

Efficient and Accurate Algorithms for Solving the Bethe-Salpeter Eigenvalue Problem for Crystalline Systems

Peter Benner[†] Carolin Penke^{†*}

[†]Max Planck Institute for Dynamics of Complex Technical Systems,
Sandtorstr. 1, 39106 Magdeburg, Germany.

*Corresponding author. Email: penke@mpi-magdeburg.mpg.de

Optical properties of materials related to light absorption and scattering are explained by the excitation of electrons. The Bethe-Salpeter equation is the state-of-the-art approach to describe these processes from first principles (*ab initio*), i.e. without the need for empirical data in the model. To harness the predictive power of the equation, it is mapped to an eigenvalue problem via an appropriate discretization scheme. The eigenpairs of the resulting large, dense, structured matrix can be used to compute dielectric properties of the considered crystalline or molecular system. The matrix always shows a 2×2 block structure. Depending on exact circumstances and discretization schemes, one ends up with a matrix structure such as

$$H_1 = \begin{bmatrix} A & B \\ -B & -A \end{bmatrix} \in \mathbb{C}^{2n \times 2n}, \quad A = A^H, B = B^H,$$

or

$$H_2 = \begin{bmatrix} A & B \\ -B^H & -A^T \end{bmatrix} \in \mathbb{C}^{2n \times 2n} \text{ or } \mathbb{R}^{2n \times 2n}, \quad A = A^H, B = B^T.$$

Additionally, certain definiteness properties typically hold. H_1 can be acquired for crystalline systems [1], H_2 is a more general form found e.g. in [2] and [3], which can for example be used to study molecules. In this work, we present new theoretical results characterizing the structure of H_1 and H_2 in the language of non-standard scalar products. These results enable us to develop a new perspective on the state-of-the-art solution approach for matrices of form H_1 . This new viewpoint is used to develop two new methods for solving the eigenvalue problem. One requires less computational effort while providing the same degree of accuracy. The other one improves the expected accuracy, compared to methods currently in use, with a comparable performance. Both methods are well suited for high performance environments and only rely on basic numerical linear algebra building blocks.

Keywords: Bethe-Salpeter, Many-body Perturbation Theory, Structured Eigenvalue Problem, Efficient Algorithms, Matrix Square Root, Cholesky Factorization, Singular Value Decomposition

AMS subject classifications: 15A18, 65F15

1 Introduction and Preliminaries

The accurate and efficient computation of optical properties of molecules and condensed matter has been an objective actively pursued in recent years [1, 4, 5]. In particular, the increasing importance of renewable en-

ergies reinforces the interest in the *in silico* prediction of optical properties of novel composite materials and nanostructures.

New theoretical and algorithmic developments need to go hand in hand with the ever advancing computer technology. In view of the ongoing massive increase in parallel computing power [6], the solution of problems that were considered almost impossible just a few years ago comes within reach. In order to unlock the full potential of a supercomputer, great attention must be paid to the development of parallelizable and reliable methods.

Ab initio spectroscopy aims to compute optical properties of materials from first principles, without the need for empirical parameters. A state-of-the-art approach is derived from many-body perturbation theory and relies on solving the Bethe-Salpeter equation for the *density fluctuation response function* $P(\omega)$. This function describes the propagation of an electron-hole pair and is used to compute optical properties such as the optical absorption spectrum. Its poles give the excitation energies of the given system (for details see [1]). Restricting the number of considered occupied and unoccupied orbitals, the propagator can be represented in (frequency-dependent) matrix form with respect to a set of resonant and antiresonant two-orbital states. The Bethe-Salpeter equation can be rewritten to show that the inverse of the matrix form of $P(\omega)$ is approximated by the matrix pencil

$$\mathcal{H} - \omega\Sigma, \quad (1)$$

where $\Sigma = \begin{bmatrix} I_n & 0 \\ 0 & -I_n \end{bmatrix}$ contains a positive and a negative identity matrix on its diagonal. $\mathcal{H} \in \mathbb{C}^{2n \times 2n}$ is a Hermitian matrix. It is computed from the Coulomb interaction and the screened interaction in matrix form (familiar from Hedin's equations [7]) and scalar energy differences between occupied and unoccupied orbitals. The periodicity of crystalline systems implies a time-inversion symmetry in the basis functions used for the discretization. This leads to \mathcal{H} having the following form:

$$\mathcal{H} = \mathcal{H}_1 = \begin{bmatrix} A & B \\ B & A \end{bmatrix}, \quad A = A^H, B = B^H. \quad (2)$$

In this paper we do not consider the generalized eigenvalue problem given in 1. Instead, we focus on the corresponding standard eigenvalue problem of the matrix

$$H_1 := \Sigma^{-1} \mathcal{H}_1 = \Sigma \mathcal{H}_1 = \begin{bmatrix} A & B \\ -B & -A \end{bmatrix}, \quad A = A^H, B = B^H, \quad (3)$$

to which we refer as a *BSE matrix of form I*.

If the time-inversion symmetry in the basis functions is not exploited or not available, i.e. when a non-crystalline system is considered, the resulting structure is slightly different. The matrix considered for a standard eigenvalue problem then turns out to have the form

$$H_2 := \Sigma \mathcal{H}_2 = \begin{bmatrix} A & B \\ -B^H & -A^T \end{bmatrix} \in \mathbb{C}^{2n \times 2n} \text{ or } \mathbb{R}^{2n \times 2n}, \quad A = A^H, B = B^T, \quad (4)$$

where H denotes the Hermitian transpose, and T denotes the regular transpose without complex conjugation. We call matrices of this form *BSE matrix of form II*.

In this paper, we contribute new methods that exploit the structure given in (3) and improve upon previous approaches in terms of performance and accuracy. They are well-suited for high performance computing as they rely on basic linear algebra building blocks for which high performance implementations are readily available.

We characterize the structure of the BSE matrices (3) and (4) by employing the concept of non-standard scalar products. We introduce the notation and concepts following [8].

A nonsingular matrix M defines a scalar product on $\mathbb{C}^n \langle \cdot, \cdot \rangle_M$, which is a bilinear or sesquilinear form, given by

$$\langle x, y \rangle_M = \begin{cases} x^T M y & \text{for bilinear forms,} \\ x^H M y & \text{for sesquilinear forms,} \end{cases}$$

for $x, y \in \mathbb{C}^n$. For a matrix $A \in \mathbb{C}^{n \times n}$, $A^{*M} \in \mathbb{C}^{n \times n}$ denotes the adjoint with respect to the scalar product defined by M . This is a uniquely defined matrix satisfying the identity

$$\langle Ax, y \rangle_M = \langle x, A^{*M}y \rangle_M$$

for all $x, y \in \mathbb{C}^n$. We call A^{*M} the M -adjoint of A and it holds

$$A^{*M} = M^{-1}A^*M, \quad (5)$$

where $*$ can refer to the transpose T or Hermitian transpose H , depending on whether a bilinear or a sesquilinear form is considered. We call the matrix M -orthogonal if $A^{*M} = A^{-1}$ (given the inverse exists), M -self-adjoint if $A = A^{*M}$ and M -skew-adjoint if $A = -A^{*M}$.

2 Results on the spectral structure of BSE matrices

We now have a language to describe the structure of the BSE matrices in a more concise way, not relying on the matrix block structure. The following two matrices, and the scalar products induced by them, play a central role:

$$J_n = \begin{bmatrix} 0 & I_n \\ -I_n & 0 \end{bmatrix}, \quad \Sigma_n = \begin{bmatrix} I_n & 0 \\ 0 & -I_n \end{bmatrix}. \quad (6)$$

We drop the index when the dimension is clear from its context. The identities $J^{-1} = -J$ and $\Sigma^{-1} = \Sigma$ are regularly used in the following.

Theorem 1. *A matrix $H \in \mathbb{C}^{2n \times 2n}$ is a BSE matrix of form I as given in (3) if and only if both of the following conditions hold.*

1. H is skew-adjoint with respect to the complex sesquilinear form induced by J , i.e. $H = -H^{*J} = JH^HJ$.
2. H is self-adjoint with respect to the complex sesquilinear form induced by Σ , i.e. $H = H^{*\Sigma} = \Sigma H^H\Sigma$.

