

Document downloaded from:

<http://hdl.handle.net/10251/176408>

This paper must be cited as:

Campos, C.; Roman, JE. (2021). NEP: A Module for the Parallel Solution of Nonlinear Eigenvalue Problems in SLEPc. ACM Transactions on Mathematical Software. 47(3):1-29.
<https://doi.org/10.1145/3447544>



The final publication is available at

<https://doi.org/10.1145/3447544>

Copyright Association for Computing Machinery

Additional Information

NEP: a module for the parallel solution of nonlinear eigenvalue problems in SLEPc*

Carmen Campos[†] Jose E. Roman[‡]

July 21, 2021

Abstract

SLEPc is a parallel library for the solution of various types of large-scale eigenvalue problems. In the last years we have been developing a module within SLEPc, called NEP, that is intended for solving nonlinear eigenvalue problems. These problems can be defined by means of a matrix-valued function that depends nonlinearly on a single scalar parameter. We do not consider the particular case of polynomial eigenvalue problems (which are implemented in a different module in SLEPc) and focus here on rational eigenvalue problems and other general nonlinear eigenproblems involving square roots or any other nonlinear function. The paper discusses how the NEP module has been designed to fit the needs of applications and provides a description of the available solvers, including some implementation details such as parallelization. Several test problems coming from real applications are used to evaluate the performance and reliability of the solvers.

1 Introduction

SLEPc, the Scalable Library for Eigenvalue Problem Computations [Hernandez et al., 2005, Roman et al., 2019], is a software package for the parallel solution of large-scale eigenvalue problems, with special focus on iterative methods to compute only part of the spectrum. It started as a collection of solvers for the linear eigenvalue problem, but its scope soon extended to other classes of related problems such as the singular value decomposition, the polynomial eigenvalue problem and the general nonlinear eigenvalue problem. SLEPc has a modular design, and this paper deals with the module called NEP, which is intended for the solution of general nonlinear eigenvalue problems. These eigenproblems depend nonlinearly on the eigenvalue parameter only, not the eigenvector. SLEPc also includes a basic nonlinear inverse iteration solver for problems that are nonlinear with respect to the eigenvector, but we do not discuss it here.

The NEP module can be used to solve rational eigenvalue problems, those defined by means of rational functions, or the more general class of nonlinear eigenproblems defined via arbitrary nonlinear functions (a formal definition will be given in §2). Note that polynomial eigenvalue problems are excluded because they belong to PEP, a different SLEPc module, whose details were described for the most part in a previous paper [Campos and Roman, 2016].

Nonlinear eigenvalue problems arise in many areas of computational science and engineering. For example, in wave phenomena eigenvalue problems become nonlinear when certain characteristic properties of the material depend on the eigenvalue in the frequency range of interest; a

*This work was supported by the Spanish Agencia Estatal de Investigación under grants TIN2016-75985-P and PID2019-107379RB-I00 (including European Commission FEDER funds).

[†]D. Sistemes Informàtics i Computació, Universitat Politècnica de València, València, Spain (mccampos@dsic.upv.es).

[‡]D. Sistemes Informàtics i Computació, Universitat Politècnica de València, València, Spain (jroman@dsic.upv.es).

typical example is found in the analysis of photonic crystals, where the relative permittivity is modeled by a rational function of the eigenvalue (the complex-valued frequency) [Brûlé et al., 2016]. Another example appears in the spectral analysis of delay differential equations [Michiels and Niculescu, 2014]. In some applications, the nonlinear character comes from using special boundary conditions for the solution of differential equations [Vandenbergh et al., 2014, van Beeumen et al., 2018].

The topic of nonlinear eigenvalue problems has been very active in the numerical linear algebra community in the last years. Many numerical methods have been proposed in the literature, which have been surveyed by Mehrmann and Voss [2004] and, more recently, by Güttel and Tisseur [2017]. The merit of some of these methods has been illustrated by solving various problems arising from real applications. However, as it often happens, these state-of-the-art algorithms are not available in the form of reliable and efficient free software implementations that can be widely used by the scientific computing community. To the best of our knowledge, apart from Matlab implementations of some algorithms, the only available package containing solvers for the general nonlinear eigenproblem is NEP-PACK [Jarlebring et al., 2018] for the Julia programming language. Our goal in SLEPc is to provide well-tested robust solvers with high performance for parallel computers that can be used from C/C++, Fortran and Python. We started the development of the NEP module in 2013, and since then it has been used in several scientific computing projects, e.g., in the simulation of photonic crystals [Demésy et al., 2020, Zolla et al., 2018], or the analysis of scattering resonances in metal-dielectric nano-structures [Araujo et al., 2020].

In this paper we will describe the methods that are currently available in SLEPc’s NEP module, and provide some relevant details of our implementation. But we will also discuss design decisions we have taken in order to organize the code. These decisions are usually geared towards making a friendly user interface. From the user point of view, defining a nonlinear eigenvalue problem is more difficult than in the linear case, since it is necessary to work with matrix-valued functions rather than with constant matrices. In SLEPc we provide ways to define these functions, and enable the solvers with capabilities to work with such functions, or their derivatives, and in some cases to compute matrix functions as well.

Rational eigenvalue problems are quite special because it is feasible to convert them to equivalent polynomial eigenvalue problems and it is also possible to solve them via linearization [Su and Bai, 2011]. But our approach is to use the same methods as for the more general nonlinear eigenproblems.

The rest of the paper is organized as follows. Section 2 gives a formal definition of the nonlinear eigenvalue problem and provides some necessary background. In section 3 we will describe the design of the NEP module and the user interface. Section 4 is devoted to the description of the different solvers available in SLEPc. A brief overview of the parallelization of the methods is included in section 5. The results of computational experiments are shown in section 6. We wrap up with some concluding remarks in section 7.

2 Mathematical background

Given a matrix-valued function $T : \Omega \rightarrow \mathbb{C}^{n \times n}$ defined in the open set $\Omega \subseteq \mathbb{C}$, the nonlinear eigenvalue problem (NEP) consists in determining eigenpairs $(x, \lambda) \in \mathbb{C}^n \times \mathbb{C}$ satisfying

$$T(\lambda)x = 0, \quad x \neq 0. \tag{1}$$

The vector x and the scalar λ are referred to as eigenvector and eigenvalue, respectively. We will assume that T is holomorphic in the domain Ω and that it is regular, that is, $\det T(z)$ is not identically zero. The eigenvalues are the solutions of the scalar equation $\det T(z) = 0$.

In addition to (right) eigenvectors, x , some applications also require the computation of left eigenvectors, y , satisfying

$$y^*T(\lambda) = 0^*, \quad y \neq 0, \quad (2)$$

where $(\cdot)^*$ denotes the conjugate transpose of a vector or matrix. Left eigenvectors can be computed as right eigenvectors of T^* .

The function T can always be expressed in *split form*

$$T(z) = \sum_{i=1}^{\ell} A_i f_i(z), \quad (3)$$

for constant coefficient matrices $A_1, \dots, A_\ell \in \mathbb{C}^{n \times n}$ and scalar holomorphic functions $f_1, \dots, f_\ell : \Omega \rightarrow \mathbb{C}$, with $\ell \leq n^2$. In many applications, T is naturally given in this form, usually with a small number of terms, $\ell \ll n^2$.

The linear eigenvalue problem is a particular case where T can be written as $T(z) = A - zB$ with $A, B \in \mathbb{C}^{n \times n}$. Compared to the linear eigenvalue problem, the solution structure of the NEP is more diverse. A NEP may have no solution at all, finitely many solutions, or infinitely many solutions, and eigenvectors belonging to distinct eigenvalues need not be linearly independent. The survey by Güttel and Tisseur [2017] reviews the solution structure of NEPs in detail. Next, we highlight the main points that are relevant for the rest of the paper.

The algebraic multiplicity of an eigenvalue λ is defined as the multiplicity of the root of $\det T(z)$ at $z = \lambda$. The geometric multiplicity of λ is the dimension of the null space of $T(\lambda)$. An eigenvalue is semisimple if its algebraic and geometric multiplicities coincide.

The Smith form is a factorization, analog of the Jordan form in linear problems, that reveals the structure of eigenvalues (including partial multiplicities) together with (generalized) right and left eigenvectors. The Smith form can also be used to work with the resolvent. In particular, Keldysh's theorem provides an explicit formula for the resolvent. In the simpler case when all eigenvalues λ_i are semisimple, the expression becomes

$$T^{-1}(z) = \sum_{i=1}^k (z - \lambda_i)^{-1} x_i y_i^* + R(z), \quad (4)$$

where k is the number of eigenvalues λ_i contained in Ω , x_i and y_i are the corresponding right and left eigenvectors, respectively, normalized so that $y_i^* T'(\lambda_i) x_i = 1$, and $R(z)$ is some function holomorphic in Ω .

Some algorithms rely on the concept of invariant pair [Kressner, 2009, Beyn et al., 2011]. A pair $(X, H) \in \mathbb{C}^{n \times k} \times \mathbb{C}^{k \times k}$ is invariant if

$$\mathbb{T}(X, H) := \frac{1}{2\pi i} \int_{\Gamma} T(z) X (zI - H)^{-1} dz = 0, \quad (5)$$

where Γ is a contour enclosing the eigenvalues of H . If T is expressed in the split form (3), then the definition can be written as

$$\mathbb{T}(X, H) := A_1 X f_1(H) + \dots + A_\ell X f_\ell(H) = 0, \quad (6)$$

where $f_i(H)$ denotes a matrix function [Higham, 2008]. Furthermore, an invariant pair is called minimal if there is an integer p such that $\text{rank } \mathcal{V}_p(X, H) = k$ for

$$\mathcal{V}_p(X, H) := \begin{bmatrix} X \\ XH \\ \vdots \\ XH^{p-1} \end{bmatrix}. \quad (7)$$

The smallest such p is called the minimality index of the pair (X, H) . If (X, H) is a minimal invariant pair for T , then the eigenvalues of H are eigenvalues of T [Kressner, 2009].

The algorithms discussed in §4 compute approximate solutions of (1). The accuracy of a computed solution $(\tilde{x}, \tilde{\lambda})$ can be assessed by means of the backward error

$$\eta(\tilde{x}, \tilde{\lambda}) = \min\{\varepsilon : (T(\tilde{\lambda}) + \Delta T(\tilde{\lambda}))\tilde{x} = 0, \|\Delta T\|_{\Omega} \leq \varepsilon\|T\|_{\Omega}\}, \quad (8)$$

where $\|T\|_{\Omega} = \sup_{z \in \Omega} \|T(z)\|$ for some matrix norm $\|\cdot\|$ on $\mathbb{C}^{n \times n}$. When $T(\lambda)$ is provided in split form (3), Güttel and Tisseur [2017] suggest using the explicit expression based on the scaled residual

$$\eta(\tilde{x}, \tilde{\lambda}) = \frac{\|T(\tilde{\lambda})\tilde{x}\|}{f(\tilde{\lambda})\|\tilde{x}\|}, \quad \text{with} \quad f(\tilde{\lambda}) = \sum_{i=1}^{\ell} |f_i(\tilde{\lambda})| \|A_i\|, \quad (9)$$

which is a straightforward generalization of [Tisseur, 2000, Theorem 1 and Lemma 3]. If the scaled residual in (9) is small then the approximate eigenpair $(\tilde{x}, \tilde{\lambda})$ is the exact eigenpair of a nearby problem. If the problem is not expressed in the split form, we will use the following upper bound instead,

$$\eta(\tilde{x}, \tilde{\lambda}) = \frac{\|T(\tilde{\lambda})\tilde{x}\|}{\|T(\tilde{\lambda})\|\|\tilde{x}\|}. \quad (10)$$

3 Design of the NEP module

In this section we describe how the NEP module is organized, including its relation with other SLEPc objects, and also how the user interacts with it, either with the programmatic interface or the command-line options.

Currently implemented methods are (details are postponed to §4):

- Successive linear problems (SLP), a Newton-type iteration where a linear eigenproblem is solved at each step to get the eigenvalue correction.
- Residual inverse iteration (RII), that computes a Newton-type correction of the eigenvector by multiplying $T(\sigma)^{-1}$ times the residual.
- Nonlinear Arnoldi, that builds an expanding subspace with the vectors generated by RII, and chooses approximate eigenpairs via Rayleigh-Ritz projection.
- Polynomial interpolation, where a matrix polynomial $P(\lambda)$ is built by evaluating $T(\cdot)$ at a few points, then a solver of class PEP is used to solve the polynomial eigenproblem.
- NLEIGS, which is based on a Krylov iteration operating on a companion-type linearization of a rational interpolant of the nonlinear function.

3.1 Overview of SLEPc

SLEPc, the Scalable Library for Eigenvalue Problem Computations [Hernandez et al., 2005, Roman et al., 2019], is a parallel library intended mainly for the solution of large-scale eigenvalue problems. The latest versions also include functionality for matrix functions. Figure 1 illustrates the relation among the different SLEPc components.

