

High Performance Solution of Skew-symmetric Eigenvalue Problems with Applications in Solving the Bethe-Salpeter Eigenvalue Problem

Carolin Penke^{a,*}, Andreas Marek^b, Christian Vorwerk^c, Claudia Draxl^{c,1}, Peter Benner^{a,1}

^aComputational Methods in Systems and Control Theory, Max Planck Institute for Dynamics of Complex Technical Systems, Germany

^bMax Planck Computing and Data Facility, Garching, Germany

^cInstitut für Physik and IRIS Adlershof, Humboldt-Universität zu Berlin, Berlin, Germany

Abstract

We present a high-performance solver for dense skew-symmetric matrix eigenvalue problems. Our work is motivated by applications in computational quantum physics, where one solution approach to solve the Bethe-Salpeter equation involves the solution of a large, dense, skew-symmetric eigenvalue problem. The computed eigenpairs can be used to compute the optical absorption spectrum of molecules and crystalline systems. One state-of-the-art high-performance solver package for symmetric matrices is the ELPA (Eigenvalue SoLvers for Petascale Applications) library. We exploit a link between tridiagonal skew-symmetric and symmetric matrices in order to extend the methods available in ELPA to skew-symmetric matrices. This way, the presented solution method can benefit from the optimizations available in ELPA that make it a well-established, efficient and scalable library. The solution strategy is to reduce a matrix to tridiagonal form, solve the tridiagonal eigenvalue problem and perform a back-transformation for eigenvectors of interest. ELPA employs a one-step or a two-step approach for the tridiagonalization of symmetric matrices. We adapt these to suit the skew-symmetric case. The two-step approach is generally faster as memory locality is exploited better. If all eigenvectors are required, the performance improvement is counteracted by the additional back transformation step. We exploit the symmetry in the spectrum of skew-symmetric matrices, such that only half of the eigenpairs need to be computed, making the two-step approach the favorable method. We compare performance and scalability of our method to the only available high-performance approach for skew-symmetric matrices, an indirect route involving complex arithmetic. In total, we achieve a performance that is up to 3.67 higher than the reference method using Intel's ScalAPACK implementation. Our method is freely available in the current release of the ELPA library.

Keywords: Distributed memory, Skew-symmetry, Eigenvalue and eigenvector computations, GPU acceleration, Bethe-Salpeter, Many-body perturbation theory

1. Introduction

A matrix $A \in \mathbb{R}^{n \times n}$ is called skew-symmetric when $A = -A^T$, where \cdot^T denotes the transposition of a matrix. We are interested in eigenvalues and eigenvectors of A .

The symmetric eigenvalue problem, i.e. the case $A = A^T$, has been studied in depth for many years. It lies at the core of many applications in different areas such as electronic structure computations. Many methods for its solution have been proposed [1] and successfully implemented. Optimized libraries for many platforms are widely available [2, 3]. With the rise of more advanced computer architectures and more powerful supercomputers, the solution of increasingly complex problems comes within reach. Parallelizability and scalability become key issues in algorithm development. The ELPA library [4] is one endeavor to tackle these challenges and provides highly competitive direct solvers for symmetric (and Hermitian) eigen-

value problems running on distributed memory machines such as compute clusters.

The skew-symmetric case [5] lacks the ubiquitous presence of its symmetric counterpart and has not received the same extensive treatment. We close this gap by extending the ELPA methodology to the skew-symmetric case.

Our motivation stems from the connection to the Hamiltonian eigenvalue problem which has many applications in control theory and model order reduction [6]. A real Hamiltonian matrix H is connected to a symmetric matrix M via the matrix $J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}$, where I denotes the identity matrix,

$$M = JH.$$

If M is positive definite, in the following denoted by $M > 0$, the Hamiltonian eigenvalue problem can be recast into a skew-symmetric eigenvalue problem using the Cholesky factorization $M = LL^T$. The eigenvalues of H are given as eigenvalues of the skew-symmetric matrix $L^T J L$ and eigenvectors can be transformed accordingly.

This situation occurs for example in [7], where a structure-preserving method for the solution of the Bethe-Salpeter eigen-

*Corresponding author

Email address: penke@mpi-magdeburg.mpg.de (Carolin Penke)

¹These authors were supported by BiGmax, the Max Planck Societys Research Network on Big-Data-Driven Materials Science.

value problem is described. Solving the Bethe-Salpeter eigenvalue problem allows a prediction of optical properties in condensed matter, a more accurate approach than currently used ones, such as time-dependent density functional theory (TDDFT) [8]. In this application context, the condition $M > 0$ ultimately follows from much weaker physical interactions represented in the off-diagonal values [9, 10]. When larger systems are of interest, the resulting matrices easily become very high-dimensional. This calls for a parallelizable and scalable algorithm. The solution of the corresponding skew-symmetric eigenvalue problem can be accelerated via the developments presented in this paper.

The remaining paper is structured as follows. Section 2 reintroduces the methods used by ELPA and points out the necessary adaptations to make them work for skew-symmetric matrices. The Bethe-Salpeter problem is presented in Section 3. Section 4 provides performance results of the ELPA extension, including GPU acceleration, and points out the speedup achieved in the context of the Bethe-Salpeter eigenvalue problem.

2. Solution Method

2.1. Solving the Symmetric Eigenvalue Problem in ELPA

The ELPA library [4, 11, 12] is a highly optimized parallel MPI-based code [13]. It shows great scalability over thousands of CPU cores and contains low-level optimizations targeting specific compute architectures [14]. When only a portion of eigenvalues and eigenvectors are needed, this is exploited algorithmically and results in performance benefits. We briefly describe the well-established procedure employed by ELPA. This forms the basis of the method for skew-symmetric matrices described in the next subsection.

ELPA contains functionality to deal with symmetric-definite generalized eigenvalue problems. In this paper, we focus on the standard eigenvalue problem for simplicity. This is reasonable as it is the most common use case and forms the basis of any method for generalized problems. We only consider real skew-symmetric problems. The reason is that any skew-symmetric problem can be transformed into a Hermitian eigenvalue problem by multiplying it with the imaginary unit i . This problem can be solved using the available ELPA functionality for complex matrices. For the real case this induces complex arithmetic which should obviously be avoided, but for complex matrices this is a viable approach.

