Document downloaded from:
http://hdl.handle.net/10251/163192
This paper must be cited as:
Campos, C.; Román Moltó, JE. (2020). Inertia-based spectrum slicing for symmetric quadratic eigenvalue problems. Numerical Linear Algebra with Applications. 27(4):1-17. https://doi.org/10.1002/nla. 2293


The final publication is available at
https://doi.org/10.1002/nla. 2293

Copyright John Wiley \& Sons

Additional Information

# Inertia-based spectrum slicing for symmetric quadratic eigenvalue problems* 

Carmen Campos ${ }^{\dagger} \quad$ Jose E. Roman ${ }^{\ddagger}$

November 7, 2019


#### Abstract

In the quadratic eigenvalue problem with all coefficient matrices symmetric, there can be complex eigenvalues. However, some applications need to compute real eigenvalues only. We propose a Lanczos-based method for computing all real eigenvalues contained in a given interval of large scale symmetric quadratic eigenvalue problems. The method uses matrix inertias of the quadratic polynomial evaluated at different shift values. In this way, for hyperbolic problems, it is possible to make sure that all eigenvalues in the interval have been computed. We also discuss the general non-hyperbolic case. Our implementation is memory-efficient by representing the computed pseudo-Lanczos basis in a compact tensor product representation. We show results of computational experiments with a parallel implementation in the SLEPc library.


## 1 Introduction

The quadratic eigenvalue problem (QEP) arises in many scientific computing applications. One such application is the analysis of damped oscillations of linear systems [27, 28], such as a springmass system, or systems that can be modeled by analogous equations, e.g., electrical networks. The discretization of the governing equation leads to the eigenproblem

$$
\begin{equation*}
\left(\lambda^{2} M+\lambda C+K\right) x=0, \tag{1}
\end{equation*}
$$

where $\lambda$ is the eigenvalue and $x \neq 0$ is the eigenvector. In this paper, we focus on the case where the coefficient matrices $M, C$ and $K$ are real symmetric (although the approach is also valid for Hermitian matrices). They are sometimes called mass, damping and stiffness matrices, respectively. Typically, $M$ and $K$ are positive definite.

In cases where damping effects can be neglected, the problem would be solved as a linear eigenproblem

$$
\begin{equation*}
A x=\mu B x . \tag{2}
\end{equation*}
$$

If the pencil $(A, B)$ is symmetric-definite then all eigenvalues $\mu_{i}$ are real and if $B$ is positive definite the eigenvectors $x_{i}$ are mutually $B$-orthogonal. For large-scale problems where the full spectrum is not required, Lanczos-type methods can be used to obtain a few eigenpairs $\left(x_{i}, \mu_{i}\right)$. If the wanted eigenvalues lie in the interior of the spectrum, then one possible strategy

[^0]is to combine Lanczos with the shift-and-invert spectral transformation, that operates with $(A-\sigma B)^{-1} B$ to compute eigenvalues closest to the shift $\sigma$. Furthermore, it is often necessary to compute all eigenvalues contained in a prescribed interval of the real line. The latter scenario is the main focus of this paper. By performing several shift-and-invert Lanczos runs at different shift locations $[10,5]$ it is possible to effectively sweep the interval and compute all eigenvalues efficiently. This type of strategy, often called spectrum slicing, relies on the matrix inertia of $(A-\sigma B)$ obtained at each shift $\sigma$, which provides the number of negative eigenvalues for this matrix and informs about the number of eigenvalues of (2) lying on the left of that particular shift value, so that a bisection scheme can guarantee that no eigenvalues have been missed. Note that alternative spectrum slicing mechanisms that do not need inertia are being developed in recent times [20], but we do not consider this kind of method here.

Our goal is to extend the inertia-based spectrum slicing scheme to the QEP (1). A common way to solve the QEP is to linearize it first to obtain an equivalent linear eigenproblem (2), then use a linear eigensolver to compute the wanted solutions, and finally extract the QEP eigenvectors from the eigenvectors of the linearization. In the case that all three coefficient matrices of the QEP are symmetric, it is convenient to use a symmetric linearization, where the resulting $A$ and $B$ matrices of the linearization are also symmetric. However, the pencil $(A, B)$ is either indefinite, or it is definite but requires to compute a certain transformation so that $B$ is positive definite $[11,23]$. In this scenario, the question is how to use the information of the matrix inertia of $A-\sigma B$ to determine the localization of eigenvalues.

The topic of inertia computation in indefinite pencils has been treated by Nakatsukasa and Noferini [22], who show that it is not possible to get the exact eigenvalue count in this general case. Our problem is a particular case since the pencil comes from the linearization of a symmetric quadratic matrix polynomial. We will see below that for some problem types inertia still gives a complete information about the location of eigenvalues in this case.

In addition to computing the correct eigenvalue count inside an interval, another difficulty is that a Krylov solver that operates on a symmetric-indefinite pencil and wants to exploit symmetry has to rely on the pseudo-Lanczos process [24]. Furthermore, the linearization duplicates the size of the matrices, that is, $A$ and $B$ have $2 n$ rows and columns, where $n$ is the dimension of the matrices in (1). A naive implementation of the solver will be inefficient in terms of computational and storage cost, because it has to operate with Lanczos vectors of length $2 n$, and also especially because a matrix of order $2 n$ must be factorized in the shift-and-invert technique. In a previous work [7] we showed that it is possible to make an efficient implementation of the pseudo-Lanczos recurrence in the case of linearization of a symmetric QEP, and do the shift-and-invert solves implicitly by only requiring the factorization of $Q(\sigma)$, where $Q$ is the matrix polynomial associated with the QEP,

$$
\begin{equation*}
Q(\lambda)=\lambda^{2} M+\lambda C+K \tag{3}
\end{equation*}
$$

Moreover, we presented a memory-efficient variant, called STOAR [7], that uses a compact representation of the Krylov basis and hence stores and operates only with vectors of length $n$.

In this paper we aim at extending STOAR with the capability to retrieve all eigenvalues contained in an interval, with inertia-based spectrum slicing. Since we are factorizing $Q(\sigma)$, rather than the linearization, we must consider how to perform bisection by using inertia information obtained from $Q(\sigma)$. This will be feasible in the particular case of hyperbolic quadratic eigenvalue problems, where all eigenvalues are real [12], but also in some non-hyperbolic problems as shown by Li and Cai [19]. This latter paper proposes a classical bisection algorithm to compute all real eigenvalues contained in a interval, using the inertia count. The advantage of our method, compared to Li and Cai's, is that the number of required factorizations is much smaller, as it is a subspace-based algorithm. Furthermore, our method computes both eigenvalues and
eigenvectors, whereas Li and Cai return eigenvalues only, so one step of inverse iteration would be necessary for each wanted eigenvector.

The main contribution of this paper is to show how all these ingredients must be combined to get a robust and efficient spectrum slicing solver for hyperbolic QEPs that can also be used to solve general symmetric QEPs under certain conditions in the selected interval. We also show results obtained with a practical implementation in SLEPc [25, 13], where we have combined our previous STOAR implementation [7] with the spectrum slicing solver for the linear case [5]. In particular, we have to consider practical issues such as how to perform deflation against basis vectors stored in a compact representation, which is required to avoid reconvergence of eigenvalues in a neighboring subinterval, or how to lock converged eigenvalues when doing the thick-restart step.

The paper is organized as follows. Preliminaries are described in section 2, including the description of inertia laws for linear and quadratic eigenproblems, as well as the restarted memoryefficient pseudo-Lanczos solver. The details of the spectrum slicing scheme for symmetric QEPs is presented in section 3 , including aspects related to parallel implementation. The performance and accuracy of the solver is assessed in section 4, and the conclusions follow in section 5 .

## 2 Previous work

In this section we describe related works which include the main ingredients used in the spectrum slicing method proposed in this work.

### 2.1 Spectrum slicing in linear eigenvalue problems

The technique of inertia-based spectrum slicing was developed in the 1980's, reaching its full potential in the paper by Grimes et al. [10]. Some years ago we implemented this technique in SLEPc, modernizing the algorithm by including features such as restart and parallel computing [5]. This implementation is being used to solve very challenging problems in computational chemistry [17]. We next give a brief overview of the method emphasizing three relevant elements: $B$-Lanczos, inertia and deflation. In this section we assume a symmetric linear eigenproblem (2) with positive definite $B$.

### 2.1.1 Solver preserving the symmetric structure

The spectrum slicing technique uses the symmetric Lanczos method and the shift-and-invert transformation to obtain interior eigenvalues. It computes the largest magnitude eigenvalues of

$$
\begin{equation*}
S x=\theta x, \quad S=(A-\sigma B)^{-1} B \tag{4}
\end{equation*}
$$

where the obtained eigenvalues $\theta=(\mu-\sigma)^{-1}$ correspond to eigenvalues $\mu$ closest to the target $\sigma \in \mathbb{R}$. The transformed matrix $S$ is not symmetric but it is self-adjoint with respect to the inner product defined by $B,\langle u, v\rangle_{B}:=u^{*} B v$. Hence, by using this non-standard inner product, the Lanczos method can still be used to generate a Lanczos decomposition

$$
\begin{equation*}
S V_{j}=V_{j} T_{j}+\beta_{j} v_{j+1} e_{j}^{*} \tag{5}
\end{equation*}
$$

where the columns of $V_{j}$ form a $B$-orthonormal basis of the Krylov subspace spanned by $S$ and an initial vector $v_{1}, V_{j}^{*} B V_{j}=I$, and $T_{j}$ is a real symmetric tridiagonal matrix.