The main solver class is EPS (Eigenvalue Problem Solver), for linear eigenvalue problems (standard or generalized). For nonlinear eigenvalue problems, there is the PEP module (specific for polynomial eigenvalue problems) and the NEP module, which is the focus of this paper. In addition, there are two more solver classes that we will not discuss further in this paper: SVD for




Figure 1: Components of the SLEPc library. Shaded modules represent the main solvers, and white blocks are auxiliary classes. Solid arrows indicate that a solver may sometimes use another solver of different type, and dashed arrows mean dependence on auxiliary objects for some computations or due to the user interface (arrows from almost all solvers to `BV`, `DS`, and `RG` have been omitted).

singular value computations and `MFN` for computing the action of a sparse matrix function on a vector.

Figure 1 also shows five auxiliary classes. The first one (`ST`) provides functionality for spectral transformations such as shift-and-invert, which are used in linear and polynomial problems but not in the `NEP` module. The other four are relevant for this paper: `BV` for functionality relative to basis vectors, `DS` for (sequential) dense solvers used in projected problems, `RG` for specifying a region of the complex plane, and `FN` for specifying a scalar function and computing dense matrix functions.

SLEPc relies on PETSc (Portable, Extensible Toolkit for Scientific Computation [Balay et al., 2020]), so in addition to the components depicted in Figure 1, SLEPc also depends on PETSc classes, most notably data objects such as vectors (`Vec`) and matrices (`Mat`), and solvers for linear systems of equations (`KSP`) and preconditioners (`PC`).

PETSc and SLEPc have an object-oriented design, where objects of the different classes interact with each other via a public interface, hiding the implementation details within them. All solver classes have several implementations that can be selected at run time, so for instance a `KSP` object may run GMRES, Bi-CGStab, or others, depending on the user’s choice. The particular solver to be used, and many other settings, may be specified within the program as well as with command-line options.

The model of parallelism is based on the message-passing standard MPI¹. Every object is bound to an MPI communicator since its creation, and some of the object’s operations are collective, meaning that the associated computation is carried out in parallel among all MPI processes in the communicator. This implies that some data structures must be distributed across those processes. Almost all functionality of PETSc and SLEPc is parallelized, one of the most notable exceptions being the matrix factorizations (`LU`, Cholesky) that are sequential in PETSc. As a workaround for this, PETSc seamlessly integrates third-party packages such as MUMPS [Amestoy et al., 2001] to perform parallel direct linear solves.

PETSc’s object-oriented model is somewhat limited, as it is implemented in the C language. Some object-oriented features such as overloading or templates are not available. Object composition is very common (for instance, a `NEP` object may *contain* an `EPS` object), and, regarding inheritance, the hierarchy of classes is essentially flat, with polymorphism achieved via *subtyping* (for instance, a `NEP` object can be of type `slp`, `rii`, `narnoldi`, `interpol`, or `nleigs`). This subtyping mechanism involves a nomenclature convention, as will be seen when describing the

¹Local computations can exploit a second level of parallelism via a multi-threaded BLAS or with GPU computing, but we do not discuss this here.

```

1 NEP      nep;      /* eigensolver context      */
2 Mat      T, J;     /* Function and Jacobian matrices */
3 Vec      xr, xi;   /* eigenvector, x                */
4 PetscScalar kr, ki; /* eigenvalue, k                 */
5 ApplCtx  ctx;     /* user-defined context          */
6 PetscInt  j, nconv;
7 PetscReal error;
8
9 NEPCreate( PETSC_COMM_WORLD, &nep );
10 /* omitted: create and preallocate T and J matrices */
11 NEPSetFunction( nep, T, T, FormFunction, &ctx );
12 NEPSetJacobian( nep, J, FormJacobian, &ctx );
13 NEPSetFromOptions( nep );
14 NEPSolve( nep );
15 NEPGetConverged( nep, &nconv );
16 for (j=0; j<nconv; j++) {
17     NEPGetEigenpair( nep, j, &kr, &ki, xr, xi );
18     NEPComputeError( nep, j, NEP_ERROR_RELATIVE, &error );
19 }
20 NEPDestroy( &nep );

```

Figure 2: Example C code for basic solution with NEP using callbacks.

user interface in §3.2 and §4. The function `XXXYYYSetZZZ` belongs to type `YYY` within class `XXX` and sets the attribute `ZZZ`. If `YYY` is omitted, it refers to an attribute of the base class. For example, `NEPSetTolerances` sets two attributes (`tol` and `max_it`) that are common to all NEP objects, and can be set also via the command line with `-nep_tol` and `-nep_max_it`, whereas `NEPRIISetHermitian` sets the hermitian flag that is specific of the `RII` type, and the syntax for setting it in the command-line is `-nep_rii_hermitian`. Objects can define a prefix for their associated command-line options, and this is used in the case of object composition. For example, a solver of type `SLP` contains an object of type `EPS`, whose options can be set, e.g., with `-nep_slp_eps_tol`. We conclude the discussion about object-oriented features noting that PETSc provides a mechanism to extend a class by adding new user-defined types via registering the constructor of the new type. In this way, a user can implement a new NEP solver and register it in the system via `NEPRegister`.

3.2 User interface

The basic usage of eigensolvers in SLEPc requires creating the solver object, passing the matrices that define the problem, and then invoking the “solve” method. After that, the solution is kept internally to the solver and can be retrieved with a few function calls. The solution process can be customized with different settings as we will discuss later.

In `EPS` and `PEP`, the user defines the problem by providing the coefficient matrices. In the case of `NEP`, the matrix-valued function $T(\cdot)$ requires also information about the nonlinear function. We will now describe how the application code can define $T(\cdot)$ as well as the derivative $T'(\cdot)$, that we will call the Jacobian and is required by some solvers. There are two ways of doing this: with callback functions and using the split form.

3.2.1 Callbacks

We start describing the callback function mechanism. Figure 2 shows basic code for solving a NEP in this way. Two variables are used to represent both the eigenvalue and the eigenvector, and the rationale for this has to do with real versus complex arithmetic as discussed below

(§3.2.3). By default, the solver attempts to compute just one eigenpair, but this and other settings can be adjusted (§3.2.4). After calling the solver, the example checks the number of converged eigenpairs, and extracts each of them from the solver object. It also computes the relative residual error, $\|T(\lambda_j)x_j\|/|\lambda_j|$.

Lines 11-12 of the example code are the relevant ones for the callback function mechanism. With `NEPSetFunction` the user provides a PETSc matrix `T`, a user-defined function `FormFunction`, and an optional context `ctx`. The user-defined function (not shown in Figure 2) must have the prototype

```
PetscErrorCode FormFunction(NEP nep, PetscScalar lambda, Mat T, Mat P, void *ctx)
```

and must contain code that computes the entries and assembles the matrix $T(\lambda)$ from the value of λ (argument `lambda`) and possibly some application-specific parameters stored in the data structure pointed by `ctx`. The computed matrix is the argument `T` (note that this matrix is created only once in the main program and reused here). Similarly, the user-defined function `FormJacobian` builds a PETSc matrix `J` that stores $T'(\lambda)$ for a given λ . The main difference between the two callback functions is that `FormFunction` has an additional argument `P` intended to build a preconditioner matrix (that for basic usage is equal to matrix `T`).

With this scheme, whenever the SLEPc solver needs to evaluate $T(\lambda)$ or $T'(\lambda)$, it invokes the user subroutines to have this done. As we will see in §4, most solvers rely on evaluating these matrices (or just $T(\lambda)$ in some cases) for different values of the parameter. Often, it is also necessary to perform linear system solves with $T(\sigma)$, usually for a fixed value σ referred to as the target (§3.2.4).

We point out that it is also possible to solve the problem in a matrix-free fashion, that is, just implementing subroutines that compute the action of $T(\lambda)$ or $T'(\lambda)$ on a vector, rather than explicitly building these two matrices. This can be done by means of PETSc *shell matrices*.

3.2.2 Split form

An alternative way to specify the nonlinear operator $T(\cdot)$ is the split form (3), where the user provides a separate set of constant matrices and nonlinear functions. This is the recommended usage because most problems coming from real applications can be written naturally in the split form with only a few terms, and also because, compared to the callback mechanism, it is simpler (the user does not have to worry about how to compute the derivative) and more robust (no chance to introduce bugs in user callbacks). Still, the callback mechanism is required for problems that cannot be described easily in the split form.

To illustrate the use of the split form in SLEPc, let us focus on a particular NEP coming from the discretization of a parabolic partial differential equation with time delay τ ,

$$(-\lambda I + A + e^{-\tau\lambda}B)x = 0. \quad (11)$$

In this case, the matrix-valued function $T(\lambda)$ has a polynomial part $(-\lambda I + A)$ and a nonlinear part $(e^{-\tau\lambda}B)$. In SLEPc we do not distinguish between the polynomial and nonlinear parts, so all terms are treated equally. Hence, for this example we need to build three matrices (the identity I , A , and B) and three analytic functions

$$f_1(\lambda) = -\lambda, \quad f_2(\lambda) = 1, \quad f_3(\lambda) = e^{-\tau\lambda}. \quad (12)$$

Figure 3 illustrates how `NEPSetSplitOperator` is used to provide an array of matrices and functions. The code for creating the matrices is not shown in the figure. The functions (12) are managed with SLEPc's auxiliary object `FN`, that can be used to represent mathematical functions from a set of predefined types and combinations thereof, in particular:


```

1  FNCreate(PETSC_COMM_WORLD,&f1); /* f1 = -lambda */
2  FNSetType(f1,FNRATIONAL);
3  coeffs[0] = -1.0; coeffs[1] = 0.0;
4  FNRationalSetNumerator(f1,2,coeffs);
5
6  FNCreate(PETSC_COMM_WORLD,&f2); /* f2 = 1 */
7  FNSetType(f2,FNRATIONAL);
8  coeffs[0] = 1.0;
9  FNRationalSetNumerator(f2,1,coeffs);
10
11 FNCreate(PETSC_COMM_WORLD,&f3); /* f3 = exp(-tau*lambda) */
12 FNSetType(f3,FNEXP);
13 FNSetScale(f3,-tau,1.0);
14
15 mats[0] = A;  funs[0] = f2;
16 mats[1] = Id; funs[1] = f1;
17 mats[2] = B;  funs[2] = f3;
18 NEPSetSplitOperator(nep,3,mats,funs,SUBSET_NONZERO_PATTERN);

```

Figure 3: Excerpt of C code illustrating the definition of the nonlinear function in split form.

- Polynomial and rational functions, by providing the list of coefficients of the numerator and (optionally) the denominator.
- The (principal) square root, $x^{\frac{1}{2}}$, and the inverse square root, $x^{-\frac{1}{2}}$.
- Transcendental functions: the exponential, the logarithm, and the φ -functions defined as

$$\varphi_0(x) = e^x, \quad \varphi_1(x) = \frac{e^x - 1}{x}, \quad \varphi_k(x) = \frac{\varphi_{k-1}(x) - 1/(k-1)!}{x}. \quad (13)$$

- Scaling factors can be specified for both the argument and the function, with `FNSetScale`, to form, e.g., $g(x) = \beta e^{\alpha x}$.
- The special type `FNCOMBINE` allows creating new functions by combining two `FN` objects with either addition, multiplication, division, or function composition.
- If the previous functionality is not enough to cover the needs of an application, the user can extend SLEPc by adding a new `FN` function implementation via `FNRegister`.

Every `FN` object, representing a function $f : \mathbb{C} \rightarrow \mathbb{C}$, provides the `FNEvaluateFunction` and `FNEvaluateDerivative` operations to evaluate $f(x)$ and $f'(x)$ for a given value x . When the nonlinear operator of the NEP is represented in split form (3), SLEPc is able to build matrix $T(\lambda)$ by evaluating $f_i(\lambda)$, $i = 1, \dots, \ell$, together with $\ell - 1$ matrix *axpy* operations, and $T'(\lambda)$ can be built similarly by evaluating $f'_i(\lambda)$ instead. Note that the *axpy* operations for sparse matrices as implemented in PETSc can be optimized in the case that the sparsity patterns are related; this is hinted with the last argument of `NEPSetSplitOperator` (see line 18 of Figure 3).

We note in passing that `FN` is also capable of evaluating matrix functions $f : \mathbb{C}^{k \times k} \rightarrow \mathbb{C}^{k \times k}$ [Higham, 2008]. This is required by some methods, as will be pointed out in §4.

3.2.3 Real and complex arithmetic

A feature of PETSc (and hence of SLEPc) is that the library is built for either real or complex scalars. This implies that the basic type `PetscScalar` (and consequently all data structures such as matrices and vectors) is either a real number or a complex number. It is not possible to mix

real and complex scalars in the same program². This has an impact in SLEPc’s user interface, because real matrices may have complex eigenpairs.

In linear eigenproblem (EPS), the solver typically works internally with a (partial) real or complex Schur form, in real and complex arithmetic, respectively. When the final computed solution is to be returned to the user, it is represented as follows. In real arithmetic, two `PetscScalar` variables are used for the eigenvalue, storing the real and imaginary parts, and similarly two `Vec` objects are employed for the eigenvector. If the eigenvalue is real, the imaginary part is set to zero. In complex arithmetic, only one variable is used for the eigenvalue and one vector for the eigenvector.