We consider the symmetric eigenvalue problem, i.e. the orthogonal diagonalization of a matrix,

$$Q^T A Q = \Lambda,$$

where $A = A^T \in \mathbb{R}^{n \times n}$ is the matrix whose eigenvalues are sought. We are looking for the orthogonal eigenvector matrix Q and the diagonal matrix Λ containing the eigenvalues. The solution is carried out in the following steps.

1. Reduce A to tridiagonal form, i.e. find an orthogonal transformation Q_{trd} s.t.

$$A_{trd} = Q_{trd}^T A Q_{trd}$$

is tridiagonal. This is done by accumulating Householder transformations

$$Q_{trd} = Q_1 Q_2 \cdots Q_{n-1},$$

where $Q_i = I - \tau_i v_i v_i^T$ represents the i -th Householder transformation that reduces the i -th column and row of the updated $Q_{i-1}^T \cdots Q_1^T A Q_1 \cdots Q_{i-1}$ to tridiagonal form. The matrices Q_i are not formed explicitly but are represented by the Householder vectors v_i . These are stored in place of the eliminated columns of A .

2. Solve the tridiagonal eigenvalue problem, i.e. find orthogonal Q_{diag} s.t.

$$\Lambda = Q_{diag}^T A_{trd} Q_{diag}.$$

In ELPA, this step employs a tridiagonal divide-and-conquer scheme.

3. Transform the required eigenvectors back, i.e. perform the computation

$$Q = Q_{trd} Q_{diag}.$$

The ELPA solver comes in two flavors which define the details of the transformation steps, i.e Steps 1 and 3. ELPA1 works as described, the reduction to tridiagonal form is performed in one step. ELPA2 splits the transformations into two parts. Step 1 becomes

1. (a) Reduce A to banded form, i.e. compute orthogonal Q_{band} s.t.

$$A_{band} = Q_{band}^T A Q_{band}$$

is a band matrix.

- (b) Reduce the banded form to tridiagonal form, i.e. compute orthogonal Q_{trd} s.t.

$$A_{trd} = Q_{trd}^T A_{band} Q_{trd}$$

is tridiagonal.

Accordingly, the back transformation step is split into two parts

3. (a) Perform the back transformation corresponding to the band-to-tridiagonal reduction

$$\tilde{Q} = Q_{trd} Q_{diag}.$$

- (b) Perform the back transformation corresponding to the full-to-band reduction

$$Q = Q_{band} \tilde{Q}.$$

The benefit of the two-step approach is that more efficient BLAS-3 procedures can be used in the tridiagonalization process and an overlap of communication and computation is possible. As a result, a lower runtime can generally be observed in the tridiagonalization, compared to the one-step approach. This comes at the cost of more operations in the eigenvector

back transformation due to the extra step that has to be performed. Therefore, ELPA2 is superior to ELPA1 in particular when only a portion of the eigenvectors is sought. In the context of skew-symmetric eigenvalue problems, this becomes pivotal as the purely imaginary eigenvalues come in pairs $\pm\lambda i$, $\lambda \in \mathbb{R}$. The eigenvectors are given as the complex conjugates of each other. It is therefore enough to compute half of the eigenvalues and eigenvectors.

Both approaches are extended to skew-symmetric matrices in this work.

2.2. Solving the Skew-symmetric Eigenvalue Problem

Like a symmetric matrix, a skew-symmetric matrix can be reduced to tridiagonal form using Householder transformations. A Householder transformation represents a reflection onto a scaled first unit vector e_1 . Let H be a transformation that acts on a vector v s.t. $Hv = \alpha e_1$. Obviously $-v$ is transformed to $H(-v) = -\alpha e_1$ by the same H . Therefore all tridiagonalization methods that work on symmetric matrices, such as the ones implemented in ELPA, can in principle work on skew-symmetric matrices as well.

A skew-symmetric tridiagonal matrix is related to a symmetric one via the following observation [5].

Lemma 1. *With the unitary matrix $D = \text{diag}\{1, i, i^2, \dots, i^{n-1}\}$, where i denotes the imaginary unit, $\alpha_j \in \mathbb{R}$, it holds*

$$-iD^H \begin{bmatrix} 0 & \alpha_1 & & & \\ -\alpha_1 & 0 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \alpha_{n-1} & 0 \\ & & & & -\alpha_{n-1} & 0 \end{bmatrix} D = \begin{bmatrix} 0 & \alpha_1 & & & \\ \alpha_1 & 0 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \alpha_{n-1} & 0 \\ & & & & \alpha_{n-1} & 0 \end{bmatrix}. \quad (1)$$

H denotes the Hermitian transpose of a matrix.

After the reduction to tridiagonal form, the symmetric tridiagonal system is solved using a divide-and-conquer method [11]. As a first step of the back transformation, the resulting (real) eigenvectors have to be multiplied by the (complex) matrix D . Then the back transformations corresponding to the tridiagonalization take place. Algorithm 1 outlines the process. It is very similar to the method employed for symmetric eigenvalue problems. The differences are the addition of step 3 and changes in the implementation, which are given in detail in Sections 2.3.1 and 2.3.2.

In ELPA2 the transformation steps (1 and 4 in Algorithm 1) are both split into two parts as described in Section 2.1.

2.3. Implementation

Extending ELPA for skew-symmetric matrices means adding the back transformation step involving D . In contrast to symmetric matrices, skew-symmetric matrices have complex eigenvectors and strictly imaginary eigenvalues. Computationally complex values are introduced in Algorithm 1 with D in step 3. Further transformations have to be performed for the real and the imaginary part individually. It is preferable to set up an array with complex data type entries representing the eigenvectors as late as possible, so that we can benefit from efficient routines in double precision. The routines for the eigenvector

Algorithm 1 Solution of a Skew-symmetric Eigenvalue Problem

Input: $A = -A^T \in \mathbb{R}^{n \times n}$

Output: Unitary eigenvectors $Q \in \mathbb{C}^{n \times n}$, $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ s.t. $Q^H A Q = \text{diag}\{\lambda_1 i, \dots, \lambda_n i\}$.

