Accepted Manuscript

Computing the density of states for optical spectra of molecules by low-rank and QTT tensor approximation

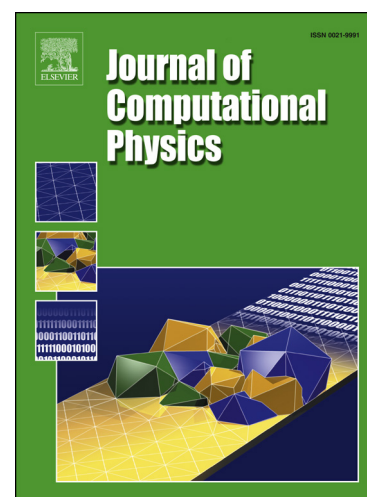Peter Benner, Venera Khoromskaia, Boris N. Khoromskij, Chao Yang

Please cite this article in press as: P. Benner et al., Computing the density of states for optical spectra of molecules by low-rank and QTT tensor approximation, *J. Comput. Phys.* (2019), https://doi.org/10.1016/j.jcp.2019.01.011

Highlights

- An economical method is introduced for calculating the traces of parametric matrix resolvents by taking advantage of the block-diagonal plus low-rank matrix structure.
- A regularized DOS can be accurately represented by a low rank quantized tensor train (QTT) interpolation.
- These techniques allow us to avoid both the need of stochastic sampling for calculating traces of matrix resolvents and the interpolation of DOS by problem independent functions like polynomials, plane waves, etc.

Computing the density of states for optical spectra of molecules by low-rank and QTT tensor approximation

Peter Benner[*]   Venera Khoromskaia[†]   Boris N. Khoromskij[‡]
Chao Yang[§]

Abstract

In this paper, we introduce a new interpolation scheme to approximate the density of states (DOS) for a class of rank-structured matrices with application to the Tamm-Dancoff approximation (TDA) of the Bethe-Salpeter equation (BSE). The presented approach for approximating the DOS is based on two main techniques. First, we propose an economical method for calculating the traces of parametric matrix resolvents at interpolation points by taking advantage of the block-diagonal plus low-rank matrix structure described in [6, 3] for the BSE/TDA problem. This allows us to overcome the computational difficulties of the traditional schemes since we avoid the construction of the matrix inverse and hence the need of stochastic sampling. Second, we show that a regularized or smoothed DOS discretized on a fine grid of size N can be accurately represented in a low rank quantized tensor train (QTT) format that can be determined through a least squares fitting procedure. The QTT tensor provides good approximation properties for strictly oscillating DOS functions with multiple gaps, in contrast to interpolation by problem independent functions like polynomials, trigonometric functions, etc. Moreover, the QTT approximant requires asymptotically much fewer (e.g., $O(\log N)$) functional calls compared with the full grid size N. Numerical tests indicate that the QTT approach yields accurate recovery of DOS associated with problems that contain relatively large spectral gaps. The QTT tensor rank only weakly depends on the size of a molecular system that paves the way for treatment of large-scale spectral problems.

[*]Max Planck Institute for Dynamics of Complex Technical Systems, Sandtorstr. 1, D-39106 Magdeburg, Germany (`benner@mpi-magdeburg.mpg.de`)

[†]Max Planck Institute for Mathematics in the Sciences, Leipzig; Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg (`vekh@mis.mpg.de`).

[‡]Max Planck Institute for Mathematics in the Sciences, Inselstr. 22-26, D-04103 Leipzig, Germany (`bokh@mis.mpg.de`).

[§]Berkeley Labs, Berkeley, USA (`cyang@lbl.gov`).

1   Introduction

Numerical approximation of the density of states (DOS) or spectral density (see §2.2) of large matrices is one of the challenging problems arising in the prediction of electronic, vibrational and thermal properties of molecules and crystals and many other applications. This topic, first developed in condensed matter physics [13, 49, 45, 12, 48], has long since attracted interest in the community of numerical linear algebra [46, 15, 44], see also a survey on commonly used methodology for approximation of DOS for large matrices of general structure [28].

Most of traditional methods are based on a polynomial or fractional-polynomial interpolation of the DOS regularized by Gaussians or Lorentzians, and computing traces of certain matrix valued functions, say matrix resolvents or polynomials, defined at a large set of interpolation points within the spectral interval of interest. The trace calculations are typically accomplished with stochastic sampling over a large number of random vectors [28], leading to Monte Carlo estimates with slow convergence rates and low accuracy. Moreover, the algorithms based on polynomial interpolants have poor approximating properties when the spectrum of a matrix exhibits gaps or highly oscillating non-regular shapes, as is the case in electronic structure calculations. Since the size of matrices resulting from real life applications is usually large (in quantum mechanics it scales as a polynomial of the molecular size), and the DOS of these matrices often exhibits very complicated shape, the above mentioned methods become prohibitively expensive.

In this paper we present a new method for efficient approximation of the DOS function for large rank-structured symmetric $n \times n$ matrices by fast evaluation of the traces of the matrix resolvent, which avoids both the solution of the eigenvalue problem and the calculation of the matrix inverse. Our main contribution is the development of a fast algorithm for evaluation of the DOS function at low cost that scales linearly in the matrix size. Furthermore, we achieve a reduction of the total number of function evaluations in the case of a fine representation grid.

We apply this approximation to the Bethe-Salpeter equation (BSE), which is a widely used model for *ab initio* estimation of the absorption spectra for molecules or surfaces of solids [38, 17, 43, 35, 30, 34]. In particular, we use the recently developed low-rank structured representation of the BSE Hamiltonian, which was introduced and analyzed in [6]. An efficient and structured eigenvalue solver for this block-diagonal plus low-rank representation of the BSE Hamiltonian as well as to its symmetric surrogate obtained by the Tamm-Dancoff approximation (TDA) is described in [3]. In the numerical tests, we confine ourselves to the construction of DOS for the symmetric positive definite matrix specified by the TDA model.

Our approach to approximating the DOS relies on the Lorentzian blurring [16]. In this case, the most computationally expensive part of the calculation amounts to the multiple evaluation of traces of shifted matrix inverses which, in general, scales cubically as $O(n^3)$ in the matrix size. The presented method reduces this cost to $O(n)$, by using the following two basic advancements.

First, we propose an economical method for calculating traces of parametric matrix resolvents at any chosen interpolation point by taking advantage of the block-diagonal plus low-rank BSE/TDA matrix structure described in [6, 3]. This becomes possible due to an explicit rank-structured representation of the matrix inverse which is evaluated efficiently

by using the Sherman-Morrison-Woodbury formula at $O(n)$ complexity. This allows us to overcome the computational difficulties of the traditional schemes since we avoid the construction of the full matrix inverse and the need of stochastic sampling. Note that the diagonal plus low-rank approximation to the BSE Hamiltonian introduced in [6] employs the low-rank approximation to the two-electron integrals tensor in the form of a Cholesky factorization developed previously [22] in the framework of a tensor-based Hartree-Fock solver [21].

The second novelty of this paper is the application of the QTT tensor approximation to the DOS sampled on a fine grid, which results in a long vector of size N displaying the spectrum of a molecule. The QTT approximation method was introduced and analyzed for function related vectors in [25], and later it proved to be useful in many applications. As a proof of concept, we first check by numerical experiments that the DOS function exhibits accurate low-rank QTT tensor approximation for strictly oscillating DOS functions with multiple gaps. The accuracy of approximation is controlled by ϵ-truncation of the corresponding tensor ranks.

These observations allow us to recover the QTT approximant of the DOS by using the TT/QTT-cross heuristic approximation that requires asymptotically much fewer (e.g., $O(\log N)$) functional calls compared with the full grid size N, thus avoiding the need for interpolation by problem independent functions like polynomials, trigonometric polynomials, etc. In this way, the QTT interpolant over $O(\log N)$ interpolation points provides a rather accurate representation of the functional N-vector of the DOS. Moreover, it resolves the positions of peaks (spikes) in DOS with good precision, which is significant from a physical point of view.

Numerical tests for moderate size molecules confirm the closeness of DOS for the TDA model to those computed on the exact BSE spectrum. We also justify that the simplified block-diagonal plus low-rank approximation of the TDA matrix recovers well the global landscape and location of spikes in the DOS curve on the whole energy interval. We demonstrate the almost linear complexity scaling of the trace calculation algorithm applied to TDA matrices of different size. Furthermore, we present the numerics on QTT tensor interpolation of DOS via the TT-cross approximation. We also sketch modifications necessary to calculate the optical absorption spectrum via a rank-structured BSE model.

The rest of the paper is structured as follows. In Section 2, we recall the main prerequisites for the description of our method including the rank-structured approximation of the BSE/TDA matrix, basic notions of the regularization of DOS by Lorentzians and a short summary on the existing methods for matrices of general structure. Section 3 discusses the main techniques of the presented method and the corresponding analysis in Theorems 3.2 and 3.3, and presents the corresponding numerics. Section 4 presents a short summary of the QTT tensor approximation of function related vectors and the analysis of the QTT tensor ranks of the DOS, see Theorem 4.1. In Section 4.4 the TT-cross based QTT interpolation is applied to the discretized DOS function, where the quality of the interpolation is illustrated numerically. The beneficial features of the new computational schemes are verified by extensive numerical experiments on the examples of various molecular systems. Section 5 outlines the possible extension of the approach to the case of full BSE system. Conclusions summarize the main results and address the application perspectives.

3

2 Main prerequisites and outline of initial applications

2.1 Rank-structured approximation to BSE matrix

In this paper we describe a method for efficient and accurate approximation of the DOS for large rank-structured symmetric matrices. Our basic application is concerned with estimating the DOS and the absorption spectrum for the Bethe-Salpeter problem describing the excitation energies of molecules.