Proof. $JH^HJ = H$ is equivalent to JH being Hermitian and $\Sigma H^H\Sigma = H$ is equivalent to ΣH being Hermitian. We observe that JH and ΣH are Hermitian, if H has BSE form I. Conversely, let $H = \begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix}$, $H_{ij} \in \mathbb{C}^{n \times n}$

and $JH = \begin{bmatrix} H_{21} & H_{22} \\ -H_{11} & -H_{12} \end{bmatrix}$ be Hermitian. It follows

$$H_{21} = H_{21}^H, \quad H_{12} = H_{12}^H, \quad H_{11} = -H_{22}^H. \quad (7)$$

Let $\Sigma H = \begin{bmatrix} H_{11} & H_{12} \\ -H_{21} & -H_{22} \end{bmatrix}$ be Hermitian. It follows

$$H_{11} = H_{11}^H, \quad H_{22} = H_{22}^H, \quad H_{21} = -H_{12}^H. \quad (8)$$

(7) and (8) give exactly BSE form I:

$$H = \begin{bmatrix} H_{11} & H_{12} \\ -H_{12}^H & -H_{11}^H \end{bmatrix} \text{ with } H_{11} = H_{11}^H, H_{12} = H_{12}^H.$$

□

Theorem 2. *A matrix H is a BSE matrix of form II as given in (4) if and only if both of the following conditions hold.*

1. H is skew-adjoint with respect to the complex bilinear form induced by J , i.e. $H = -H^{*J} = JH^TJ$.

2. H is self-adjoint with respect to the complex sesquilinear form induced by Σ , i.e. $H = H^{*\Sigma} = \Sigma H^H \Sigma$.

Proof. The proof works exactly as the proof of Theorem 1, but here, the complex transpose \cdot^\top is associated with J instead of the Hermitian transpose \cdot^H . \square

Using this new characterization, we see that eigenvalues and eigenvectors also exhibit special structures. Matrices, that are skew-adjoint with respect to the sesquilinear form induced by J are called *Hamiltonian*, and play an important role in control theory and model order reduction (see e.g. [9]). The same property with respect to the bilinear form is called *J-symmetric* in [8] and explored further in [10].

The first two propositions of the following theorem are well known facts about Hamiltonian [9] and *J*-symmetric matrices [10]. Here, $\bar{\lambda}$ denotes the complex conjugate of λ .

Theorem 3. Let $H \in \mathbb{C}^{2n \times 2n}$.

1. If H is skew-adjoint with respect to the sesquilinear form induced by J , i.e. $JH = -H^H J$, then its eigenvalues come in pairs $(\lambda, -\bar{\lambda})$. If x is a right eigenvector of H corresponding to λ , then $x^H J$ is the left eigenvector of H corresponding to $-\bar{\lambda}$.
2. If H is skew-adjoint with respect to the bilinear form induced by J , i.e. $JH = -H^\top J$, then its eigenvalues come in pairs $(\lambda, -\lambda)$. If x is a right eigenvector of H corresponding to λ , then $x^\top J$ is the left eigenvector of H corresponding to $-\lambda$.
3. If H is self-adjoint with respect to the sesquilinear form induced by Σ , i.e. $\Sigma H = H^H \Sigma$, then eigenvalues come in pairs $(\lambda, \bar{\lambda})$. If x is a right eigenvector of H corresponding to λ , then $x^H \Sigma$ is the left eigenvector of H corresponding to $\bar{\lambda}$.

Proof. 1. Using $H^H = JHJ$ and $J^{-1} = -J$, we see that

$$Hx = \lambda x \quad \Leftrightarrow \quad x^H JH = -\bar{\lambda} x^H J.$$

2. Using $H^\top = JHJ$ and $J^{-1} = -J$, we see that

$$Hx = \lambda x \quad \Leftrightarrow \quad x^\top JH = -\lambda x^\top J.$$

3. Using $H^H = \Sigma H \Sigma$ and $\Sigma^{-1} = \Sigma$, we see that

$$Hx = \lambda x \quad \Leftrightarrow \quad x^H \Sigma H = \bar{\lambda} x^H \Sigma.$$

\square

Theorem 3 reveals that symmetries defined by the matrices J or Σ are reflected in connections between left and right eigenvectors of the considered matrix. The BSE matrices show a symmetry with respect to two scalar products (Theorem 1 and 2). This double-structure leads to eigenvalues that show up not only in pairs but in quadruples if they have a real and an imaginary component. Additionally, it yields a connection between right eigenvectors, clarified in the following theorem.

Theorem 4. Let $H \in \mathbb{C}^{2n \times 2n}$ be self-adjoint with respect to Σ and skew-adjoint with respect to (a) the sesquilinear scalar product or (b) the bilinear scalar product induced by J . Then

1. The eigenvalues of H come in pairs $(\lambda, -\lambda)$ if $\lambda \in \mathbb{R}$ or $\lambda \in i\mathbb{R}$, or in quadruples $(\lambda, \bar{\lambda}, -\lambda, -\bar{\lambda})$.
2. a) If v is an eigenvector of H with respect to λ , then $J\Sigma v$ is an eigenvector of H with respect to $-\lambda$.
b) If v is an eigenvector of H with respect to λ , then $J\Sigma \bar{v}$ is an eigenvector of H with respect to $-\bar{\lambda}$.

Proof. 1. The quadruple property comes from combining the propositions given in Theorem 3 (1. and 3. or 2. and 3., respectively). The pair property for real eigenvalues comes (a) from Theorem 3, proposition 1, or (b) from Theorem 3, proposition 2. The pair property for imaginary eigenvalues follows from Theorem 3, proposition 3 in both cases.

2. a) With $\Sigma JHJ\Sigma = H$ we have

$$Hv = \lambda v \Leftrightarrow HJ\Sigma v = -\lambda J\Sigma v.$$

b) With $\Sigma JHJ\Sigma = \bar{H}$ we have

$$Hv = \lambda v \Leftrightarrow \bar{H}J\Sigma v = -\lambda J\Sigma v \Leftrightarrow HJ\Sigma \bar{v} = -\bar{\lambda} J\Sigma \bar{v}.$$

□

The special case of (b) (i.e. for BSE matrices of form II) has been proven in [10]. Our proof does not rely on the particular block structure of the matrix, but works with the given symmetries and is therefore more concise and easily extendable to other double-structured matrices.

In the practice of computing excitation properties of materials, there is even more structure available. It can be exploited for devising efficient algorithms. The Hermitian matrices

$$\mathcal{H}_1 := \Sigma H_1 = \begin{bmatrix} A & B \\ B & A \end{bmatrix}, \quad \mathcal{H}_2 := \Sigma H_2 = \begin{bmatrix} A & B \\ B^H & A^T \end{bmatrix}, \quad (9)$$

introduced in (1), are called BSE Hamiltonians. They are typically positive definite [1, 11]. The term ‘‘Hamiltonian’’ might cause confusion in this context, as it has a different meaning in numerical linear algebra and in electronic structure theory. We have used the numerical linear algebra meaning, by which a matrix H is called Hamiltonian if it holds $JH = -JH^H$. In electronic structure theory, the term ‘‘Hamiltonian’’ is inspired by the Hamiltonian operator from basic quantum mechanics. It refers to the left matrix \mathcal{H} of a generalized, Schrödinger-like, eigenvalue problem, such as (1), which is typically Hermitian.

The definiteness property has consequences for the structure of the eigenvalue spectrum. To study these, we consider the more general class of Σ -Hermitian matrices, by which we mean matrices that are self-adjoint with respect to the scalar product induced by a signature matrix Σ .

Theorem 5. *Let $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$, $\sigma_i \in \{+1, -1\}$ be a signature matrix with p positive and $n - p$ negative diagonal entries. Let $H \in \mathbb{C}^{n \times n}$ be given such that ΣH is Hermitian positive definite. Then H is diagonalizable and its eigenvalues are real, of which p are positive and $n - p$ are negative.*

Proof. As ΣH is positive definite, and Σ is symmetric, they can be diagonalized simultaneously (see [12], Corollary 8.7.2), i.e. there is a nonsingular $X \in \mathbb{C}^{n \times n}$, s.t

$$X^H \Sigma H X = I_n, \quad (10)$$

$$X^H \Sigma X = \Lambda \in \mathbb{R}^{n \times n}, \quad (11)$$

where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ gives the eigenvalues of the matrix pencil $\Sigma x - \lambda \Sigma H$. It follows from (11) and Sylvester’s law of inertia that Λ has p positive and $n - p$ negative values. We have

$$X^{-1} H X = \Lambda^{-1}, \quad (12)$$

i.e. H is diagonalizable and Λ^{-1} contains the eigenvalues of H . □

The spectral structure of the BSE matrices given in practice follows immediately from the presented theorems and is summarized in the following lemma.