The Krylov subspace contains increasingly accurate Ritz approximations ( $\tilde{x}_{i}, \tilde{\theta}_{i}$ ) of eigenpairs of (4), where $T_{j} y_{i}=\tilde{\theta}_{i} y_{i}$ and $\tilde{x}_{i}=V_{j} y_{i}$. A cheap convergence criterion for the obtained
approximations is $\beta_{j}\left|y_{j i}\right| \leq \varepsilon\left|\tilde{\theta}_{i}\right|$ for a given tolerance $\varepsilon$, which involves a bound for the $B$-norm of the residual

$$
\begin{equation*}
r_{i}=S \tilde{x}_{i}-\tilde{\theta}_{i} \tilde{x}_{i} . \tag{6}
\end{equation*}
$$

In a spectrum slicing strategy we compute all eigenvalues in a given interval by chunks, using the shift-and-invert transformation at a sequence of shifts, $\sigma_{k}, k=1,2, \ldots$, and validating the number of eigenvalues obtained using the information supplied by the inertia of the matrices $\left(A-\sigma_{k} B\right)$.

### 2.1.2 Counting eigenvalues using matrix inertias

We denote as $\nu(M)$ the number of negative eigenvalues of a symmetric matrix $M$, and we refer to it as the inertia of $M$. In virtue of Sylvester's inertia law [16], this value remains unchanged under congruence transformations. In generalized linear eigenproblems (2) with $B$ positive definite, Grimes et al. [10] consider the Cholesky decomposition $B=L_{B} L_{B}^{*}$ and the reduction to standard form $L_{B}^{-1} A L_{B}^{-*} x=\lambda x$ to see that $\nu(A-\sigma B)=\nu\left(L_{B}^{-1} A L_{B}^{-*}-\sigma I\right)$ gives the number of negative eigenvalues for the shifted eigenproblem. Thus, in this case the number of eigenvalues of $A-\lambda B$ to the left of a shift $\sigma$ can be obtained as the inertia of the matrix $(A-\sigma B)$.

In the context of a Lanczos iteration for computing eigenvalues closest to a target value, to expand the Krylov subspace associated to $S$ (4), we use an indefinite (block-)triangular factorization $A-\sigma_{k} B=L_{k} D_{k} L_{k}^{*}$ where the matrix $D_{k}$ is block-diagonal with $1 \times 1$ or $2 \times 2$ diagonal blocks. As a byproduct of this factorization the number of eigenvalues on the left of $\sigma_{k}$ can be obtained, $\nu_{k}:=\nu\left(A-\sigma_{k} B\right)=\nu\left(D_{k}\right)$. In this way, running Lanczos on two different shifts $\sigma_{k}<\sigma_{k+1}$, we know that the number of eigenvalues, counted with their multiplicities, in the interval $\left[\sigma_{k}, \sigma_{k+1}\right)$ is $\nu_{k+1}-\nu_{k}$.

### 2.1.3 Avoiding reconvergence via deflation

Computing eigenvalues from different Lanczos runs may entail some difficulties that spectrum slicing has to address. One of them is ascertaining whether equal (in machine precision) values computed from different shifts are spurious duplicates or genuine multiples of a particular eigenvalue. Another problem faced by spectrum slicing is to ensure $B$-orthogonality between eigenvectors associated to a multiple eigenvalue (or a cluster of eigenvalues) when calculated from more than one shift. The tool used for solving these problems is deflation, which avoids reconvergence by forcing the Lanczos method to work in the orthogonal complement of a set of selected vectors. Since eigenvectors of (2) are mutually $B$-orthogonal, in the context of the $B$ Lanczos method, such vectors can be easily deflated incorporating them directly in the Lanczos basis against which $B$-orthogonalization is performed. This fact makes the $B$-Lanczos procedure a suitable Krylov method when considering the deflation of eigenvector sets.

Although selected deflation reduces the possibility of recalculating eigenvalues already obtained, it also involves an additional cost since at each iteration of the Lanczos method one vector has to be orthogonalized against all vectors selected for deflation. Therefore, the amount of deflation to be applied for each shift must be adjusted. The strategy used in our slicing technique for linear eigenproblems [5] selects for deflation the set of computed eigenvectors associated to eigenvalues inside the smallest subinterval $\left[\sigma_{\ell}, \sigma_{r}\right]$ containing the current shift $\sigma$ and having endpoints at the set of previous shifts. This scheme accepts any value computed from $\sigma$ in $\left[\sigma_{\ell}, \sigma_{r}\right]$, where reconvergence is not possible, and discards any eigenvalue computed outside this interval.

Figure 1 depicts the spectrum slicing process for three consecutive shifts. First, it shows a trust subinterval (between $a$ and $\sigma_{k-1}$ ) with no missing eigenvalues, obtained after $k-1$ shifts


Figure 1: Three steps of a sample spectrum slicing computation, corresponding to three consecutive shift values $\sigma_{k-1}, \sigma_{k}$ and $\sigma_{k+1}$. The picture illustrates a situation of backtracking, because $\sigma_{k+1}$ is placed on the left of $\sigma_{k}$ while normal movement is to the right. Vertical lines represent eigenvalues, with taller and thicker strokes for already computed eigenvalues (violet color for those being used for deflation) and smaller strokes for eigenvalues that have not been found by the algorithm yet. The shaded area is the trust subinterval, where computed inertia matches the number of eigenvalues found. The interval $\left[\sigma_{\ell}, \sigma_{r}\right]$ is used to choose eigenvectors for deflation.
have been processed; second, it shows the subinterval selected for deflation at the $k$ th Lanczos run, and the newly computed eigenvalues at $\sigma_{k}$; and third, it illustrates the deflation subinterval when backtracking to search for missing eigenvalues, and the extension of the trust subinterval (between $a$ and $\sigma_{k}$ ) once the number of computed eigenvalues matches the inertia count.

### 2.2 Symmetric quadratic eigenproblems

We now turn our attention to the quadratic eigenvalue problem (1), where we assume that all matrices are symmetric and $M$ is non-singular. These problems have special spectral properties, such as for example that eigenvalues $\lambda$ are real or come in complex conjugate pairs $\lambda$ and $\bar{\lambda}$, even if the matrices defining the QEP are complex Hermitian. Also, in some cases we can ensure that all the eigenvalues are real, as discussed below. Under some assumptions [29], real eigenvalues of symmetric QEPs present a variational characterization allowing results to be exploited in a spectrum slicing method relying on Sylvester's law of inertia, similar to those mentioned in section 2.1.

A standard way of solving the quadratic eigenvalue problem is via linearization, that is, via a linear eigenvalue problem whose eigenvalues coincide with those of the QEP. For the case of a symmetric QEP, the pencil

$$
L(\lambda)=A-\lambda B, \quad \text { with } \quad A=\left[\begin{array}{cc}
K & 0  \tag{7}\\
0 & -M
\end{array}\right], \quad B=\left[\begin{array}{cc}
-C & -M \\
-M & 0
\end{array}\right],
$$

is a linearization that preserves symmetry. The eigenvectors of $L(\lambda) y=0$ have the form $y=\left[\begin{array}{c}x \\ \lambda x\end{array}\right]$, where $(x, \lambda)$ are the eigenpairs of $Q(\lambda) x=0$.

For our purpose, we will consider the general case of symmetric quadratic eigenvalue problems as well as the special subclass of hyperbolic quadratic eigenproblems. A symmetric QEP is hyperbolic if $M>0$ and $\left(x^{*} C x\right)^{2}>4\left(x^{*} M x\right)\left(x^{*} K x\right)$ for all $x \in \mathbb{C}^{n} \backslash\{0\}$ [27]. In this case, all eigenvalues are guaranteed to be real, forming two disjoint groups of $n$ eigenvalues, each having linearly independent eigenvectors. Hyperbolic QEPs are also characterized by the symmetric


Figure 2: Pictorial representation of the two groups of eigenvalues of a hyperbolic QEP. The line $\nu(Q)$ illustrates how the value of the inertia of matrix $Q(\sigma)$ changes for varying $\sigma$.
pencil (7) being definite [15]. A particular case of hyperbolic QEP, often called overdamped, is when $C>0$ and $K \geq 0$, in which case all its eigenvalues are nonpositive. On the other hand, a more general class of symmetric QEPs where all eigenvalues are real is the so-called definite QEP, characterized by the fact that a symmetric-definite linearization exists [15, 1]. Symmetric QEPs of both hyperbolic and definite type can be transformed in a way that the $B$ matrix of the symmetric linearization pencil (7) is positive definite. A procedure for this has been proposed by Niendorf and Voss [23]. Then it is possible to apply the spectrum slicing technique of section 2.1 without modification. However, we advocate for modifying the spectrum slicing method so that it can be used without the need of the Niendorf-Voss transformation, resulting in a cheaper computation as will be illustrated in section 4. Furthermore, our approach may also be suitable for some general symmetric QEPs (not definite). For the modified spectrum slicing method, we are not exploiting the overdamped and definite structures, we will only distinguish between hyperbolic QEP and general symmetric QEP.