- 1: Reduce A to tridiagonal form, i.e. generate Q_{trd} s.t.

$$Q_{trd}^T A Q_{trd} = A_{trd} = \begin{bmatrix} 0 & \alpha_1 & & & \\ -\alpha_1 & 0 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \alpha_{n-1} & 0 \\ & & & & -\alpha_{n-1} & 0 \end{bmatrix}.$$

- 2: Solve the eigenvalue problem for the symmetric tridiagonal matrix $-iD^H A_{trd} D$, where $D = \text{diag}\{1, i, i^2, \dots, i^{n-1}\}$, i.e. generate Q_{diag} s.t.

$$Q_{diag}^T \begin{bmatrix} 0 & \alpha_1 & & & \\ \alpha_1 & 0 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \alpha_{n-1} & 0 \\ & & & & \alpha_{n-1} & 0 \end{bmatrix} Q_{diag} = \begin{bmatrix} \lambda_1 & & & & \\ & \lambda_2 & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & \lambda_n \end{bmatrix}.$$

- 3: Back transformation corresponding to symmetrization (see Lemma 1), i.e. compute

$$Q \leftarrow D Q_{diag} \in \mathbb{C}^{n \times n}.$$

- 4: Back transformation corresponding to band-to-tridiagonal reduction, i.e. compute

$$Q \leftarrow Q_{trd} Q.$$

back transformation corresponding to tridiagonalization do not change, because all they do is to apply Householder transformations to non-symmetric (and non-skew-symmetric) matrices. They are applied on the real and imaginary part independently, realizing the complex back transformation in real arithmetic. The symmetric tridiagonal eigensolver can be used as is. Making it aware of the zeros on the diagonal might turn out to be numerically or computationally beneficial.

We now examine the implementation of the two tridiagonalization approaches in ELPA1 and ELPA2 in more detail. At many points in the original implementation, symmetry of the matrix is assumed in order to avoid unnecessary computations and to efficiently reuse data available in the cache. In this section we recollect some details of the tridiagonal reduction in order to point out these instances. Here, the implicit assumptions can be changed from “symmetric” to “skew-symmetric” by simple sign changes.

ELPA is based on the well established and well documented 2D block-cyclic data layout introduced by ScaLAPACK for load balancing reasons. It is therefore compatible to ScaLAPACK and can act as a drop-in replacement while no ScaLAPACK routines are used by ELPA itself. In general, each process works on the part of the matrix that was assigned to it. This chunk of data resides in the local memory of the process. Com-

munication between processes is realized via MPI. Each process calls serial BLAS routines. Additional CUDA and OpenMP support is available.

2.3.1. Tridiagonalization in ELPA1

In ELPA1, the tridiagonalization is realized in one step using Householder transformations. The computation of the Householder vectors is not affected by the symmetry of a matrix. Essentially, the tridiagonalization of a matrix comes down to a series of rank-2 updates [15], described in the following. Given a Householder vector v , the update of the trailing submatrix is performed as

$$A \leftarrow (I - \tau vv^T)A(I - \tau vv^T) \quad (2)$$

$$= A + v \underbrace{(0.5\tau^2 v^T A v v^T - \tau v^T A)}_{u_1^T} + \underbrace{(0.5\tau^2 v v^T A v - \tau A v)}_{u_2} v^T \quad (3)$$

$$= A + v u_1^T + u_2 v^T \quad (4)$$

$$= A + \begin{bmatrix} v & u_2 \end{bmatrix} \begin{bmatrix} u_1 & v \end{bmatrix}^T. \quad (5)$$

For symmetric matrices it holds $u_1 = u_2$. This is assumed in the original ELPA implementation. For skew-symmetric matrices it holds $u_1 = -u_2$. In ELPA1, the two matrices $\begin{bmatrix} v & u_2 \end{bmatrix}$ and $\begin{bmatrix} u_1 & v \end{bmatrix}^T$ are stored explicitly. Actual updates are performed using GEMM and GEMV routines. The matrices differ in the data layout, i.e. which process owns which part of the matrix. After the vector u_1 is computed, it is transposed and redistributed to represent u_2 in $\begin{bmatrix} v & u_2 \end{bmatrix}$. Here, for the skew-symmetric variant, a sign change is introduced. The skew-symmetric update now reads

$$A \leftarrow A + \begin{bmatrix} v & -u_1 \end{bmatrix} \begin{bmatrix} u_1 & v \end{bmatrix}^T. \quad (6)$$

During the computation of u_1 , symmetry is assumed in the computation of $A^T v$. In particular, the code assumes that an off-diagonal matrix tile is the same as in the transposed matrix. Another sign change corrects this assumption for skew-symmetric matrices.

2.3.2. Tridiagonalization in ELPA2

In ELPA2, the tridiagonalization is split into two parts. First, the matrix is reduced to banded form, then to tridiagonal form. For the reduction to banded form, the Householder vectors are computed by the process column owning the diagonal block. They are accumulated in a triangular matrix $T \in \mathbb{R}^{nb \times nb}$, where nb is the block size. The product of Householder matrices is stored via its storage-efficient representation [16]

$$Q = H_1 \cdots H_{nb} = I - VT V^T, \quad (7)$$

where $V = \begin{bmatrix} v_1 & \cdots & v_{nb} \end{bmatrix}$ contains the Householder vectors. $H_i = I - \tau_i v_i v_i^T$ is the Householder matrix corresponding to the i -th Householder transformation.

In this context, the update of the matrix A takes the following shape, analogous to the direct tridiagonalization described

in Section 2.3.1.

$$A \leftarrow (I - VT V^T)^T A (I - VT V^T) \quad (8)$$

$$= A + V \underbrace{(0.5T^T V^T A V T V^T - T^T V^T A)}_{u_1^T} + \underbrace{(0.5V T^T V^T A V T - A V T)}_{u_2} V^T \quad (9)$$

$$= A + \begin{bmatrix} V & U_2 \end{bmatrix} \begin{bmatrix} U_1 & V \end{bmatrix}^T. \quad (10)$$

It holds $U_1 = U_2$ if A is symmetric, and $U_1 = -U_2$ if A is skew-symmetric. Each process computes the relevant parts of U_1 in a series of (serial) matrix operations and updates the portion of A that resides in its memory. Here, the symmetry of A is assumed and exploited at various points in the implementation. Sign changes have to be applied at these instances.

For the banded-to-tridiagonal reduction, the matrix is redistributed in the form of a 1D block cyclic data layout. Each process owns a diagonal and a subdiagonal block. The reduction of a particular column introduces fill-in in the neighboring block. The ‘‘bulge-chasing’’ is realized as a pipelined algorithm where computation and communication can be overlapped by reordering certain operations [11, 17].