The 2×2-block matrix representation of the Bethe-Salpeter Hamiltonian (BSH) leads to the following eigenvalue problem.

$$H \begin{pmatrix} \mathbf{x}_k \\ \mathbf{y}_k \end{pmatrix} \equiv \begin{pmatrix} A & B \\ -B^* & -A^* \end{pmatrix} \begin{pmatrix} \mathbf{x}_k \\ \mathbf{y}_k \end{pmatrix} = \omega_k \begin{pmatrix} \mathbf{x}_k \\ \mathbf{y}_k \end{pmatrix}, \tag{2.1}$$

where the matrix blocks of size $n \times n$, with $n = N_{ov} = N_o(N_b - N_o)$, are defined by

$$A = \mathbf{\Delta \varepsilon} + V - \widehat{W}, \quad B = V - \widetilde{W}, \tag{2.2}$$

and eigenvalues ω_k correspond to the excitation energies. Here $\mathbf{\Delta \varepsilon}$ is a diagonal matrix and

$$V = [v_{ia,jb}] \quad a, b \in \mathcal{I}_v := \{N_o + 1, \ldots, N_b\}, \quad i, j \in \mathcal{I}_o := \{1, \ldots, N_o\},$$

is the rank-R_B two-electron integrals (TEI) matrix projected onto the Hartree-Fock molecular orbital basis, where N_b is the number of Gaussian type orbital (GTO) basis functions and N_o denotes the number of occupied orbitals [6].

The method for solving the Bethe-Salpeter equation (BSE) using low-rank factorizations of the generating matrices has been introduced in [6]. It is based on a tensor-structured grid-based Hartree-Fock (HF) solver which provides not only the full set of eigenvalues and HF orbitals, but also the two-electron integrals tensor in the form of a low-rank Cholesky factorization, see [21] and references therein.

The matrix V inherits its low rank from the two-electron integrals tensor, and $\widetilde{W}$ is also proven to have a small ϵ-rank (see [6]). In particular, there holds

$$V \approx L_V L_V^T, \quad L_V \in \mathbb{R}^{n \times R_V}, \quad R_V \leq R_B, \tag{2.3}$$

with the rank estimates $R_V = R_V(\varepsilon) = \mathcal{O}(N_b |\log \varepsilon|)$, and $\text{rank}(\widetilde{W}) \leq \text{rank}(V)$.

In [3], it was shown that the matrix $\widehat{W}$, which does not exhibit an accurate low rank representation, can be well approximated by a block diagonal matrix

$$\widehat{W} \approx \text{blockdiag}[\widehat{B}, D],$$

where $\widehat{B}$ is a $N_W \times N_W$ dense block with $N_W = O(n^\alpha)$, $\alpha < 1$. In our numerical examples below the block-size N_W is almost of the same order as the rank parameter of L_V. As a result, the TDA matrix A can be approximated by a sum of a block-diagonal matrix and a low rank matrix shown in Figure 2.1, i.e.,

$$A \approx \widehat{A} = \mathbf{\Delta \varepsilon} + QQ^T - \text{blockdiag}[\widehat{B}, D] \equiv \text{blockdiag}[B_0, D_0] + QQ^T,$$

4

Figure 2.1: Diagonal plus low-rank plus reduced-block structure of the matrix $\widehat{A}$.

where $Q = L_V$, see also Section 3.

An efficient structured solver designed to calculate a number of minimal eigenvalues of the block-diagonal plus low-rank representation of the BSE/TDA matrices is described in [3]. It is based on an efficient subspace iteration of the matrix inverse, which for rank-structured matrix formats can be evaluated efficiently by using the Sherman-Morrison-Woodbury formula, thus reducing the numerical expense of the direct diagonalization down to $\mathcal{O}(N_b^2)$ in the size of the atomic orbitals basis set, N_b[1]. Furthermore, this solver also includes a QTT-based compression scheme, where both eigenvectors and the rank-structured BSE matrix blocks are represented by block-QTT tensors. The block-QTT representation of the eigenvector is determined by an alternating least squares (ALS) iterative algorithm. The overall asymptotic complexity for computing several smallest in modulo eigenvalues in the BSE spectral problem by using the QTT approximation is estimated by $\mathcal{O}(\log(N_o)N_o^2)$, where N_o is the number of occupied orbitals.

Matrices in the form (2.1) are called J-symmetric or Hamiltonian, see [5] for implications on the algebraic properties of the BSE matrix. In particular, solutions of equation (2.1) come in pairs: excitation energies ω_k with eigenvectors $(\mathbf{x}_k, \mathbf{y}_k)$, and de-excitation energies $-\omega_k$ with eigenvectors $(\mathbf{y}_k^*, \mathbf{x}_k^*)$.

The simplification in the BSH, H, defined by the $n \times n$ symmetric diagonal block A is called the Tamm-Dancoff (TDA) approximation. In what follows, we are interested in the TDA spectral problem,

$$A\mathbf{u}_k = \lambda_k \mathbf{u}_k, \quad k = 1, \ldots, n,$$

providing good approximations to $\omega_k, \mathbf{x}_k$.

In general, methods for solving partial eigenvalue problems for matrices with a special structure as in the BSE setting are conceptually related to the approaches for Hamiltonian matrices [4, 7, 14, 9], particularly to those based on minimization principles [1, 2]. A structured Lanczos algorithm for estimation of the optical absorption spectrum was described in [41]. Various structured eigensolvers tailored for electronic structure calculations are discussed in [36, 37, 10, 29, 28, 42].

---

[1]Notice that a more accurate cost estimate is $\mathcal{O}(N_o N_b)$, where $N_o \sim N_b/10$ for typical basis sets we used so far. In general, the above estimate only reduces the constant in the basic complexity bound $\mathcal{O}(N_b^2)$.

2.2 Density of states for symmetric matrices

To fix the idea, we first consider the case of symmetric matrices. Following [28], we use the simple definition of the DOS for symmetric matrices

$$\phi(t) = \frac{1}{n} \sum_{j=1}^{n} \delta(t - \lambda_j), \quad t, \lambda_j \in [0, a], \tag{2.4}$$

where δ is the Dirac distribution and the λ_j's are the eigenvalues of $A = A^T$ ordered as $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_n \leq a$. Here $[0, a]$ is the energy interval of interest.

Several classes of blurring approximations to $\phi(t)$ are used in the literature. One can replace each Dirac-δ by a Gaussian function with width $\eta > 0$, i.e.,

$$\delta(t) \rightsquigarrow g_\eta(t) = \frac{1}{\sqrt{2\pi}\eta} \exp\left(-\frac{t^2}{2\eta^2}\right), \tag{2.5}$$

where the choice of the regularization parameter η depends on the particular problem setting. As a result, (2.4) can be approximated by

$$\phi(t) \approx \phi_\eta(t) := \frac{1}{n} \sum_{j=1}^{n} g_\eta(t - \lambda_j), \tag{2.6}$$

on the whole energy interval $t \in [0, a]$.

We may also replace each Dirac-δ by a Lorentzian, i.e.,

$$\delta(t) \rightsquigarrow L_\eta(t) := \frac{1}{\pi} \frac{\eta}{t^2 + \eta^2} = \frac{1}{\pi} \text{Im}\left(\frac{1}{t - i\eta}\right), \tag{2.7}$$

so that an approximate DOS can be written as

$$\phi(t) \approx \phi_\eta(t) := \frac{1}{n} \sum_{j=1}^{n} L_\eta(t - \lambda_j). \tag{2.8}$$

When $\eta \to 0_+$, both Gaussians and Lorentzians converge to the Dirac distribution, i.e.,

$$\lim_{\eta \to 0_+} g_\eta(t) = \lim_{\eta \to 0_+} L_\eta(t) = \delta(t).$$

However, they exhibit different features of the approximant for small $\eta > 0$. In the case of Gaussians, one expects a sharp resolution of the spectral peaks, while the Lorentzian based representation aims to resolve better the global landscape of $\phi(t)$.

Both functions $\phi_\eta(t)$ and $L_\eta(t)$ are continuous, hence, they can be discretized by sampling on a fine grid Ω_h over $[0, a]$. In the following, we use the uniform cell-centered N-point grid with the mesh size $h = a/N$.

In what follows, we focus on the case of Lorentzian blurring, which will be motivated later on, and apply it to the TDA approximation of the BSE problem (see §2.1 below). We use the simplified block-diagonal plus low-rank approximation to the matrix A, see [6, 3], which allows efficient explicit representation of the shifted inverse matrix.
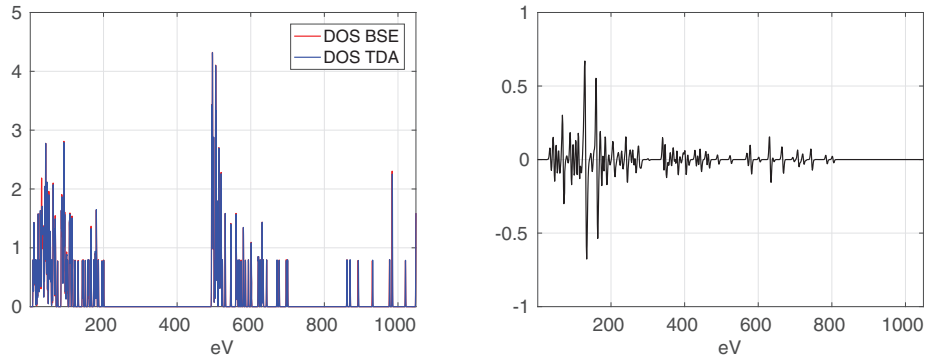
6

Figure 2.2: DOS for H_2O, $\eta = 0.5$: exact BSE vs. TDA on the full spectrum (left), the corresponding error (right).