The spectrum slicing technique described in section 2.1 relies on the eigenvalue count inside an interval, which is computed using the matrix inertia of $A-\sigma B$. To extend this scheme to the symmetric QEP we review similar results available for some particular QEP types, such as the hyperbolic case.

### 2.2.1 Matrix inertia in quadratic symmetric matrix polynomials

In the sequel, we will use $n_{\ell}(\sigma)$ to denote the number of real eigenvalues to the left of $\sigma \in \mathbb{R}$ in a given symmetric eigenproblem (linear or quadratic).

In symmetric linear eigenvalue problems (2), specific properties such as $B$ being positive definite give us the possibility of computing $n_{\ell}(\sigma)$, which in this case is equal to the value of the matrix inertia $\nu(A-\sigma B)$. For the more general case of symmetric nonlinear eigenproblems, Kostić and Voss [18] give sufficient conditions that allow using the matrix inertia information to locate real eigenvalues. In the particular case of hyperbolic QEPs, they show that it is always possible to compute $n_{\ell}(\sigma)$ from the matrix inertia $\nu(Q(\sigma))$, in a similar way as in the linear case. Focusing particularly on the hyperbolic case, we now summarize results that have been used in the eigensolver implementation (section 3) for computing $n_{\ell}(\sigma)$.

As mentioned before, the eigenvalues of a hyperbolic QEP form two disjoint groups of $n$ real eigenvalues each,

$$
\begin{equation*}
\lambda_{1}^{-} \leq \lambda_{2}^{-} \leq \cdots \leq \lambda_{n}^{-}<\lambda_{1}^{+} \leq \lambda_{2}^{+} \leq \cdots \leq \lambda_{n}^{+}, \tag{8}
\end{equation*}
$$

as illustrated in Figure 2. The intervals $J_{-}:=\left[\lambda_{1}^{-}, \lambda_{n}^{-}\right]$and $J_{+}:=\left[\lambda_{1}^{+}, \lambda_{n}^{+}\right]$correspond to the two real solutions $\rho_{-}(x)$ and $\rho_{+}(x)$, respectively, of the scalar equation

$$
\begin{equation*}
x^{*} Q(\lambda) x=\lambda^{2} x^{*} M x+\lambda x^{*} C x+x^{*} K x=0, \quad x \in \mathbb{C}^{n} \backslash\{0\} . \tag{9}
\end{equation*}
$$

Another property of hyperbolic problems is that the matrix $Q(\sigma)$ is negative definite for all $\sigma$ in the open interval $\left(\lambda_{n}^{-}, \lambda_{1}^{+}\right)$, and positive definite in the intervals $\left(-\infty, \lambda_{1}^{-}\right)$and $\left(\lambda_{n}^{+},+\infty\right)$.

To compute $n_{\ell}(\sigma)$, for a given shift $\sigma \in \mathbb{R}$, from the matrix inertia $\nu(Q(\sigma))$, we need to know the relative position of $\sigma$ with respect to the intervals $J_{-}$and $J_{+}$. We distinguish three cases [18]:

- If $\nu(Q(\sigma))=n$, then $\sigma$ is located between the two intervals $J_{-}$and $J_{+}$, and hence there are $n$ eigenvalues on the left of $\sigma, n_{\ell}(\sigma)=n$.
- If $\nu(Q(\sigma))=0$, then $\sigma$ is either on the left of $J_{-}$or on the right of $J_{+}$. In that case, we compute $t:=x^{*} Q^{\prime}(\sigma) x$ for a random vector $x$ to determine which side is $\sigma$ on. It holds that $n_{\ell}(\sigma)=2 n$ if $t>0$, and $n_{\ell}(\sigma)=0$ otherwise.
- If $0<\nu(Q(\sigma))<n$, then $\sigma$ is either inside $J_{+}$or $J_{-}$. Similarly to the previous case, to discriminate between the two intervals we need to compute $t:=x^{*} Q^{\prime}(\sigma) x$, in this case for a vector $x$ such that $x^{*} Q(\sigma) x>0$. One such $x$ is for example a Ritz vector approximating an eigenvector associated to a positive eigenvalue of $Q(\sigma)$. Then,

$$
n_{\ell}(\sigma)= \begin{cases}\nu(Q(\sigma)), & \text { if } t<0  \tag{10}\\ 2 n-\nu(Q(\sigma)), & \text { if } t>0\end{cases}
$$

The fact that $n_{\ell}(\sigma)$ can always be computed in the hyperbolic case makes the spectrum slicing technique applicable for these problems. The general case of symmetric QEPs is trickier. There may be both real and complex eigenvalues, and there is no general formula to compute $n_{\ell}(\sigma)$ from $\nu(Q(\sigma))$, even if all eigenvalues are real. As a consequence, the described technique cannot be applied in general for all symmetric QEPs. Still, we can find conditions that allow us to extend the spectrum slicing technique to certain non-hyperbolic problems. For that case, it is relevant to define eigenvalue types of symmetric matrix polynomials [8]. A real eigenvalue $\lambda_{i}$ of $Q(\lambda) x=0$ is of positive type if

$$
\begin{equation*}
x^{*} Q^{\prime}\left(\lambda_{i}\right) x=2 \lambda_{i} x^{*} M x+x^{*} C x>0, \quad \forall x \in \operatorname{ker}\left(Q\left(\lambda_{i}\right)\right) \backslash\{0\} . \tag{11}
\end{equation*}
$$

Likewise, it is of negative type if $x^{*} Q^{\prime}\left(\lambda_{i}\right) x<0$. We say that a real eigenvalue is of definite type if it is either of positive or negative type, and otherwise it is of mixed type. In the case of hyperbolic QEPs, all eigenvalues are of definite type, more precisely those contained in $J_{-}$are of negative type and those in $J_{+}$of positive type.

Given an interval $[a, b]$, if all real eigenvalues contained in such interval are of definite type and, furthermore, all of them are of the same type (either positive or negative) and semi-simple, then it is possible to show [19] that $n_{\ell}(b)-n_{\ell}(a)=|\nu(Q(b))-\nu(Q(a))|$. Hence, under these assumptions, the spectrum slicing technique of section 3 can be used to find all real eigenvalues contained in the interval in general symmetric QEPs. Unfortunately, it is not possible to know a priori if a given problem satisfies the assumptions or not.

### 2.3 Memory-efficient Krylov solvers for the symmetric QEP

Our solver relies on a Krylov method operating on the pencil of the linearization (7), which is not definite in general. As hinted in section 2.1, we need to use an iterative solver that preserves the structure of the linearization, in a way that the computed eigenvectors can be readily used in the deflation mechanism required for spectrum slicing. This can be accomplished with pseudoLanczos $[24,7]$. Furthermore, due to the structure of the linearization it is very convenient to use specific memory-efficient methods such as STOAR [7].

### 2.3.1 The pseudo-Lanczos method to solve symmetric indefinite linear eigenproblems

The pseudo-Lanczos recurrence preserves the symmetric-indefinite structure of the problem by using an indefinite inner product $[9],\langle u, v\rangle_{B}:=u^{*} B v$, for which $\langle z, z\rangle_{B}$ can be negative. The computed pseudo-Lanczos decomposition, analogous to (5), is

$$
\begin{equation*}
S V_{j}=V_{j} \Omega_{j}^{-1} T_{j}+t_{j+1, j} \omega_{j+1}^{-1} v_{j+1} e_{j}^{*} \tag{12}
\end{equation*}
$$

where $V_{j}$ is the Krylov basis satisfying $V_{j}^{*} B V_{j}=\Omega_{j}, T_{j}$ is a real symmetric tridiagonal matrix, and $\Omega_{j}=\operatorname{diag}( \pm 1)$ is a signature matrix. Algorithm 1 shows the required steps to compute the pseudo-Lanczos decomposition.

```
Algorithm 1 Pseudo-Lanczos iteration
Input: Matrices \(B\) and \(S=(A-\sigma B)^{-1} B\), initial vector \(v_{1}\), number of steps \(k\)
Output: Pseudo-Lanczos basis \(V_{k+1}\), tridiagonal matrix \(T_{k}\), signature matrix \(\Omega_{k+1}\)
    Normalize: \(\omega_{1}=\operatorname{sign}\left(v_{1}^{*} B v_{1}\right), \quad v_{1}=v_{1} \omega_{1} / \sqrt{\left|v_{1}^{*} B v_{1}\right|}\)
    for \(j=1,2, \ldots, k\) do
        \(z=S v_{j}\)
        \(t_{1: j, j}=V_{j}^{*} B z\)
        \(z=z-V_{j} \Omega_{j}^{-1} t_{1: j, j}\)
        \(t_{j+1, j}=\sqrt{\left|z^{*} B z\right|}\)
        \(\omega_{j+1}=\operatorname{sign}\left(z^{*} B z\right)\)
        \(v_{j+1}=z \omega_{j+1} / t_{j+1, j}\)
    end for
```

Our implementation also incorporates a thick restart based on the eigenvectors of the projected matrix $\Omega_{j}^{-1} T_{j}[7]$.