This strategy carries over to nonlinear eigenproblems (PEP and NEP). However, the NEP case needs further clarification because, as described above, the solvers need to evaluate $T(\lambda)$ (or its derivative) and the question is what happens if λ is complex when the computation is done in real arithmetic. This case would oblige to represent $T(\lambda)$ with two `Mat` objects, one for the real part and one for the imaginary part, with the consequent complication of the user interface: for instance the user callback `FormFunction` would require two scalar arguments `lambdar` and `lambdai`, and four matrix arguments `Tr`, `Ti`, `Pr`, `Pi`. Furthermore, this design would involve a high complication in real arithmetic to use PETSc’s linear solvers in KSP.

To avoid the complication described above, we decided to forbid the evaluation of $T(\lambda)$ for a complex λ in case of real arithmetic. If this happens, the solver would fail with an error message. Even though this may seem a severe limitation, in practice this implies that solvers will not be able to compute complex eigenvalues in real arithmetic, and for real eigenvalues the search region (described next) must be restricted to an interval of the real line.

3.2.4 Settings

We now describe general settings of the NEP object. The user can change them within the code or with a command-line option. For instance, calling `NEPSetTarget(nep, 1.4)` is equivalent to running the program adding the command-line argument `-nep_target 1.4` (command-line options are processed in the `NEPSetFromOptions` operation, see line 13 of Figure 2).

Basic settings include the following:

- `NEPSetDimensions` allows specifying the number of eigenvalues to compute (`nev`) and the maximum allowed dimension of the search subspace (`ncv`, only relevant in some solvers).
- `NEPSetTolerances` indicates the tolerance for the stopping criterion and the maximum number of iterations.
- `NEPSetWhichEigenpairs` is used to indicate which eigenvalues are of interest, e.g., those with largest magnitude, with largest real part, or closest to a given target value σ .
- `NEPSetTarget` is used to provide the value of the target, σ .
- `NEPSetType` selects the solver to be used. Available solvers are described in §4.
- `NEPSetProblemType` can be used to indicate that the problem is `rational` (as opposed to `general`). The aim of problem types in SLEPc is to allow users to give a hint to the solver so that it can make optimizations in some cases.

As other eigensolver classes, NEP can be used with a region defined via an auxiliary object `RG`. The `RG` object is a simple mechanism in SLEPc to define regions of the complex plane, such

²This design decision is sometimes criticized as a weakness of PETSc, but it eliminates many complications/inefficiencies.

ALGORITHM 1: Successive Linear Problems (SLP)

Input: Initial eigenvalue approximation $\lambda^{(0)}$
Output: Computed eigenpair $(x^{(k)}, \lambda^{(k)})$

- 1 **for** $k=0,1,2,\dots$ **do**
- 2 Evaluate $A = T(\lambda^{(k)})$, $B = T'(\lambda^{(k)})$
- 3 **if** $k > 0$ and $\eta(x^{(k)}, \lambda^{(k)}) < \mathit{tol}$ **then** exit
- 4 Compute $(x^{(k+1)}, \mu^{(k)})$, the smallest magnitude eigenpair of $Ax = \mu Bx$
- 5 Update $\lambda^{(k+1)} = \lambda^{(k)} - \mu^{(k)}$
- 6 **end**

as an interval, a polygon, an ellipse, or a ring. In NEP, the meaning of the region is different depending on the solver type (see details in §4):

- In interpolation methods, the region is used to select the interpolation nodes, and also as a filtering mechanism (eigenvalue approximations outside the region are discarded).
- In the other methods (Newton-based) the region is not taken into account.

3.2.5 Left eigenvectors and the resolvent

In order to compute left eigenvectors (2), the user must select a two-sided variant of the solver with `NEPSetTwoSided`. Currently, only the NLEIGS solver (§4.5) supports this option. Then, after the solver has finished, left eigenvectors can be retrieved with `NEPGetLeftEigenvector`. Furthermore, the operation `NEPApplyResolvent` is available to compute the action of the resolvent (4) on a vector.

4 Solvers

This section describes the solvers that are available in the NEP module. The solvers are very different from each other and it is convenient that the user knows how they work in order to decide which one is the most appropriate for a particular problem. In general, NLEIGS (§4.5) is often our most competitive solver, but other solvers may be better suited for special cases.

4.1 Classical Iterations

Ruhe [1973] proposed the method of successive linear problems (SLP). This method is based on linearization via Taylor's formula,

$$T(\lambda + \mu) = T(\lambda) + \mu T'(\lambda) + \frac{\mu^2}{2} R(\lambda, \mu), \quad (14)$$

from which high-order terms R are discarded. This results in an iteration where at each step the correction μ to the eigenvalue is obtained by computing the smallest magnitude eigenvalue of a linear eigenvalue problem, see Algorithm 1. Several comments are in order:

- The convergence criterion shown in line 3 of Algorithm 1 is based on the scaled residual $\eta(x, \lambda)$ defined in (10) and (9), but the user can also choose the absolute residual, the relative residual, or even a custom convergence test.
- The eigenproblem in line 4 is solved via an EPS solver. In particular, it is solved with shift-and-invert Krylov-Schur with target $\sigma = 0$, which in practice is equivalent to computing

largest magnitude eigenvalues of $A^{-1}Bx = \mu^{-1}x$. Therefore, matrix $T(\lambda^{(k)})$ is factorized at each step k .

- When solving the eigenproblem in line 4, $x^{(k)}$ is set as initial guess for the wanted eigenvector, $x^{(k+1)}$.
- Algorithm 1 computes just one eigenpair. Computation of several eigenpairs requires adding a deflation mechanism, described in §4.3.

The main drawback of SLP is that a factorization of $T(\lambda^{(k)})$ is required at each iteration. The same problem is present in nonlinear inverse iteration, also discussed in [Ruhe, 1973]. We have implemented the residual inverse iteration (RII), a modification proposed by Neumaier [1985] that replaces $T(\lambda^{(k)})$ with $T(\sigma)$, for a constant value σ , with the consequent reduction of computational cost. To derive these methods, start with Newton's method applied to the system of $n + 1$ nonlinear equations consisting of $T(\lambda)x = 0$ together with a normalization condition. The Newton iteration can be written as

$$T(\lambda^{(k)})x^{(k+1)} = -(\lambda^{(k+1)} - \lambda^{(k)})T'(\lambda^{(k)})x^{(k)}, \quad (15)$$

followed by normalization of $x^{(k+1)}$. The Newton update (15) can be expressed in the form of nonlinear inverse iteration, where the new eigenvector approximation is computed by solving the linear system

$$T(\lambda^{(k)})x^{(k+1)} = T'(\lambda^{(k)})x^{(k)}, \quad (16)$$

and then the eigenvalue approximation $\lambda^{(k+1)}$ is updated with a correction. Neumaier's idea is to rewrite (15) as

$$\begin{aligned} x^{(k)} - x^{(k+1)} &= x^{(k)} + (\lambda^{(k+1)} - \lambda^{(k)})T(\lambda^{(k)})^{-1}T'(\lambda^{(k)})x^{(k)} \\ &= T(\lambda^{(k)})^{-1} \left(T(\lambda^{(k)}) + (\lambda^{(k+1)} - \lambda^{(k)})T'(\lambda^{(k)}) \right) x^{(k)}. \end{aligned} \quad (17)$$

The expression in parenthesis is a first-order approximation of $T(\lambda^{(k+1)})$, as in (14), resulting in the iteration

$$x^{(k+1)} = x^{(k)} - T(\lambda^{(k)})^{-1}T(\lambda^{(k+1)})x^{(k)}. \quad (18)$$

The main difference of residual inverse iteration (18), with respect to Newton, is that the new eigenvalue approximation $\lambda^{(k+1)}$ must be computed in advance. The algorithm in [Neumaier, 1985] replaces $\lambda^{(k)}$ with a fixed shift σ , showing that convergence is not affected.

Our implementation is shown in Algorithm 2. Lines 3–6 represent a Newton iteration to solve the scalar nonlinear equation

$$x^{(k)*}T(\sigma)^{-1}T(z)x^{(k)} = 0, \quad (19)$$

where the computed solution is taken as $\lambda^{(k+1)}$. As suggested by Neumaier, this equation can be replaced with the cheaper version $x^{(k)*}T(z)x^{(k)} = 0$ if $T(\lambda)$ is Hermitian. In our implementation, the user can choose between the two versions with `NEPRIISetHermitian`. The stopping criterion for this loop uses a tolerance equal to the square root of the machine epsilon, ε . In line 9, we give the user the possibility of updating σ with the value of $\lambda^{(k+1)}$. The user interface function `NEPRIISetLagPreconditioner` allows to do this, by specifying the number of iterations after which the preconditioner $T(\sigma)$ must be updated (0 means a fixed σ , and 1 implies updating at every iteration as in Newton). The linear solve of line 10 is carried out via a PETSc's `KSP` object, that can also be customized by the user (by default, GMRES with block Jacobi is used). In case of iterative linear solves, we offer the possibility of using a constant tolerance or an exponentially decreasing tolerance (`NEPRIISetConstCorrectionTol`).

ALGORITHM 2: Residual Inverse Iteration (RII)

Input: Target value σ , initial guess $x^{(0)}$
Output: Computed eigenpair $(x^{(k)}, \lambda^{(k+1)})$

- 1 $\lambda^{(0)} = \sigma$
- 2 **for** $k=0, 1, 2, \dots$ **do**
- 3 $\lambda^{(k+1)} = \lambda^{(k)}$
- 4 **repeat**
- 5 $\lambda^{(k+1)} = \lambda^{(k+1)} - \mu$, with $\mu = \frac{x^{(k)*} T(\sigma)^{-1} T(\lambda^{(k+1)}) x^{(k)}}{x^{(k)*} T(\sigma)^{-1} T'(\lambda^{(k+1)}) x^{(k)}}$
- 6 **until** $|\mu|/|\lambda^{(k+1)}| < \sqrt{\varepsilon}$
- 7 Compute residual $r = T(\lambda^{(k+1)})x^{(k)}$
- 8 **if** $\eta(x^{(k)}, \lambda^{(k+1)}) < \mathit{tol}$ **then** exit
- 9 **if** *required* **then** update σ , evaluate $T(\sigma)$
- 10 Solve $T(\sigma)v = r$
- 11 Update and normalize $x^{(k+1)} = \tilde{x}^{(k+1)}/\|\tilde{x}^{(k+1)}\|$, with $\tilde{x}^{(k+1)} = x^{(k)} - v$
- 12 **end**

ALGORITHM 3: Nonlinear Arnoldi (N-Arnoldi)

Input: Target value σ , initial guess $x^{(0)}$
Output: Computed eigenpair $(x^{(k)}, \lambda^{(k)})$

- 1 Initialize basis $V_1 = [x^{(0)}/\|x^{(0)}\|]$
- 2 **for** $k=1, 2, \dots$ **do**
- 3 Compute $(y, \tilde{\lambda})$ satisfying projected problem $V_k^* T(\tilde{\lambda}) V_k y = 0$
- 4 Compute Ritz pair $(x^{(k)} = V_k y, \lambda^{(k)} = \tilde{\lambda})$ and residual $r = T(\lambda^{(k)})x^{(k)}$
- 5 **if** $\eta(x^{(k)}, \lambda^{(k)}) < \mathit{tol}$ **then** exit
- 6 Solve $T(\sigma)v = r$
- 7 Expand subspace: $\tilde{v} = v - V_k V_k^* v$, $V_{k+1} = [V_k, \tilde{v}/\|\tilde{v}\|]$
- 8 **end**

4.2 Nonlinear Arnoldi

Both the SLP and RII methods presented in the previous section are single-vector iterations. However, in the area of linear eigenproblems it is well-known that single-vector iterations can be less effective than algorithms based on projection onto an expanding subspace. This idea can also be applied to the NEP. Voss [2004] suggests a projection method with RII iterates, that is, correction vectors obtained in step 10 of Algorithm 2 are used to expand a subspace whose orthogonal basis is V_k . Then, instead of solving a one-dimensional nonlinear equation (19), the problem is projected onto this k -dimensional subspace, resulting in a NEP of size k

$$V_k^* T(\tilde{\lambda}) V_k y = 0. \quad (20)$$

An eigenvalue $\tilde{\lambda}$ satisfying (20) is the Ritz approximate eigenvalue $\lambda^{(k)}$ of the original NEP, and the corresponding eigenvector y is related to the Ritz vector $x^{(k)} = V_k y$. Voss called this method Nonlinear Arnoldi (N-Arnoldi). Algorithm 3 sketches our implementation of this method in SLEPc.

Next, we discuss how the N-Arnoldi solver in SLEPc manages the projected problem (20). With the split form, (20) turns into

$$\left(\sum_{i=1}^{\ell} B_i f_i(\tilde{\lambda}) \right) y = 0, \quad B_i = V_k^* A_i V_k. \quad (21)$$

The projected matrices B_i are computed incrementally, that is, their k th rows and columns are computed at step k using the last column of V_k . The projected NEP (21) is solved with **DS**, which is one of SLEPc's auxiliary objects depicted in Figure 1. **DS** stands for Direct Solver (or Dense System), and is intended to solve small-sized eigenproblems whose coefficient matrices are dense. For linear eigenproblems, **DS** involves calling one or more LAPACK subroutines. But for NEP dense systems, we need to provide a nonlinear eigensolver at the **DS** level. Currently, we have implemented a dense version of the SLP method of Algorithm 1. This dense implementation receives the B_i matrices and a handle to the f_i functions, and performs all steps of the algorithm with dense computations sequentially, with LAPACK's `_ggev` to compute the eigenvalue correction.