Lemma 6. *Let H be a BSE matrix of form I (see (3)) or form II (see (4)), such that the BSE Hamiltonian (9) is positive definite. Then the eigenvalues are real and come in pairs $\pm \lambda$. If v is an eigenvector associated with λ , then*

$$1. \ y = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix} v \text{ is an eigenvector associated with } -\lambda \text{ if } H \text{ is of BSE form I,}$$

2. $y = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix} \bar{v}$ is an eigenvector associated with $-\lambda$ if H is of BSE form II.

In the remaining part of this section we focus on BSE matrices of form I, paving the way for new, efficient algorithms. Two essential observations to make the problem more tractable are the following, which can e.g. be found in [2], albeit for real matrices.

Lemma 7. Let H be a BSE matrix of form I (3).

1. With the matrix $Q = \frac{1}{2} \begin{bmatrix} I & I \\ -I & I \end{bmatrix}$ we have

$$Q^{-1}HQ = \begin{bmatrix} 0 & A-B \\ A+B & 0 \end{bmatrix}. \quad (13)$$

2. ΣH is positive definite if and only if $A+B$ and $A-B$ are positive definite.

The following theorem plays the central role in this paper. We use it as a new tool to derive an existing solution approach. Within this new framework, other solution approaches will become apparent, yielding significant benefits compared to the existing approach.

Theorem 8. Let H be a BSE matrix of form I (see (3)) with a positive definite BSE Hamiltonian (9). Then $M_1 := A+B$, $M_2 = A-B$ are Hermitian positive definite. Let

$$M_1 M_2 v_1 = \mu v_1, \quad v_2^H M_1 M_2 = \mu v_2^H, \quad v_1^H v_2 = 1, \quad (14)$$

define a pair of right and left eigenvectors of the matrix product $M_1 M_2$. Then $\mu \in \mathbb{R}$, $\mu > 0$. With $Q := \frac{1}{2} \begin{bmatrix} I & I \\ -I & I \end{bmatrix}$ and scaling factors $\lambda_1 := v_2^H M_1 v_2 > 0$, $\lambda_2 := v_1^H M_2 v_1 > 0$, an eigenpair of H is given by

$$\lambda = \sqrt{\mu}, \quad v_\lambda = Q \begin{bmatrix} v_1 \lambda_1^{\frac{1}{4}} \lambda_2^{-\frac{1}{4}} \\ v_2 \lambda_1^{-\frac{1}{4}} \lambda_2^{\frac{1}{4}} \end{bmatrix}, \quad (15)$$

i.e. $Hv_\lambda = \lambda v_\lambda$. If γ is another eigenvalue of $M_1 M_2$ and v_γ is the corresponding constructed vector, it holds

$$v_\lambda^H \Sigma v_\gamma = \begin{cases} 1 & \text{if } \lambda = \gamma \\ 0 & \text{if } \lambda \neq \gamma. \end{cases} \quad (16)$$

Proof. We observe $(Q^{-1}HQ)^2 = \text{diag}(M_1 M_2, M_2 M_1)$. Therefore the eigenvalues of $M_1 M_2$ must be a subset of the eigenvalues of H^2 . These are positive real according to Theorem 5. Note that v_2 being a left eigenvector of $M_1 M_2$ is equivalent to v_2 being a right eigenvector of $M_2 M_1$, because M_1 and M_2 are Hermitian. It follows from (14) that

$$M_1 M_2 M_1 v_2 = \mu M_1 v_2,$$

i.e. $M_1 v_2$ lies in the eigenspace of $M_1 M_2$ corresponding to μ . If $M_1 v_2$ is not a multiple of v_1 it must hold $(M_1 v_2)^H v_2 = 0$, which contradicts the fact, that M_1 is positive definite. So there is $\lambda_1 \in \mathbb{C}$, s.t.

$$M_1 v_2 = \lambda_1 v_1. \quad (17)$$

Similarly it follows from (14) that there is $\lambda_2 \in \mathbb{C}$, s.t.

$$M_2 v_1 = \lambda_2 v_2. \quad (18)$$

Following from (17) and (18) and using $v_1^H v_2 = 1$, λ_1 and λ_2 can be computed as

$$\lambda_1 = v_2^H M_1 v_2, \quad \lambda_2 = v_1^H M_2 v_1. \quad (19)$$

It follows that $\lambda_1, \lambda_2 \in \mathbb{R}_+$ because M_1 and M_2 are positive definite. Inserting (18) in (17) we get

$$M_1 M_2 v_1 = \lambda_1 \lambda_2 v_1. \quad (20)$$

With (14) it follows $\lambda_1 \lambda_2 = \mu$. For $\lambda \in \mathbb{C}$,

$$\begin{bmatrix} 0 & M_1 \\ M_2 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \lambda \begin{bmatrix} x \\ y \end{bmatrix} \quad (21)$$

holds if and only if

$$M_1 y = \lambda x \quad \text{and} \quad M_2 x = \lambda y. \quad (22)$$

This is achieved by $\lambda := \sqrt{\lambda_1 \lambda_2} = \sqrt{\mu}$ and $x := v_1 \lambda_1^{\frac{1}{4}} \lambda_2^{-\frac{1}{4}}$ and $y := v_2 \lambda_1^{-\frac{1}{4}} \lambda_2^{\frac{1}{4}}$, following from (17) and (18).

In conclusion we have

$$H v_\lambda = H Q \begin{bmatrix} x \\ y \end{bmatrix} = Q \begin{bmatrix} 0 & M_1 \\ M_2 & 0 \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} = \lambda Q \begin{bmatrix} x \\ y \end{bmatrix} = \lambda v_\lambda.$$

The Σ -orthogonality condition (16) remains to be shown. We observe $Q^H \Sigma Q = \frac{1}{2} \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}$ and with

$$v_\lambda = Q \begin{bmatrix} x_\lambda \\ y_\lambda \end{bmatrix}, \quad x_\lambda = v_{1,\lambda} \lambda_1^{\frac{1}{4}} \lambda_2^{-\frac{1}{4}}, \quad y_\lambda = v_{2,\lambda} \lambda_1^{-\frac{1}{4}} \lambda_2^{\frac{1}{4}}, \quad (23)$$

$$v_\gamma = Q \begin{bmatrix} x_\gamma \\ y_\gamma \end{bmatrix}, \quad x_\gamma = v_{1,\gamma} \gamma_1^{\frac{1}{4}} \gamma_2^{-\frac{1}{4}}, \quad y_\gamma = v_{2,\gamma} \gamma_1^{-\frac{1}{4}} \gamma_2^{\frac{1}{4}}, \quad (24)$$

we see

$$v_\lambda^H \Sigma v_\gamma = \frac{1}{2} (v_{2,\lambda}^H v_{1,\gamma} \bar{\lambda}_1^{-\frac{1}{4}} \bar{\lambda}_2^{\frac{1}{4}} \gamma_1^{\frac{1}{4}} \gamma_2^{-\frac{1}{4}} + v_{1,\lambda}^H v_{2,\gamma} \bar{\lambda}_1^{\frac{1}{4}} \bar{\lambda}_2^{-\frac{1}{4}} \gamma_1^{-\frac{1}{4}} \gamma_2^{\frac{1}{4}}). \quad (25)$$

This expression is equal to 0, if $v_{2,\lambda}$ and $v_{1,\gamma}$ are left and right eigenvectors corresponding to different eigenvalues $\lambda \neq \gamma$ of $M_1 M_2$. If $\lambda = \gamma$ and if v_λ and v_γ were constructed in the same way, we have $\lambda_1 = \gamma_1$ and $\lambda_2 = \gamma_2$. Using that λ_1 and λ_2 are real, (25) simplifies to

$$v_\lambda^H \Sigma v_\lambda = \frac{1}{2} (v_{2,\lambda}^H v_{1,\lambda} + v_{1,\lambda}^H v_{2,\lambda}) = 1, \quad (26)$$

where we used the normalization of the vectors given in (14). \square

3 Algorithms for crystalline systems (Form I)

We have seen in Theorem 8 that the Bethe-Salpeter eigenvalue problem of form I with size $2n \times 2n$ can be interpreted as a product eigenvalue problem with two Hermitan factors of size $n \times n$. In practice, the complete set of eigenvectors provides additional insight to excitonic effects. To compute them, left and right eigenvectors of the smaller product eigenvalue problem are needed. Product eigenvalue problems are well studied (see e.g. [13]). A general way to solve these problems, taking the product structure into account to improve numerical properties, is the periodic QR algorithm. This tool can be used for solving general Hamiltonian eigenvalue problems [14]. In this work we focus on non-iterative methods that work for Hermitian factors and transform the problem such that it can be treated using available HPC libraries.