The pseudo-Lanczos method can be seen as non-symmetric Lanczos particularized for the case of a symmetric generalized eigenproblem. As such, it may suffer from breakdown (or near breakdown) in the case that an isotropic vector is generated, that is, if $\langle z, z\rangle_{B}=0$ for $z \neq 0$. Also, a solver based on pseudo-Lanczos may become unstable, either due to the pseudo-Lanczos iteration itself (involving indefinite Gram-Schmidt orthogonalization [26]) or the resolution of the symmetric-indefinite projected problem. A careful implementation of these two parts can reduce the risk of instability, and furthermore it is possible to detect instability by monitoring the loss of pseudo-symmetry in the projected problem [7].

In the case of shift-and-invert transformation, the matrix used to expand the Krylov subspace is

$$
S=(A-\sigma B)^{-1} B=\left[\begin{array}{cc}
-K-\sigma C & -\sigma M  \tag{13}\\
-\sigma M & M
\end{array}\right]^{-1}\left[\begin{array}{cc}
C & M \\
M & 0
\end{array}\right] .
$$

A practical implementation should avoid the explicit computation of the inverse of $A-\sigma B$, and instead perform a linear solve whenever a matrix-vector product with $S$ is required. Rather than factoring this $2 n \times 2 n$ matrix, a Schur complement approach is preferred, which involves operating with the (implicit) inverse of $Q(\sigma)$, that is, the quadratic polynomial (3) evaluated at the current shift $\sigma$. Since $Q(\sigma)$ is symmetric (for a real $\sigma$ ), this can be done with an indefinite (block-)triangular factorization $Q(\sigma)=L D L^{*}$, as described in section 2.1 for the linear case. Then the inertia $\nu(Q(\sigma))=\nu(D)$ is readily available to determine the number of eigenvalues in a subinterval.

### 2.3.2 The STOAR method

The dimension of the matrices of the linearization (7) is $2 n$, where $n$ is the size of the coefficient matrices of $Q$. This increased dimension not only implies double storage requirements for the Krylov basis, $V \in \mathbb{C}^{2 n \times j}$ (for some dimension $j$ ), but also a high computational cost. Memoryefficient Krylov solvers [21] try to exploit the block structure of the linearization in a way that the memory requirements are restricted to a basis with vectors of dimension $n$ only, $U \in \mathbb{C}^{n \times(j+1)}$, plus some additional coefficients, $G \in \mathbb{C}^{2(j+1) \times j}$. These bases are related as

$$
\begin{equation*}
V=\left(I_{2} \otimes U\right) G \tag{14}
\end{equation*}
$$

These algorithms are also adapted to carry out much less operations than would be done in a naive implementation. Details of how these techniques have been implemented in SLEPc for general matrix polynomials of arbitrary degree are given in a previous paper [6].

Here we focus on the particular case of symmetric matrix polynomials of degree 2. The details of the memory-efficient Krylov solver for this case have already been worked out in a previous work [7]. The STOAR method, that we will use in the spectrum slicing method of section 3, is a variant of TOAR (two-level orthogonal Arnoldi) [21] that exploits symmetry by using the pseudo-Lanczos recurrence for the matrix

$$
\check{S}=\check{B}^{-1} \check{A}=\left[\begin{array}{cc}
0 & I  \tag{15}\\
-M^{-1} K & -M^{-1} C
\end{array}\right]
$$

associated with the linearization

$$
\check{L}(\lambda)=\check{A}-\lambda \check{B}, \quad \text { with } \quad \check{A}=\left[\begin{array}{cc}
0 & K  \tag{16}\\
K & C
\end{array}\right], \quad \check{B}=\left[\begin{array}{cc}
K & 0 \\
0 & -M
\end{array}\right]
$$

It generates a Krylov relation (12) where the pseudo-Lanczos vectors are built in a compact form (14),

$$
V_{j}=\left[\begin{array}{l}
V_{j}^{0}  \tag{17}\\
V_{j}^{1}
\end{array}\right]=\left[\begin{array}{cc}
U_{j+1} & 0 \\
0 & U_{j+1}
\end{array}\right]\left[\begin{array}{l}
G_{j}^{0} \\
G_{j}^{1}
\end{array}\right]
$$

where $G_{j}^{0}$ and $G_{j}^{1}$ are the $U_{j+1}$ coordinates of the blocks $V_{j}^{0}$ and $V_{j}^{1}$, respectively. We now summarize the main steps carried out by the STOAR method [7].

The first step is the expansion of the Krylov subspace $z=\check{S} v_{j}$ (line 3 of Algorithm 1). Since the last pseudo-Lanczos vector $v_{j}=\left[\begin{array}{c}v_{j}^{0} \\ v_{j}^{1}\end{array}\right]$ is expressed in the compact representation, $v_{j}^{0}=U_{j+1} g_{j}^{0}, v_{j}^{1}=U_{j+1} g_{j}^{1}$, a compact representation of $z=\left[\begin{array}{c}z^{0} \\ z^{1}\end{array}\right]$, where

$$
\left\{\begin{array}{l}
z^{0}=v_{j}^{1}  \tag{18}\\
z^{1}=-M^{-1}\left(K v_{j}^{0}+C v_{j}^{1}\right)=-M^{-1}\left(K U_{j+1} g_{j}^{0}+C U_{j+1} g_{j}^{1}\right)
\end{array}\right.
$$

is computed taking $U_{j+2}=\left[U_{j+1}, u_{j+2}\right]$, with orthonormal columns, and coefficients $g=\left[\begin{array}{c}g^{0} \\ g^{1}\end{array}\right]$ so that $z^{0}=U_{j+2} g^{0}, z^{1}=U_{j+2} g^{1}$. The vector $u_{j+2}$ and coefficients $g^{1}$ are obtained via GramSchmidt orthonormalization of $z^{1}$ against the columns of $U_{j+1}$, whereas $g^{0}=\left[\begin{array}{c}g_{j}^{1} \\ 0\end{array}\right]$.

The orthogonalization and normalization steps (lines 4-8 of Algorithm 1) start with vector $z=\left(I_{2} \otimes U_{j+2}\right) g$. Using the fact that $U_{j+2}$ has orthonormal columns, the Gram-Schmidt orthogonalization turn into the orthogonalization of $g$ against the columns of $G_{j}$ with the indefinite inner product defined by

$$
\left[\begin{array}{cc}
U_{j+2}^{*} K U_{j+2} & 0  \tag{19}\\
0 & -U_{j+2}^{*} M U_{j+2}
\end{array}\right]
$$

The normalization proceeds in the same way, by normalizing the resulting vector to obtain $g_{j+1}$ and $\omega_{j+1}$.

STOAR is a restarted method that compresses the Krylov basis when it reaches a specified maximum size. While reducing the Krylov basis from size $m$ to $p$ by a transformation of the form $\tilde{V}_{p}=V_{m} Q=\left(I_{2} \otimes U_{m+1}\right) G_{m} Q$, for $Q \in \mathbb{C}^{m \times p}$, this method performs (in the same way as TOAR), an additional step which involves an SVD decomposition of the matrix $\left[G_{m}^{0}, G_{m}^{1}\right]$ to adjust the size of the basis $U_{m+1}$ to the new Krylov basis dimension.

## 3 Spectrum slicing for quadratic eigenproblems

In this section we show how the spectrum slicing technique of section 2.1 can be adapted to symmetric QEPs solved with the STOAR method of section 2.3 using the information provided by inertia as discussed in section 2.2 .

### 3.1 Basic scheme

Algorithm 2 shows a scheme of the spectrum slicing technique particularized for symmetric QEPs. The algorithm computes all real eigenvalues contained in a given interval $[a, b]$. The overall computation is similar to the case of linear eigenproblems, but with important differences that will be highlighted below.

```
Algorithm 2 Spectrum slicing for symmetric QEPs
Input: Matrices \(A\) and \(B\), interval \([a, b]\)
Output: Computed eigenpairs ( \(x_{i}, \lambda_{i}\) ) with \(\lambda_{i} \in[a, b]\)
    Set \(k=1\)
    Compute \(n_{\ell}(a), n_{\ell}(b)\)
    \(\Sigma=\left\{\left(a, a, b, n_{\ell}(a), n_{\ell}(b)\right)\right\}\)
    while \(\Sigma \neq \emptyset\) do
        Extract \(\left(\sigma_{k}, \sigma_{\ell}, \sigma_{r}, n_{\ell}\left(\sigma_{\ell}\right), n_{\ell}\left(\sigma_{r}\right)\right)\) from \(\Sigma\)
        Compute \(n_{\ell}\left(\sigma_{k}\right)\)
        Select eigenvectors computed at \(\left[\sigma_{\ell}, \sigma_{r}\right]\) as deflation vectors
        Run STOAR at \(\sigma_{k}\) to obtain \(m\) eigenpairs ( \(x_{i}, \lambda_{i}\) )
        Discard \(\left(x_{i}, \lambda_{i}\right)\) for \(\lambda_{i} \notin\left[\sigma_{\ell}, \sigma_{r}\right]\)
        Split eigenvalues in two groups \(\left\{\lambda_{i}^{\ell}\right\}_{i=1}^{m_{\ell}}\) and \(\left\{\lambda_{i}^{r}\right\}_{i=1}^{m_{r}}\) with \(\lambda_{i}^{\ell}<\sigma_{k}<\lambda_{i}^{r}\)
        if \(m_{\ell} \neq n_{\ell}\left(\sigma_{k}\right)-n_{\ell}\left(\sigma_{\ell}\right)\) then
            Choose new shift \(\sigma\) in \(\left[\sigma_{\ell}, \sigma_{k}\right]\)
            Insert \(\left(\sigma, \sigma_{\ell}, \sigma_{k}, n_{\ell}\left(\sigma_{\ell}\right), n_{\ell}\left(\sigma_{k}\right)\right)\) into \(\Sigma\)
        end if
        if \(m_{r} \neq n_{\ell}\left(\sigma_{r}\right)-n_{\ell}\left(\sigma_{k}\right)\) then
            Choose new shift \(\sigma\) in \(\left[\sigma_{k}, \sigma_{r}\right]\)
            Insert \(\left(\sigma, \sigma_{k}, \sigma_{r}, n_{\ell}\left(\sigma_{k}\right), n_{\ell}\left(\sigma_{r}\right)\right)\) into \(\Sigma\)
        end if
        \(k=k+1\)
    end while
```