The numerical illustrations in §2.2 represent the DOS for the H_2O molecule and H_2 chains broadened by Gaussians (2.6). The data corresponds to the reduced basis approach via rank-structured approximation applied to the symmetric TDA model [6, 3] described by the matrix block A of the full BSE system matrix.

It was numerically demonstrated in [6] that the spectrum of the TDA model provides a good approximation to the spectrum of the full BSE Hamiltonian. The difference between the two is on the order of 10^{-2} for molecules of moderate size.


Figure 2.3: DOS for H_2O on the energy sub-interval $[0, 40]$: exact BSE vs. TDA (left), and the error (right).

Figure 2.2, left, compares the DOS for the H_2O molecule calculated via the eigenvalues of the full BSE Hamiltonian and those of the TDA approximation, while on the right we display the corresponding error.

Figure 2.3, left, compares the same DOS calculations but zoomed on the first compact energy interval $[0, 40]$ eV. The red curve corresponds to the full BSE data, and the blue one represents the TDA case. The figure on the right displays the corresponding error. We observe that the maximal error amplitude corresponds to the model error between the full BSE and TDA systems, see Figure 2.3, right.

Figure 2.4, left, represents the DOS for H_2O computed by using the exact TDA spectrum

7

(blue) and its approximation based on a simplified model obtained via low-rank approximation to A (red), while the right figure shows the zoom in the energy interval $[3, 25]$ eV. We observe that the locations of spikes are well recovered. For example, the first two spikes for H_2O molecule are positioned at 8.7 eV and 10.8 eV, respectively, that are close to quantities known in the literature on the optical spectrum of water (cf. [18]).

Figures 2.5 presents the DOS for H_{16} (left) and H_{32} (right) chains of Hydrogen atoms. We observe the essential similarity in the shapes (only the amplitude is changing) which is apparently a consequence of quasi-periodicity of the system.
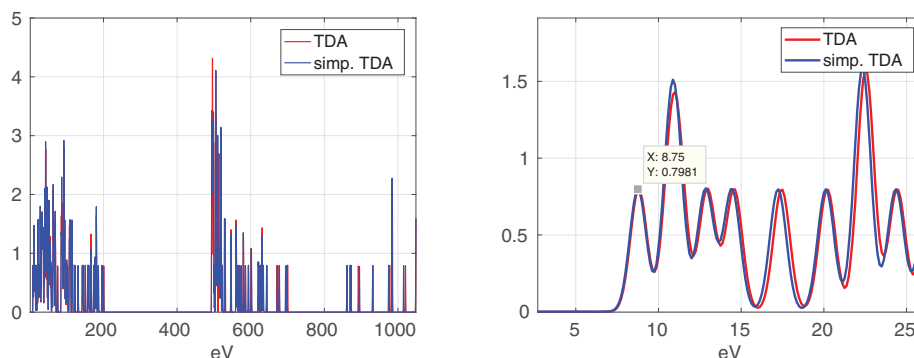

Figure 2.4: DOS for H_2O. Exact TDA vs. simplified TDA (left), zoom of the small spectral interval (right).
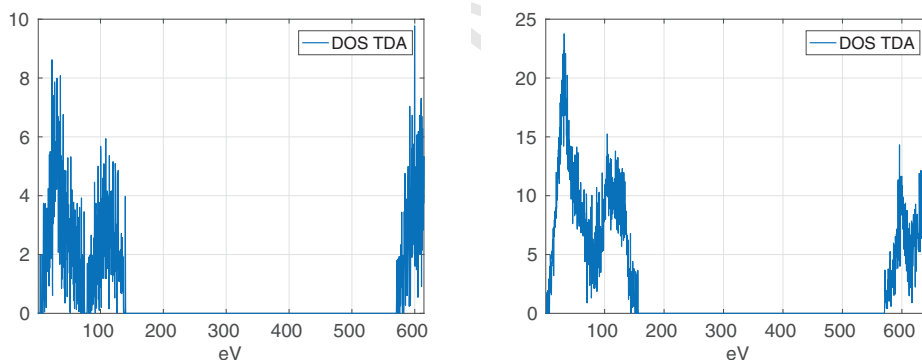
Figure 2.5: DOS for H_{16} (left) and H_{32} (right) chains of Hydrogen atoms.

The rank-structured approach to calculation of the molecular absorption spectrum in the case of full BSE is sketched in §5. This topic will be addressed elsewhere.

2.3 General description of the existing computational schemes

One of the commonly used approaches to the numerical approximation of both functions $g_\eta(t)$ and $L_\eta(t)$ is based on the construction of certain polynomial or fractional polynomial interpolants whose evaluation at each sampling point t_k requires the solution of a large linear system with the BSE/TDA matrix, i.e., remains expensive.

8

In the case of Lorentzian broadening (2.8) the regularized DOS takes the form

$$\phi(t) \approx \phi_\eta(t) := \frac{1}{n\pi} \sum_{j=1}^{n} \mathrm{Im}\left(\frac{1}{(t-\lambda_j)-i\eta}\right) = \frac{1}{n\pi}\mathrm{Im}\,\mathrm{Trace}[(tI - A - i\eta I)^{-1}]. \qquad (2.9)$$

To keep real-valued arithmetics, likewise, we can write the latter in the form

$$\phi_\eta(t) := \frac{1}{n\pi} \sum_{j=1}^{n} \frac{\eta}{(t-\lambda_j)^2 + \eta^2} = \frac{1}{n\pi}\mathrm{Trace}[((tI - A)^2 + \eta^2 I)^{-1}]. \qquad (2.10)$$

In both cases the task of computing the approximate DOS $\phi_\eta(t)$ reduces to approximating the trace of the matrix resolvent

$$(tI - A - i\eta I)^{-1} \quad \text{or} \quad ((tI - A)^2 + \eta^2 I)^{-1}.$$

Here, the price to pay for real-valued arithmetics is to address the more complicated low-rank structure in $(tI - A)^2$.

The traditional approach [28] to approximately computing the traces of the matrix-valued analytic function $f(A)$ reduces this task to the estimation of the mean of $v_m^T f(A) v_m$ over a sequence of random vectors v_m, $m = 1, \ldots, m_r$, satisfying certain orthogonality conditions (see [28], Theorem 3.1 for the detailed discussion of stochastic sampling in such a computational scheme). That is, $\mathrm{Trace}[f(A)]$ is approximated by

$$\mathrm{Trace}[f(A)] \approx \frac{1}{m_r} \sum_{m=1}^{m_r} v_m^T f(A) v_m. \qquad (2.11)$$

The calculation of (2.11) for

$$f_1(A) = (tI - A - i\eta I)^{-1} \quad \text{or} \quad f_2(A) = ((tI - A)^2 + \eta^2 I)^{-1} \qquad (2.12)$$

reduces to solving linear systems in the form of

$$(tI - i\eta I - A)x = v_m \quad \text{for} \quad m = 1, \ldots, m_r, \qquad (2.13)$$

or

$$(\eta^2 I + (tI - A)^2)x = v_m \quad \text{for} \quad m = 1, \ldots, m_r. \qquad (2.14)$$

These linear systems need to be solved for many target points $t = t_k \in [a, b]$ in the course of a chosen interpolation scheme.

In the case of rank-structured matrices A, the solution of equations (2.13) or (2.14) can be implemented with a lower cost. However, even in this favorable situation one requires a relatively large number m_r of stochastic realizations to obtain satisfactory mean value approximation. The convergence rate is expected to be of the order of $O(1/\sqrt{m_r})$. On the other hand, with the limited number of interpolation points, the polynomial type of interpolation schemes applied to highly non-regular shapes as shown, say, in Figure 2.4 (left), can only provide limited resolution and is unlikely to reveal spectral gaps and many local spikes of interest.

In what follows, we propose the new approach for fast calculation of the DOS function in the form of (2.9) and (2.10) that is based on evaluating the trace terms directly (without stochastic sampling).

9

3 Fast evaluation of DOS for rank-structured matrices

In this section we describe fast algorithms for evaluation of the DOS for rank-structured matrices at a fixed point in the energy interval $t \in [0, a]$. This approach relies on the explicit expression for the trace of a matrix resolvent, which we introduce below for the class of block-diagonal plus low-rank matrices arising in the reduced model approach for the BSE problem [6, 3].

3.1 DOS by the trace of rank-structured matrix inverse

We consider the rank-structured matrices of the form (see §2.1 for more details)

$$A = E_0 + PQ^T, \quad \text{with} \quad P, Q \in \mathbb{R}^{n \times R}, \quad E_0 = \text{blockdiag}\{B_0, D_0\}, \tag{3.1}$$

where the rank parameter R is small compared to n, and the full $N_W \times N_W$ matrix block B_0 is of small size. Here D_0 is a diagonal matrix of size $n - N_W$.

Remark 3.1 *In our applications to DOS calculation for the BSE/TDA matrix, we normally have the relation $N_W = O(n^\alpha)$, with some $0 < \alpha < 1$ and constant depending on R. In particular, for the large molecular systems to be considered in the following numerical examples we have the relation $N_W \approx C_0(n\,R^2)^\alpha$ with $\alpha = 1/3$ and $C_0 = 2$, which implies the almost linear complexity scaling of our algorithm in the matrix size n.*

Notice that even in the case of structured matrices in (3.1) the traditional approach by (2.11) leads to a sequence of linear systems (2.13) to be solved many times in the course of stochastic sampling, for each of many interpolation points $t \in [0, a]$.