4.3 Deflation

The methods presented so far (SLP, RII and N-Arnoldi) approximate just one eigenpair. However, SLEPc's implementation of these methods support the computation of more than one eigenpair. Subsequent eigenpairs are computed by implementing a deflation mechanism, to avoid unwanted convergence to previously computed eigenpairs. Deflation can be achieved via a non-equivalence transformation to the matrix-valued function $T(\cdot)$ that maps already computed eigenvalues to infinity [Güttel and Tisseur, 2017]. We have opted for implementing the deflation technique developed by Effenberger [2013], as described below.

Effenberger's deflation is based on the concept of minimal invariant pair described in §2. Suppose a minimal invariant pair $(X, H) \in \mathbb{C}^{n \times k} \times \mathbb{C}^{k \times k}$ of the NEP has already been obtained, with minimality index p . We want to compute an extended pair

$$(\tilde{X}, \tilde{H}) = \left([X \quad x], \begin{bmatrix} H & t \\ 0 & \lambda \end{bmatrix} \right) \quad (22)$$

such that it contains a new eigenpair (x, λ) and is also a minimal invariant pair of the same NEP. As we will see below, imposing the conditions of invariant pair and minimality on the extended pair (22) results in an extended nonlinear eigenvalue problem (of size $n+k$),

$$\tilde{T}(\lambda) \begin{bmatrix} x \\ t \end{bmatrix} = \begin{bmatrix} T(\lambda) & U(\lambda) \\ A(\lambda) & B(\lambda) \end{bmatrix} \begin{bmatrix} x \\ t \end{bmatrix} = 0, \quad (23)$$

where $U(\lambda)$, $A(\lambda)$ and $B(\lambda)$ are matrix-valued functions of dimensions $n \times k$, $k \times n$ and $k \times k$, respectively. Solving the extended NEP (23) provides the required data (λ, x, t) for the wanted extension (22). Hence, in SLEPc the implementations of SLP, RII and N-Arnoldi apply the respective algorithm to the extended NEP (23), operating by blocks as discussed below, to compute one eigenpair at a time.

To enforce the definition of invariant pair (5) on the extended pair, Effenberger uses the fact that

$$(zI - \tilde{H})^{-1} = \begin{bmatrix} (zI - H)^{-1} & (zI - H)^{-1}t(z - \lambda)^{-1} \\ 0 & (z - \lambda)^{-1} \end{bmatrix}, \quad (24)$$

and writes

$$\mathbb{T}(\tilde{X}, \tilde{H}) = [\mathbb{T}(X, H) \quad T(\lambda)x + U(\lambda)t], \quad (25)$$

with

$$U(\lambda) = \frac{1}{2\pi i} \int_{\Gamma} T(z)X(zI - H)^{-1}(z - \lambda)^{-1}dz, \quad (26)$$

where the contour Γ encloses the eigenvalues of H and λ . Since (X, H) is invariant, the condition $\mathbb{T}(\tilde{X}, \tilde{H}) = 0$ holds if and only if $T(\lambda)x + U(\lambda)t = 0$. This relation constitutes the first block

equation of the extended NEP (23). The second block equation, $A(\lambda)x + B(\lambda)t = 0$, results from imposing the condition that (\tilde{X}, \tilde{H}) is minimal, (7), with minimality index not exceeding $p+1$. Effenberger [2013] proves this and gives the expressions

$$A(\lambda) = \sum_{i=0}^p \lambda^i (XH^i)^*, \quad B(\lambda) = \sum_{i=1}^p (XH^i)^* X q_i(\lambda), \quad (27)$$

with

$$q_i(\lambda) = \sum_{j=0}^{i-1} \lambda^j H^{i-j-1}. \quad (28)$$

$A(\lambda)$ and $B(\lambda)$ are matrix polynomials, which are easy to evaluate. However, evaluating $U(\lambda)$ is more involved. It can be shown [Effenberger, 2013] that if λ is not an eigenvalue of H ,

$$U(\lambda) = T(\lambda)X(\lambda I - H)^{-1}. \quad (29)$$

Alternatively, if $T(\cdot)$ is expressed in the split form (3), then

$$U(\lambda) = \sum_{i=1}^{\ell} A_i X \phi_i(\lambda), \quad (30)$$

where $\phi_i(\lambda)$ is the $k \times k$ leading principal submatrix of the matrix function $f_i \left(\begin{bmatrix} H & I \\ 0 & \lambda I \end{bmatrix} \right)$. For this, our solvers make use of matrix functions as implemented in the FN module, as was mentioned at the end of §3.2.2. If any of the provided f_i functions lacks a matrix function implementation, we fall back to (29) hoping that λ is not an eigenvalue of H .

We now discuss how we have implemented the main operations in a block fashion, that is, without explicitly building the operator of the extended NEP (23). According to Algorithms 1, 2 and 3, we need to perform matrix-vector multiplications with $\tilde{T}(\lambda)$ and $\tilde{T}'(\lambda)$, and linear solves with $\tilde{T}(\sigma)$. The general strategy is that we have **Vec** objects representing vectors $\begin{bmatrix} x \\ t \end{bmatrix}$ of length $n+k$, from which we can easily extract two subvectors x and t , also **Vec** objects, of length n and k , respectively, to perform the operation by blocks, the first vector being a parallel vector (distributed across all processes) while the second one is sequential (stored redundantly in all processes to minimize communication, see §5).

Matrix-vector products with the extended operator, $\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \tilde{T}(\lambda) \begin{bmatrix} z_1 \\ z_2 \end{bmatrix}$, can be accomplished with

$$\begin{aligned} y_1 &= T(\lambda)z_1 + \sum_{i=1}^{\ell} A_i X \phi_i(\lambda)z_2, \\ y_2 &= \sum_{i=0}^p \lambda^i (XH^i)^* z_1 + \sum_{i=1}^p (XH^i)^* X q_i(\lambda)z_2. \end{aligned} \quad (31)$$

For the linear system solves $\begin{bmatrix} T(\sigma) & U(\sigma) \\ A(\sigma) & B(\sigma) \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = \begin{bmatrix} b_1 \\ b_2 \end{bmatrix}$ we define the Schur complement

$$S(\sigma) := B(\sigma) - A(\sigma)T(\sigma)^{-1}U(\sigma), \quad (32)$$

that in the case of callback representation is computed as $B(\sigma) - A(\sigma)X(\sigma I - H)^{-1}$ using (29) and assuming that σ is not an eigenvalue of H . The operation is divided in three steps:

$$v = T(\sigma)^{-1}b_1 \quad \begin{cases} x_2 = S(\sigma)^{-1}(b_2 - A(\sigma)v), \\ x_1 = v - T(\sigma)^{-1}U(\sigma)x_2 = v - X(\sigma I - H)^{-1}x_2, \end{cases} \quad (33)$$

where we have used (29) again for x_1 . The computation of v amounts to a standard linear solve of size n with PETSc's KSP.

Finally, N-Arnoldi requires an additional operation (step 3 of Algorithm 3) to perform the projection of the extended operator onto a basis of length $n+k$,

$$\begin{aligned} \begin{bmatrix} V_1^* & V_2^* \end{bmatrix} \begin{bmatrix} T(\lambda) & U(\lambda) \\ A(\lambda) & B(\lambda) \end{bmatrix} \begin{bmatrix} V_1 \\ V_2 \end{bmatrix} &= V_1^* T(\lambda) V_1 + V_1^* U(\lambda) V_2 + V_2^* A(\lambda) V_1 + V_2^* B(\lambda) V_2 = \\ \sum_{i=1}^{\ell} V_1^* A_i V_1 f_i(\lambda) + \sum_{i=1}^{\ell} V_1^* A_i X \phi_i(\lambda) V_2 + \sum_{i=0}^p \lambda^i V_2^* (X H^i)^* V_1 + \sum_{i=1}^p V_2^* (X H^i)^* X q_i(\lambda) V_2. \end{aligned} \quad (34)$$

The first term contains the B_i matrices of (21), that are updated for each new basis vector.

In some problems, using the deflation may produce slightly inaccurate eigenpairs. The reason is that deflating with respect to an eigenpair that has been computed with a certain error (given by the user's tolerance) introduces a perturbation that limits the attainable accuracy of subsequent eigenpairs. This means that, for instance, the second eigenpair has about one significant digit less than the first one. As a remedy, we have introduced a threshold parameter in SLP and RII (`NEPRIISetDeflationThreshold`) that works as follows. At the beginning, the solver operates with the extended operator so that previous eigenpairs are deflated. When the relative residual error of the current eigenpair is below the threshold, the solver switches to the simpler variant without deflation. In that way, the last part of the computation proceeds without the perturbation, and better accuracy can be attained. We assume that once the threshold has been reached, the iterates will be already in the convergence basin of the current eigenpair. Still, there is risk of misconvergence if this option is used. In the experiments of §6, this option was needed only in one problem (turned off by default).

4.4 Polynomial Interpolation

The methods presented in §4.1 and §4.2 derive from Newton's method in one way or another. A completely different strategy for solving the NEP consists in approximating $T(\lambda)$ with a matrix polynomial $P(\lambda)$ (or a rational matrix-valued function $R(\lambda)$, see §4.5) and then computing the solution of this alternative eigenproblem, hoping that its solution is in good agreement with the solution of the original problem.

In this section we discuss SLEPc's `interp` solver, that implements the method proposed by Effenberger and Kressner [2012] based on a Chebyshev interpolation of $T(\lambda)$. This solver can be used for the particular case when the wanted eigenvalues of the NEP are real and lie within a prescribed closed interval of the real line, $\mathcal{I} = [a, b]$. The user specifies the interval with an `RG` object of type `interval`. The method has two stages. First, it computes the coefficient matrices of $P_d(\lambda)$, the Chebyshev polynomial (of the first kind) of degree d that interpolates $T(\lambda)$ at the Chebyshev nodes

$$\nu_i := \cos \left(\left(i + \frac{1}{2} \right) \frac{\pi}{d+1} \right), \quad i = 0, \dots, d, \quad (35)$$

in the interval $[-1, 1]$. The mapping from \mathcal{I} to $[-1, 1]$ usually requires a change of variable,

$$\lambda = \frac{b-a}{2} \theta + \frac{b+a}{2}. \quad (36)$$

If P_d is expressed in the basis of Chebyshev polynomials $\tau_k(\cdot)$,

$$P_d(\lambda) = \frac{C_0}{2} \tau_0(\lambda) + \sum_{k=1}^d C_k \tau_k(\lambda), \quad (37)$$

then $C_k = \frac{2}{d+1} \sum_{i=0}^d T(\nu_i) \tau_k(\nu_i)$, for ν_i defined in (35), or equivalently,

$$C_k = \frac{2}{d+1} \sum_{i=0}^d T(\cos \omega_i) \cos(k\omega_i), \quad \omega_i = \left(i + \frac{1}{2}\right) \frac{\pi}{d+1}. \quad (38)$$

Secondly, the polynomial eigenvalue problem $P_d(\lambda)x = 0$ is solved and the eigensolutions (x, λ) are returned as solutions of the NEP. The nice thing of this approach is that the polynomial eigenproblem is managed with a SLEPc solver of class `PEP`, so here we leverage all the functionality available in that module. By default, the polynomial eigenproblem is solved via linearization and a memory-efficient Krylov solver. The details of how this linearization is done when the polynomial is expressed in the Chebyshev basis can be found in [Campos and Roman, 2016]. An alternative to linearization is to employ a polynomial Jacobi-Davidson solver with support for Chebyshev basis [Campos and Roman, 2020]. In any case, the solver must be used in a combination with a region filter (`RG`) in order to discard computed eigenvalue approximations that do not lie in the interval (or its immediate vicinity).

Effenberger and Kressner [2012] show that the approximation error $\|T(\lambda) - P_d(\lambda)\|_F$ decreases exponentially, uniformly in $[-1, 1]$, when the interpolation degree is increased. They also show that, for sufficiently large d , the eigenvalues obtained in the associated polynomial eigenproblem are good approximations to eigenvalues of the NEP whenever those are close to the interval \mathcal{I} . The degree of the polynomial, d , is a user setting that can be specified at run time (`NEPInterpolSetInterpolation`). It must be noted that the solver will compute the eigenpairs of the polynomial eigenproblem up to the accuracy requested by the user (tolerance), but the actual error corresponding to the NEP may be larger if the value d is not large enough. There is no automatic procedure to choose an appropriate value of d .

4.5 Rational Interpolation

As mentioned in the previous section, the polynomial interpolation solver is restricted to the case of real eigenvalues. Another limitation is that the approximation error gets worse if the interpolation points are not sufficiently far away from the singularities of $T(\cdot)$. In such cases it is preferable to have recourse to a rational interpolation that takes into account the singularities. In SLEPc we provide a solver based on the NLEIGS method [Güttel et al., 2014], consisting in a (rational) Krylov method operating on the linearization pencil resulting from the rational interpolation. We provide a memory-efficient implementation in the sense that a compact (tensor-product) representation of the Krylov basis is employed throughout the computation (we call it the TOAR representation by analogy with the polynomial eigensolvers of [Campos and Roman, 2016]). In this way, our solver fits the framework of the CORK method [van Beeumen et al., 2015]. Furthermore, we have enriched the solver with other features such as automatic determination of singularities and a two-sided variant to compute left eigenvectors. All the details are described next.