The algorithms presented in this section compute the positive part of the spectrum of a BSE matrix and the corresponding eigenvectors. If the eigenvectors corresponding to the negative mirror eigenvalues are of interest, they can easily be computed by employing Lemma 6.

3.1 Square root approach

A widely used approach for solving the Bethe-Salpeter eigenvalue problem of form I (e.g. in [1]) relies on the computation of the matrix square root of $M_2 = A - B$. We present it in the following, relating it to the framework given by Theorem 8.

Starting with the first equation of (14), we see

$$M_1 M_2 V_1 = V_1 \Lambda^2 \quad (27)$$

$$\Leftrightarrow M_2^{\frac{1}{2}} M_1 M_2^{\frac{1}{2}} \hat{V} = \hat{V} \Lambda^2, \quad (28)$$

where $\hat{V} = M_2^{\frac{1}{2}} V_1$ contains the eigenvectors of the Hermitian matrix $M_2^{\frac{1}{2}} M_1 M_2^{\frac{1}{2}}$.

On the other hand, with the second equation of (14), we see

$$V_2^H M_1 M_2 = \Lambda^2 V_2^H \quad (29)$$

$$\Leftrightarrow \hat{V}^H M_2^{\frac{1}{2}} M_1 M_2^{\frac{1}{2}} = \Lambda^2 \hat{V}^H \quad (30)$$

where $\hat{V} = M_2^{-\frac{1}{2}} V_2$ contains the left eigenvectors of $M_2^{\frac{1}{2}} M_1 M_2^{\frac{1}{2}}$. Here, note that because $M_2^{\frac{1}{2}} M_1 M_2^{\frac{1}{2}}$ is Hermitian, left and right eigenvectors coincide, we can denote both left and right eigenvector matrices by \hat{V} .

If the generalized eigenvalue problem (14) is solved in this particular way, we can say more about the resulting scalar factors λ_1 and λ_2 in (17) and (18).

Lemma 9. *Let M_1 and M_2 be given as in Theorem 8, Λ^2 be a diagonal matrix containing the eigenvalues of $M_1 M_2$ and \hat{V} contain the eigenvectors of $M_2^{\frac{1}{2}} M_1 M_2^{\frac{1}{2}}$ and $\hat{V}^H \hat{V} = I$. Then $V_1 := M_2^{-\frac{1}{2}} \hat{V}$, $V_2 := M_2^{\frac{1}{2}} \hat{V}$ fulfill*

$$V_1^H V_2 = I, \quad M_2 V_1 = V_2, \quad M_1 V_2 = V_1 \Lambda^2. \quad (31)$$

Proof. The first statement is immediately obvious from the normalization of \hat{V} and because M_2 is Hermitian. Starting with (28) we see

$$M_2^{\frac{1}{2}} M_1 M_2^{\frac{1}{2}} \hat{V} = \hat{V} \Lambda^2 \Leftrightarrow M_1 V_2 = V_1 \Lambda^2.$$

Using $M_1 V_2 = V_1 \Lambda^2$, (28) also yields

$$M_2^{\frac{1}{2}} M_1 M_2^{\frac{1}{2}} \hat{V} = \hat{V} \Lambda^2 \Leftrightarrow M_2 V_1 = V_2.$$

□

Lemma 9 states that the scaling factors in (17) and (18) are given by $\lambda_1 = \lambda^2$ and $\lambda_2 = 1$ in this case. Theorem 8 therefore suggests to scale the acquired eigenvectors in order to obtain eigenvectors of the full matrix as $v := Q \begin{bmatrix} v_1 \lambda^{\frac{1}{2}} \\ v_2 \lambda^{-\frac{1}{2}} \end{bmatrix}$.

These observations suggest Algorithm 1, where the left and right eigenvectors needed in Theorem 8 are computed from the eigenvectors of the Hermitian matrix $M_2^{\frac{1}{2}} M_1 M_2^{\frac{1}{2}}$.

The essential work of this algorithm is the computation of the matrix square root (Step 1) and the solution of a Hermitian eigenvalue problem (Step 2). Computing the (principal) square root of a matrix is a nontrivial task (see e.g. [15], Chapter 6) with a (perhaps surprisingly) high computational demand. Its efficient computation has been an active area of research. Given a Hermitian positive definite matrix C , its principal square root S , s.t. $S^2 = C$, can be computed by diagonalizing $M = V_C D_C V_C^H$, and taking the square roots of the diagonal entries of D_C . Then $S := V_C \sqrt{D_C} V_C^H$.

The main computational effort of Algorithm 1 is therefore the subsequent solution of two Hermitian eigenvalue problems.

Algorithm 1 Compute eigenvectors of BSE matrix of form I, using the matrix square root.

Input: $A = A^H \in \mathbb{C}^{n \times n}$, $B = B^H \in \mathbb{C}^{n \times n}$, defining a BSE matrix of form I: $H = \begin{bmatrix} A & B \\ -B & -A \end{bmatrix}$, s.t. $A + B$ and

$A - B$ are positive definite.

Output: $V \in \mathbb{C}^{2n \times n}$, $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n) \in \mathbb{R}_+^{n \times n}$ s.t. $HV = V\Lambda$.

1: $S \leftarrow (A - B)^{\frac{1}{2}}$

2: Compute eigendecomposition of Hermitian positive definite matrix $M := S(A + B)S$, i.e.

$$MV_M = V_M D,$$

where V_M contains the normalized eigenvectors of M and $D = \text{diag}(d_1, \dots, d_n) \in \mathbb{R}_+^{n \times n}$ the eigenvalues of M .

3: $\Lambda \leftarrow D^{\frac{1}{2}}$

4: $V_1 \leftarrow S^{-1}V_M\Lambda^{\frac{1}{2}}$

5: $V_2 \leftarrow SV_M\Lambda^{-\frac{1}{2}}$

6: $V \leftarrow \begin{bmatrix} \frac{1}{\sqrt{2}}(V_1 + V_2) \\ \frac{1}{\sqrt{2}}(V_2 - V_1) \end{bmatrix}$.

3.2 Cholesky factorization approach

We now lay out how the product eigenvalue problem given in (14) can be solved by using Cholesky factorizations. Let the Cholesky factorization $M_2 = LL^H$ be given.

Starting with (14) we see

$$M_1 M_2 V_1 = V_1 \Lambda^2 \quad (32)$$

$$\Leftrightarrow L^H M_1 L \hat{V} = \hat{V} \Lambda^2, \quad (33)$$

where $\hat{V} = L^H V_1$ contains the eigenvectors of the Hermitian matrix $L^H M_1 L$.

On the other hand, with the second equation of (14), we have

$$V_2^H M_1 M_2 = \Lambda^2 V_2^H \quad (34)$$

$$\Leftrightarrow \hat{V}^H L^H M_1 L = \Lambda^2 \hat{V}^H, \quad (35)$$

with $\hat{V} = L^{-1} V_2$ as the eigenvectors of $L^H M_1 L$.

The analogy to Lemma 9 is given in the following, which can be proven in the same way.