The update of the diagonal blocks takes the same form as in ELPA1 (Equations (2) to (5)). Here, no matrix multiplication is employed but BLAS-2 routines are used working directly with the Householder vectors. It holds $u_1 = u_2$ for symmetric A and $u_1 = -u_2$ for skew-symmetric A . In the symmetric case, the update is realized via a symmetric rank-2 update (SYR2). We implemented a skew-symmetric variant of this routine which realizes the skew-symmetric rank-2 update $A \leftarrow A - v u^T + u v^T$. For the setup of u , a skew-symmetric variant of the BLAS routine performing a symmetric matrix vector product (SYMV) is necessary.

The other parts of Algorithm 1 are adopted from the symmetric implementation without changes. The computation of Householder vectors, the accumulation of the Householder transformations in a triangular matrix and the update of the local block during reduction to banded form do not have to be changed compared to symmetric ELPA. This is because they act on the lower part of the matrix so that possible (skew-)symmetry has no effect.

3. The Bethe-Salpeter Eigenvalue Problem

Ab initio spectroscopy aims to describe the excitations in condensed matter from first principles, *i.e.* without the input of any empirical parameters. For light absorption and scattering, the Bethe-Salpeter Equation (BSE) approach is the state-of-the-art methodology for both crystalline systems [18, 19, 20, 21, 8] as well as condensed molecular systems [22, 23, 24, 25]. This approach takes its name from the *Bethe-Salpeter Equation* [26], the equation of motion of the electron-hole correlation function, as derived from many-body perturbation theory [27, 8]. In practice, the problem of solving the BSE is mapped to an effective eigenvalue problem. Specifically, its eigenvalues and -states are

employed to construct dielectric properties, such as the spectral density, absorption spectrum, and the loss function [7, 28]. An appropriate discretization scheme leads to a finite-dimensional representation in matrix form H_{BS} that shows a particular block structure [29]:

$$H_{BS} = \begin{bmatrix} A & B \\ -\bar{B} & -\bar{A} \end{bmatrix} = \begin{bmatrix} A & B \\ -B^H & -A^\top \end{bmatrix}, \quad (11)$$

$$A = A^H, \quad B = B^\top \in \mathbb{C}^{n \times n}.$$

Note that the Hermitian transpose H as well as the regular transpose without complex conjugation $^\top$ play a role in this structure.

In general, we are interested in all eigenpairs of the Hamiltonian, as they contain valuable information on the excitations of the system. Specifically, they describe the *bound excitons*, localized electron-hole pairs that form due to correlation between an excited electron and a hole. The BSE eigenstates are used to reconstruct the excitonic wavefunction and obtain the excitonic binding energy.

In this paper, we present a solution strategy for the most general formulation of the BSE problem. As such, A and B are generally dense and complex-valued, which holds in the case of excitations in condensed matter.

H_{BS} belongs to the slightly more general class of J -symmetric matrices [30]. This class of matrices display a symmetry $(\lambda, -\lambda)$ in the spectrum. The additional structure in H_{BS} leads to an additional symmetry $(\lambda, -\lambda, \bar{\lambda}, -\bar{\lambda})$ and a relation between the corresponding eigenvectors. Following [7], we consider the definite Bethe-Salpeter eigenvalue problem. H_{BS} is called definite when the property

$$\begin{bmatrix} I_n & 0 \\ 0 & -I_n \end{bmatrix} H_{BS} = \begin{bmatrix} A & B \\ \bar{B} & \bar{A} \end{bmatrix} > 0 \quad (12)$$

is fulfilled, which often holds in practice. In this case, the eigenvalues are real and therefore come in pairs $(\lambda, -\lambda)$. The method presented in this work relies on this assumption.

We aim for a solution method that preserves this structure under the influence of inevitable numerical errors, i.e. that guarantees that the eigenvalues come in pairs or quadruples, respectively. General methods for eigenvalue problems, such as the QR/QZ algorithm, destroy this property. In this case it is not clear anymore which eigenpairs correspond to the same excitation state.

A structure-preserving method running in parallel on distributed memory systems is developed in [7] and has been made available as BSEPACK. It relies on assumption (12) and exploits a connection to a Hamiltonian eigenvalue problem given in the following Theorem.

Theorem 2. Let $Q = \frac{1}{\sqrt{2}} \begin{bmatrix} I & -iI \\ I & iI \end{bmatrix}$, then Q is unitary and

$$Q^H \begin{bmatrix} A & B \\ -\bar{B} & -\bar{A} \end{bmatrix} Q = i \begin{bmatrix} \text{Im}(A+B) & -\text{Re}(A-B) \\ \text{Re}(A+B) & \text{Im}(A-B) \end{bmatrix} =: iH,$$

where H is real Hamiltonian, i.e. $JH = (JH)^\top$ with

$$J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}.$$

Let

$$M = JH = \begin{bmatrix} \text{Re}(A+B) & \text{Im}(A-B) \\ -\text{Im}(A+B) & \text{Re}(A-B) \end{bmatrix} \quad (13)$$

be the symmetric matrix associated with the Hamiltonian matrix H . Its positive definiteness follows from property (12), which can be seen in the following way. Let the matrices S and Ω be given as

$$S = \begin{bmatrix} I_n & \\ & -I_n \end{bmatrix}, \quad \Omega = \begin{bmatrix} A & B \\ \bar{B} & \bar{A} \end{bmatrix}, \quad (14)$$

i.e. $H_{BS} = S\Omega$. With the matrix Q from Theorem 2 we have

$$M = -iJQ^H S \Omega Q. \quad (15)$$

It is easily verified that

$$-iJQ^H S Q = I_n, \quad (16)$$

i.e. $-iJQ^H S$ is the inverse of Q . The construction of M (15) can therefore be seen as a similarity transformation of Ω . If Ω is positive definite (12), so is M . The method described in [7] relies on this property in order to guarantee the existence of the Cholesky factorization of M .

It performs the following steps.

1. Construct M as in (13).
2. Compute a Cholesky factorization $M = LL^\top$.
3. Compute eigenpairs of the skew-symmetric matrix $L^\top J L$, where $J = \begin{bmatrix} 0 & I \\ -I & 0 \end{bmatrix}$.
4. Perform the eigenvector back transformation associated with Cholesky factorization and transformation to Hamiltonian form (Theorem 2).

The eigenvalues and eigenvectors can be used to compute the optical absorption spectrum of the material in a postprocessing step.

The main workload is given as the solution of a skew-symmetric eigenvalue problem (Step 3). As a proof of concept, solution routines for the symmetric eigenvalue problem from the ScaLAPACK reference implementation [3] were adapted to the skew-symmetric setting. The matrix is reduced to tridiagonal form using Householder transformations. The tridiagonal eigenvalue problem is solved via bisection and inverse iteration.