4.5.1 Overview of NLEIGS

Following the notation of [Güttel et al., 2014], the goal is to find eigenvalues located in a compact target set $\Sigma \subset \Omega$, in which $T(\cdot)$ is analytic, by approximating T with a rational matrix-valued function whose poles are selected from the set of singularities of T , denoted by Ξ . Considering the degree-graded rational Newton basis functions defined by the recursion

$$b_0(\lambda) = 1, \quad b_j(\lambda) = \frac{\lambda - \sigma_{j-1}}{\beta_j(1 - \lambda/\xi_j)} b_{j-1}(\lambda), \quad j = 1, 2, \dots \quad (39)$$

with nonzero poles at $\xi_j \in \Xi$, the rational matrix

$$R_d(\lambda) := \sum_{j=0}^d b_j(\lambda) D_j \quad (40)$$

is built in such a way that it interpolates T at the nodes $\sigma_j \in \partial\Sigma$ (the boundary of Σ). The scaling factors β_j in (39) are chosen so that $\max_{\lambda \in \partial\Sigma} |b_j(\lambda)| = 1$. If all interpolation nodes σ_j are pairwise distinct, the coefficient matrices D_j of (40) can be computed easily, from the interpolation conditions $R_j(\sigma_j) = T(\sigma_j)$, by means of the recurrence

$$D_0 = \beta_0 T(\sigma_0), \quad D_j = \frac{T(\sigma_j) - R_{j-1}(\sigma_j)}{b_j(\sigma_j)}, \quad j = 1, 2, \dots \quad (41)$$

These matrices are referred to as rational divided difference matrices. In case T is expressed in the split form (3), they can be written as

$$D_j = \sum_{i=0}^{\ell} d_i^j A_i, \quad j \geq 0, \quad (42)$$

where d_i^j denotes the j th rational divided difference corresponding to the scalar function f_i . As shown in [Güttel et al., 2014], these scalar divided differences can be obtained in a numerically stable way using matrix functions, through the expression

$$d_i^j = \beta_0 e_j^T f_i(H_{d+1} K_{d+1}^{-1}) e_1, \quad (43)$$

where H_{d+1} and K_{d+1} are bidiagonal matrices given by

$$H_{d+1} := \begin{bmatrix} \sigma_0 & & & & \\ \beta_1 & \sigma_1 & & & \\ & \ddots & \ddots & & \\ & & & \beta_d & \sigma_d \end{bmatrix}, \quad K_{d+1} := \begin{bmatrix} 1 & & & & \\ \beta_1/\xi_1 & 1 & & & \\ & \ddots & \ddots & & \\ & & & \beta_d/\xi_d & 1 \end{bmatrix}. \quad (44)$$

In SLEPc's implementation of NLEIGS, the D_j matrices are computed via (41) if the NEP is represented with callback functions, and using the expression (43) if the split form (3) is being used. The latter case relies on the matrix function evaluation capabilities of FN, as discussed in §3.2.2. Furthermore, in the split case the D_j matrices (42) are not computed explicitly, and instead the solver works with them implicitly, that is, operating directly with the A_i matrices that appear in the definition of operator T .

The degree d of the interpolant (40) is determined at run time, by evaluating the norms of the rational divided difference matrices as they are generated, until the relation $\|D_d\|/\|D_0\| < \mathbf{tol}$ holds for a given tolerance \mathbf{tol} . The user can provide this tolerance and the maximum allowed degree with `NEPNLEIGSSetInterpolation`. In problems with split representation, for which the D_j matrices are not available explicitly, we use the estimates of their norms proposed in [Güttel et al., 2014], that make use of the divided differences of the f_i functions: $\delta^d := \max\{|d_i^d| : i = 0, \dots, \ell\}$.

The interpolation nodes and poles that determine the approximation $R_d(\lambda)$ of (40) are obtained as a sequence of Leja–Bagby points for (Σ, Ξ) [Güttel et al., 2014]: starting from an arbitrary $\sigma_0 \in \partial\Sigma$, select nodes $\sigma_j \in \partial\Sigma$ and poles $\xi_j \in \Xi$ recursively such that they satisfy the conditions

$$\max_{z \in \partial\Sigma} |s_j(z)| = |s_j(\sigma_{j+1})| \quad \text{and} \quad \min_{z \in \Xi} |s_j(z)| = |s_j(\xi_{j+1})|, \quad (45)$$

for

$$s_j(\lambda) := \frac{\prod_{k=0}^j (\lambda - \sigma_k)}{\prod_{k=1}^j (1 - \lambda/\xi_k)}, \quad j = 0, 1, \dots \quad (46)$$

The fact that the interpolation points σ_j lie at the boundary $\partial\Sigma$ is a direct consequence of the maximum module theorem. In our solver, the target set Σ is specified by the user via a region object (`RG`, §3.2.4), and hence its boundary $\partial\Sigma$ is discretized automatically using the functionality offered by `RG`. In contrast, defining the singularity set Ξ is more subtle and we offer several possibilities:

- The user may provide a callback function (with `NEPNLEIGSSetSingularitiesFunction`) that returns a sequence of points as a discretization of Ξ . This function may return less points than requested by the method, in which case the remaining values are set to infinity. In the extreme case that all $\xi_j = \infty$, then the basis functions of (39) become polynomials and a Newton polynomial interpolation at nodes σ_j is effectively used.
- If the callback function is not provided, but the NEP is of type `rational` (§3.2.4) and has been defined in the split form (3) using only rational functions or a simple combination (sum or product) of quotients of polynomials, then the solver is able to automatically determine the poles. The poles of a quotient of polynomials are obtained by computing the roots of the denominator via the eigenvalues of the associated companion matrix.
- Otherwise, a discretization of Ξ is computed via the AAA algorithm [Nakatsukasa et al., 2018]. In particular, we follow the procedure developed by Elsworth and Güttel [2019]. Note that the AAA algorithm is used only to obtain the ξ_j points, not to approximate the NEP.

It is worth mentioning the special case when $T(\cdot)$ is a rational matrix-valued function of type (p, q) . Then the interpolant R_d of (40) with $d = \max\{p, q\}$ will be exact for any choice of the sampling points σ_j . As a consequence, our implementation will be able to solve rational eigenvalue problems (REP) via linearization in an alternative way compared to the method of Su and Bai [2011], with the advantage that the REP does not need to be expressed in any specific form.

Once the rational approximation (40) has been built, a linearization is carried out resulting in a linear eigenvalue problem

$$\mathcal{A}y = \lambda\mathcal{B}y, \quad (47)$$

whose eigenvalues λ are the same as for the rational eigenproblem $R_d(\lambda)x = 0$ and whose eigenvectors have the form

$$y = \begin{bmatrix} b_0(\lambda)x \\ \vdots \\ b_{d-1}(\lambda)x \end{bmatrix}. \quad (48)$$

For the matrices \mathcal{A} and \mathcal{B} of the linearization (47), we use the simplification proposed in [Güttel et al., 2014] of setting the last pole $\xi_d = \infty$, resulting in

$$\mathcal{A} = \begin{bmatrix} D_0 & D_1 & \dots & D_{d-2} & (D_{d-1} - \frac{\sigma_{d-1}}{\beta_d} D_d) \\ \sigma_0 I & \beta_1 I & & & \\ & \ddots & \ddots & & \\ & & \ddots & \beta_{d-2} I & \\ & & & \sigma_{d-2} I & \beta_{d-1} I \end{bmatrix}, \quad \mathcal{B} = \begin{bmatrix} 0 & 0 & \dots & 0 & -\frac{D_d}{\beta_d} \\ I & \frac{\beta_1}{\xi_1} I & & & \\ & \ddots & \ddots & & \\ & & \ddots & \frac{\beta_{d-2}}{\xi_{d-2}} I & \\ & & & I & \frac{\beta_{d-1}}{\xi_{d-1}} I \end{bmatrix}. \quad (49)$$

To solve the linear eigenproblem (47), Güttel et al. [2014] propose two main variants of the rational Krylov method [Ruhe, 1984]. In the dynamic variant, the approximation and the linearization are built incrementally as the Krylov subspace grows, with rational Krylov shifts equal to the interpolation nodes σ_j , while in the static variant the linearization is created a priori and then solved. In SLEPc we have implemented the static variant only, where the user can optionally provide the list of rational Krylov shifts (`NEPNLEIGSSetRKShifts`), otherwise only one shift is used (the target), which is equivalent to the shift-and-invert Krylov–Schur method [Stewart, 2001]. In our experience, the rational Krylov method tends to be less efficient for large problems, since it requires several matrix factorizations. We provide details of the Krylov–Schur case in the next subsections.

4.5.2 Expansion of the Krylov basis

In our NLEIGS solver, the linearization matrices are never built explicitly, and instead the Krylov–Schur method proceeds by exploiting their block structure. To generate a new Krylov vector, we need to multiply the last vector of the basis with matrix

$$\mathcal{S} = (\mathcal{A} - \sigma\mathcal{B})^{-1}\mathcal{B}, \quad (50)$$

where \mathcal{A} and \mathcal{B} are given in (49), and $\sigma \in \Sigma$ is the shift or target value (a value around which the eigenvalues are sought). To derive the recurrences that will apply matrix \mathcal{S} implicitly, we consider the factorization $(\mathcal{A} - \sigma\mathcal{B})\Pi = U_\sigma L_\sigma$, where $\Pi = \begin{bmatrix} 0 & I_{(d-1)n} \\ I_n & 0 \end{bmatrix}$ is a permutation matrix, and

$$L_\sigma = \begin{bmatrix} \frac{1}{b_{d-1}(\sigma)}I & & & & \\ -\frac{b_0(\sigma)}{b_{d-1}(\sigma)}I & I & & & \\ \vdots & & \ddots & & \\ -\frac{b_{d-2}(\sigma)}{b_{d-1}(\sigma)}I & & & & I \end{bmatrix}, \quad U_\sigma = \begin{bmatrix} R_d(\sigma) & D_0 & D_1 & \dots & D_{d-2} \\ & (\sigma_0 - \sigma)I & \beta_1(1 - \frac{\sigma}{\xi_1})I & & \\ & & (\sigma_1 - \sigma)I & \ddots & \\ & & & \ddots & \beta_{d-2}(1 - \frac{\sigma}{\xi_{d-2}})I \\ & & & & (\sigma_{d-2} - \sigma)I \end{bmatrix}. \quad (51)$$

Using this factorization, the product $w = \mathcal{S}x = \Pi L_\sigma^{-1} U_\sigma^{-1} \mathcal{B}x$ is carried out by first solving the block upper triangular system $U_\sigma y = \mathcal{B}x$, by means of the recurrence

$$\begin{aligned} y^{d-1} &= \frac{1}{\sigma_{d-2} - \sigma} x^{d-2} + \frac{\beta_{d-1}}{(\sigma_{d-2} - \sigma)\xi_{d-1}} x^{d-1}, \\ y^j &= \frac{1}{\sigma_{j-1} - \sigma} x^{j-1} + \frac{\beta_j}{(\sigma_{j-1} - \sigma)\xi_j} x^j - \frac{\beta_j}{\sigma_{j-1} - \sigma} \left(1 - \frac{\sigma}{\xi_j}\right) y^{j+1}, \quad j = d-2, \dots, 1, \\ y^0 &= R_d(\sigma)^{-1} \left(-D_0 y^1 - D_1 y^2 - \dots - D_{d-2} y^{d-1} - \frac{1}{\beta_d} D_d x^{d-1} \right). \end{aligned} \quad (52)$$

In the above expressions, x^i denotes the i th block of vector $x = \text{vec}(x^0, \dots, x^{d-1})$. Secondly, the block lower triangular system $L_\sigma \tilde{w} = y$ is solved with

$$\begin{aligned} \tilde{w}^0 &= b_{d-1}(\sigma) y^0, \\ \tilde{w}^j &= y^j + b_{j-1}(\sigma) y^0, \quad j = 1, \dots, d-1. \end{aligned} \quad (53)$$

Finally, the solution vector is obtained by applying the permutation $w = \Pi \tilde{w}$.

The operations in (52)–(53) involve vector *axy* operations, matrix-vector products with the divided difference matrices, and the construction of $R_d(\sigma)$ and its inverse. Of course, the inverse is not built explicitly, but instead a (sparse) linear solve is done (using a `KSP` linear solver

object) which usually implies a factorization. It is noteworthy that, in our implementation, when the NEP has been defined in the split form (3), the last expression of (52) operates with the divided difference matrices D_j without forming them explicitly. Since the D_j consist in linear combinations of the problem matrices A_i , those operations are rearranged in a way that they are expressed in terms of the A_i matrices only.

4.5.3 Krylov basis representation

We provide two different implementations of NLEIGS, one that operates with explicitly stored Krylov vectors (full basis) and another one with a compact representation (TOAR basis). The advantage of the full basis is that it is possible to use any of the EPS eigensolvers (not only Krylov–Schur) to solve the linear eigenvalue problem (47), whereas the TOAR basis allows a significant reduction of both computational and storage cost. By default, the compact representation is used, and the other one can be selected with the user option `NEPNLEIGSSetFullBasis`.