The computation proceeds by placing shifts $\sigma_{k}$ at different points of the interval. Normally the method sweeps the interval either from from left to right or right to left (depending on whether the interval is open in one end, for instance), but it does not mean that the values of the shifts increase (or decrease) monotonically, since in some cases the algorithm must get back to look for missing eigenvalues. There is a bag $\Sigma$ of unprocessed shifts, where each shift
is represented by a tuple $\left(\sigma_{k}, \sigma_{\ell}, \sigma_{r}, n_{\ell}\left(\sigma_{\ell}\right), n_{\ell}\left(\sigma_{r}\right)\right)$ indicating the current shift value $\sigma_{k}$, the interval for deflation $\left[\sigma_{\ell}, \sigma_{r}\right]$ and the number of (real) eigenvalues on the left (inertia) of both ends of this interval. At the beginning of the algorithm, $\Sigma$ must be filled with the initial shift, which is one of the ends of the interval, together with the whole interval $[a, b]$ (which can be open on one side) and the eigenvalue count at both $a$ and $b$ (see section 3.4 for details about how $n_{\ell}(\cdot)$ is obtained). Next, the algorithm proceeds by extracting one shift at a time from the bag and processing it, until the bag is empty (at each shift, at most two new shifts may be inserted into the bag).

New shifts are added to the bag in the case that the number of computed eigenvalues is still less than the eigenvalue count provided by inertia (computed at line 6), either on the left or on the right of $\sigma_{k}$ (or both). The values of the new shifts added to the bag (lines 12 and 16 of Algorithm 2) are determined in two ways. The normal case is when either $\sigma_{\ell}$ or $\sigma_{r}$ coincide with one end of the global interval; in that case, the distance of the new shift with respect of the previous one takes into account the separation of computed eigenvalues [5]. In the case that the algorithm gets back to compute missing eigenvalues, then the new shift is just the midpoint of $\left[\sigma_{\ell}, \sigma_{k}\right]$ (or $\left[\sigma_{k}, \sigma_{r}\right]$ ).

Since eigenvalues are computed by chunks, from different pseudo-Lanczos factorizations at different shifts, it may happen that some of the eigenvalues are computed more than once. To avoid that, we perform explicit deflation, but only against a limited set of eigenvectors, as was mentioned in section 2.1. Eigenvectors corresponding to computed eigenvalues located in $\left[\sigma_{\ell}, \sigma_{r}\right]$ are selected for deflation (line 7 of the algorithm), and this prevents reconvergence of these eigenvalues in the next pseudo-Lanczos run. In the event that an eigenvalue outside $\left[\sigma_{\ell}, \sigma_{r}\right]$ is computed we must discard it (line 9 of the algorithm) because we cannot tell if it is a genuine or spurious copy.

### 3.2 Linearization

The computation associated with each shift consists essentially in a run of the STOAR method to compute a fixed number, $m$, of eigenpairs around $\sigma_{k}$. In our implementation, the employed linearization is a general one, that includes (7) and (16) as particular cases. We consider symmetric pencils proposed by Higham et al. [14] belonging to the space $\mathbb{H}(Q)$,

$$
\hat{L}(\lambda)=\hat{A}-\lambda \hat{B}, \quad \text { with } \quad \hat{A}=\left[\begin{array}{cc}
\beta K & \alpha K  \tag{20}\\
\alpha K & \alpha C-\beta M
\end{array}\right], \quad \hat{B}=\left[\begin{array}{cc}
\alpha K-\beta C & -\beta M \\
-\beta M & -\alpha M
\end{array}\right],
$$

with $\alpha, \beta \in \mathbb{R}$. A pencil of this type is a linearization for the matrix polynomial (3) whenever the matrix

$$
\begin{equation*}
R_{\alpha, \beta}:=\beta^{2} M-\alpha \beta C+\alpha^{2} K \tag{21}
\end{equation*}
$$

is non-singular [14]. The eigenvectors of the linearization have the form $y=\left[\begin{array}{c}x \\ \lambda x\end{array}\right]$, where $(x, \lambda)$ is an eigenpair of the QEP. The linearizations (7) and (16) are particular cases of (20) taking $(\alpha, \beta)=(1,0)$ and $(0,1)$, respectively. By leaving $(\alpha, \beta)$ as free parameters, we give the user the possibility to improve the properties of the indefinite inner product matrix $\hat{B}$ for a more stable Gram-Schmidt orthogonalization [26].

For the spectrum slicing scheme we need to perform the shift-and-invert transformation at each shift. Contrary to the approach used in our previous paper [7], where the transformation was done on the quadratic problem, the deflation technique, explained in section 3.3, requires the shift-and-invert transformation to be performed on the linearization. The rationale is that in this latter case the transformed problems generated from different shifts have the same eigenvectors, which can be directly used from one shift to another for deflation.

To compute the action of $(\hat{A}-\sigma \hat{B})^{-1} \hat{B}$ on some vector $v \in \mathbb{C}^{n}$ without explicitly building and factorizing the matrix $(\hat{A}-\sigma \hat{B})$ of size $2 n$, we use a block factorization. Assuming $\beta \neq 0$, the factorization is

$$
\begin{align*}
\hat{A}-\sigma \hat{B} & =\left[\begin{array}{cc}
\beta K-\sigma \alpha K+\sigma \beta C & \alpha K+\sigma \beta M \\
\alpha K+\sigma \beta M & \alpha C-\beta M+\sigma \alpha M
\end{array}\right] \\
& =\left[\begin{array}{cc}
I & 0 \\
\frac{\alpha}{\beta} I & I
\end{array}\right]\left[\begin{array}{cc}
I & 0 \\
0 & \frac{1}{\beta} R_{\alpha, \beta}
\end{array}\right]\left[\begin{array}{cc}
I & -(\alpha K+\sigma \beta M) \\
0 & I
\end{array}\right]\left[\begin{array}{cc}
\beta Q(\sigma) & 0 \\
\sigma I & -I
\end{array}\right], \tag{22}
\end{align*}
$$

with $R_{\alpha, \beta}$ defined above. From decomposition (22) we get

$$
(\hat{A}-\sigma \hat{B})^{-1} \hat{B}=\left[\begin{array}{cc}
\beta Q(\sigma) & 0  \tag{23}\\
\sigma I & -I
\end{array}\right]^{-1}\left[\begin{array}{cc}
I & -(\alpha K+\sigma \beta M) \\
0 & I
\end{array}\right]^{-1} \tilde{B},
$$

where

$$
\tilde{B}:=\left[\begin{array}{cc}
I & 0  \tag{24}\\
0 & \beta R_{\alpha, \beta}^{-1}
\end{array}\right]\left[\begin{array}{cc}
I & 0 \\
-\frac{\alpha}{\beta} I & I
\end{array}\right] \hat{B}=\left[\begin{array}{cc}
\alpha K-\beta C & -\beta M \\
-I & 0
\end{array}\right] .
$$

Expressions (23) and (24) allow us to make computations for $z=(\hat{A}-\sigma \hat{B})^{-1} \hat{B} v$ by blocks, operating with matrices and vectors of size $n$, with the formula

$$
\left\{\begin{array}{l}
z^{0}=-Q(\sigma)^{-1}\left(\sigma M v^{0}+C v^{0}+M v^{1}\right)  \tag{25}\\
z^{1}=\sigma z^{0}+v^{0}
\end{array}\right.
$$

If $\beta=0$, it is possible to derive the same formula assuming $\alpha \neq 0$. Considering that $v$ is expressed in compact representation, vector $z$ must also be computed in this representation as explained in section 2.3.

As in the linear case, the computation of (25) must not evaluate the inverse of $Q(\sigma)$ explicitly. Instead, a symmetric-indefinite triangular factorization of $Q(\sigma)$ is used to expand the basis. Furthermore, this factorization will provide the inertia of this matrix, which informs about the localization of eigenvalues, as discussed in section 2.