In our approach, for the class of rank-structured matrices (3.1), we propose to avoid stochastic sampling in (2.11) by introducing a direct scheme that allows us to evaluate the trace of matrices $f_1(A)$ or $f_2(A)$ defined in (2.12), corresponding to the matrix resolvent in (2.9) and (2.10), respectively, by one-step straightforward matrix calculation.

To that end, let us first construct the reduced-model approximation to the matrix resolvent $S(t)^{-1} := f_1(A)$ for the matrix in (3.1), where $S(t)$ denotes the diagonal shift of A depending on the parameter t,

$$S(t) = tI - E_0 + PQ^T - i\eta I =: E(t) + PQ^T. \tag{3.2}$$

Here the block-diagonal part E_0 is modified by the diagonal shift,

$$E(t) = E_0 + tI - i\eta I \equiv \text{blockdiag}\{B(t), D(t)\}$$

corresponding to the case of (2.9), i.e.,

$$B(t) = tI_B - i\eta I_B + B_0, \quad D(t) = tI_D - i\eta I_D + D_0. \tag{3.3}$$

Here B_0 and D_0 denote the corresponding matrix blocks in the representation of the diagonal block A in the initial BSE matrix, see (3.1), and I_B, I_D denote the identity matrices corresponding to the respective index subsets. For the ease of exposition, we further assume that the matrix size of the block B in (3.3) is bounded by $N_W = O(n^\alpha)$ with $\alpha \le 1/3$, see

10

Remark 3.1. This assumption on the block size ensures the linear complexity scaling of our algorithm in the matrix size n.

In what follows, we use the notion $\mathbf{1}_m$ for a length-m vector of all ones, and $\odot$ for the Hadamard product of matrices.

The following result asserts that the cost of trace calculations is estimated to be $O(nR)$.

Theorem 3.2 *Let the matrix family $S(t)$, $t \in [0, a]$, be given by (3.2), with*

$$E(t) = \text{blockdiag}\{B(t), D(t)\},$$

where $B(t), D(t)$ are defined in (3.3). Then the trace of the matrix inverse $S(t)^{-1}$ can be calculated explicitly by

$$\text{trace}[S(t)^{-1}] = \text{trace}[B(t)^{-1}] + \text{trace}[D(t)^{-1}] - \mathbf{1}_n^T(U(t) \odot V(t))\mathbf{1}_R, \qquad (3.4)$$

where $U(t) = E(t)^{-1}PK(t)^{-1} \in \mathbb{R}^{n \times R}$, $V(t) = E(t)^{-1}Q \in \mathbb{R}^{n \times R}$, and

$$K(t) = I_R + Q^T E(t)^{-1}(t)P$$

is a small $R \times R$ matrix. For fixed $t \in [0, a]$, the numerical cost is estimated by $O(N_W^3 + nR^2)$ provided that N_W and R are of the same order.

Proof. The analysis relies on the particular structure of the matrix blocks. Indeed, we use the direct trace representation for both rank-R and block-diagonal matrices. Our argument is based on the observation that the trace of a rank-R matrix $U(t)V(t)^T$, where $U(t), V(t) \in \mathbb{R}^{n \times R}$, $U(t) = [\mathbf{u}_1, \ldots, \mathbf{u}_R]$, $V(t) = [\mathbf{v}_1, \ldots, \mathbf{v}_R]$, $\mathbf{u}_k, \mathbf{v}_k \in \mathbb{R}^n$, can be calculated in terms of skeleton vectors by

$$\text{trace}[U(t)V(t)^T] = \sum_{k=1}^R \langle \mathbf{u}_k, \mathbf{v}_k \rangle = \mathbf{1}_n^T(U(t) \odot V(t))\mathbf{1}_R,$$

at the expense $O(Rn)$. For fixed t, define the rank-R matrices by

$$U(t) = E(t)^{-1}PK(t)^{-1}, \quad V(t) = E(t)^{-1}Q,$$

then the Sherman-Morrison-Woodbury scheme leads to the representation, see [3],

$$S(t)^{-1} = \text{blockdiag}\{B(t)^{-1}, D(t)^{-1}\} - E(t)^{-1}PK(t)^{-1}Q^T E(t)^{-1}, \qquad (3.5)$$

where the last term simplifies to

$$E(t)^{-1}PK(t)^{-1}Q^T E(t)^{-1} = U(t)V(t)^T.$$

Now we apply the above formula for the trace of a rank-R matrix to obtain the desired representation.

The complexity estimate follows taking into account the bound on the size of matrix block B. Indeed, forming $U(t)$ involves solving the linear system $P_1(t) = U(t)K(t)$, for $U(t)$, where $P_1(t)$ is the pre-computed $E(t)^{-1}P$, which can be evaluated at the cost $O(N_W^3 + N_W^2 R + nR)$. Here $P_1(t)$ would be re-used to compute $K(t)$ itself, and thus stored. The cost for solving this

11

system of equations is $2/3R^3$ (LU factorization of $K(t)$), plus $2nR^2$ for backward/forward solves. This completes the proof. ∎

The above procedure, see (3.4), has to be applied many times for calculating the trace of $E(t_m)^{-1}PK(t_m)^{-1}Q^T E(t_m)^{-1}$ at each fixed interpolation point t_m, $m = 1, \ldots, M$, $M \leq N$, to represent the DOS function on a fine N-point grid. The interpolation points $\{t_m\}$ are predefined by the chosen interpolating scheme.

Here, we notice that the price to pay for the real arithmetics in equation (2.14) is that we compute with squared matrices which, however, do not increase the asymptotic complexity since there is no increase of the rank in the rank-structured representation of the system matrix, see the following Theorem 3.3. In what follows we denote by $[U, V]$ the concatenation of two matrices of compatible size. We confine ourselves to the symmetric case $P = Q$, resulting from the TDA model in which the matrix A is symmetric. The DOS calculation for non-symmetric matrices like in the case of BSE problem is beyond the scope of this paper.

Theorem 3.3 *Given matrix $S(t) = (tI - A)^2 + \eta^2 I$, where A is defined by (3.1) with $P = Q$, then the trace of the real-valued matrix resolvent $S^{-1}(t)$ can be calculated explicitly by*

$$\text{trace}[S(t)^{-1}] = \text{trace}[\overline{E}(t)^{-1}] - \mathbf{1}_n^T (\overline{U}(t) \odot \overline{V}(t))\mathbf{1}_{2R}, \tag{3.6}$$

with $\overline{U}(t) = \overline{E}(t)^{-1}\overline{P}(t)K(t)^{-1} \in \mathbb{R}^{n \times 2R}$, and $\overline{V}(t) = \overline{E}(t)^{-1}\overline{Q} \in \mathbb{R}^{n \times 2R}$, where the real-valued block-diagonal matrix $\overline{E}(t)$ is given by

$$\overline{E}(t)(t) = \eta^2 I + t^2 I - 2tE_0 + E_0^2 = (\eta^2 + t^2)I + \text{blockdiag}[B^2 - 2tB, D^2 - 2tD],$$

and the rank-$2R$ matrices $\overline{P}(t), \overline{Q}$ are represented via concatenation

$$\overline{P}(t) = [-2tQ + E_0Q + QE_0 + Q(Q^TQ), Q] \in \mathbb{R}^{n \times 2R}, \quad \overline{Q} = [Q, E_0Q] \in \mathbb{R}^{n \times 2R},$$

such that the small core matrix $K(t) \in \mathbb{R}^{2R \times 2R}$ takes the form $K(t) = I_R + \overline{Q}^T \overline{E}(t)^{-1}\overline{P}(t)$.

For fixed $t \in [0, a]$, the numerical cost is estimated by $O(N_W^3 + nR^2)$ provided that N_W and R are of the same order.

Proof. Indeed, given the block-diagonal plus low-rank matrix A in the form (3.1), we obtain

$$S(t) = (tI - A)^2 + \eta^2 I = \overline{E}(t) + \overline{P}(t)\overline{Q}^T,$$

where the block-diagonal matrix $\overline{E}(t)$ and the rank-$2R$ matrix $\overline{P}(t)\overline{Q}^T$ are defined as above. Applying the Sherman-Morrison-Woodbury scheme as above to the block-diagonal plus rank-$2R$ matrix structure in S, the representation result follows. Now we take into account that

$$\text{trace}[\overline{E}(t)^{-1}] = \text{trace}[(B^2 - 2tB)^{-1}] + \text{trace}[(D^2 - 2tD)^{-1}],$$

then the restriction on the size of the block B proves the complexity bound by the similar argument as in the proof of Theorem 3.2. ∎

Based on Theorems 3.2 and 3.3, the calculations of traces for the parametric matrix resolvent can be implemented efficiently in both complex and real arithmetics. The following numerics demonstrates the efficiency of the DOS calculation for the rank-structured TDA matrix implemented in real arithmetics as described by (3.6) in Theorem 3.3.

Figure 3.1: Left: DOS for H$_2$O vs. its recovering by using the trace of matrix resolvents; Right: zoom on the small energy interval.


Figure 3.2: Left: DOS for Ethanol molecule vs. its recovering by using the trace of matrix resolvents; Right: zoom on the small energy interval.

Figures 3.1 and 3.2 demonstrate that using only the structure-based trace representation (3.6) in Theorem 3.3, we obtain the approximation which resolves well the shape of DOS function and positions of spikes on the examples of H$_2$O and C$_2$H$_2$OH (Ethanol) molecules. The exact DOS is shown by the blue line, while the results of structure-based DOS calculation is indicated by the red line (we use the acronym "SMW" for the Sherman-Morrison-Woodbury scheme).

