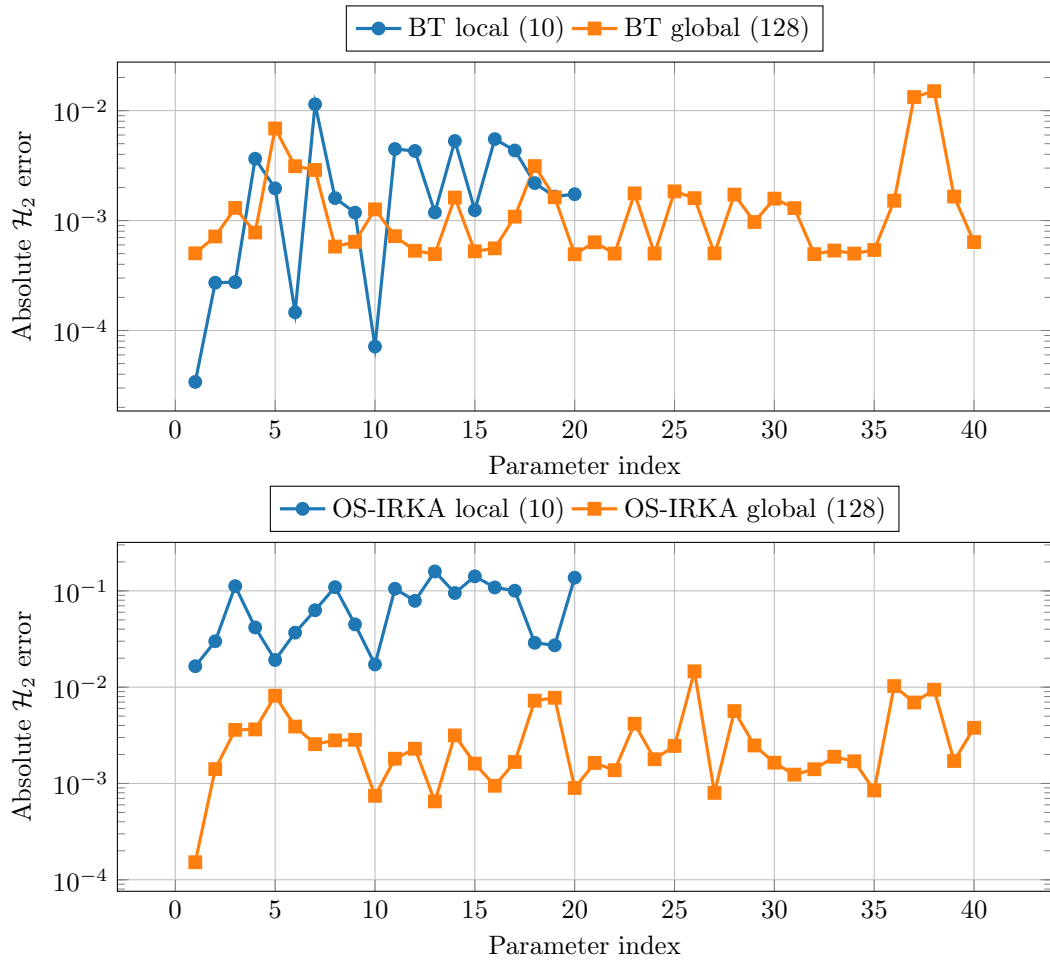


Figure 5: Balanced truncation (BT) and one-sided iterative rational Krylov algorithm (OS-IRKA) for four-parameter model

C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, and P. van Mulbregt. SciPy 1.0: fundamental algorithms for scientific computing in python. *Nat. Methods*, February 2020.

- [15] P. Wittmuess, C. Tarin, A. Keck, E. Arnold, and O. Sawodny. Parametric model order reduction via balanced truncation with Taylor series representation. *IEEE Trans. Autom. Control*, 61(11):3438–3451, November 2016.

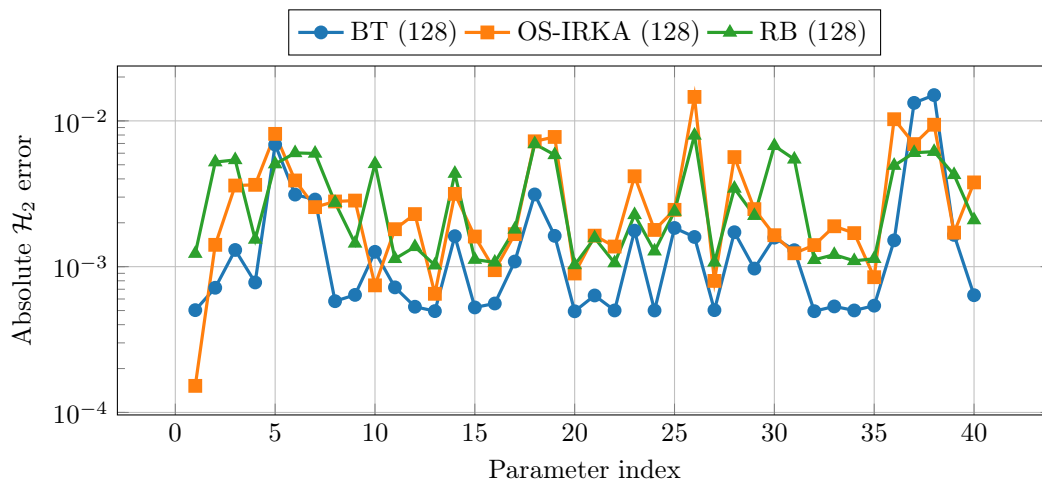


Figure 6: Comparison of methods for the four-parameter model