The ScaLAPACK reference implementation is not regarded as a state-of-the-art solver library. When performance and scalability are issues, one generally turns to professionally maintained and optimized libraries such as ELPA [4] or vendor-specific implementations such as Intel's MKL. Within BSEPACK, ScaLAPACK can be substituted by ELPA working on skew-symmetric matrices. The resulting performance benefits are discussed in Section 4.2.

4. Numerical Experiments

4.1. ELPA Benchmarks

In this section we present performance results for the skew-symmetric ELPA extension. All test programs are run on the




Figure 1: Scaling of the ELPA solver for skew-symmetric matrices. For comparison the runtimes for the alternative solution method via complex Hermitian solvers is included. Here, ELPA and Intel’s MKL 2018 routines `pzhheevd` and `pzhheevr` are used. The matrix has a size of $n = 20000$.

mechthild compute cluster, located at the Max Planck Institute for Dynamics of Complex Technical Systems in Magdeburg, Germany. Up to 32 nodes are used, which consist of 2 Intel Xeon Silver 4110 (Skylake) processors with 8 cores each, running at 2.1 GHz. The Intel compiler, MPI library and MKL in the 2018 version are used in all test programs. The computations use randomly generated skew-symmetric matrices in double-precision.

Figure 1 shows the resulting performance and the scaling properties of ELPA for a medium sized skew-symmetric matrix ($n = 20000$). As an alternative to the approach described in this work, the skew-symmetric matrix can be multiplied with the imaginary unit i . The resulting complex Hermitian matrix can be diagonalized using available methods in ELPA or Intel’s ScaLAPACK implementation shipped with the MKL. This represents the only previously available approach to solve skew-symmetric eigenvalue problems in a massively parallel high-performance setting.

For skew-symmetric matrices, only 50% of eigenvalues and eigenvectors need to be computed, as they are purely imaginary and come in pairs $\pm\lambda i, \lambda \in \mathbb{R}$. The runtime measurements for 100% are included for reference.

Figure 1 shows that all approaches display good scalability in the examined setting. Skew-symmetric ELPA runs 2.76 to 3.28 times faster than the complex MKL based solver, where both only compute 50% of eigenpairs. The data gives further insight into how this improvement is achieved. Table 1 compares the runtimes for different solvers and presents the achieved speedups. When we compare complex 100% solvers, ELPA already improves performance by a factor of 1.1 to 1.29 (column 2

Table 1: Execution time speedups achieved by different aspects of the solution approach.

#Cores	Compl. ELPA2 100% vs. Compl. MKL 100 %	Compl. ELPA2 50% vs. Compl. MKL 50%	Skew-Sym. ELPA2 50% vs. Compl. ELPA2 50%	Skew-Sym. ELPA2 50% vs. Compl. MKL 50%
16	1.10	1.41	2.33	3.28
32	1.29	1.41	2.30	3.24
64	1.11	1.40	2.32	3.25
128	1.18	1.33	2.20	2.93
256	1.17	1.28	2.16	2.76
512	1.21	1.51	1.87	2.82

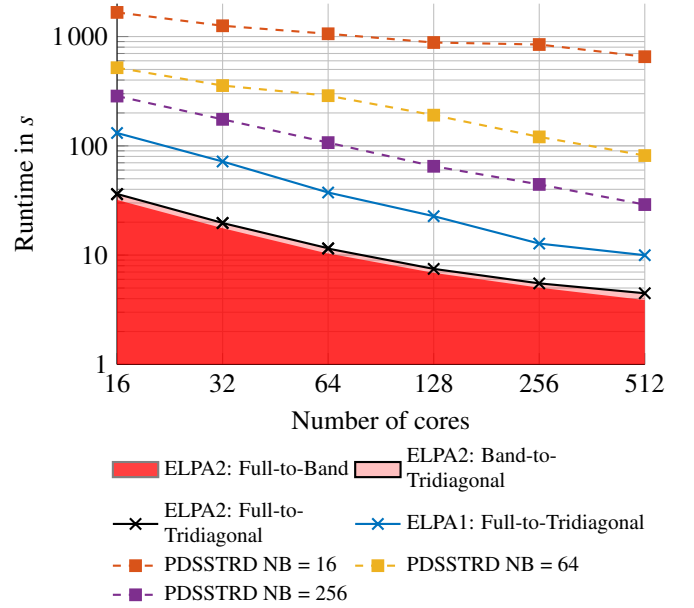


Figure 2: Scaling of the tridiagonalization in two steps (ELPA2) and one step (ELPA1). We compare it to the runtimes of the tridiagonalization routine for skew-symmetric matrices PDSSTRD available in BSEPACK [7] for different block sizes NB . The matrix size is $n = 20000$.

in Table 1). When all eigenpairs are computed, ELPA1 and ELPA2 yield very similar runtime results which is why only ELPA2 is considered in Table 1. The two-step approach employed by ELPA2 pays off in particular when not all eigenpairs are sought, which is the case here. When complex 50% solvers are compared (ELPA2 vs. MKL, column 3 in Table 1), the achieved speedup increases to a value between 1.28 and 1.51. The largest impact on the performance is caused by avoiding complex arithmetic. This is represented by the speedup achieved by the skew-symmetric 50% ELPA2 implementation compared to the complex 50% ELPA2 implementation (column 4 of Table 1). This accounts for an additional speedup of 1.87 to 2.33.

The tridiagonalization is an essential step in every considered solution scheme and contributes a significant portion of the execution time. The fewer eigenpairs are sought, the more dominant it becomes with respect to computation time. Figure 2 displays the runtimes and scalability of available tridiagonaliza-




Figure 3: Runtimes for solving eigenvalue problems of larger sizes. 256 CPU cores were used, i.e. 16 nodes on the *mechthild* compute cluster.

tion techniques for skew-symmetric matrices. As an alternative implementation to the presented approaches there is a tridiagonalization routine PDSSTRD shipped in BSEPACK [7]. It is an adapted version of the ScaLAPACK reference implementation.