We now describe the specialized version of Krylov–Schur that handles the compact representation of the Krylov basis, in a similar way as the TOAR solver for polynomial eigenvalue problems in PEP [Campos and Roman, 2016]. The idea is to exploit the block structure of the matrices \mathcal{A} and \mathcal{B} of the linearization, and derive dependency relations among the d blocks of the generated Krylov vectors, whose global dimension is dn . In this way, it is possible to define a basis U_{k+d} of vectors of length n , from which all blocks of the Krylov basis V_k can be reconstructed, resulting in the relation

$$V_k = (I_d \otimes U_{k+d})G_k, \quad (54)$$

for some matrix of coefficients G_k .

The Krylov–Schur method applied to the linear eigenproblem (47) produces an Arnoldi relation of order k for matrix \mathcal{S} (50), that can be written as

$$\mathcal{B}V_k = (\mathcal{A} - \sigma\mathcal{B})(V_k H_k + h_{k+1,k} v_{k+1} e_k^*). \quad (55)$$

Considering that the Krylov basis V_k is divided in d blocks of n consecutive rows, $\{V_k^i\}_{i=0}^{d-1}$, and equating the last $d-1$ block-rows of (55), we get $d-1$ relations that reveal the linear dependency between the vectors of the set formed by the columns of blocks $\{V_k^i\}_{i=0}^{d-1}$ and the vectors $\{v_{k+1}^i\}_{i=0}^{d-1}$. This implies the equivalence of subspaces

$$\text{span}(\cup_{i=0}^{d-1} \{V_k^i e_j\}_{j=1}^k \cup \{v_{k+1}^i\}_{i=0}^{d-1}) = \text{span}(\{V_k^0 e_j\}_{j=1}^k \cup \{v_{k+1}^i\}_{i=0}^{d-1}), \quad (56)$$

meaning that each of the d parts of the $k+1$ Arnoldi vectors stored in the columns of V_{k+1} can be obtained from linear combinations of $k+d$ vectors, that is, V_k can be expressed as

$$[V_k^i \quad v_{k+1}^i] = U_{k+d} [G_k^i \quad g_{k+1}^i] \quad i = 0, \dots, d-1, \quad (57)$$

where $U_{k+d} \in \mathbb{C}^{n \times (k+d)}$, $\{G_k^i\}_{i=0}^{d-1} \subset \mathbb{C}^{(k+d) \times k}$ and $\{g_{k+1}^i\}_{i=0}^{d-1} \subset \mathbb{C}^{k+d}$. Both U_{k+d} and G_k must have orthonormal columns by construction. This is the TOAR representation of (54), that has been described in [Campos and Roman, 2016] in the context of polynomial eigenvalue problems.

The recurrences (52)–(53) that generate a new Arnoldi vector must be adapted to operate with vectors in the TOAR representation. Assuming that the current Arnoldi basis is (57), the function

$$[u_{k+d+1}, g] = \text{expand}(\{D_i\}_{i=0}^d, U_{k+d}, g_{k+1}), \quad (58)$$

generates a vector u_{k+d+1} to extend the TOAR basis (if necessary), and the coefficients g used to represent the expansion vector $w = \mathcal{S}v_{k+1}$ in terms of such basis by the following steps. Firstly,

ALGORITHM 4: Arnoldi for NLEIGS with TOAR representation

Input: Divided difference matrices $\{D_i\}_{i=0}^d \subset \mathbb{C}^{n \times n}$, vectors $\{w^i\}_{i=0}^{d-1} \subset \mathbb{C}^n$ (d parts of initial vector), number of iterations $k \in \mathbb{N}$

Output: $H_k \in \mathbb{C}^{k \times k}$, $h_{k+1,k} \in \mathbb{R}$, $U_{k+d} \in \mathbb{C}^{n \times (k+d)}$, $\{G_k^i\}_{i=0}^d \subset \mathbb{C}^{(k+d) \times k}$, $\{g_{k+1}^i\}_{i=0}^{d-1} \subset \mathbb{C}^{k+d}$ satisfying the Arnoldi relation in TOAR representation

```

1  $G_0 \leftarrow [ \quad ], H_0 \leftarrow [ \quad ]$ 
2  $[Q, R] = \text{qr}([w^0, \dots, w^{d-1}], 0)$ 
3  $U_d \leftarrow Q$ 
4  $g_1 \leftarrow \text{vec } R / \|\text{vec } R\|$ 
5 for  $j = 1, \dots, k$  do
6    $[u_{j+d}, g] = \text{expand}(\{D_i\}_{i=0}^d, U_{j-1+d}, g_j)$ 
7    $U_{j+d} \leftarrow [U_{j-1+d} \quad u_{j+d}]$ 
8    $G_j^i \leftarrow \begin{bmatrix} G_{j-1}^i & g_j^i \\ 0 & 0 \end{bmatrix}, i = 0, \dots, d-1$ 
9    $h_j = G_j^* g, \hat{g} \leftarrow g - G_j h_j$ 
10   $h_{j+1,j} = \|\hat{g}\|, g_{j+1} \leftarrow \hat{g} / h_{j+1,j}$ 
11   $H_j \leftarrow \begin{bmatrix} H_{j-1} & \\ & h_j \end{bmatrix}$ 
12 end

```

we compute the TOAR coefficients of the y^j vectors in (52),

$$\begin{cases} \tilde{g}^{d-1} = \frac{1}{\sigma_{d-2} - \sigma} g_{k+1}^{d-2} + \frac{\beta_{d-1}}{(\sigma_{d-2} - \sigma)\xi_{d-1}} g_{k+1}^{d-1}, \\ \tilde{g}^j = \frac{1}{\sigma_{j-1} - \sigma} g_{k+1}^{j-1} + \frac{\beta_j}{(\sigma_{j-1} - \sigma)\xi_j} g_{k+1}^j - \frac{\beta_j}{\sigma_{j-1} - \sigma} \left(1 - \frac{\sigma}{\xi_j}\right) \tilde{g}^{j+1}, \quad j = d-2, \dots, 1. \end{cases} \quad (59)$$

To compute y^0 , it is necessary to reconstruct the preceding vectors using the expressions $y^i = U_{k+d} \tilde{g}^i$, for $i = 1, \dots, d-1$, and the last block of the input vector as $v_{k+1}^{d-1} = U_{k+d} g_{k+1}^{d-1}$. Then

$$y^0 = R_d(\sigma)^{-1} \left(-D_0 y^1 - D_1 y^2 - \dots - D_{d-2} y^{d-1} - \frac{1}{\beta_d} D_d v_{k+1}^{d-1} \right), \quad (60)$$

which is used to generate vector u_{k+d+1} that extends the orthogonal basis U_{k+d} (unless it is linearly dependent). Then, from the orthogonalization coefficients, we obtain the coordinates \tilde{g}^0 that express y^0 with respect to the extended basis,

$$y^0 = U_{k+d+1} \tilde{g}^0. \quad (61)$$

Finally, the coefficients g^i of (58) can be obtained from the coefficients \tilde{g} using (53) and taking into account the permutation Π , as

$$\begin{cases} g^j = \tilde{g}^{j+1} + b_j(\sigma) \tilde{g}^0, \quad j = 0, \dots, d-2, \\ g^{d-1} = b_{d-1}(\sigma) \tilde{g}^0. \end{cases} \quad (62)$$

Algorithm 4 presents the Arnoldi method that is used by default within the NLEIGS solver, where step 6 performs the expansion of the Krylov subspace as explained above. The remaining part (Gram-Schmidt orthogonalization in lines 8–9 and normalization in line 10) are done in exactly the same way as in the polynomial eigensolver, as detailed in [Campos and Roman, 2016]. On the other hand, we have implemented the Krylov–Schur restart on top of Algorithm 4, including locking of converged eigenvalues. To keep the presentation short, we again direct interested readers to [Campos and Roman, 2016] where relevant details can be found.

4.5.4 Two-sided variant

The NLEIGS solver in SLEPc provides a two-sided variant that allows computing left eigenvectors when enabled as explained in §3.2.5. For this, we implement a method that is similar to the two-sided Krylov–Schur of Zwaan and Hochstenbach [2017].

The two-sided variant works with two Krylov bases, to approximate right and left eigenvectors, respectively. The left basis is built from matrix-vector products with the conjugate transpose, \mathcal{S}^* . Hence, we need to complement (52)–(53) with analogue recurrences that perform this operation by blocks. To obtain $w = \mathcal{S}^*x = \mathcal{B}^*U_\sigma^{-*}L_\sigma^{-*}\Pi^*x$ we proceed as follows. Let $y = L_\sigma^{-*}\Pi^*x$, then $L_\sigma^*y = \Pi^*x$ and we have

$$\begin{cases} y^i = x^{i-1}, & i = 1, \dots, d-1, \\ y^0 = \bar{b}_{d-1}(\sigma)x^{d-1} + \sum_{i=1}^{d-1} \bar{b}_{i-1}(\sigma)y^i = \bar{b}_{d-1}(\sigma)x^{d-1} + \sum_{i=0}^{d-2} \bar{b}_i(\sigma)x^i, \end{cases} \quad (63)$$

where the bar denotes complex conjugation. For $z = U_\sigma^{-*}y$ the following relations hold,

$$\begin{aligned} R_d(\sigma)^*z^0 &= y^0, \\ D_0^*z^0 + (\bar{\sigma}_0 - \bar{\sigma})z^1 &= y^1, \\ D_1^*z^0 + \beta_1 \left(1 - \frac{\bar{\sigma}}{\bar{\xi}_1}\right) z^1 + (\bar{\sigma}_1 - \bar{\sigma})z^2 &= y^2, \\ &\dots \\ D_{d-2}^*z^0 + \beta_{d-2} \left(1 - \frac{\bar{\sigma}}{\bar{\xi}_{d-2}}\right) z^{d-2} + (\bar{\sigma}_{d-2} - \bar{\sigma})z^{d-1} &= y^{d-1}, \end{aligned} \quad (64)$$

from which we obtain the recurrence

$$\begin{cases} z^0 = R_d(\sigma)^{-*}y^0, \\ z^1 = \frac{1}{\bar{\sigma}_0 - \bar{\sigma}} (y^1 - D_0^*z^0), \\ z^i = \frac{1}{\bar{\sigma}_{i-1} - \bar{\sigma}} \left(y^i - D_{i-1}^*z^0 - \beta_{i-1} \left(1 - \frac{\bar{\sigma}}{\bar{\xi}_{i-1}}\right) z^{i-1} \right), & i = 2, \dots, d-1. \end{cases} \quad (65)$$

Finally, $w = \mathcal{B}^*z$ is computed with

$$\begin{cases} w^0 = z^1, \\ w^i = \frac{\beta_i}{\bar{\xi}_i} z^i + z^{i+1}, & i = 1, \dots, d-2, \\ w^{d-1} = -\frac{1}{\beta_d} D_d^*z^0 + \frac{\beta_{d-1}}{\bar{\xi}_{d-1}} z^{d-1}. \end{cases} \quad (66)$$

This computation can be carried out efficiently, the most computationally expensive step being the application of $R_d(\sigma)^{-*}$, which implies a linear solve with the same KSP object as in (52) (the factorization is reused).

Unfortunately, the above recurrences cannot be adapted to the TOAR representation, as was done for the right Krylov basis in (59)–(62). The reason is that the linearization (47) allows the compact representation of right eigenvectors only. There exist alternative linearizations that are appropriate for representing both right and left eigenvectors in a compact way [Lietaert et al., 2018], and this will be a topic of future research. In our current implementation, we restrict the two-sided variant to the case of using the full basis setting (`NEPNLEIGSSetFullBasis`, see §4.5.3).

When this option is enabled, the solver is more expensive in terms of memory and computation, but it allows to generate the left basis required in the two-sided variant. In this case, there is no need to implement a specialized Arnoldi (Algorithm 4), so we solve the linear eigenproblem using an EPS object (which implements two-sided Krylov–Schur).

5 Parallel Implementation

The model of parallelism of PETSc and SLEPc has been briefly introduced in §3.1. We now provide a few details about how the algorithms described in preceding sections have been implemented in parallel. The way to proceed is very similar to other SLEPc solvers.

The parallelization is based on distributing large-scale data objects such as matrices and vectors across the available MPI processes, whereas small-scale objects such as the Hessenberg matrix H_j in Algorithm 4 are normally stored redundantly (all processes hold a copy).

Matrices in PETSc are usually distributed by blocks of rows, that is, every process owns a contiguous range of rows. Vectors are distributed in the same way, with global and local dimensions being consistent for all objects involved in an algorithm. For example, to evaluate the residual as in step 7 of Algorithm 2, $r = T(\lambda)x$, in the case that T is given in the split form (3), the computations will be done as $r = \sum_{i=1}^{\ell} f_i(\lambda)A_i x$, which implies ℓ matrix-vector products and $\ell - 1$ vector *axy* operations, all of them performed in parallel provided that all matrices and vectors have a compatible distribution.