Lemma 10. Let M_1 and M_2 be given as in Theorem 8, $M_2 = LL^H$ be a Cholesky decomposition, Λ^2 be a diagonal matrix containing the eigenvalues of $M_1 M_2$ and \hat{V} contain eigenvectors of $L^H M_1 L$ and $\hat{V}^H \hat{V} = I$. Then $V_2 := L \hat{V}$, $V_1 := V_2^{-H} = L^{-H} \hat{V}$ fulfill

$$V_1^H V_2 = I, \quad M_2 V_1 = V_2, \quad M_1 V_2 = V_1 \Lambda^2. \quad (36)$$

This suggests the same scaling factors as in Algorithm 1, leading to Algorithm 2. The same key idea is used in the standard approach for solving generalized symmetric-definite eigenvalue problems [12], implemented in various libraries [16, 17].

Comparing Algorithm 1 to 2, we see that the essential work in both algorithms is solving Hermitian $n \times n$ eigenvalue problems. The Cholesky variant (Algorithm 2) solves one explicitly at Step 2. The square root variant (Algorithm 1) solves one for computing the matrix square root, which is then used to set up the matrix for the second eigenvalue problem. Then both left and right eigenvectors of the product eigenvalue problem can be inferred from the computed ones.

Algorithm 2 Compute eigenvectors of BSE matrix of form I, using a Cholesky factorization.

Input: $A = A^H \in \mathbb{C}^{n \times n}$, $B = B^H \in \mathbb{C}^{n \times n}$, defining a BSE matrix of form $IH = \begin{bmatrix} A & B \\ -B & -A \end{bmatrix}$, s.t. $A + B$ and $A - B$

are positive definite.

Output: $V \in \mathbb{C}^{2n \times n}$, $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n) \in \mathbb{R}_+^{n \times n}$ s.t. $HV = V\Lambda$.

- 1: Compute Cholesky factorization $LL^H = A - B$.
- 2: Compute eigendecomposition for $M = L^H(A + B)L$

$$MV_M = V_M D$$

- 3: $\Lambda \leftarrow D^{\frac{1}{2}}$
 - 4: $V_1 \leftarrow L^{-H} V_M \Lambda^{\frac{1}{2}}$
 - 5: $V_2 \leftarrow L V_M \Lambda^{-\frac{1}{2}}$
 - 6: $V \leftarrow \begin{bmatrix} \frac{1}{\sqrt{2}}(V_1 + V_2) \\ \frac{1}{\sqrt{2}}(V_2 - V_1) \end{bmatrix}$.
-

3.3 Singular Value Decomposition approach

Both the square root approach discussed in Section 3.1 and the Cholesky approach discussed in Section 3.2 compute the squared eigenvalues of the original problem. In the numerical linear algebra community this procedure is well known to limit the attainable accuracy [18].

The methods essentially work on the (transformed) matrix product $M_1 M_2$. It corresponds to the squared matrix H^2 , as

$$Q^H H^2 Q = \begin{bmatrix} M_1 M_2 & \\ & M_2 M_1 \end{bmatrix}, \quad \text{where } Q := \frac{1}{\sqrt{2}} \begin{bmatrix} I & I \\ -I & I \end{bmatrix}.$$

See also Lemma 7 and the proof of Theorem 8. Now the scaling factor of Q is different to ensure its orthogonality. H belongs to the class of Hamiltonian matrices (see Section 2). When the eigenvalues are computed from the squared matrix H^2 , employing a backward-stable method, the computational error can be approximated using first-order perturbation theory [18–20]. It is given as

$$|\lambda - \hat{\lambda}| \approx \varepsilon \frac{\|H\|_2}{s(\lambda)} \min \left\{ \frac{\|H\|_2}{\lambda}, \frac{1}{\sqrt{\varepsilon}} \right\}, \quad (37)$$

where λ denotes an exact eigenvalue of H , $\hat{\lambda}$ the corresponding computed value, $s(\lambda)$ the condition number of the eigenvalue, and ε the machine precision. Unless λ is very large, the expression is dominated by $\frac{\sqrt{\varepsilon}\|H\|_2}{s(\lambda)}$. Essentially, the number of significant digits of the eigenvalues is halved, compared to direct backward-stable methods. For example, applying the QR algorithm on the original matrix H would yield an approximate error of $\frac{\varepsilon\|H\|_2}{s(\lambda)}$. It fails however, to preserve and exploit the structure of the problem and is undesirable from a numerical as well as from a performance point of view.

A remedy is given in the approach discussed in this section, making use of the singular value decomposition (SVD).

Given the Cholesy factorizations $L_1 L_1^H = M_1$, $L_2 L_2^H = M_2$ and the SVD $U \Lambda V^H = L_1 L_2^H$, we observe that Λ contains the eigenvalues of the BSE matrix, i.e. the square roots of the eigenvalues of the matrix product $M_1 M_2$. The details of the eigenvector computation are given in the following Lemma.

Lemma 11. *Let M_1 and M_2 be given as in Theorem 8, $L_1 L_1^H = M_1$, $L_2 L_2^H = M_2$ be Cholesky factorizations, and $L_1^H L_2 = U \Lambda V^H$ be a singular value decomposition. Then $V_1 := L_1 U \Lambda^{-\frac{1}{2}}$, $V_2 := L_2 V \Lambda^{-\frac{1}{2}}$ fulfill*

$$V_1^H V_2 = I, \quad M_2 V_1 = V_2 \Lambda, \quad M_1 V_2 = V_1 \Lambda. \quad (38)$$

Proof. It holds

$$V_1^H V_2 = \Lambda^{-\frac{1}{2}} U^H L_1^H L_2 V \Lambda^{-\frac{1}{2}} = \Lambda^{-\frac{1}{2}} U^H U \Lambda V^H V \Lambda^{-\frac{1}{2}} = I$$

and

$$M_2 V_1 = M_2 L_1 U \Lambda^{-\frac{1}{2}} = L_2 L_2^H L_1 U \Lambda^{-\frac{1}{2}} = L_2 V \Lambda U^H U \Lambda^{-\frac{1}{2}} = L_2 V \Lambda^{\frac{1}{2}} = V_2 \Lambda.$$

$M_1 V_2 = V_1 \Lambda$ is proved in the same way. \square

Lemma 11 states that the scaling described in (15) boils down to a scaling factor of 1 as $\lambda_1 = \lambda_2$. This suggests Algorithm 3.

Algorithm 3 Compute eigenvectors of BSE matrix of form I, using the singular value decomposition.

Input: $A = A^H \in \mathbb{C}^{n \times n}$, $B = B^H \in \mathbb{C}^{n \times n}$, defining a BSE matrix of form I $H = \begin{bmatrix} A & B \\ -B & -A \end{bmatrix}$, s.t. $A + B$ and $A - B$ are positive definite.

Output: $V \in \mathbb{C}^{2n \times n}$, $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n) \in \mathbb{R}_+^{n \times n}$ s.t. $HV = V\Lambda$.

- 1: Compute Cholesky factorization $L_1 L_1^H = M_1$.
 - 2: Compute Cholesky factorization $L_2 L_2^H = M_2$.
 - 3: Compute singular value decomposition $U_{SVD} \Lambda V_{SVD}^H = L_1^H L_2$
 - 4: $V_1 \leftarrow L_1 U_{SVD} \Lambda^{-\frac{1}{2}}$
 - 5: $V_2 \leftarrow L_2 V_{SVD} \Lambda^{-\frac{1}{2}}$
 - 6: $V \leftarrow \begin{bmatrix} \frac{1}{2}(V_1 + V_2) \\ \frac{1}{2}(V_2 - V_1) \end{bmatrix}$.
-

The main difference between the SVD-based algorithm (Algorithm 3) and the other ones, from a numerical point of view, is that the eigenvalue matrix Λ is computed directly by the SVD and not as a square root of another diagonal matrix D .

The way real BSE matrices of form II (4) are treated in [2] is based on the same idea.

We can expect to see a higher accuracy in the eigenvalues than in the square root and the Cholesky approach, because the eigenvalues are computed directly, using a backward-stable method for the singular value decomposition. Perturbation theory [21] yields an approximate error of

$$|\lambda - \hat{\lambda}| \approx \varepsilon \frac{\|H\|_2}{s(\lambda)}. \quad (39)$$

In the other approaches, a similar approximation only holds for the error of the squared eigenvalues λ^2 , and translates in form of (37) to the non-squared ones.