All discussed implementations are based on the 2D-block-cyclic data distribution established by ScaLAPACK. Here, the matrix is divided into blocks of a certain size NB . The blocks are distributed to processes organized in a 2D grid in a cyclic manner. Typically, the block size is a parameter chosen once in a software project. The data redistribution to data layouts defined by other block sizes is avoided as this involves expensive all-to-all communication. The main disadvantage of the PDSSTRD routine is that it is very susceptible to the chosen block size, both with regard to scalability and overall performance. This makes it less suitable to be included in larger software projects, where the block size is a parameter predefined by other factors. ELPA (both the one and two-step version) on the other hand does not have this problem and performs equally well for all data layouts [31].

Figure 2 also displays the advantage of the two-step tridiagonalization over the one-step approach. Here the performance is dominated by the first step, i.e. the reduction to banded form. In the context of electronic structure computations, the matrices of interest can become extremely large. Figure 3 displays the achieved runtime improvements for larger matrices up to a size of $n = 125000$. The individual speedups are presented in Table 2. For large matrices we achieve a speedup of up to 3.67 compared to the available MKL routine.

4.1.1. GPU Acceleration

For the 1-step tridiagonalization approach (ELPA1), there is a GPU-accelerated version available that gets shipped with

Table 2: Execution time speedups achieved by different aspects of the solution approach.

Matrix size	Compl. ELPA2 100% vs. Compl. MKL 100%	Compl. ELPA2 50% vs. Compl. MKL 50%	Skew-Sym. ELPA2 50% vs. Compl. ELPA2 50%	Skew-Sym. ELPA2 50% vs. Compl. MKL 50%
50 000	1.17	1.45	2.32	3.35
75 000	1.16	1.46	2.39	3.50
100 000	1.17	1.47	2.42	3.57
125 000	1.17	1.49	2.46	3.67

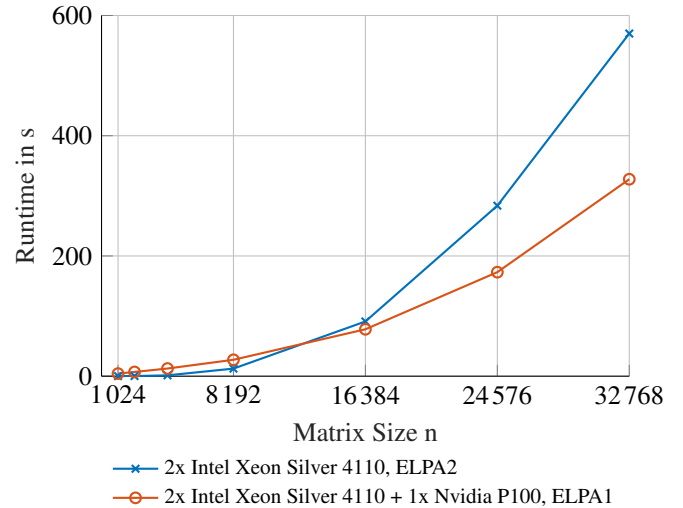


Figure 4: Runtimes for solving eigenvalue problems on one node on the *mechthild* compute cluster employing a GPU.

the ELPA library [32]. The design approach is to stick with the same code base as the CPU-only version, and offload compute-intensive parts, such as BLAS-3 operations, to the GPU in order to benefit from its massive parallelism. This is done using the CUBLAS library provided by NVIDIA. Because ELPA2 employs more fine-grained communication patterns, this approach works best for ELPA1. Here, the performance can benefit when the computational intensity is high enough, i.e. when big chunks of data are being worked on by the GPU.

Figure 4 shows the performance that can be achieved on one node of the *mechthild* compute cluster, that is equipped with an NVIDIA P100 GPU as an accelerator device. The GPU version is based on ELPA1 and therefore does not benefit from the faster tridiagonalization in ELPA2 (see Figure 2 and the discussion in the previous section). Despite this fact, the GPU-accelerated ELPA1 version eventually outperforms the ELPA2 CPU-only version, if the matrix is large enough. In our case the turning point is at around $n = 15000$. For smaller matrices the additional work of setting up the CUDA environment and transferring the matrix counteracts any possible performance benefits and results in a larger runtime. For matrices of size $n = 32768$ employing the GPU can reduce the runtime from 570 seconds to 328 seconds, i.e. by 41%.

The take-away message of these results is the following. If nodes equipped with GPUs are available and to be utilized, it is




Figure 5: Scaling of the direct, complex BSEPACK eigenvalue solver for computing the optical absorption spectrum of hexagonal boron nitride. The Bethe-Salpeter matrix (11) has a size of 51200.

important to make sure each node has enough data to work on. This way, the available resources are used most efficiently.

4.2. Accelerating BSEPACK

We consider the performance improvements that can be achieved by using the newly developed skew-symmetric eigenvalue solver in the BSEPACK [7] software, described in Section 3. In this procedure, Step 3, the computation of eigenpairs of the skew-symmetric matrix $L^T J L$, is now performed by the ELPA library.

To demonstrate the speedup, we consider the example of hexagonal boron nitride at a fixed size of the BSE Hamiltonian. The excitations in hexagonal boron nitride are widely studied both experimentally and theoretically [33, 34, 35, 36, 37, 38, 39], as its wide band gap and the layered geometrical structure yield strong effects of electron-hole correlation, such as the formation of bound excitons. Previous studies have shown that the BSE approach yields the optical absorption and excitonic properties with high accuracy. In our calculations, the BSE Hamiltonian is constructed on a $16 \times 16 \times 4$ \mathbf{k} -grid in the 1st Brillouin zone, the 5 highest valence and 5 lowest conduction bands are employed to construct the transition space, leading to a matrix size of $2 \times 16 \times 16 \times 4 \times 5 \times 5 = 51200$. In the calculation of the BSE Hamiltonian, single-particle wavefunctions and the static dielectric function are expanded in plane waves with a cut-off of 387 eV and 132 eV, respectively. The static dielectric function is obtained from ABINIT [40], while the BSE Hamiltonian is constructed using the EXC code [41].

Figure 5 displays the achieved runtimes of BSEPACK for this fixed-size matrix for different core counts. We compare the original version and a version that employs ELPA. The performance of the original solver is highly dependent on the chosen block size (see also Figure 2). This parameter determines how the matrix is distributed to the available processes in the form of a 2D block-cyclic data layout. The default is given as

$NB = 64$, but choosing a larger block size can increase the performance dramatically, as can be seen in Figure 5 for $NB = 256$. Typically, software packages (e.g. [42, 28]) developed for electronic structure computations are large and contain many features, implementing methods for different quantities of interest. The block size is typically predetermined by other considerations. It would mean a serious effort to change it, in order to optimize just one building block of the software. Furthermore the optimal block size of the original BSEPACK is probably dependent on the given hardware and the given matrix size. Autotuning frameworks could help, but are also very costly and impose an additional implementation effort. A software, that does not show this kind of runtime dependency is greatly preferable. Employing ELPA for the main computational task in BSEPACK fulfills this requirement. The performance of ELPA is independent of the chosen NB , because the block size on the node level for optimal cache use is decoupled from the block size defining the multi-node data layout.