A key operation in SLP (Algorithm 1) is the explicit evaluation of the function and its derivative, $T(\lambda)$ and $T'(\lambda)$, for some scalar value λ . When working in split form, both operations are very similar and consist in $\ell - 1$ matrix *axy* operations, $A = A + \alpha B$. This operation, which is implemented in PETSc, is trivially parallelizable, and its main complication is the correct management of the sparsity pattern of the involved matrices (in the `NEPSetSplitOperator` function mentioned in §3.2.2 the user indicates whether the sparsity pattern of all A_i matrices is equal or not). On the other hand, if T has been specified via the callback mechanism of §3.2.1, then it is the user’s responsibility to provide functions that compute $T(\lambda)$ and $T'(\lambda)$ in parallel.

The most costly step in SLP is the computation of the correction (step 4 of Algorithm 1) via the linear eigenproblem $A^{-1}Bx = \mu^{-1}x$. This is done with an EPS solver, whose parallelization involves various operations that appear in NEP as well and we discuss below: matrix-vector product (including linear solve), vector orthogonalization, and projected problem.

Parallel matrix-vector products involve neighbour-wise communication, and hence they should have good scalability (assuming that the problem matrices A_i have an appropriate sparse pattern, e.g., coming from a mesh-based discretization). On the other hand, orthogonalization of vectors (e.g., step 7 of Algorithm 3) and other operations such as inner products and norms require global communication, which may limit the scalability to large number of processes. In SLEPc, we strive to minimize the impact of global communication [Hernandez et al., 2007].

Often it is necessary to solve linear systems of equations at each eigensolver iteration, for instance in the linear eigensolve within SLP, or in RII (step 10 of Algorithm 2), N-Arnoldi and NLEIGS. As pointed out previously, this is handled via a PETSc’s KSP linear solver. In terms of parallelization, we repeat here that parallel factorization methods are obtained by means of external packages such as MUMPS. Note that direct linear solvers usually offer better robustness, compared to iterative ones, but may not scale well to many MPI processes.

We now discuss the solution of the projected problem, that is, the part of the algorithm that operates with small-scale matrices such as step 3 of Algorithm 3. As mentioned before, these small matrices are stored redundantly and the associated computation is also redundant (all processes perform the computation sequentially and independently of each other). Since this part of the algorithm is sequential, its cost must not be too high, otherwise parallel efficiency will

deteriorate. In normal usage, this cost is negligible because the size of the projected problem is very small, but for the case of large number of requested eigenvalues it could become larger, so we include a user-defined parameter in `NEPSetDimensions` to keep the dimension of the projected problem bounded.

In the implementation of NLEIGS, all relevant operations have been mentioned already. The only additional comment is that the coefficient matrices G_j of the compact basis are stored redundantly, so the second-level orthogonalization (steps 9-10 of Algorithm 4) is also a redundant (sequential) operation.

We conclude this section stating that the most difficult operations in terms of parallel implementation are those associated with the deflation technique described in §4.3. We have opted for increasing the memory usage to store auxiliary quantities whenever this leads to a reduction of the required MPI communication in the main operations (31), (33) and (34). For example, in the split variant, we maintain data structures for storing $A_i X$, or $X^* X$, that are updated every time an eigenpair has converged. In this way, the matrix-vector product (31) only needs to perform communication among the processes in the operations $T(\lambda)z_1$ and $X^* z_1$ when computing y_1 and y_2 , respectively (recall that z_2 is stored redundantly in all processes). Similarly, we compute $T(\sigma)^{-1} X$ every time X is extended (or σ updated) to reduce the amount of communication in linear solves with the extended operator \tilde{T} .

6 Performance Evaluation

We now present results of some computational experiments to assess the performance of the solvers in terms of accuracy, convergence, efficiency, and scalability. The executions are carried out on the Tirant III computer, which consists of 336 computing nodes, each of them with two Intel Xeon SandyBridge E5-2670 processors (16 cores each) running at 2.6 GHz with 32 GB of memory, linked with an Infiniband network. We allocated 4 MPI processes per node at most. The presented results correspond to SLEPc version 3.13, together with PETSc 3.13 and MUMPS 5.2.1. All software has been compiled with Intel C and Fortran compilers (version 18) and Intel MPI.

Table 1 lists the problems used in the experiments, summarizing some properties and parameters. Here is a short description of the problems:

- The `delay` problem arises from the discretization of a parabolic partial differential equation with time delay τ [Jarlebring, 2008], with $T(\lambda) = -\lambda I + A + e^{-\tau\lambda} B$. The nonlinear character comes from the exponential, so in this case there are no singularities ($\Xi = \{\infty\}$). For the time delay parameter we have used the value $\tau = 0.001$.
- The `loaded_string` problem is a rational eigenproblem from a discretization of a boundary value problem describing the vibration of a string with a load of mass m attached by an elastic spring of stiffness κ [Solov'ev, 2006]. Here, the rational matrix $T(\lambda) = A - \lambda B + \frac{\lambda}{\lambda - \sigma} C$ has a single pole $\sigma = \kappa/m$, i.e., $\Xi = \{\kappa/m\}$. In the results below we take $m = 1$ and $\kappa = 1$.
- The `gun` problem models a radio-frequency gun cavity [Liao et al., 2010]. Here, $T(\lambda) = K - \lambda M + i\sqrt{\lambda - \kappa_1^2} W_1 + i\sqrt{\lambda - \kappa_2^2} W_2$, where we have used $\kappa_1 = 0$ and $\kappa_2 = 108.8774$. The nonlinearity comes from the square roots, and in this case $\Xi = (-\infty, \kappa_2^2)$. Figure 4 shows a pictorial representation of this problem.
- The `dimer` problem is taken from [Araujo et al., 2020] and corresponds to the analysis of scattering resonances (TM polarization) of a dimer nano-structure (two disks coated with gold). The differential equation is expressed as $\Delta u + \varepsilon(x, \omega)\omega^2 u = 0$, where ω is the eigenvalue. A Drude-Lorentz model is used for the relative permittivity $\varepsilon(x, \omega)$, which

Table 1: Description of the test problems used in the computational experiments. The problem dimension is given, as well as the region where eigenvalues are sought (it is used only by the Interpol and NLEIGS solvers). In the tests, nev eigenvalues are requested around the value σ with a tolerance tol .

name	dim	region	nev	σ	tol
delay 4M	2^{22}	$[-100, 50]$	10	0	1×10^{-8}
delay 100K	100000	$[-100, 50]$	5	1	1×10^{-6}
loaded_string 4M	2^{22}	$[4, 800]$	9	10	1×10^{-13}
loaded_string 200K	200000	$[4, 800]$	9	10	1×10^{-8}
gun	9956	see Fig. 4	10	$65000 + 500i$	1×10^{-8}
dimer	173725	$[-1, 20] \times [-2, 0]$	20	$5.3 - 0.25i$	1×10^{-8}




Figure 4: Graphical representation of the **gun** problem. The shaded region is the target set Σ . The nodes (blue points) and poles (orange points) used in the NLEIGS solver are Leja-Bagby pairs, picked from the discretization of the boundary $\delta\Sigma$ and the singularity set Ξ , respectively. Red crosses are the eigenvalues lying inside the region.

approximates it as a rational function. Hence, the nonlinear eigenproblem to be solved is rational, with $T(\omega) = A_0 - \omega^2 A_1 - \omega^2 \varepsilon(x, \omega) A_2$. In the case of gold, the rational representation of ε has 12 poles, which are depicted in Figure 5.

Both **delay** and **loaded_string** have real eigenvalues, so the computation has been performed in real arithmetic. In the case of **gun** and **dimer** the problem matrices are complex, as well as the eigenvalues, and hence the solvers have been configured to use complex arithmetic.

All problems have been represented in the split form (3). The accuracy of the computed eigenpairs is assessed with the scaled residual $\eta(x, \lambda)$, (9), using ∞ -norms for practical computation of matrix norms. Table 2 shows the maximum value of the scaled residual obtained by the different solvers on the problems of Table 1. A few comments are required to understand the results:

- The large problem size 4M for **delay** and **loaded_string** has been used for the scalability tests



Figure 5: Graphical representation of the **dimer** problem. The orange points are the 12 poles of the rational function. The shaded part is the region used in the NLEIGS solver, while the blue circle is the target value used in NLEIGS as well as the other solvers. Red crosses are the 20 eigenvalues computed by NLEIGS (there are more eigenvalues inside the region, and other solvers may not return exactly the same ones).

Table 2: Results for the different solvers corresponding to executions using 16 MPI processes. Apart from the method name, we show the number of iterations, the number of required linear solves, the time required to build the preconditioner, the total time, and the scaled residual error.

test problem	method	iter	lin. solves	PC build	time	$\eta(x, \lambda)$
delay 100K	SLP	2	16	0.7	3.0	5×10^{-9}
	RII	1	15	0.7	2.9	4×10^{-9}
	N-Arnoldi	1	15	0.7	2.7	1×10^{-8}
	Interpol	1	20	0.7	3.6	4×10^{-11}
	NLEIGS	1	20	0.7	3.6	6×10^{-14}
loaded_string 200K	SLP	6	96	4.3	31.7	4×10^{-10}
	RII	1	45	1.4	14.5	2×10^{-7}
	N-Arnoldi	1	45	1.4	14.6	5×10^{-6}
	Interpol	10	96	1.4	31	2×10^{-9}
	NLEIGS	3	41	1.4	13.4	2×10^{-10}
gun	SLP	7	121	1.5	2.1	5×10^{-7}
	RII	40	88	4.7	5.7	5×10^{-6}
	N-Arnoldi	42	123	7.0	3.0	2×10^{-6}
	NLEIGS	3	39	0.4	0.6	1×10^{-11}
dimer	SLP	18	401	21.6	55.5	1×10^{-9}
	RII	89	249	49	329	1×10^{-8}
	N-Arnoldi	115	571	90	249	9×10^{-9}
	Interpol	3	75	2.0	8.9	2×10^{-5}
	NLEIGS	2	58	2.0	5.4	2×10^{-14}

(see below). In these cases, the computation uses an iterative linear solver, in particular, Bi-CGStab with AMG preconditioning (Hypre). The rest of executions use an LU factorization computed with MUMPS.

- In the gun problem, the RII solver is run with deflation threshold activated, cf. §4.3.
- In RII, we use the Hermitian variant of the scalar equation (19) (`NEPRIISetHermitian`) whenever the problem is Hermitian, i.e., in `loaded_string` and `delay`.
- The Newton-type solvers (SLP, RII, N-Arnoldi) compute *nev* eigenvalues at most, while Krylov-type solvers (Interpol, NLEIGS) may return more than *nev* eigenvalues in some cases. In particular, for the `delay 100K` problem Interpol and NLEIGS compute 6 eigenvalues, and for the `gun` problem NLEIGS computes 13.
- In Krylov-type solvers (Interpol, NLEIGS), the `ncv` parameter (maximum dimension of the subspace, `NEPSetDimensions`) is set to the default $\max\{2 \cdot nev, nev + 15\}$, and the number of iterations listed in Table 2 refers to the number of *restarts*. For instance, in the `delay` problem a single restart cycle was enough, so the solver built 20 Krylov vectors (see the column ‘lin. solves’ in the table).

From the table, we can see that Newton-based methods (SLP, RII, N-Arnoldi) are always slower than the polynomial or rational interpolation methods. In the case of the largest problem size (4M), they do not succeed to give the solution in a reasonable time. Among the classical solvers, we can see that N-Arnoldi is significantly faster than RII in the `gun` and `dimer` problems,

Table 3: Results for different variants of the NLEIGS solver with 16 MPI processes. The information shown is the same as in Table 2.

test problem	NLEIGS variant	iter	lin. solves	PC build	time	$\eta(x, \lambda)$
gun	TOAR basis	3	39	0.4	0.6	1×10^{-11}
	full basis	3	39	0.4	0.6	9×10^{-13}
	two-sided	4	102	0.4	1.1	6×10^{-13}
dimer	TOAR basis	2	58	2.0	5.4	2×10^{-14}
	full basis	2	58	2.0	7.2	2×10^{-14}
	two-sided	2	126	2.0	21	2×10^{-14}

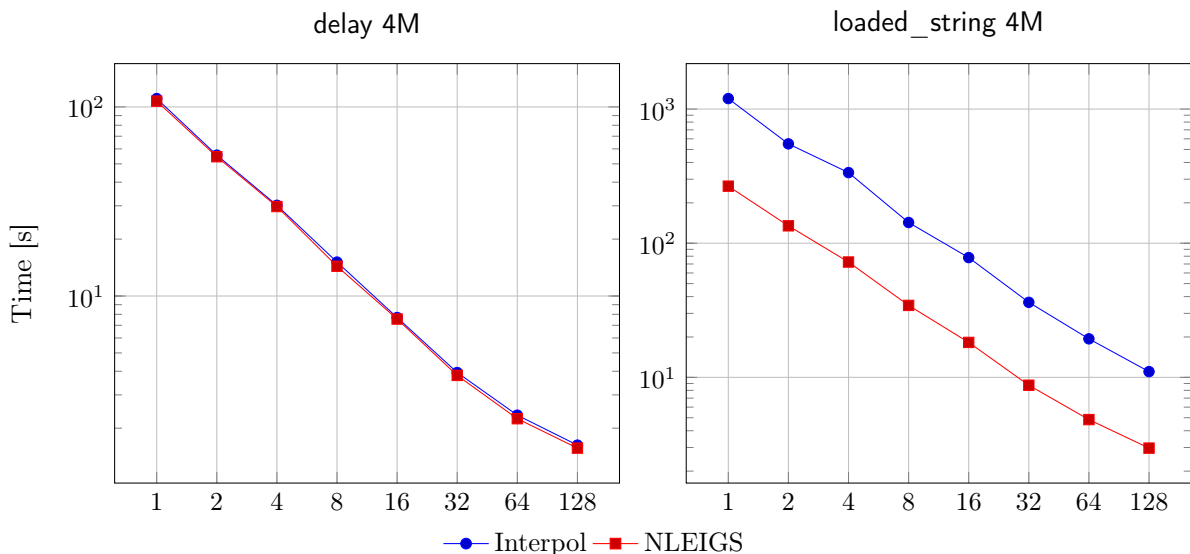


Figure 6: Execution time (in seconds) with up to 128 MPI processes for the problems `delay` (left) and `loaded_string` (right), solved with Interpol and NLEIGS. The execution parameters are shown in Table 1.

which is to be expected from the accelerating effect of the expanding subspace in N-Arnoldi. Still, the SLP method beats both of them in these two problems.