### 3.3 Deflating eigenvectors in compact representation

As mentioned in section 2.1, the deflation mechanism for spectrum slicing consists in augmenting the Krylov subspace with already available eigenvectors. These eigenvectors have been computed in a previous shift $\sigma$ and deflation prevents their re-computation. In the case of Algorithm 2, that uses the STOAR method, we augment the pseudo-Lanczos basis, which is expressed in the compact representation (17), with eigenvectors of the linearization $\left[\begin{array}{c}x \\ \lambda x\end{array}\right]$, which are also expressed in compact representation. This has several implications:

- During the spectrum slicing loop it is necessary to keep eigenvectors of the linearization $\left[\begin{array}{c}x \\ \lambda x\end{array}\right]$, and only at the end of the computation we extract the relevant part $x$. Note that this does not mean an increase of memory storage, since $\left[\begin{array}{c}x \\ \lambda x\end{array}\right]$ needs to store about $n$ elements in compact representation.
- The management of the subspace basis is much more involved than in the linear case, because augmenting the basis is not as simple as appending new columns, due to the compact representation, and similarly for other operations such as extracting a number of columns from the subspace basis.

To illustrate the issues associated with management of the basis, let us consider several situations. The simplest case is when after step 7 of Algorithm 2 we have a basis of $p$ deflation vectors expressed in compact representation $V_{p}=\left(I_{2} \otimes U_{p}\right) G_{p}$. Deflation in pseudo-Lanczos is effected by including these vectors in the orthogonalization of steps $4-5$ of Algorithm 1. A simple way to achieve this is to initialize the pseudo-Lanczos basis to $V_{p}$ and then start the loop at $j=p$. This will extend the basis naturally to order $k$, and we just have to consider only the trailing principal submatrix of order $(k-p)$ of the projected matrices $T_{k}$ and $\Omega_{k}$ (the leading part contains the already known eigenvalues).

The situation is somewhat more complicated when the set $V$ of vectors selected for deflation is just a part of the eigenvectors obtained from a previous shift, or when such set includes vectors computed from two distinct shifts. In the first case, $V$ would be stored in compact representation $V=\left(I_{2} \otimes U\right) G$ with $G$ having less columns than $U$, making it necessary to perform an operation to adjust and level off both dimensions. In the second case, where we have two sets of vectors $V_{1}=\left(I_{2} \otimes U_{1}\right) G_{1}$ and $V_{2}=\left(I_{2} \otimes U_{2}\right) G_{2}$, obtained from two different STOAR runs, we would need to orthogonally extend one of the two bases, for instance $\tilde{U}=\left[U_{1}, \tilde{U}_{2}\right]$, in such a way that the new basis $\tilde{U}$ spans both $U_{1}$ and $U_{2}$, and we would get in this way a unified compact representation for the two groups of eigenvectors $\left[V_{1}, V_{2}\right]=\left(I_{2} \otimes \tilde{U}\right)\left[G_{1}, \tilde{G}_{2}\right]$.

The latter two cases get simplified if after every STOAR run the basis of computed eigenvectors (that includes those used for deflation) is divided in two groups (corresponding to eigenvalues lying on the left and right of the shift), each of them represented independently in compact form. In this way, the vectors associated with each of the subintervals generated at steps 11 and 15 of Algorithm 2, get stored directly in the appropriate way for their use in deflation of subsequent subintervals.

The basis split is accomplished by means of a small computation involving the SVD of a matrix as follows. The first $q<p$ vectors of the computed basis are $V_{q}=\left(I_{2} \otimes U_{p}\right) G_{q}$ with $U_{p} \in \mathbb{C}^{n \times p}$ and $G_{q} \in \mathbb{C}^{2 p \times q}$, so we need to adjust the number of columns of $U_{p}$ and rows of $G_{q}$. The rank of $\left[V_{q}^{0}, V_{q}^{1}\right]$ is $q$ and similarly for $\left[G_{q}^{0}, G_{q}^{1}\right] \in \mathbb{C}^{p \times 2 q}$. We then compute the compact singular value decomposition

$$
\left[\begin{array}{ll}
G_{q}^{0} & G_{q}^{1} \tag{26}
\end{array}\right]=\check{U} \check{\Sigma} \check{V}^{*},
$$

where $\check{\Sigma}$ is a $q \times q$ diagonal matrix of singular values, and $\check{U}, \check{V}$ have dimensions $p \times q$ and $q \times 2 q$, respectively. This decomposition allows us to express the $V_{q}$ vectors in the form

$$
\begin{equation*}
V_{q}^{i}=U_{p} G_{q}^{i}=U_{p} \check{U} \check{\Sigma} \check{V^{i^{*}}}, \quad \text { for } i=0,1, \tag{27}
\end{equation*}
$$

where we have written $\check{V}^{*}=\left[\check{V}^{0^{*}}, \check{V}^{1^{*}}\right]$. Updating $U_{p} \leftarrow U_{p} \check{U}$ and $G_{q}^{i} \leftarrow \check{\Sigma} \check{V}^{i^{*}}$, for $i=0$, 1, we obtain the representation for $V_{q}$ we are looking for.

### 3.4 Additional details

To complement the description, we provide some additional implementation details.
A fundamental value in Algorithm 2 is $n_{\ell}(\sigma)$, the number of real eigenvalues to the left of $\sigma \in \mathbb{R}$. In a practical implementation this value is explicitly computed only in the case of hyperbolic QEPs, but not in the general symmetric case that we will discuss below.

In the hyperbolic case, the computation of $n_{\ell}(\sigma)$ has been described in section 2.2. Apart from computing the inertia of $Q(\sigma)$ (which implies a factorization), the case $0<\nu(Q(\sigma))<n$ involves an additional cost because it requires computing a value $t:=x^{*} Q^{\prime}(\sigma) x$ to determine the relative position of $\sigma$ with respect to the intervals $J_{-}$and $J_{+}$. In our implementation, $x$ is an approximate Ritz vector of $Q(\sigma)$ corresponding to a positive eigenvalue, and we compute it via a linear eigensolver from SLEPc's EPS module. This computation is required unconditionally at
the beginning of Algorithm 2 (step 2), but inside the loop we can avoid it unless the endpoints of the interval satisfy $a \in J_{-}, b \in J_{+}$.

As we have mentioned in section 2.2 , in the general symmetric case there is no known way to explicitly compute $n_{\ell}(\sigma)$. Still, the spectrum slicing technique can be implemented if the interval $[a, b]$ contains only eigenvalues of the same definite type. The algorithm in this case works similarly to Algorithm 2, but storing the inertias $\nu(Q(\sigma))$ instead of $n_{\ell}(\sigma)$, and using $\left|\nu\left(Q\left(\sigma_{2}\right)\right)-\nu\left(Q\left(\sigma_{1}\right)\right)\right|$ as a count of the number of eigenvalues in $\left[\sigma_{1}, \sigma_{2}\right]$. If the assumption that the interval contains only eigenvalues of the same type does not hold, there is no guarantee of a proper functioning and it may happen that the number of returned eigenvalues is smaller than the actual number of real eigenvalues in $[a, b]$. To minimize the occurrence of this latter scenario, we have implemented several mechanisms to check situations where the required assumption does not hold:

- The first check is done at the beginning of the algorithm on both endpoints of the initial interval, $a$ and $b$. For each of them, we compute the eigenpair $(x, \lambda)$ satisfying the QEP $Q(\lambda) x=0$ for $\lambda$ inside the interval $[a, b]$ and closest to the endpoint. The type of this eigenvalue is given by the sign of $x^{*} Q^{\prime}(\lambda) x$. If the sign is different in both eigenvalues, then the solver aborts because the assumption does not hold.
- If the previous check did not fail, the eigenvalue type must also be checked for all eigenvalues as they are computed, because in principle eigenvalues of different type may be mixed arbitrarily. In this case, instead of explicitly computing the eigenvector $x$ of the quadratic eigenproblem, the type of the eigenvalue $\lambda$ can be determined by checking the type of $\lambda$ directly in the linearization. From the relation

$$
\begin{equation*}
y^{*} \hat{L}^{\prime}(\lambda) y=-y^{*} \hat{B} y=(\alpha \lambda+\beta) x^{*} Q^{\prime}(\lambda) x \tag{28}
\end{equation*}
$$

where $y=\left[\begin{array}{c}x \\ \lambda x\end{array}\right]$, we see that the type of $\lambda$ in the QEP can be easily determined from the type of $\lambda$ in the linearization $\hat{L}$. Furthermore, the sign of $y^{*} \hat{B} y$ is readily available at the end of each STOAR run. If the solver detects eigenvalues of different type within the interval, the execution is aborted with an informative error message.

- The algorithm will also abort the execution if it detects a mismatch between the inertia count and the number of found eigenvalues in a given subinterval.

Still, in the general case there could be cases in which the solver does not detect that the input interval does not satisfy the required conditions. In that case, the number of returned eigenvalues would be less than the actual number of eigenvalues contained in the interval.

We conclude this section with a brief description of the parallel implementation. SLEPc relies on PETSc [3] for linear algebra operations such as sparse matrix-vector products or linear solves. These operations are parallelized for distributed memory parallel computers with a messagepassing paradigm (MPI). Oversimplifying, parallelizing an algorithm essentially amounts to deciding which objects (matrices and vectors) must be parallel (stored in a distributed way), and then PETSc (and SLEPc) internally manage the data exchange necessary to perform the different operations in parallel. Regarding Algorithm 2, the main cost lies in the computation of the inertia via a triangular factorization, as well as the STOAR runs for each shift. The inertia involves a triangular factorization, and parallel implementations of such factorizations are available in PETSc via third party packages such as MUMPS. For the STOAR runs we can distinguish two main building blocks: on one hand the operations in (25), including triangular solves and matrix-vector products with the matrix polynomial coefficients, and on the other hand the orthogonalization of vectors. For the latter operation, we have implemented parallel subroutines in SLEPc that operate with vectors expressed in a compact representation (14).