Figure 3.3 shows the rescaled CPU time, i.e. $T_0 = T/R^2$, where T denotes the total CPU time for computing the DOS by the algorithm implied by Theorem 3.3. This demonstrates almost linear complexity scaling of the algorithm in n, $O(R^2 n)$. The dashed line shows linear scaling $T_0 = C n$ with $C = 2.3$. We applied the algorithm to molecules of different system size n (i.e. the size of TDA matrix) varying from $n = 180$ to $n = 4488$ (see Table 3.1 for more details). In all cases the N-point representation grid with fixed $N = 2^{14}$ was used. We point out that in the particular numerical tests for the BSE matrix, the technical condition on the size of the fully populated matrix block, $N_W = O(n^\alpha)$ with $\alpha \le 1/3$, implies linear complexity scaling of the algorithm in the matrix size n, see Remark 3.1. We control the performance by verifying the linear computational cost in the matrix size n, see Figure 3.3.

We conclude that the algorithm based on representation (3.6) demonstrates the good

13

Figure 3.3: Algorithm based on Theorem 3.3: the rescaled CPU time $T_0 = T/R^2$ versus n. The dashed line shows linear scaling $T_0 = C\,n$ with $C = 2.3$.

resolution of the DOS function at linear complexity in the system size n which allows to treat large molecules.

3.2 Calculating multiple traces of A^{-1} with lower cost

Any particular interpolation scheme usually requires DOS function evaluation for many different parameters $t_m \in \tau = \{t_1, \ldots, t_M\} \subset [0, a]$, $M \leq N$, in the matrix resolvent. Finer resolution of the spectrum for large molecular systems leads to a considerable increase of the number of samples M that might be practically equal to the grid size, $M \approx N$. Hence, the total cost $O(MnR^2)$ may become prohibitively expensive since the trace computation for each fixed value of t_m still requires complicated matrix operations (see Theorems 3.2 and 3.3).

In this section, we describe a further enhancement scheme for fast multiple calculation of traces on the large set of interpolation points. We outline how it is possible to reduce the complexity of these calculations (reduced model) by using a certain smoothness in t in the parametric matrix resolvent by introducing the low rank approximation of the large $n^2 \times M$ matrices

$$\mathbb{E}_M = [vec(E(t_1)^{-1}), \ldots, vec(E(t_M)^{-1})] \in \mathbb{R}^{n^2 \times M}$$

Molecule	H_2O	NH_3	H_2O_2	N_2H_4	C_2H_5OH	$C_2H_5\ NO_2$	$C_3H_7\ NO_2$
$n = N_{ov}$	180	215	531	657	1430	3000	4488
Rank R	36	30	68	54	74	129	147
Total time T (s)	6.7	7.7	33	47	219	1084	2223
Scaled time T/R^2 (s)	0.005	0.008	0.007	0.017	0.041	0.065	0.103

Table 3.1: Scaled times for the Algorithm in Theorem 3.3.

14

and

$$\mathbb{K}_M = [vec(K(t_1)^{-1}), \ldots, vec(K(t_M)^{-1})] \in \mathbb{R}^{R^2 \times M},$$

obtained by concatenation of vectorized matrices $E(t_m)^{-1}$ and $K(t_m)^{-1}$, $m = 1, \ldots, M$, respectively. The idea is that

$$E(t)^{-1} = \text{blockdiag}[P(t)^{-1}, D(t)^{-1}]$$

defines an analytic matrix family on the spectral interval $t \in [0, a]$, and so is the family of core matrices $\{K^{-1}(t)\}$. This property allows the model reduction via low rank approximation of the matrices $\mathbb{E}_M$ and $\mathbb{K}_M$, $t_m \in \tau$. Suppose that the approximate representations

$$K(t_m)^{-1} = \sum_{k=1}^{R_K} c_k(t_m) K_k, \quad m = 1, \ldots, M,$$

and

$$E(t_m)^{-1} = \text{blockdiag}[P(t_m)^{-1}, D(t_m)^{-1}] = \sum_{\ell=1}^{R_E} p_\ell(t_m) E_\ell, \quad m = 1, \ldots, M$$

are precomputed with small rank parameters R_K and R_E, where matrices $E_\ell = \text{blockdiag}[P_\ell, D_\ell] \in \mathbb{R}^{n \times n}$ and $K_k \in \mathbb{R}^{R \times R}$ do not depend on t_m, and E_ℓ inherits the block-diagonal structure that $E(t)^{-1}$ obeys. This can be understood as a low-rank approximation procedure which separates the parameter t_m and the matrix index. This may reduce the cost of trace calculation in the case of a large number of sampling points M.

To that end we take into account that Q does not depend on t, and plug the above decompositions in the representation of $\text{Trace}[S(t_m)^{-1}]$ in (3.5) to obtain

$$\text{Trace}[E^{-1}QK^{-1}Q^T E^{-1}(t_m)] = \text{Trace}\left[\sum_{\ell=1}^{R_E} p_\ell(t_m) E_\ell\, Q \left(\sum_{k=1}^{R_K} c_k(t_m) K_k \right) Q^T \sum_{\ell'=1}^{R_E} p_{\ell'}(t_m) E_{\ell'} \right]$$

for all $t_m \in \tau$. Now it follows that

$$\text{Trace}[E^{-1}QK^{-1}Q^T E^{-1}(t_m)] = \sum_{\ell=1}^{R_E} p_\ell(t_m) \sum_{k=1}^{R_K} c_k(t_m) \sum_{\ell'=1}^{R_E} p_{\ell'}(t_m) \text{Trace}[E_\ell Q K_k Q^T E_{\ell'}],$$

where $K_k \in \mathbb{R}^{R \times R}$ is a small matrix, $Q \in \mathbb{R}^{n \times R}$, $E_\ell = \text{blockdiag}[P_\ell, D_\ell]$ with diagonal D_ℓ and the full $n_P \times n_P$ matrix P_ℓ, such that n_P is much smaller than n, see Remark 3.1 for further details.

With these prerequisites, we pre-compute a set of "time-independent" traces

$$T_{\ell k \ell'} = \text{Trace}[E_\ell Q K_k Q^T E_{\ell'}], \quad \ell, \ell' = 1, \ldots, R_E, \quad k = 1, \ldots, R_K, \tag{3.7}$$

and store the $\frac{1}{2}R_E^2 R_K$ numbers $T_{\ell k \ell'}$ (symmetric tensor) to obtain the cheap representation of the trace in terms of only a scalar sum for $t_m \in \tau$

$$\text{Trace}[E^{-1}QK^{-1}Q^T E^{-1}](t_m) = \sum_{\ell=1}^{R_E} \sum_{k=1}^{R_K} \sum_{\ell'=1}^{R_E} p_\ell(t_m) c_k(t_m) p_{\ell'}(t_m) T_{\ell k \ell'}.$$

15

The cost of pre-computing each trace-value $T_{mkm'}$ is estimated by $O(n_P^3 + n_P^2 R)$ as justified in the proof of Theorem 3.2. The number of coefficients to be stored is about $O(R_E R_K)$ and it is expected to be small or moderate. With these data at hand, the evaluation of the required trace for the particular $t_\nu \in \tau$ takes $O(R_E^2 R_K)$ scalar operations independent of n.

Notice that the computations in (3.7) are intrinsically parallel, which can be exploited on modern computing hardware using multi-threading or distributed computing.

4   QTT tensor approximation of DOS

In the recent decade, the tensor-structured numerical methods are becoming ubiquitous in solving the multidimensional problems in scientific computing [26, 21]. The success of the tensor approach lies in the ability to reduce the numerical solution of the multivariate integral-differential equations to essentially one-dimensional operations. This was motivated by the previous nonlinear approximation theory (see [26, 24] for the detailed discussion), where for some classes of function related tensors the exponentially fast convergence of their rank-structured approximation was proven.

4.1   Tensor decompositions, tensor-train format

A real-valued tensor $\mathbf{A}$ of order d is defined as an element of the linear finite dimensional Hilbert space $\mathbb{W}_{\mathbf{n}} = \mathbb{R}^{n_1 \times \dots \times n_d}$ such that its entry-wise representation reads

$$\mathbf{A} = [a_{i_1,\dots,i_d}] \equiv [a(i_1,\dots,i_d)] \equiv [a_{\mathbf{i}}], \quad i_\ell = \{1,\dots,n_\ell\} \in I_\ell,$$

with the index set $\mathbf{i} \in \mathcal{I} = I_1 \times \dots \times I_d$, and the Euclidean scalar product, defined by

$$\langle \mathbf{A}, \mathbf{B} \rangle := \sum_{\mathbf{i} \in \mathcal{I}} a_{\mathbf{i}} b_{\mathbf{i}}, \quad \mathbf{A}, \mathbf{B} \in \mathbb{W}_{\mathbf{n}}.$$

Storage size for a dth order tensor scales exponentially in d, as n^d (for simplicity, let $n_\ell = n$), that causes the so-called "curse of dimensionality". The rank-structured parametrization of a tensor provides a mechanism to avoid or reduce the curse of dimensionality. The canonical and Tucker (additive) tensor formats developed in multilinear algebra (see the survey paper [27]) are constructed by linear combination of the simplest separable representations given by rank-1 tensors,

$$\mathbf{U} = \mathbf{u}^{(1)} \otimes \dots \otimes \mathbf{u}^{(d)} \in \mathbb{R}^{n_1 \times \dots \times n_d}, \quad \mathbf{u}^{(\ell)} \in \mathbb{R}^n,$$

with entries given by $u_{i_1,\dots i_d} = u_{i_1}^{(1)} \cdots u_{i_d}^{(d)}$, which can be stored using dn real numbers.

Important motivation for tensor numerical methods was the renewal of the factorized (multiplicative) representation of dth order tensors in a form of a tensor train (TT) format [32], which is a particular case of the *matrix product states* (MPS) decomposition [50, 47]. The MPS tensor format was introduced in 1992, see [50], for solving spin system modeling problems. It has since been successfully applied in the physics and quantum chemistry community [40].