The ELPA-accelerated version is up to 9.22 times as fast as the original code with the default block size. Even when the block size is increased, using the new solver always yields a better performance. In the case of $NB = 256$, the ELPA-version still performs up to 2.76 times as fast. Choosing even larger block sizes has in general no further positive effect on the performance of the original BSEPACK. Employing ELPA also leads to an improved scalability over the number of cores.

5. Conclusions

We have presented a strategy to extend existing solver libraries for symmetric eigenvalue problems to the skew-symmetric case. Applying these ideas to the ELPA library, makes it possible to compute eigenvalues and eigenvectors of large skew-symmetric matrices in parallel with a high level of efficiency and scalability. We benefit from the maturity of the ELPA software project, where many optimizations have been realized over the years. All of these, including GPU support, find their way into the presented skew-symmetric solver. As far as we know, no other solvers dedicated to the skew-symmetric eigenvalue problem exist in an HPC setting. It is always possible to solve a complex Hermitian eigenvalue problem instead of a skew-symmetric one. Our newly developed solver outperforms this strategy, implemented via Intel MKL ScaLAPACK, by a factor of 3. We also observe an increase in performance concerning the Bethe-Salpeter eigenvalue problem. Here we improve the runtime of available routines by a factor of almost 10, making the BSEPACK library with ELPA a viable choice as a building block for larger electronic structure packages.

6. Acknowledgment

We thank Francesco Sottile for fruitful discussion and his support in generating the BSE Hamiltonian for hexagonal BN.