Table 1: Computational results for the spectrum slicing solver with hyperbolic quadratic matrix polynomials. For each problem, the table shows the matrix dimensions, $n$, the interval of interest, and the number of eigenvalues contained in the interval, $n_{\text {eig }}$. Results include the number of processed shifts, $\# \sigma$, the maximum backward error, $\eta_{\text {max }}$, and the run time (in seconds) with 16 MPI processes.

| problem | $n$ | interval | $n_{\text {eig }}$ | $\# \sigma$ | $\eta_{\max }$ | $t(\mathrm{~s})$ |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| spring | 1.5 mill. | $(-\infty,-49.494891]$ | 411 | 29 | $3 \cdot 10^{-14}$ | 181 |
| spring | 20000 | $[-9.7,-0.5277]$ | 1423 | 99 | $1 \cdot 10^{-11}$ | 41.3 |
| loaded_string | 20000 | $[4,100000]$ | 101 | 11 | $1 \cdot 10^{-10}$ | 2.7 |

Recall that, in this representation, the $U$ basis consists of a set of long vectors (of length $n$ ) that will be distributed in the parallel code, whereas the $G$ factor is in comparison very small and hence it must be stored redundantly in all MPI processes.

## 4 Experimental results

We now show how our new spectrum slicing solver behaves in terms of accuracy as well as computational performance (including parallel efficiency). We first focus on hyperbolic problems and defer the general symmetric case until the end of this section.

The computer used for the executions is Tirant 3, consisting 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, interconnected with an Infiniband FDR10 network. All runs placed at most 4 MPI process per node. The presented results correspond to SLEPc version 3.10, together with PETSc 3.10 and MUMPS 5.1.2. All software has been compiled with Intel C and Fortran compilers (version 18) and Intel MPI.

For the hyperbolic QEPs, we make use of the NLEVP collection [4]. In particular, we consider two problems:

- spring is a QEP obtained from the discretization of a linearly damped mass-spring system. It is a monic polynomial $Q(\lambda)=\lambda^{2} I+\lambda C+K$, where both $C$ and $K$ are tridiagonal. We consider the simplest case where all masses are equal, and all dampers and springs have the same damping constants. In particular, we use the default values $\mu=1$ (masses), $\tau=10$ (dampers) and $\kappa=5$ (springs), which result in a hyperbolic (and overdamped) problem.
- loaded_string is a rational eigenvalue problem arising from the discretization of a loaded vibrating string, with a load of mass $m$ attached by an elastic spring of stiffness $k$. Considering $\mu=k / m$, the problem is formulated as

$$
\begin{equation*}
\left(A-\lambda B+\frac{\lambda}{\lambda-\mu} C\right) x=0 \tag{29}
\end{equation*}
$$

with $A, B$ tridiagonal and positive definite, and $C=k e_{n} e_{n}^{T}$. This problem can be easily transformed to a QEP by multiplication with $\lambda-\mu$. We consider $m=1$ and $k=1$, in which case the resulting QEP is hyperbolic.
Table 1 shows information and results for these two problems. The spring problem has been solved for two different sizes and two different intervals, one of them is half-bounded and

Table 2: Computational results for the spectrum slicing solver with hyperbolic quadratic matrix polynomials, where the Niendorf-Voss transformation has been previously applied. Results include the number of processed shifts, $\# \sigma$, the maximum backward error, $\eta_{\max }$, the run time (in seconds) with 16 MPI processes for the initial transformation, $t_{\text {transf }}$, and for the spectrum slicing computation, $t_{\text {slice }}$.

| problem | $n$ | $\# \sigma$ | $\eta_{\max }$ | $t_{\text {slice }}(\mathrm{s})$ | $t_{\text {transf }}(\mathrm{s})$ |
| :--- | :---: | :---: | :---: | :---: | :---: |
| spring | 1.5 mill. | 27 | $5 \cdot 10^{-14}$ | 175 | 36.6 |
| spring | 20000 | 100 | $4 \cdot 10^{-12}$ | 43.4 | 0.42 |

the other one intersects both $J^{-}$and $J^{+}$and hence contains eigenvalues of both negative and positive type. In this way, we show that the solver is able to cope with all cases that may appear in hyperbolic problems.

The accuracy of a computed solution $(x, \lambda)$ is assessed by means of its relative backward error

$$
\begin{equation*}
\eta(x, \lambda)=\frac{\|Q(\lambda) x\|}{\left(\left|\lambda^{2}\right|\|M\|+|\lambda|\|C\|+\|K\|\right)\|x\|}, \tag{30}
\end{equation*}
$$

where we use $\infty$-norms for practical computation of matrix norms. Note that the solver does not evaluate this backward error in the stopping criterion. Instead, the stopping criterion relies on the residual of the linearization, that is readily available during the pseudo-Lanczos iteration. By default, we check $\left\|r_{i}\right\| / /\left|\tilde{\theta}_{i}\right|<$ tol, where $\tilde{\theta}_{i}$ is the approximate eigenvalue and $r_{i}$ is the residual (6). The tolerance used for the experiments in this section was $10^{-10}$.

The results of Table 1 show that all computed eigenvalues satisfy the requested tolerance. Apart from the reported cases, we have tested the solver with many other problem sizes and intervals, serially and in parallel, and we can conclude that it is remarkably robust in all cases for hyperbolic problems. It is worth mentioning that even in the case that one of the STOAR runs quits earlier due to loss of symmetry (as mentioned in section 2.3), the overall algorithm tolerates this situation because it can recover by computing missing eigenvalues from a different shift. In terms of performance, the run times largely depend on how many shifts are processed by the spectrum slicing algorithm, since each shift implies a sparse factorization. In any case, the total computation time is generally much smaller compared to a naive solution of the linearized problem, apart from the fact that this latter approach does not guarantee that all eigenpairs are returned.

Table 2 shows the results for the same problems as in Table 1 but solved by previously applying the Niendorf-Voss transformation (cf. section 2.2), consisting in shifting the QEP with a value $\mu$ such that matrix $Q(\mu)$ is negative definite. Then, the symmetric linearization pencil has a definite $B$ matrix, and the spectrum slicing method will operate always with a definite inner product. In the spring problem, both approaches provide a similar accuracy, and the computation time corresponding to the spectrum slicing method is also similar. However, the Niendorf-Voss transformation has an additional cost that in the case of $n=1.500 .000$ is nonnegligible. The loaded_string problem was not included in Table 2 since we could not carry out the Niendorf-Voss transformation because the required linear eigenproblems did not converge.

Results for parallel performance when the solver is applied to the spring problem of size 1.5 million are reported in Figure 4 (left), with up to 128 MPI processes. We can see that the cost associated with orthogonalization of vectors (expressed in compact representation) scales almost linearly. However, the overall scalability is hindered by the factorization and linear solves (red line) that have worse parallel efficiency. It is well known that parallel direct solvers do not scale

Table 3: Computational results for the spectrum slicing solver with symmetric quadratic matrix polynomials (non-hyperbolic). See Table 1 for a description of the columns.

| problem | $n$ | interval | $n_{\text {eig }}$ | $\# \sigma$ | $\eta_{\max }$ | $t(\mathrm{~s})$ |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| spring $(\mathrm{n} / \mathrm{o})$ | 20000 | $[-2,-1.55]$ | 215 | 14 | $6 \cdot 10^{-12}$ | 1.8 |
| sleeper | 1.5 mill. | $[-0.99,-0.97]$ | 234 | 21 | $9 \cdot 10^{-12}$ | 144 |
| atmos $(0.5 \mathrm{~Hz})$ | 40000 | $[0.00628,0.00973]$ | 230 | 18 | $2 \cdot 10^{-13}$ | 3 |
| atmos $(2 \mathrm{~Hz})$ | 40000 | $[0.02513,0.03779]$ | 844 | 67 | $2 \cdot 10^{-13}$ | 11 |



Figure 3: Full spectrum of the three non-hyperbolic problems: non-overdamped spring (left), sleeper (center), and atmos (right). In all cases, the matrix size used to generate the plots is 400 .
well beyond a few tens of processes. For intervals containing many eigenvalues, it would be possible to implement a two-level parallelization scheme, where the interval is split in several subintervals, each of them being processed by a subset of MPI processes. In this way, the number of processes participating in a single factorization is limited. We have implemented this in SLEPc in the spectrum slicing algorithm for linear problems [17], but not yet for quadratic eigenproblems.

We now turn our attention to symmetric QEPs that are not hyperbolic. We consider three problems, the first two belonging to the NLEVP collection. The information associated with these problems is shown in Table 3.