For a given rank parameter $\mathbf{r} = (r_1, \dots, r_{d-1})$, and the respective index sets $J_\ell = \{1,\dots,r_\ell\}$ $(\ell = 1,\dots,d-1)$, the rank-$\mathbf{r}$ TT format represents a tensor $\mathbf{A} = [a_{i_1,\dots,i_d}] \in \mathbb{W}_{\mathbf{n}}$ as the

16

contracted products of 3-dimensional tensors over the d-fold product index set $\mathcal{J} := \times_{\ell=1}^{d-1} J_\ell$, such that

$$\mathbf{A} = \sum_{(\alpha_1,\ldots,\alpha_{d-1})\in\mathcal{J}} \mathbf{a}_{1,\alpha_1}^{(1)} \otimes \mathbf{a}_{\alpha_1,\alpha_2}^{(2)} \otimes \cdots \otimes \mathbf{a}_{\alpha_{d-1},1}^{(d)}, \quad \text{with } \mathbf{a}_{\alpha_{\ell-1},\alpha_\ell}^{(\ell)} \in \mathbb{R}^{n_\ell}$$

or entry-wise (see [31, 32])

$$a_{i_1,\ldots,i_d} = A^{(1)}(i_1)A^{(2)}(i_2)\cdots A^{(d)}(i_d),$$

with $r_{\ell-1} \times r_\ell$ matrices $A^{(\ell)}(i_\ell) = [a_{\alpha_{\ell-1},\alpha_\ell}^{(\ell)}(i_\ell)]$, $(\ell = 1,\ldots,d)$ and $r_0 = r_d = 1$. The TT representation reduces the storage cost to $\mathcal{O}(dr^2 n)$, $r = \max r_\ell$, $n = \max n_\ell$.

4.2 Quantized-TT approximation of function related vectors

The quantized-TT (QTT) approximation method was introduced and analyzed for the class of function related vectors in [25].

The QTT-type approximation of a vector of size n with $n = q^L$, $L \in \mathbb{N}$, $q = 2, 3, \ldots$, is defined as the tensor decomposition (approximation) in the TT or canonical format applied to a tensor obtained by the folding (reshaping) of the initial vector to an L-dimensional $q \times \cdots \times q$ data array. The latter is considered as an element of the L-dimensional quantized tensor space $\mathbb{Q}_{q,L} = \bigotimes_{j=1}^{L} \mathbb{K}^q$, $\mathbb{K} \in \{\mathbb{R},\mathbb{C}\}$, and L is the auxiliary dimension parameter that measures the depth of the quantization transform. A vector $\mathbf{x} = [x_i]_{i\in I} \in \mathbb{R}^n$, is reshaped to its multi-dimensional quantized image in $\mathbb{Q}_{q,L}$ by q-adic folding,

$$\mathcal{F}_{q,L} : \mathbf{x} \to \mathbf{Y} = [y(\mathbf{j})] \in \mathbb{Q}_{q,L}, \quad \mathbf{j} = \{j_1,\ldots,j_L\},$$

with $j_\nu \in \{1,\ldots,q\}$ for $\nu = 1,\ldots,L$. Here, for fixed i, we have $y(\mathbf{j}) := x_i$, and $j_\nu = j_\nu(i)$ is defined via q-coding, $j_\nu - 1 = C_{-1+\nu}$, such that the coefficients $C_{-1+\nu}$ are found from the q-adic representation of $i-1$ (binary coding for $q = 2$),

$$i - 1 = C_0 + C_1 q^1 + \cdots + C_{L-1}q^{L-1} \equiv \sum_{\nu=1}^{L}(j_\nu - 1)q^{\nu-1}.$$

Assuming that for the rank-$\mathbf{r}$ TT approximation of the quantized image $\mathbf{Y}$ there holds $r_k \leq r$, $k = 1,\ldots,L$, the complexity of such representation for the tensor $\mathbf{Y}$ reduces to logarithmic scaling

$$qr^2 \log_q n \ll n.$$

It was proven in [25] that for a vector of size $n = q^L$ (say, for $q = 2$) obtained by discretization of a certain classical function, its QTT image in the L-dimensional tensor space with $L = \log_2 n$ exhibits an amazingly low separation rank r_{qtt} independent of of the size of the original vector. In particular, $r_{qtt} = 1$ for exponential functions, $r_{qtt} = 2$ for trigonometric functions, $r_{qtt} = p+1$ for polynomials of degree p, etc. Thus the QTT tensor compresses the amount of numbers for a vector representation from n to $O(\log_2 n)$. In cases when the exact low-rank QTT representation is not known, an ε-approximation in the QTT format can be computed by using the standard TT multi-linear approximation tools [31].

17

Similar low rank QTT representations were proven for a wide class of functions [26], including strongly oscillating functions of nontrivial shape, see for example [20, 23] and the new results in §4.5 below. We also refer to [11, 20, 19, 26] for further results on QTT approximation of functional vectors and various applications of the QTT approximation tecniques. A more detailed discussion of the QTT approximation for function related vectors can be found in [25, 26]. For general classes of functional vectors, the small ε-rank QTT approximation also leads to storage size that scales logarithmically in n.

In estimating the numerical complexity we use the average QTT rank further denoted by r_{qtt} and calculated as follows,

$$r_{qtt} = \sqrt{\frac{1}{L-1} \sum_{k=1}^{L-1} r_k^2}, \tag{4.1}$$

where the QTT rank parameters r_k are the TT ranks of the quantized image $\mathbf{Y}$ of a vector.

As a first illustration, we consider the QTT approximation of the DOS for the 1D finite difference Laplacian operator in $[0, \pi]$ with Dirichlet boundary conditions, $A = -\mathrm{tridiag}\{1, -2, 1\} \in \mathbb{R}^{n \times n}$, discretized on the uniform grid of size $h = \pi/(n+1)$ with $n = 2047$. The corresponding eigenvalues are given by

$$\lambda_k = 4\sin^2(\frac{\pi k}{2n}), \quad k = 1, \dots, n.$$

Figure 4.1 represents the Lorentzian-DOS for discrete Laplacian (left) and the corresponding approximation error for its rank-5 QTT approximation computed on the representation grid of size $N = 2^{14}$ (right).


Figure 4.1: DOS for Laplacian (left), and the error ($\approx 10^{-5}$) of its QTT approximation with $r_{qtt} = 5$ (right).

In what follows, we apply the QTT approximation method to the DOS regularized by Lorentzians and sampled on a fine representation grid of size $N = 2^L$. The QTT approximant can be viewed as the rank structured ε-interpolant of the highly non-regular function ϕ_η regularizing the exact DOS. In this case, the application of traditional polynomial or trigonometric type interpolation is inefficient. The QTT approach provides a good approximation

18

to ϕ_η on the whole spectral interval and requires only a moderate number of representation parameters $r_{qtt}^2 \log N \ll N$, where the average QTT rank r_{qtt} defined in (4.1) is a small rank parameter adaptively depending on the truncation error $\epsilon > 0$.

Based on the existence of the low-rank QTT approximation for DOS functions, we describe a tensor based heuristic QTT approximation of the DOS by using only an incomplete set of sampling points, i.e., QTT representation by adaptive cross approximation (TT-cross) [33, 39]. Furthermore, we derive an upper bound on the QTT ranks of the DOS by the Gaussians broadening.

In what follows, we describe the low-rank QTT tensor interpolation of the function $L_\eta(t)$ sampled on a fine uniform grid $\{t_1, \ldots, t_N\}$ in the whole spectral interval $[0, a]$ or on some subinterval of $[0, a]$. We show that QTT parametrization allows asymptotically fewer interpolation points (functional calls) $M \leq N$, than the size N of the fine representation grid. This might be beneficial in the limit of a large number of representation points N since each functional evaluation of the DOS is highly expensive requiring computation of some matrix valued functions.

4.3   QTT approximation of DOS via Lorentzians: proof of concept

In this section we demonstrate the efficiency of the QTT approximation applied to the DOS via both Gaussian and Lorentzian blurring. By various numerical experiments we verify that the low-rank QTT approximant resolves well the exact DOS, thereby providing a proof of concept for practical use of the TT/QTT-cross approximation tools.
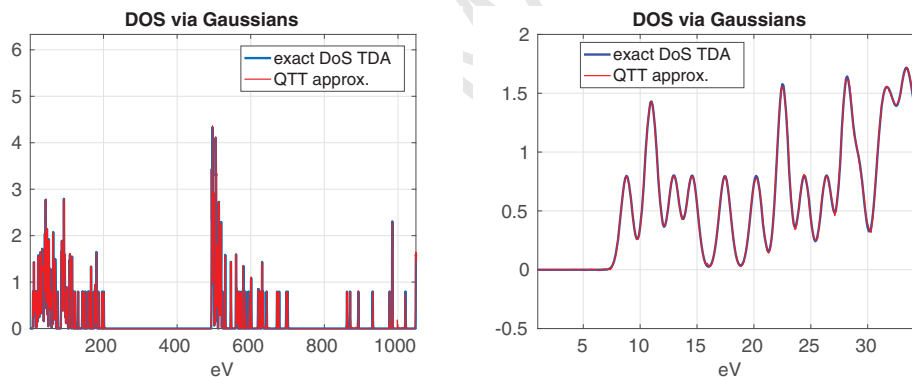

Figure 4.2: DOS (in eV) for the H_2O molecule via Gaussians (left), and zoom on the left most part of the spectrum. Here $r_{QTT} = 9.4$, $\eta = 0.4$

In the following numerical examples, we use a sampling vector defined on a grid of size $N \approx 2^{14}$. We set the QTT truncation error to $\epsilon_{QTT} = 0.04$, if not explicitly indicated[2]. For ease of interpretation, we set the pre-factor in (2.4) to 1. It is worth noting that the QTT-approximation scheme is applied to the full TDA spectrum. Our results demonstrate that it renders good resolution in the whole range of energies (in eV) including large "zero gaps".