References

- [1] G. H. Golub, C. F. Van Loan, *Matrix Computations*, 4th Edition, Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, 2013.
- [2] E. Anderson, Z. Bai, C. Bischof, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, D. Sorensen, *LAPACK Users' Guide*, SIAM, Philadelphia, PA, 2nd Edition (1995). doi:10.1137/1.9780898719604.
- [3] L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, R. C. Whaley, *ScaLAPACK User's Guide*, Vol. 4 of Software, Environments and Tools, SIAM Publications, Philadelphia, PA, USA, 1997. doi:10.1137/1.9780898719642.
- [4] 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.
- [5] R. C. Ward, L. J. Gray, Eigensystem computation for skew-symmetric and a class of symmetric matrices, *ACM Trans. Math. Softw.* 4 (3) (1978) 278–285. doi:10.1145/355791.355798.
- [6] P. Benner, D. Krefßner, V. Mehrmann, Skew-Hamiltonian and Hamiltonian eigenvalue problems: Theory, algorithms and applications, in: Z. Drmač, M. Marusic, Z. Tutek (Eds.), *Proc. Conf. Appl. Math. Scientific Comp.*, Springer-Verlag, Dordrecht, 2005, pp. 3–39. doi:10.1007/1-4020-3197-1_1.
- [7] 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 (Supplement C) (2016) 148 – 167. doi:10.1016/j.laa.2015.09.036.
- [8] G. Onida, L. Reining, A. Rubio, Electronic excitations: density-functional versus many-body Greens-function approaches, *Rev. Mod. Phys.* 74 (2) (2002) 601. doi:10.1103/RevModPhys.74.601.
- [9] F. Furche, On the density matrix based approach to time-dependent density functional response theory, *The Journal of Chemical Physics* 114 (14) (2001) 5982–5992. doi:10.1063/1.1353585.
- [10] J. Čížek, J. Paldus, Stability conditions for the solutions of the Hartree-Fock equations for atomic and molecular systems. application to the Pi-electron model of cyclic polyenes, *The Journal of Chemical Physics* 47 (10) (1967) 3976–3985. doi:10.1063/1.1701562.
- [11] T. Auckenthaler, V. Blum, H.-J. Bungartz, T. Huckle, R. Johanni, L. Krämer, B. Lang, H. Lederer, P. Willems, Parallel solution of partial symmetric eigenvalue problems from electronic structure calculations, *Parallel Computing* 37 (12) (2011) 783 – 794, 6th International Workshop on Parallel Matrix Algorithms and Applications (PMAA'10). doi:10.1016/j.parco.2011.05.002.
- [12] A. Alvermann, A. Basermann, H.-J. Bungartz, et al., Benefits from using mixed precision computations in the ELPA-AEO and ESSEX-II eigensolver projects, *Japan Journal of Industrial and Applied Mathematics* 36 (2) (2019) 699–717. doi:10.1007/s13160-019-00360-8.
- [13] Message Passing Interface Forum, *MPI: A message-passing interface standard*, Tech. rep., Knoxville, TN, USA (1994).
- [14] P. Küs, A. Marek, S. Köcher, H.-H. Kowalski, C. Carbogno, C. Scheurer, K. Reuter, M. Scheffler, H. Lederer, Optimizations of the eigensolvers in the ELPA library, *Parallel Computing* 85 (2019) 167 – 177. doi:10.1016/j.parco.2019.04.003.
- [15] R. S. Martin, C. Reinsch, J. H. Wilkinson, Householder's tridiagonalization of a symmetric matrix, *Numerische Mathematik* 11 (3) (1968) 181–195. doi:10.1007/BF02161841.
- [16] R. S. Schreiber, C. Van Loan, A storage-efficient WY representation for products of Householder transformations, *SIAM J. Sci. Statist. Comput.* 10 (1989) 53–57.
- [17] T. Auckenthaler, H.-J. Bungartz, T. Huckle, L. Krmer, B. Lang, P. Willems, Developing algorithms and software for the parallel solution of the symmetric eigenvalue problem, *Journal of Computational Science* 2 (3) (2011) 272 – 278. doi:10.1016/j.jocs.2011.05.002.
- [18] M. Rohlfing, S. G. Louie, Electron-hole excitations in semiconductors and insulators, *Phys. Rev. Lett.* 81 (1998) 2312–2315. doi:10.1103/PhysRevLett.81.2312.
- [19] L. X. Benedict, E. L. Shirley, R. B. Bohn, Optical absorption of insulators and the electron-hole interaction: An ab initio calculation, *Phys. Rev. Lett.* 80 (1998) 4514–4517. doi:10.1103/PhysRevLett.80.4514.
- [20] S. Albrecht, L. Reining, R. Del Sole, G. Onida, Ab initio calculation of excitonic effects in the optical spectra of semiconductors, *Phys. Rev. Lett.* 80 (1998) 4510–4513. doi:10.1103/PhysRevLett.80.4510.
- [21] 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.
- [22] J. C. Grossman, M. Rohlfing, L. Mitas, S. G. Louie, M. L. Cohen, High accuracy many-body calculational approaches for excitations in molecules, *Phys. Rev. Lett.* 86 (3) (2001) 472. doi:10.1103/PhysRevLett.86.472.
- [23] C. Faber, P. Boulanger, C. Attaccalite, I. Duchemin, X. Blase, Excited states properties of organic molecules: from density functional theory to the GW and Bethe-Salpeter Green's function formalisms, *Phil. Trans. R. Soc. A* 372 (2011) (2014) 20130271. doi:10.1098/rsta.2013.0271.
- [24] C. Cocchi, C.-D. Draxl, Optical spectra from molecules to crystals: Insight from many-body perturbation theory, *Phys. Rev. B* 92 (2015) 205126. doi:10.1103/PhysRevB.92.205126.
- [25] D. Hirose, Y. Noguchi, O. Sugino, All-electron G W+ Bethe-Salpeter calculations on small molecules, *Phys. Rev. B* 91 (20) (2015) 205111. doi:10.1103/PhysRevB.91.205111.
- [26] E. E. Salpeter, H. A. Bethe, A relativistic equation for bound-state problems, *Phys. Rev.* 84 (1951) 1232–1242. doi:10.1103/PhysRev.84.1232.
- [27] G. Strinati, Application of the Green's functions method to the study of the optical properties of semiconductors, *Riv. Nuovo Cimento* 11 (12) (1988) 1–86. doi:10.1007/BF02725962.
- [28] 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.
- [29] 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.
- [30] 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.
- [31] P. Benner, A. Marek, C. Penke, Improving the performance of numerical algorithms for the Bethe-Salpeter eigenvalue problem, *Proc. Appl. Math. Mech.* 18 (1). doi:10.1002/pamm.201800255.
- [32] P. Küs, H. Lederer, A. Marek, GPU optimization of large-scale eigenvalue solver, in: F. A. Radu, K. Kumar, I. Berre, J. M. Nordbotten, I. S. Pop (Eds.), *Numerical Mathematics and Advanced Applications ENUMATH 2017*, Springer International Publishing, Cham, 2019, pp. 123–131. doi:10.1007/978-3-319-96415-7_9.
- [33] G. Cappellini, G. Satta, M. Palumbo, G. Onida, Optical properties of BN in cubic and layered hexagonal phases, *Phys. Rev. B* 64 (2001) 035104. doi:10.1103/PhysRevB.64.035104.
- [34] X. Blase, A. Rubio, S. G. Louie, M. L. Cohen, Quasiparticle band structure of bulk hexagonal boron nitride and related systems, *Phys. Rev. B* 51 (1995) 6868–6875. doi:10.1103/PhysRevB.51.6868.
- [35] S. Galambosi, L. Wirtz, J. A. Soininen, J. Serrano, A. Marini, K. Watanabe, T. Taniguchi, S. Huotari, A. Rubio, K. Hämäläinen, Anisotropic excitonic effects in the energy loss function of hexagonal boron nitride, *Phys. Rev. B* 83 (2011) 081413. doi:10.1103/PhysRevB.83.081413.
- [36] G. Fugallo, M. Aramini, J. Koskela, K. Watanabe, T. Taniguchi, M. Hakala, S. Huotari, M. Gatti, F. Sottile, Exciton energy-momentum map of hexagonal boron nitride, *Phys. Rev. B* 92 (2015) 165122. doi:10.1103/PhysRevB.92.165122.
- [37] P. Cudazzo, L. Sponza, C. Giorgetti, L. Reining, F. Sottile, M. Gatti, Exciton band structure in two-dimensional materials, *Phys. Rev. Lett.* 116 (2016) 066803. doi:10.1103/PhysRevLett.116.066803.
- [38] J. Koskela, G. Fugallo, M. Hakala, M. Gatti, F. Sottile, P. Cudazzo, Excitons in van der Waals materials: From monolayer to bulk hexagonal boron nitride, *Phys. Rev. B* 95 (2017) 035125. doi:10.1103/PhysRevB.95.035125.

- [39] W. Aggoune, C. Cocchi, D. Nabok, K. Rezouali, M. A. Belkhir, C. Draxl, Dimensionality of excitons in stacked van der Waals materials: The example of hexagonal boron nitride, *Phys. Rev. B* 97 (2018) 241114. doi:10.1103/PhysRevB.97.241114.
- [40] X. Gonze, F. Jollet, F. Abreu Araujo, D. Adams, B. Amadon, T. Applencourt, C. Audouze, J. M. Beuken, J. Bieder, A. Bokhanchuk, E. Bousquet, F. Bruneval, D. Caliste, M. Côté, F. Dahm, F. Da Pieve, M. Delaveau, M. Di Gennaro, B. Dorado, C. Espejo, G. Geneste, L. Genovese, A. Gerossier, M. Giantomassi, Y. Gillet, D. R. Hamann, L. He, G. Jomard, J. Laflamme Janssen, S. Le Roux, A. Levitt, A. Lherbier, F. Liu, I. Lukačević, A. Martin, C. Martins, M. J. T. Oliveira, S. Poncé, Y. Pouillon, T. Rangel, G. M. Rignanese, A. H. Romero, B. Rousseau, O. Rubel, A. A. Shukri, M. Stankovski, M. Torrent, M. J. Van Setten, B. Van Troeye, M. J. Verstraete, D. Waroquiers, J. Wiktor, B. Xu, A. Zhou, J. W. Zwanziger, Recent developments in the ABINIT software package, *Computer Physics Communications* 205 (2016) 106–131. doi:10.1016/j.cpc.2016.04.003.
- [41] Exc webpage, www.bethe-salpeter.org, accessed: 2019-11-26.
- [42] 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, *J. Phys. Condens. Matter.* 26 (36) (2014) 363202. doi:10.1088/0953-8984/26/36/363202.