- spring is the same problem considered before, but with modified parameters $\tau=0.6202$ and $\kappa=0.4807$ such that the QEP becomes non-hyperbolic. This problem is referred to as non-overdamped mass-spring system by Li and Cai [19].
- sleeper is a QEP modelling a railtrack resting on sleepers. It is a monic polynomial $Q(\lambda)=$ $\lambda^{2} I+\lambda C+K$, where both $C$ and $K$ are pentadiagonal.
- atmos is the name we use to refer to a QEP from modal analysis of atmospheric infrasound propagation [2]. It corresponds to a novel formulation valid for wide-angle propagation in atmospheres with high Mach number flow, summarized in this equation

$$
\begin{equation*}
\left[k_{H}^{2}\left(\frac{v_{0}^{2}}{c^{2}}-1\right)-k_{H} \frac{2 \omega v_{0}}{c^{2}}+\frac{d^{2}}{d z^{2}}+\frac{\omega^{2}}{c^{2}}\right] \psi=0 . \tag{31}
\end{equation*}
$$

A finite difference discretization is used along a vertical grid (z), resulting in a symmetric QEP with $M, C$ diagonal, and $K$ tridiagonal.


Figure 4: Execution times (in seconds) with up to 128 MPI processes for the spectrum slicing solver on the spring (left) and sleeper (right) problems of dimension 1.5 million. The problem parameters are shown in Table 1 and Table 3, respectively.

Figure 3 shows the full spectrum of these problems, corresponding to a problem of size $n=400$. We can see that there are both real and complex eigenvalues. We have checked that, in these problems, real eigenvalues located on the left of a particular value are of negative type, and those on the right are of positive type. This implies that our spectrum slicing solver should work without difficulties, provided that the requested interval does not extend over eigenvalues of the two groups. The intervals shown in Table 3 satisfy this restriction ${ }^{1}$, and in the case of the atmos problem they represent the range of values with physical interest for the application at the given frequency.

Results in Table 3 show that our solver provides good accuracy also for the general symmetric case, with very fast computation times. Regarding the performance, the same comments in previous paragraphs also apply in this case. In particular, the scalability plot of Figure 4 for the non-hyperbolic case (sleeper, right plot) shows a very similar trend as before.

## 5 Conclusions

We have designed a spectrum slicing method for symmetric (or Hermitian) quadratic eigenvalue problems that, under certain conditions, is able to compute all real eigenvalues contained in a given interval. This method extends the well-known inertia-based spectrum slicing technique for linear symmetric-definite eigenproblems, by relying upon different ingredients such as the pseudo-Lanczos recurrence, the symmetric linearization of quadratic matrix polynomials and the inertia count. Our implementation is also memory-efficient as it stores the basis vectors in a compact representation, which enables significant memory savings at the cost of a more complicated implementation of deflation during the iteration.

Our solver is particularly effective for the special case of hyperbolic QEPs. In this case, the user can specify any interval of interest and the solver returns all the eigenvalues lying inside it. The more general case of symmetric QEPs can also be addressed, but only if the user-provided interval contains eigenvalues of the same definite type (positive or negative). Our solver is also compatible with the Niendorf-Voss transformation [23], so if the QEP has been transformed with this method the spectrum slicing computation will work all the way using a definite inner product, and still use a memory-efficient representation of the basis.

[^1]Regarding the performance of the solver, we have presented results of both serial and parallel executions on various test cases. The parallel scalability is limited by the required factorizations at each shift. As a future work we will consider a two-level parallelism (based on subdivision of the interval) in such a way that the number of processes can be increased further without losing efficiency.

## Acknowledgments

The computational experiments of section 4 were carried out on the supercomputer Tirant 3 belonging to Universitat de València. The matrices from the atmospheric infrasound propagation problem were kindly provided by Jelle Assink [2].

## References

[1] M. Al-Ammari and F. Tisseur. Hermitian matrix polynomials with real eigenvalues of definite type. Part I: Classification. Linear Algebra Appl., 436(10):3954-3973, 2012.
[2] J. Assink, R. Waxler, and D. Velea. A wide-angle high Mach number modal expansion for infrasound propagation. J. Acoust. Soc. Am., 141(3):1781-1792, 2017.
[3] S. Balay, S. Abhyankar, M. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, 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 3.10, Argonne National Laboratory, 2018.
[4] T. Betcke, N. J. Higham, V. Mehrmann, C. Schröder, and F. Tisseur. NLEVP: a collection of nonlinear eigenvalue problems. ACM Trans. Math. Software, 39(2):7:1-7:28, 2013.
[5] C. Campos and J. E. Roman. Strategies for spectrum slicing based on restarted Lanczos methods. Numer. Algorithms, 60(2):279-295, 2012.
[6] 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.
[7] C. Campos and J. E. Roman. Restarted Q-Arnoldi-type methods exploiting symmetry in quadratic eigenvalue problems. BIT, 56(4):1213-1236, 2016.
[8] I. Gohberg, P. Lancaster, and L. Rodman. Spectral analysis of selfadjoint matrix polynomials. Annals of Mathematics, 112(1):33-71, 1980.
[9] I. Gohberg, P. Lancaster, and L. Rodman. Indefinite Linear Algebra and Applications. Birkhaüser, Basel, Switzerland, 2005.
[10] Roger G. Grimes, John G. Lewis, and Horst D. Simon. A shifted block Lanczos algorithm for solving sparse symmetric generalized eigenproblems. SIAM J. Matrix Anal. Appl., 15(1):228-272, 1994.
[11] C.-H. Guo, N. J. Higham, and F. Tisseur. An improved arc algorithm for detecting definite Hermitian pairs. SIAM J. Matrix Anal. Appl., 31(3):1131-1151, 2009.
[12] C.-H. Guo and P. Lancaster. Algorithms for hyperbolic quadratic eigenvalue problems. Math. Comp., 74(252):1777-1791, 2005.
[13] 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.
[14] N. J. Higham, D. S. Mackey, N. Mackey, and F. Tisseur. Symmetric linearizations for matrix polynomials. SIAM J. Matrix Anal. Appl., 29(1):143-159, 2006.
[15] N. J. Higham, D. S. Mackey, and F. Tisseur. Definite matrix polynomials and their linearization by definite pencils. SIAM J. Matrix Anal. Appl., 31(2):478-502, 2009.
[16] Roger A. Horn and Charles R. Johnson. Matrix Analysis. Cambridge University Press, second edition, 2013.
[17] M. Keçeli, F. Corsetti, C. Campos, J. E. Roman, H. Zhang, Á. Vázquez-Mayagoitia, P. Zapol, and A. F. Wagner. SIESTA-SIPs: Massively parallel spectrum-slicing eigensolver for an ab initio molecular dynamics package. J. Comput. Chem., 39(22):1806-1814, 2018.
[18] A. Kostić and H. Voss. On Sylvester's law of inertia for nonlinear eigenvalue problems. Electron. Trans. Numer. Anal., 40:82-93, 2013.
[19] H. Li and Y. Cai. Solving the real eigenvalues of Hermitian quadratic eigenvalue problems via bisection. Electron. J. Linear Algebra, 30(1):721-743, 2015.
[20] R. Li, Y. Xi, E. Vecharynski, C. Yang, and Y. Saad. A thick-restart Lanczos algorithm with polynomial filtering for Hermitian eigenvalue problems. SIAM J. Sci. Comput., 38(4):A2512-A2534, 2016.
[21] D. Lu, Y. Su, and Z. Bai. Stability analysis of the two-level orthogonal Arnoldi procedure. SIAM J. Matrix Anal. Appl., 37(1):195-214, 2016.
[22] Y. Nakatsukasa and V. Noferini. Inertia laws and localization of real eigenvalues for generalized indefinite eigenvalue problems. Technical report, 2017. Preprint arXiv:1711.00495.
[23] V. Niendorf and H. Voss. Detecting hyperbolic and definite matrix polynomials. Linear Algebra Appl., 432(4):1017-1035, 2010.
[24] B. N. Parlett and H. C. Chen. Use of indefinite pencils for computing damped natural modes. Linear Algebra Appl., 140(1):53-88, 1990.
[25] J. E. Roman, C. Campos, E. Romero, and A. Tomas. SLEPc users manual. Technical Report DSIC-II/24/02-Revision 3.9, D. Sistemes Informàtics i Computació, Universitat Politècnica de València, 2018.
[26] M. Rozloz̆ník, F. Okulicka-Dluzewska, and A. Smoktunowicz. Cholesky-like factorization of symmetric indefinite matrices and orthogonalization with respect to bilinear forms. SIAM J. Matrix Anal. Appl., 36(2):727-751, 2015.
[27] F. Tisseur and K. Meerbergen. The quadratic eigenvalue problem. SIAM Rev., 43(2):235286, 2001.
[28] K. Veselić. Damped Oscillations of Linear Systems. Springer-Verlag, Berlin Heidelberg, 2011.
[29] H. Voss, B. Werner, and K. P. Hadeler. A minimax principle for nonlinear eigenvalue problems with applications to nonoverdamped systems. Math. Meth. Appl. Sci., 4(1):415424, 1982.


[^0]:    *This work was supported by the Spanish Agencia Estatal de Investigación (AEI) under project SLEPc-HS (TIN2016-75985-P), which includes European Commission ERDF funds.
    ${ }^{\dagger}$ D. Sistemes Informàtics i Computació, Universitat Politècnica de València, València, Spain (mccampos@dsic.upv.es).
    ${ }^{\ddagger}$ D. Sistemes Informàtics i Computació, Universitat Politècnica de València, València, Spain (jroman@dsic.upv.es).

[^1]:    ${ }^{1}$ We have also considered tests where the requested interval does not satisfy the assumption of having only eigenvalues of the same type, and our solver has detected this situation in all cases.