---

[2]The typical accuracy in calculations of the optical absorption spectra in a physics community is of the order of 0.1eV (for shifts in the positions of spikes), see for example [18]. Our choice $\epsilon_{QTT} = 0.04$ for the error control is within this accuracy.

19

Figure 4.2, left, represents the TDA DOS (blue line) for H_2O computed by Gaussian blurring with the parameter $\eta = 0.4$, and the corresponding rank-9.4 QTT tensor approximation (red line) to the discretized function $\phi_\eta(t)$. For this example, the number of eigenvalues is given by $n = N_{BSE}/2 = 180$. Figure 4.2, right, provides a zoom of the corresponding DOS and its QTT approximant within the small energy interval $[0, 40]$eV.
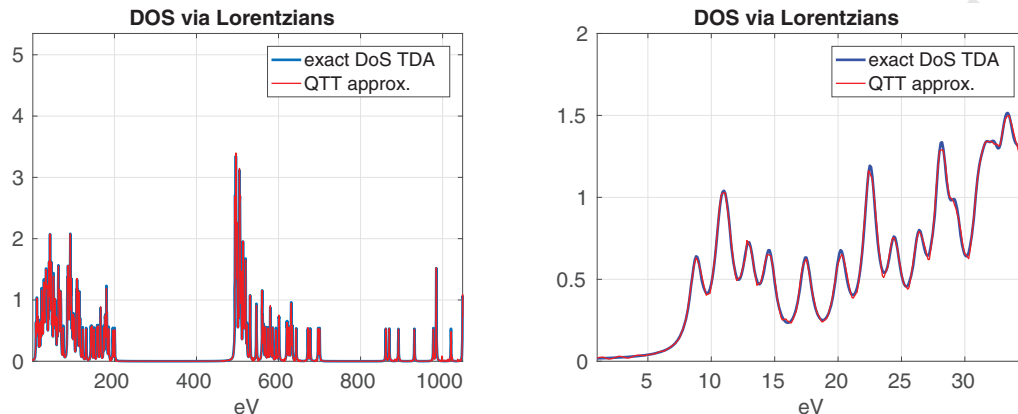

Figure 4.3: DOS for H_2O molecule via Lorentzians (blue) and its QTT approximation (red) (left). Zoom on the left most part of the spectrum (right). We have ε=0.04, $r_{QTT} = 10.5$.

Figure 4.3 demonstrates the resolution of the QTT approximation to the DOS via the Lorentzian blurring indicating similar QTT-ranks as in the case of the Gaussians regularization.
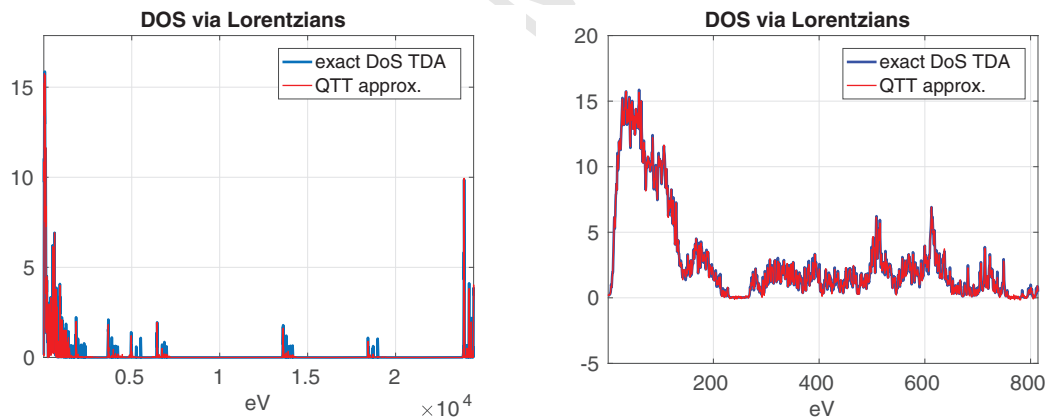

Figure 4.4: Left: DOS for Glycine amino acid via Lorentzians (blue) and its QTT approximation (red); Right: zoom of the low energy part of the spectrum. We have ε=0.04, $r_{QTT} = 16$.

Figure 4.4 (Lorentzian blurring) represents similar data, but for the large Glycine amino acid with $n = N_{TDA} = 3000$. It is worth noting that the average QTT rank of $\phi_\eta(t)$ sampled on $N = 2^{14}$ grid points is about $r_{QTT} = 16$, $(\epsilon_{QTT} = 0.04)$ though the number of eigenvalues n in this case is about 20 times larger than for the water molecule. This means that for a fixed η, the QTT-rank remains rather modest relative to the molecular size. This observation confirms Theorem 4.1 in Section 4.5.

20

A comparison of Figures 4.2 and 4.3 indicates that the Lorentzian based DOS blurring is slightly smoother than Gaussian blurring. The moderate size of the QTT ranks in Figures 4.3 and 4.4 clearly shows the potentials of the QTT ε-interpolation for modeling the DOS of large lattice type clusters.

We observe several gaps in the spectral densities and their highly oscillating form, see Figure 4.2, 4.3 and 4.4, indicating that polynomial, rational or trigonometric interpolation can be applied only to some very small energy sub-intervals, but not in the whole interval $[a, b]$. Remarkably, the QTT approximant resolves well the DOS function in the whole energy interval including nearly zero values within the spectral gaps (hardly possible for polynomial/rational based interpolation).

4.4   Numerics for the QTT interpolation of the DOS function

In the previous section we demonstrated that the QTT tensor approximation provides good resolution for the DOS function calculated for a number of molecules. In what follows, we describe a tensor based heuristic QTT approximation of the DOS by using only an incomplete set of sampling points, i.e., QTT representation by the QTT-cross approximation (via the TT-cross algorithm) [33, 39]. This allows us to recover the spectral density with M functional calls in controllable accuracy, where M asymptotically scales logarithmically in the grid size N. It is worth to note that the positions and hights of picks in the DOS function are weel resolved.

This heuristic approach can be viewed as a kind of "adaptive QTT ε-interpolation". In particular, we show by numerical experiments that the low-rank QTT-cross approximation provides a good resolution of the target DOS function with the number of functional calls that asymptotically scales logarithmically, i.e., $O(\log N)$, in the size N of the representation grid.

In the case of large N, the QTT interpolant can be computed by the TT-cross tensor approximation procedure (see [33, 39] for the detailed description) that, in general, does not require the full set of functional values over the N-grid. In the case of large N this beneficial feature allows to compute the QTT approximation by requiring less than N computationally expensive functional evaluations of $\phi_\eta(t)$.

The QTT interpolation via TT-cross tensor approximation serves to recover the representation parameters of the QTT tensor approximant and normally requires about

$$M = C_s r_{qtt}^2 \log_2 N \tag{4.2}$$

samples of the target N-vector[3] with a small pre-factor C_s, usually satisfying $C_s \leq 10$, that is independent of the fine interpolation grid size $N = 2^L$, see, for example, [23]. This cost estimate seems promising in the perspective of extended or lattice type molecular systems, requiring large spectral intervals and, as a result, a large interpolation grid of size N. Here the QTT rank parameter r_{qtt} naturally depends on the required truncation threshold $\varepsilon > 0$, characterizing the L_2-error between the exact DOS and its QTT interpolant. The QTT tensor interpolation reduces the number of functional calls, i.e., $M < N$, *if the QTT rank*

---

[3]In our application, this is the DOS functional N-vector corresponding to representations via matrix resolvents in (2.9) or (2.10).

parameters (or threshold $\varepsilon > 0$) are chosen to satisfy the condition

$$M = C_s r_{qtt}^2 \log_2 N \leq N. \tag{4.3}$$

The expression on the left-hand side provides a rather accurate estimate on the number of functional evaluations.

To complete this discussion, we present numerical tests for the low-rank QTT tensor interpolation applied to the long vector discretizing the Lorentzian-DOS on a fine representation grid of size $N = 2^L$.
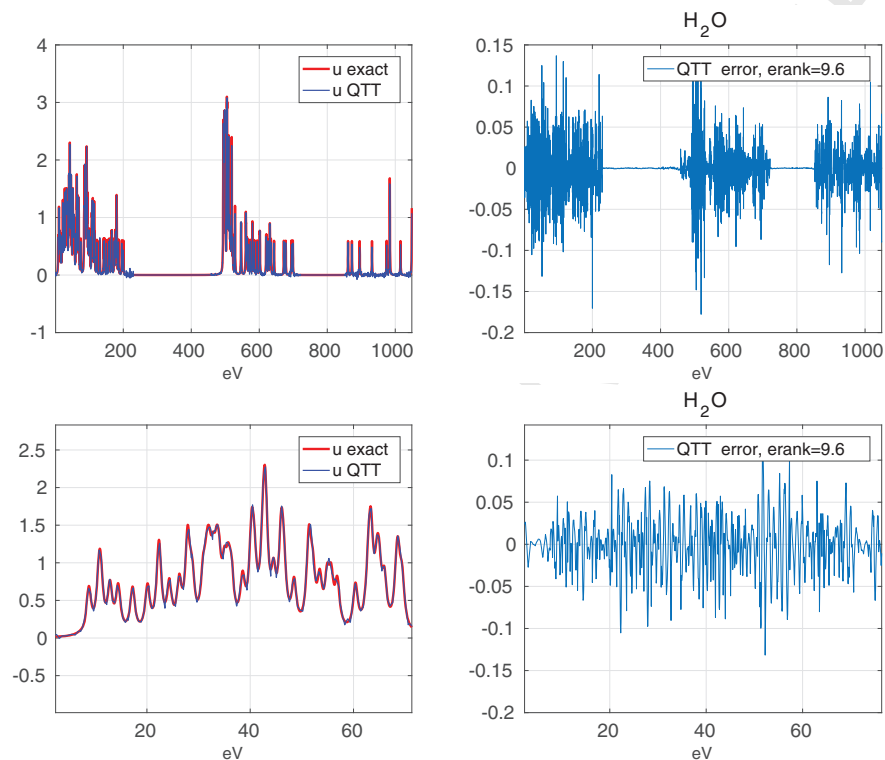

Figure 4.5: QTT-cross interpolation of the DOS for H_2O (top) and zoom in to a small spectral interval (bottom).