As pointed out in §4.4, the Interpol solver is restricted to real eigenvalues. In the case of `delay` we can hardly see any difference between Interpol and NLEIGS, as expected since NLEIGS with no singularities (poles set to infinity) behaves like a polynomial interpolation method. In contrast, `loaded_string` already displays a significant gain of NLEIGS compared to Interpol.

We have carried out a few additional runs to compare the different NLEIGS variants: TOAR basis (the default), full basis, and two-sided. Table 3 shows the results. In the `gun` problem, the overhead of the full basis version is negligible (even though the number of terms in the linearization is quite large, $d = 15$) because the orthogonalization time is small. In contrast, for the `dimer` problem we see that there is a significant benefit in using the TOAR basis variant instead of working with the full basis. In this case both the number of terms ($d = 12$) and the problem size are large. Regarding the two-sided variant, we observe that in addition to the overhead attributable to using full basis, there is an extra cost needed for the approximation of the left eigenspace. A related comment is that parallel operations involving matrix transposes are always less efficient than the non-transposed counterparts, due to the fact that matrices are stored by rows.

We conclude this section by analyzing the parallel scalability of the two interpolation solvers. Figure 6 shows the execution time of both Interpol and NLEIGS when solving the large problem sizes (4M) with increasing number of MPI processes. The plots show a very good scaling up to 128 processes.

7 Conclusions

This paper collects in one place all the relevant details related to the solvers for the nonlinear eigenvalue problem that we have been developing in SLEPc in the last years. These solvers constitute the NEP module of SLEPc, which provides a flexible user interface to specify the nonlinear problem in different application contexts. We have discussed the details of SLP, RII, N-Arnoldi, Interpol and NLEIGS, including important issues such as deflation or parallelization. The available user options have been mentioned as well, as a way of brief users manual. We have also illustrated the performance of the implemented methods when solving a set of representative problems. The results show that good accuracy can be attained with all the methods, with more or less efficiency, and that very large scale problems can be addressed thanks to the good parallel scalability.

Apart from a mere description of the implementations, this paper includes two new developments: (1) a version of NLEIGS that stores the Krylov basis in compact format (à la TOAR), in the same spirit as previous works such as [Campos and Roman, 2016] for polynomial eigenproblems and [Dopico and González-Pizarro, 2018] for rational eigenproblems; and (2) a two-sided version of NLEIGS that can be used to compute left eigenvectors, which are required, e.g., to approximate the resolvent $T^{-1}(z)$. Our two-sided NLEIGS method operates without explicitly building the linearization matrices, but in contrast it stores the Krylov basis explicitly, as opposed to the variant of [Lietaert et al., 2018]. We have also derived the details for a practical implementation of Effenberger deflation in the context of classical iterations (SLP, RII, and N-Arnoldi).

Research on nonlinear eigenvalue problems is still very active, and we foresee that novel methods will be proposed in the next years. Our intention is to continue adding solvers to the NEP module, so that users have a wide range of choices from which to select the most appropriate one for each particular application. As an example, we recently finished the implementation of the two-sided variant of SLP with non-equivalence deflation, that was included in version 3.13 of SLEPc. Examples of recently proposed methods that could be incorporated in the future are [Gavin et al., 2018], [Mele and Jarlebring, 2018], [Xue, 2018]. On the other hand, the area of nonlinear eigenvector problems (that is, problems that depend nonlinearly on the eigenvector) remains as a topic of future investigation.

During the review process, the preprint [Güttel et al., 2020] appeared, which proposes several improvements over the rational interpolation solvers of §4.5, such as repeating the poles cyclically in Leja-Bagby order instead of adding extra poles at infinity. We will adapt these ideas to our NEP solvers in forthcoming versions of SLEPc.

Acknowledgements

The authors would like to thank the anonymous referees whose constructive comments helped improve the presentation. The computational experiments of section 6 were carried out on the supercomputer Tirant III belonging to Universitat de València.

References

- P. R. Amestoy, I. S. Duff, J.-Y. L'Excellent, and J. Koster. A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J. Matrix Anal. Appl.*, 23(1):15–41, 2001.
- J. C. Araujo, C. Campos, C. Engström, and J. E. Roman. Computation of scattering resonances in absorptive and dispersive media with applications to metal-dielectric nano-structures. *J. Comput. Phys.*, 407:109220, 2020.
- S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. Gropp, D. Karpeyev, D. Kaushik, M. Knepley, D. May, L. Curfman McInnes, R. Mills, T. Munson, K. Rupp, P. Sanan, B. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.13, Argonne National Laboratory, 2020.
- W.-J. Beyn, C. Effenberger, and D. Kressner. Continuation of eigenvalues and invariant pairs for parameterized nonlinear eigenvalue problems. *Numer. Math.*, 119(3):489–516, 2011.
- Y. Brûlé, B. Gralak, and G. Demésy. Calculation and analysis of the complex band structure of dispersive and dissipative two-dimensional photonic crystals. *J. Opt. Soc. Am. B*, 33(4):691–702, 2016.
- C. Campos and J. E. Roman. Parallel Krylov solvers for the polynomial eigenvalue problem in SLEPc. *SIAM J. Sci. Comput.*, 38(5):S385–S411, 2016.
- C. Campos and J. E. Roman. A polynomial Jacobi-Davidson solver with support for non-monomial bases and deflation. *BIT*, 60(2):295–318, 2020.
- G. Demésy, A. Nicolet, B. Gralak, C. Geuzaine, C. Campos, and J. E. Roman. Non-linear eigenvalue problems with GetDP and SLEPc: Eigenmode computations of frequency-dispersive photonic open structures. *Comput. Phys. Commun.*, 257:107509, 2020.
- F. M. Dopico and J. González-Pizarro. A compact rational Krylov method for large-scale rational eigenvalue problems. *Numer. Linear Algebra Appl.*, 26(1):e2214, 2018.
- C. Effenberger. Robust successive computation of eigenpairs for nonlinear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 34(3):1231–1256, 2013.
- C. Effenberger and D. Kressner. Chebyshev interpolation for nonlinear eigenvalue problems. *BIT*, 52(4):933–951, 2012.
- S. Elsworth and S. Güttel. Conversions between barycentric, RKFUN, and Newton representations of rational interpolants. *Linear Algebra Appl.*, 576:246–257, 2019.
- B. Gavin, A. Międlar, and E. Polizzi. FEAST eigensolver for nonlinear eigenvalue problems. *J. Comput. Sci.*, 27:107–117, 2018.
- S. Güttel and F. Tisseur. The nonlinear eigenvalue problem. *Acta Numerica*, 26:1–94, 2017.
- S. Güttel, R. van Beeumen, K. Meerbergen, and W. Michiels. NLEIGS: A class of fully rational Krylov methods for nonlinear eigenvalue problems. *SIAM J. Sci. Comput.*, 36(6):A2842–A2864, 2014.
- S. Güttel, G. M. Negri Porzio, and F. Tisseur. Robust rational approximations of nonlinear eigenvalue problems. Technical Report MIMS EPrint 2020.24, The University of Manchester, 2020. Available at <http://eprints.maths.manchester.ac.uk/2796/>.

- V. Hernandez, J. E. Roman, and V. Vidal. SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans. Math. Software*, 31(3):351–362, 2005.
- V. Hernandez, J. E. Roman, and A. Tomas. Parallel Arnoldi eigensolvers with enhanced scalability via global communications rearrangement. *Parallel Comput.*, 33(7–8):521–540, 2007.
- N. J. Higham. *Functions of Matrices: Theory and Computation*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2008.
- E. Jarlebring. *The spectrum of delay-differential equations: numerical methods, stability and perturbation*. PhD thesis, TU Carolo-Wilhelmina zu Braunschweig, 2008.
- E. Jarlebring, M. Benedich, G. Mele, E. Ringh, and P. Upadhyaya. NEP-PACK: A Julia package for nonlinear eigenproblems - v0.2. *arXiv:1811.09592 : retrieved 13 Oct 2019*, 2018.
- D. Kressner. A block Newton method for nonlinear eigenvalue problems. *Numer. Math.*, 114:355–372, 2009.
- B.-S. Liao, Z. Bai, L.-Q. Lee, and K. Ko. Nonlinear Rayleigh-Ritz iterative method for solving large scale nonlinear eigenvalue problems. *Taiwan. J. Math.*, 14(3A):869–883, 2010.
- P. Lietaert, K. Meerbergen, and F. Tisseur. Compact two-sided Krylov methods for nonlinear eigenvalue problems. *SIAM J. Sci. Comput.*, 40(5):A2801–A2829, 2018.
- V. Mehrmann and H. Voss. Nonlinear eigenvalue problems: a challenge for modern eigenvalue methods. *GAMM Mitt.*, 27(2):121–152, 2004.
- G. Mele and E. Jarlebring. On restarting the tensor infinite Arnoldi method. *BIT*, 58(1):133–162, 2018.
- W. Michiels and S.-I. Niculescu. *Stability, Control, and Computation for Time-Delay Systems. An Eigenvalue-Based Approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, second edition, 2014.
- Y. Nakatsukasa, O. Sète, and L. N. Trefethen. The AAA algorithm for rational approximation. *SIAM J. Sci. Comput.*, 40(3):A1494–A1522, 2018.
- A. Neumaier. Residual inverse iteration for the nonlinear eigenvalue problem. *SIAM J. Numer. Anal.*, 22(5):914–923, 1985.
- J. E. Roman, C. Campos, E. Romero, and A. Tomas. SLEPc users manual. Technical Report DSIC-II/24/02–Revision 3.12, D. Sistemes Informàtics i Computació, Universitat Politècnica de València, 2019.
- A. Ruhe. Algorithms for the nonlinear eigenvalue problem. *SIAM J. Numer. Anal.*, 10(4):674–689, 1973.
- A. Ruhe. Rational Krylov sequence methods for eigenvalue computation. *Linear Algebra Appl.*, 58:391–405, 1984.
- S. I. Solov’ev. Preconditioned iterative methods for a class of nonlinear eigenvalue problems. *Linear Algebra Appl.*, 415(1):210–229, 2006.
- G. W. Stewart. A Krylov–Schur algorithm for large eigenproblems. *SIAM J. Matrix Anal. Appl.*, 23(3):601–614, 2001.

- Y. Su and Z. Bai. Solving rational eigenvalue problems via linearization. *SIAM J. Matrix Anal. Appl.*, 32(1):201–216, 2011.
- F. Tisseur. Backward error and condition of polynomial eigenvalue problems. *Linear Algebra Appl.*, 309(1–3):339–361, 2000.
- R. van Beeumen, K. Meerbergen, and W. Michiels. Compact rational Krylov methods for nonlinear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 36(2):820–838, 2015.
- R. van Beeumen, O. Marques, E. G. Ng, C. Yang, Z. Bai, L. Ge, O. Kononenko, Z. Li, C.-K. Ng, and L. Xiao. Computing resonant modes of accelerator cavities by solving nonlinear eigenvalue problems via rational approximation. *J. Comput. Phys.*, 374:1031–1043, 2018.
- W. G. Vandenberghe, M. V. Fischetti, R. van Beeumen, K. Meerbergen, W. Michiels, and C. Effenberger. Determining bound states in a semiconductor device with contacts using a nonlinear eigenvalue solver. *J. Comput. Electron.*, 13(3):753–762, 2014.
- H. Voss. An Arnoldi method for nonlinear eigenvalue problems. *BIT*, 44(2):387–401, 2004.
- F. Xue. A block preconditioned harmonic projection method for large-scale nonlinear eigenvalue problems. *SIAM J. Sci. Comput.*, 40(3):A1809–A1835, 2018.
- F. Zolla, A. Nicolet, and G. Demésy. Photonics in highly dispersive media: the exact modal expansion. *Opt. Lett.*, 43(23):5813–5816, 2018.
- I. N. Zwaan and M. E. Hochstenbach. Krylov–Schur-type restarts for the two-sided Arnoldi method. *SIAM J. Matrix Anal. Appl.*, 38(2):297–321, 2017.