3.4 Comparison

In recent years, various packages have been developed to facilitate the computation of the electronic structure of materials. See e.g. [22–24], or <https://www.nomad-coe.eu/externals/codes> for an overview. In particular, computing excited states via methods based on many-body-perturbation theory has come into focus, as powerful computational resources become more widely available. Here, the Bethe-Salpeter approach constitutes a state-of-the-art method for computing optical properties such as the optical absorption spectrum. To this end, Algorithm 1 is typically used to solve the resulting eigenvalue problem after the matrices A and B have been set up [1].

The main contribution of the previous section was to provide a unified frame of reference, which can be used to derive the existing approach (Algorithm 1) as well as two new ones (Algorithm 2 and Algorithm 3). Due to this unified framework, the similarities between the realizations of the different approaches become apparent. In all algorithms we clearly see four steps.

1. Preprocessing: Setup a matrix M .
2. Decomposition: Compute spectral, respectively, singular value decomposition of M .
3. Postprocessing: Transform resulting vectors to (left and right) eigenvectors of matrix $(A + B)(A - B)$.
4. Final setup: Form eigenvectors of original BSE matrix.

A detailed compilation is given in Table 1.

	SQRT (Alg. 1)	CHOL (Alg. 2)	CHOL+SVD (Alg. 3)
1. Preprocessing	$S = (A - B)^{\frac{1}{2}}, \quad (11n^3)$ $M = S(A + B)S \quad (4n^3)$	$LL^H = A - B, \quad (\frac{1}{3}n^3)$ $M = L(A + B)L^H \quad (2n^3)$	$L_1L_1^H = A - B, \quad (\frac{1}{3}n^3)$ $L_2L_2^H = A + B, \quad (\frac{1}{3}n^3)$ $M = L_1^H L_2 \quad (n^3)$
2. Decomposition	$M = V_M \Lambda^2 V_M^H \quad (9n^3)$	$M = V_M \Lambda^2 V_M^H \quad (9n^3)$	$M = U_{SVD} \Lambda V_{SVD}^H \quad (21n^3)$
3. Postprocessing	$V_1 := S^{-1} V_M \Lambda^{\frac{1}{2}}, \quad (\frac{8}{3}n^3)$ $V_2 = S V_M \Lambda^{-\frac{1}{2}} \quad (2n^3)$	$V_1 = L^{-H} V_M \Lambda^{\frac{1}{2}}, \quad (n^3)$ $V_2 = L V_M \Lambda^{-\frac{1}{2}} \quad (n^3)$	$V_1 = L_1 U_{SVD} \Lambda^{-\frac{1}{2}}, \quad (n^3)$ $V_2 = L_2 V_{SVD} \Lambda^{-\frac{1}{2}} \quad (n^3)$
4. Final setup	$V = \begin{bmatrix} \frac{1}{2}(V_1 + V_2) \\ \frac{1}{2}(V_2 - V_1) \end{bmatrix}.$		

Table 1: Algorithmic steps of the different methods. The number in brackets estimates the number of flops, where lower-order terms are neglected.

Seeing the algorithms side by side enables a direct comparison. The amount of flops is based on estimates for sequential, non blocked implementations [12], and lower order terms i.e. $\mathcal{O}(n^2)$ and $\mathcal{O}(n)$, are neglected. The preprocessing step is most expensive in the square root approach. Computing the square root of a Hermitian matrix involves the solution of a Hermitian eigenvalue problem. Additionally, the matrices S and M need to be set up, using 3 matrix-matrix products. This makes the preprocessing step even more expensive as the following “main” eigenvalue computation. The CHOL and the CHOL+SVD approach, on the other hand, only rely on one or two Cholesky factorizations and matrix multiplications, which are comparatively cheap to realize. The computational effort in the decomposition step is the highest in the CHOL+SVD step. The post-processing step again is most expensive in the SQRT approach, because the matrix S is a general square matrix, while the L matrices in CHOL and CHOL+SVD are triangular. In total, SQRT takes an estimated amount of $28\frac{2}{3}n^3$ flops, CHOL takes $13\frac{1}{3}n^3$ flops and CHOL+SVD takes $24\frac{2}{3}n^3$ flops. The classical QR algorithm applied to the full, non-Hermitian matrix takes about $25(2n)^3 = 200n^3$ flops (not including the computation of eigenvectors from the Schur vectors). Solving the Hermitian-definite eigenvalue problem (1) can exploit symmetry, but still acts on the large problem and can be expected to perform $14(2n)^3 = 112n^3$ flops.

According to this metric, we expect both new approaches to perform faster than the square root approach. The actual performance of algorithms on modern architectures is not simply determined by the number of operations performed, but by their parallelizability and communication costs. All presented approaches have a high computational intensity of $\mathcal{O}(n^3)$, such that the memory bandwidth is not likely to be a bottleneck. All methods rely on the same standard building blocks from numerical linear algebra, for which optimized versions (e.g. blocked variants for cache-efficiency) are available. This setting makes a fair comparison possible where the arithmetic complexity has a high explanatory power.

To summarize, we expect CHOL to be about twice as fast as SQRT, while keeping the same accuracy. CHOL+SVD performs more computations than CHOL, and will take more time, but could improve the accuracy of the computations. It might be faster than SQRT, depending on how efficient the diagonalizations in SQRT and the SVD in CHOL+SVD are implemented.

Method	Relative Error				Runtime
	$\kappa = 10$	$\kappa = 10^3$	$\kappa = 10^6$	$\kappa = 10^9$	
eig	1.28e-14	5.08e-14	3.82e-11	1.26e-08	62.7 ms
generalized eig	7.89e-15	6.67e-15	1.89e-11	1.97e-09	10.7 ms
haeig	4.73e-15	7.82e-15	4.32e-11	2.23e-08	50.9 ms
SQRT	5.45e-15	3.11e-12	4.64e-06	1.39e+00	5.87 ms
CHOL	4.23e-15	2.17e-12	1.32e-06	1.19e-05	3.09 ms
CHOL + SVD	1.23e-15	2.20e-14	2.53e-11	2.38e-09	4.28 ms

Table 2: Comparison of different methods for eigenvalue computation for Bethe-Salpeter matrix of form I of size 400×400 .

The comparison in Table 1 is helpful when implementing the new approaches in codes that already use the square root approach. For the Cholesky approach we need to substitute the computation of the matrix square root with the computation of a Cholesky factorization (LAPACK routine `zpotrf`), compute the matrix M using triangular matrix multiplications (`ztrmm`), and use a triangular solve (`ztrsml`) and a triangular matrix multiplication (`ztrmm`) in the post-processing step. For the CHOL+SVD approach, an additional Cholesky factorization is necessary and the Hermitian eigenvalue decomposition is substituted by a singular value decomposition (`zgesvd`). The post-processing involves two triangular matrix products instead of a matrix inversion and two general matrix products.

4 Numerical Experiments

We implemented and compared serial versions of Algorithms 1, 2 and 3 in MATLAB. They compute positive eigenvalues and associated eigenvectors of a BSE matrix $H \in \mathbb{C}^{2n \times 2n}$ of form I (3), which fulfills the definiteness property $\Sigma H > 0$ discussed in Section 2. The eigenvalues are given as a diagonal matrix $D \in \mathbb{R}^{n \times n}$. The eigenvectors $V \in \mathbb{C}^{2n \times n}$ are scaled s.t. Σ -orthogonality holds, i.e. $V^H \Sigma V = I_n$. The Σ -orthogonality is an important property in the application. It is exploited in order to construct the polarizability operator ultimately used for the computation of the absorption spectrum.

We also include the MATLAB eigensolver `eig` for comparison. `eig` can either work on the BSE matrix H or solve the generalized eigenvalue problem for the matrix pencil $(\Sigma H, \Sigma)$. In this formulation, both matrices are Hermitian and one is positive definite, which allows for a faster computation.

The experiments were performed on a laptop with an Intel(R) Core(TM) i7-8550U processor using MATLAB R2018a.

The first experiments aim to assess the accuracy of the computed eigenvalues. The matrices A and B are of size $n = 200$ and are created in the following way for a given value $\kappa \in \mathbb{R}$. Let $d = [1, \dots, \frac{1}{3}\kappa] \in \mathbb{R}^n$ be a vector with elements equally spaced between 1 and $\frac{1}{3}\kappa$. The BSE matrix is constructed as

$$H = \begin{bmatrix} A & B \\ -B & -A \end{bmatrix} := \begin{bmatrix} Q & 0 \\ 0 & Q \end{bmatrix}^H \begin{bmatrix} \text{diag}(d) & \frac{1}{2} \text{diag}(d) \\ -\frac{1}{2} \text{diag}(d) & -\text{diag}(d) \end{bmatrix} \begin{bmatrix} Q & 0 \\ 0 & Q \end{bmatrix},$$

where $Q \in \mathbb{C}^{n \times n}$ is a randomly generated, unitary matrix. It can be shown, that $\text{cond}(H) = \kappa$ and the eigenvalues are given as $\frac{\sqrt{3}}{2}d$.

Table 2 shows the relative error in the smallest eigenvalue $\lambda = \frac{\sqrt{3}}{2}$, using the methods discussed in Section 3. We also included the routine `haeig` from the SLICOT package [25, 26]. Because `haeig` can only compute eigenvalues, not eigenvectors, we also only compute eigenvalues in the other methods in order to make the runtimes comparable.

The MATLAB `eig` function has the largest runtime. `haeig` is slightly faster, because it exploits the available Hamiltonian structure. However, the routine is not optimized for cache-reuse, which is why this effect can

not be observed more clearly and vanishes for larger matrices. The generalized eigenvalue problem can be solved much faster, because it can be transformed to a Hermitian eigenvalue problem of size $2n \times 2n$. The other methods ultimately act on Hermitian matrices of size $n \times n$, which explains the much lower runtimes.

The observed eigenvalue errors also comply with the error analysis given in Section 3.3. The state-of-the-art square root approach performs even worse than expected, yielding a completely wrong eigenvalue for matrices with a condition number $\kappa = 10^9$. In the application context, the small eigenvalues are of special interest. They correspond to bound exciton states, representing a strong electron-hole interaction. They are the reason why the Bethe-Salpeter approach is used instead of simpler schemes based on time-dependent density functional theory [5]. The smallest eigenvalues suffer the most from this numerical inaccuracy.

The second experiment aims to assess the runtime of the sequential implementations, including the eigenvector computation in the measurement. The matrices A and B are setup as random matrices, where the diagonal of A has been scaled up in order to guarantee the definiteness property $\Sigma H > 0$. The measured runtimes are found in Figure 1 and serve as a rough indicator of computational effort.

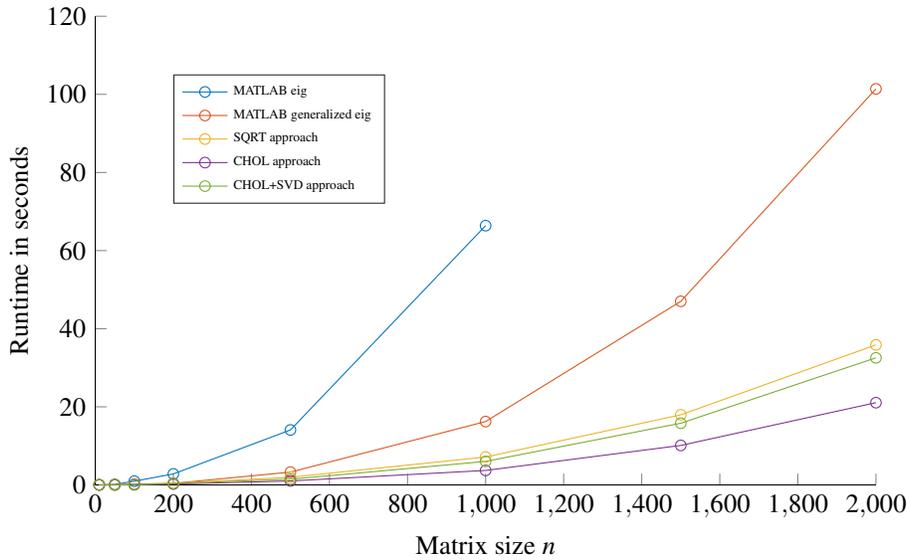


Figure 1: Runtimes for different methods, $A, B \in \mathbb{C}^{n \times n}$ with varying matrix sizes.

As expected, the Cholesky approach yields the fastest runtime of all approaches. The SVD approach also performs better than the square root approach. However, this picture could easily look different in another computational setup. An approach based on the `eig` command becomes prohibitively slow, when larger matrices are considered. Matrices in real applications become extremely large, up to dimensions of order 100 000, in order to get reasonable results. The effect would be even more drastic in a parallel setting, as the solution of a nonsymmetric dense eigenvalue problem is notoriously difficult to parallelize.

Figure 2 shows the achieved Σ -orthogonality of the eigenvector matrices for matrices with certain condition numbers. To this end, we manipulate the diagonal of the randomly generated matrix A such that badly conditioned BSE matrices H are generated. For the square root and the Cholesky approach, the Σ -orthogonality breaks down completely for badly conditioned matrices. This can have dramatic consequences and lead to completely wrong results, when further computations rely on this property.

To show the applicability to real life examples, we extracted a Bethe-Salpeter matrix corresponding to the excitation of Lithium-Fluoride from the `exciting` software package [22]. Computational details on how the matrix is generated can be found in the documentation¹. Here, it is pointed out that a Tamm-Dancoff approximation, i.e. setting the off-diagonal block B to zero, already yields satisfactory results. The resulting 2560×2560

¹<http://exciting-code.org/carbon-excited-states-from-bse>

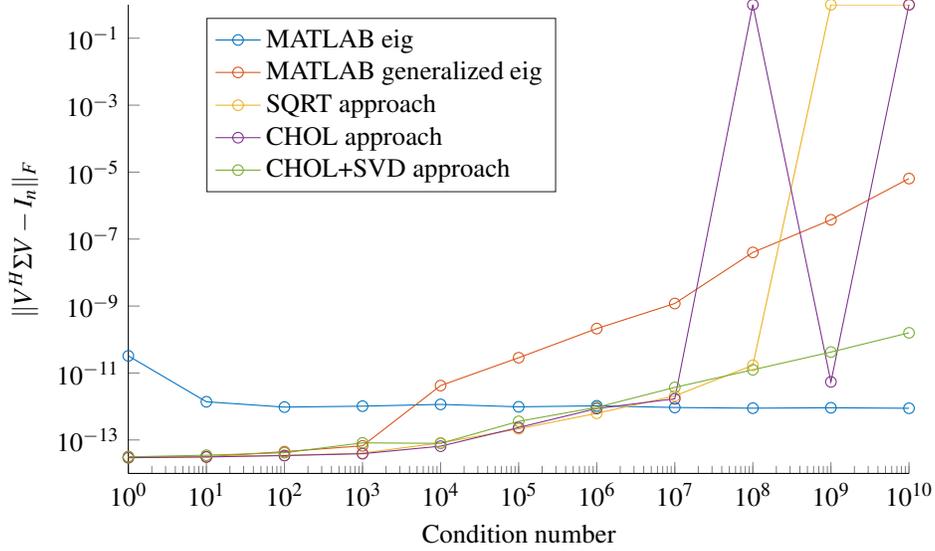


Figure 2: Deviation from Σ -orthogonality for different methods, $A, B \in \mathbb{C}^{200 \times 200}$ with a certain condition number.

	λ_1	λ_2	λ_3	Runtime
eig	4.6423352497493209e-01	4.6524229149750918e-01	4.6872644706731720e-01	32.49 s
generalized eig	4.6423352497493126e-01	4.6524229149750407e-01	4.6872644706732447e-01	10.62 s
haeig	4.6423352497493725e-01	4.6524229149750940e-01	4.6872644706732514e-01	71.43 s
SQRT	4.6423352497493120e-01	4.6524229149750490e-01	4.6872644706732541e-01	3.44 s
CHOL	4.6423352497493031e-01	4.6524229149750573e-01	4.6872644706732414e-01	2.06 s
CHOL + SVD	4.6423352497493092e-01	4.6524229149750473e-01	4.6872644706732453e-01	3.41 s
TDA	4.6427305979874345e-01	4.6528180480128906e-01	4.6877150201685513e-01	0.88 s

Table 3: Computed eigenvalues for Lithium Fluoride example.

BSE matrix has a condition number (computed using `cond` in MATLAB) of 5.33. We do not expect the algorithms to suffer from the numerical difficulties observed in the first example.

The three smallest eigenvalues computed by different methods are found in Table 3. Indeed, all approaches coincide in the first 14 significant digits. The Tamm-Dancoff approximation (TDA) applies MATLAB `eig` on the diagonal Block A and provides eigenvalues, that are correct up to 4 significant digits which is sufficient for practical applications. The measured runtimes reflect the results of the other experiments. Now the lack of low-level optimization in the `haeig` routine becomes apparent and leads to the lowest performance of all approaches.

5 Conclusions

We presented two new approaches for solving the Bethe-Salpeter eigenvalue problem as it appears in the computation of optical properties of crystalline systems. The presented methods are superior to the one currently used, which is based on the computation of a matrix square root. Computing the matrix square root constitutes a high computational effort for nondiagonal matrices. Our first proposed method substitutes the matrix square root with a Cholesky factorization which can be computed much easier. The total runtime is reduced by about 40% in preliminary experiments, while the same accuracy is achieved. In order to achieve a higher accuracy we proposed a second method, which also relies on Cholesky factorizations and uses a singular value

decomposition instead of an eigenvalue decomposition.

We also gave new theoretical results on structured matrices, which served as a foundation of the proposed algorithms.

References

- [1] T. Sander, E. Maggio, G. Kresse, Beyond the Tamm-Dancoff approximation for extended systems using exact diagonalization, *Phys. Rev. B* 92 (2015) 045209. doi:10.1103/PhysRevB.92.045209.
- [2] M. Shao, F. H. da Jornada, C. Yang, J. Deslippe, S. G. Louie, Structure preserving parallel algorithms for solving the Bethe-Salpeter eigenvalue problem, *Linear Algebra and its Applications* 488 (2016) 148–167. doi:10.1016/j.laa.2015.09.036.
- [3] C. Penke, A. Marek, C. Vorwerk, C. Draxl, P. Benner, High performance solution of skew-symmetric eigenvalue problems with applications in solving the Bethe-Salpeter eigenvalue problem, *Parallel Computing* 96 (2020) 102639. doi:10.1016/j.parco.2020.102639.
- [4] C. Vorwerk, B. Aurich, C. Cocchi, C. Draxl, Bethe-Salpeter equation for absorption and scattering spectroscopy: implementation in the exciting code, *Electronic Structure* 1 (3) (2019) 037001. doi:10.1088/2516-1075/ab3123.
- [5] S. Sagmeister, C. Ambrosch-Draxl, Time-dependent density functional theory versus Bethe-Salpeter equation: an all-electron study, *Phys. Chem. Chem. Phys.* 11 (2009) 4451–4457. doi:10.1039/B903676H.
- [6] The top500 list, available at <http://www.top500.org>.
- [7] L. Hedin, S. Lundqvist, Effects of electron-electron and electron-phonon interactions on the one-electron states of solids, Vol. 23 of *Solid State Physics*, Academic Press, 1970, pp. 1 – 181. doi:10.1016/S0081-1947(08)60615-3.
- [8] D. S. Mackey, N. Mackey, F. Tisseur, Structured factorizations in scalar product spaces, *SIAM J. Matrix Anal. Appl.* 27 (3) (2005) 821–850. doi:10.1137/040619363.
- [9] P. Benner, D. Kressner, V. Mehrmann, Skew-Hamiltonian and Hamiltonian eigenvalue problems: Theory, algorithms and applications, in: *Proc. Conf. Appl Math. Scientific Comp.*, Springer-Verlag, 2005, pp. 3–39. doi:10.1007/1-4020-3197-1_1.
- [10] P. Benner, H. Faßbender, C. Yang, Some remarks on the complex J-symmetric eigenproblem, *Linear Algebra and its Applications* 544 (2018) 407 – 442. doi:10.1016/j.laa.2018.01.014.
- [11] G. Onida, L. Reining, A. Rubio, Electronic excitations: density-functional versus many-body Green’s-function approaches, *Rev. Mod. Phys.* 74 (2002) 601–659. doi:10.1103/RevModPhys.74.601.
- [12] G. H. Golub, C. F. Van Loan, *Matrix Computations*, 4th Edition, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, 2013.
- [13] D. Kressner, *Numerical Methods for General and Structured Eigenvalue Problems*, Vol. 46 of *Lecture Notes in Computational Science and Engineering*, Springer-Verlag, Berlin, 2005.
- [14] P. Benner, V. Mehrmann, H. Xu, A numerically stable, structure preserving method for computing the eigenvalues of real Hamiltonian or symplectic pencils, *Numer. Math.* 78 (3) (1998) 329–358. doi:10.1007/s002110050315.
- [15] N. J. Higham, *Functions of Matrices: Theory and Computation*, Applied Mathematics, SIAM Publications, Philadelphia, PA, 2008. doi:10.1137/1.9780898717778.

-
- [16] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, D. Sorensen, *LAPACK Users' Guide*, SIAM, Philadelphia, PA, 3rd Edition (1999).
- [17] A. Marek, V. Blum, R. Johanni, V. Havu, B. Lang, T. Auckenthaler, A. Heinecke, H.-J. Bungartz, H. Lederer, The ELPA library: scalable parallel eigenvalue solutions for electronic structure theory and computational science, *Journal of Physics: Condensed Matter* 26 (21) (2014) 213201. doi:10.1088/0953-8984/26/21/213201.
- [18] C. F. Van Loan, A symplectic method for approximating all the eigenvalues of a Hamiltonian matrix, *Linear Algebra Appl.* 61 (1984) 233–251. doi:10.1016/0024-3795(84)90034-X.
- [19] J. H. Wilkinson, *The Algebraic Eigenvalue Problem*, Oxford University Press, Oxford, 1965.
- [20] P. Benner, R. Byers, E. Barth, Algorithm 800. Fortran 77 subroutines for computing the eigenvalues of Hamiltonian matrices I: The square-reduced method, *ACM Trans. Math. Software* 26 (1) (2000) 49–77. doi:10.1145/347837.347852.
- [21] G. W. Stewart, J.-G. Sun, *Matrix Perturbation Theory*, Academic Press, New York, 1990.
- [22] A. Gulans, S. Kontur, C. Meisenbichler, D. Nabok, P. Pavone, S. Rigamonti, S. Sagmeister, U. Werner, C. Draxl, exciting: a full-potential all-electron package implementing density-functional theory and many-body perturbation theory, *Journal of Physics: Condensed Matter* 26 (36) (2014) 363202. doi:10.1088/0953-8984/26/36/363202.
- [23] J. Deslippe, G. Samsonidze, D. A. Strubbe, M. Jain, M. L. Cohen, S. G. Louie, BerkeleyGW: A massively parallel computer package for the calculation of the quasiparticle and optical properties of materials and nanostructures, *Computer Physics Communications* 183 (6) (2012) 1269 – 1289. doi:10.1016/j.cpc.2011.12.006.
- [24] D. Sangalli, A. Ferretti, H. Miranda, C. Attaccalite, I. Marri, E. Cannuccia, P. Melo, M. Marsili, F. Paleari, A. Marrazzo, G. Prandini, P. Bonfà, M. O. Atambo, F. Affinito, M. Palumbo, A. Molina-Sánchez, C. Hogan, M. Gruning, D. Varsano, A. Marini, Many-body perturbation theory calculations using the yambo code, *Journal of Physics: Condensed Matter* 31 (32) (2019) 325902. doi:10.1088/1361-648x/ab15d0.
- [25] P. Benner, V. Mehrmann, H. Xu, A note on the numerical solution of complex Hamiltonian and skew-Hamiltonian eigenvalue problems, *Electron. Trans. Numer. Anal.* 8 (1999) 115–126.
- [26] P. Benner, D. Kressner, V. Sima, A. Varga, Die SLICOT-Toolboxen für MATLAB, *at-Automatisierungstechnik* 58 (1) (2010) 15–25. doi:10.1524/auto.2010.0814.