Figure 4.5 represents the results of the QTT based interpolating approximation to the discretized DOS function (H_2O molecule). We use the QTT-cross approximation algorithm based on [25, 33, 39] and implemented in MATLAB TT-toolbox[4]. Here we set $\varepsilon = 0.08$, $\eta = 0.1$ and $N = 2^{14}$, providing $r_{QTT} = 9.8$. The top two figures display the results on the whole spectral interval, while the bottom figures show the zoom of the same data in the small spectral interval $[0, 55]$eV.

Figure 4.6 illustrates the logarithmic increase in the number of samples required for the QTT interpolation of the DOS (for the H_2O molecule) represented on the grid of size $N = 2^L$, where $L = 11, 12, \ldots, 16$, provided that the rank truncation threshold is chosen by $\epsilon = 0.05$ and the regularization parameter is $\eta = 0.2$. In this example, the effective pre-factor in

---

[4]https://github.com/oseledets/TT-Toolbox

Figure 4.6: DOS for H_2O via Lorentzians: the number of functional calls for QTT-cross approximation (blue) vs. the full grid size N.

(4.2) is estimated by $C_s \leq 10$. This pre-factor characterizes the average number of samples required for the recovery of each of the $r_{qtt}^2 \log N$ representation parameters involved in the QTT tensor ansatz.

We observe that the QTT tensor interpolant recovers the exact DOS with a good precision. The positions of spikes and their heights are well resolved. The logarithmic scaling $O(\log N)$ vs. the grid size for the number of functional calls requested by the QTT-cross routine can be observed in Figure 4.6 (blue line).

4.5 Upper bounds on the QTT ranks of DOS

In this section we analyze the upper bounds on the QTT ranks of the discretized DOS obtained by Gaussian broadening. Our numerical tests indicate that Lorentzian blurring leads to a similar QTT rank compared with Gaussians blurring when both are applied to the same grid and the same truncation threshold $\varepsilon > 0$ is used in the QTT approximation. We consider the more general case of a symmetric interval, i.e. $t, \lambda_j \in [-a, a]$.

Assume that the function $\phi_\eta(t) = \frac{1}{n} \sum_{j=1}^{n} g_\eta(t - \lambda_j)$, $t \in [-a, a]$, in equation (2.6) is discretized by sampling over the uniform N-grid Ω_h with $N = 2^d$, where the generating Gaussian is given by $g_\eta(t) = \frac{1}{\sqrt{2\pi}\eta} \exp\left(-\frac{t^2}{2\eta^2}\right)$, see (2.5). Denote the corresponding N-vector by $\mathbf{g} = \mathbf{g}_\eta$, and the resulting discretized density vector by

$$\phi_\eta(t) \mapsto \mathbf{p} = \mathbf{p}_\eta = \frac{1}{n} \sum_{j=1}^{n} \mathbf{g}_{\eta,j} \in \mathbb{R}^N,$$

where the shifted Gaussian is assigned by the vector $g_\eta(t - \lambda_j) \mapsto \mathbf{g}_j = \mathbf{g}_{\eta,j}$.

Without loss of generality, we suppose that all eigenvalues are situated within the set of grid points, i.e. $\lambda_j \in \Omega_h$. Otherwise, we can slightly relax their positions provided that the mesh size h is small enough. This is not a severe restriction for the QTT approximation of functional vectors since storage and complexity requests depend only logarithmically on N.

23

Theorem 4.1 *Assume that the effective support of the shifted Gaussians $g_\eta(t - \lambda_j)$, $j = 1, \ldots, n$, is included in the computational interval $[-a, a]$. Then the QTT ε-rank of the vector $\mathbf{p}_\eta$ is bounded by*

$$rank_{QTT}(\mathbf{p}_\eta) \leq C \log^{3/2}(|\log \varepsilon|),$$

with a constant $C = O(\eta^{-1}) > 0$.

Proof. The main argument of the proof is similar to that in [20, 11]: the sum of discretized shifted Gaussians with the same exponential η^{-1}, each represented in Fourier basis, can be expanded with merely the same number of Fourier harmonics (uniform basis) as each individual Gaussian. Taking into account that the Fourier transform of a Gaussian is a Gaussian with reciprocal exponent, we now apply the estimate on the number of essential Fourier coefficients for the Gaussian vectors $\mathbf{g}_{\eta,j}$ with a fixed exponent parameter η (see [11] or Lemma 4.12 in [26] for more details)

$$m_0 = O(\eta^{-1} \log^{3/2}(|\log \varepsilon|)),$$

up to lower order terms. Here, $\varepsilon > 0$ denotes the rank truncation threshold which is also applied as a threshold for the essential Fourier coefficients. Since each Fourier harmonic has exact rank-2 QTT representation (see Section 4.2), we arrive at the claimed bound. ∎

Notice that the Fourier transform of the Lorentzian in (2.7) is given by

$$e^{-|k|\eta},$$

thus a similar QTT rank bound can be derived for the case of Lorentzian blurred DOS.

Molecule	H_2O	NH_3	H_2O_2	N_2H_4	C_2H_5OH	$C_2H_5\ NO_2$	$C_3H_7\ NO_2$
$n = N_{ov}$	180	215	531	657	1430	3000	4488
QTT ranks	11	11	12	11	15	16	13

Table 4.1: QTT ranks of Lorentzians-DOS for some molecules; $\varepsilon = 0.04$, $\eta = 0.4$, $N = 2^{14}$.

Table 4.1 shows that the average QTT tensor rank remains almost independent of the molecular size, which confirms Theorem 4.1. A weak dependence of the rank parameter on the molecular geometry can be observed.

5 Towards calculation of the BSE absorption spectrum

In this section we describe the generalization of our approach to the case of the full BSE system. Within the BSE framework, the optical absorption spectrum of a molecule is defined by

$$\epsilon(\omega) \equiv d_r^{\mathsf{H}} \delta(\omega I_{2n} - H) d_l = \sum_{j=1}^{2n} \frac{(d_r^{\mathsf{H}}(z_r)_j)((z_l)_j^{\mathsf{H}} d_l)}{(z_l)_j^{\mathsf{H}}(z_r)_j} \delta(\omega - \lambda_j), \tag{5.1}$$

where

$$d_r = \begin{bmatrix} d \\ -\bar{d} \end{bmatrix} \quad \text{and} \quad d_l = \begin{bmatrix} d \\ \bar{d} \end{bmatrix}$$

24

are the right and left *optical transition vectors*, respectively, and d is a vector reshaped from a transition matrix T of dimension $N_o \times (N_b - N_o)$. The (i, a)th element of T is given by $\langle \psi_i | \vec{x} | \psi_a \rangle$, where $\vec{x}$ is a position operator in the direction of x and ψ_i and ψ_a are a pair of occupied and unoccupied molecular orbitals [8]. Here z_r and z_l are the right and the left eigenvectors of the BSE Hamiltonian H in (2.1).

Similar to the DOS, the function $\epsilon(\omega)$ is a sum of Dirac-δ peaks centered at eigenvalues of the BSH. However, the height of each peak, which is often referred to as the oscillator strength, is determined by the projection of the corresponding left and right eigenvectors of H onto the optical transition vectors d_l and d_r.

A smooth approximation of (5.1) can be obtained by replacing the Dirac-δ function with either a Gaussian or a Lorentzian with an appropriate broadening width. If we choose to smooth Dirac-δ by a Lorentzian, we then need to compute

$$\epsilon(\omega) \approx \frac{1}{\pi} \mathrm{Im} \left[d_r^{\mathsf{H}} \left(\omega I_{2n} - H - i\eta I_{2n} \right)^{-1} d_l \right], \tag{5.2}$$

where η is related to the width of broadening.

For a fixed frequency ω, (5.2) can be evaluated by solving a linear system of the form

$$\left(\omega I_{2n} - H - i\eta I_{2n} \right) x = d_l.$$

The block sparse and low-rank structure of H can be used to reduce the cost for solving such a linear system.

The detailed numerical analysis of this scheme for the BSE system is a topic of a forthcoming paper.

6   Conclusions

The new approach to approximating the DOS of the TDA approximation for the BSE Hamiltonian is based on two main techniques. First, we developed an economical method for evaluating the trace of the parametric matrix resolvent by taking advantage of the block-diagonal plus low rank structure of the TDA matrix. The presented algorithm provides an efficient way to calculate the DOS function, regularized by Lorentzians, at each point on the fine representation grid, thus avoiding commonly used stochastic sampling. The numerical cost scales linearly with respect to the matrix size. Second, the QTT-cross tensor ε-interpolation scheme is used to approximate the DOS function discretized on large representation grids. This approximation scheme allows us to estimate the DOS (and especially the positions of spikes) with M function evaluations, where M scales logarithmically with respect to the size of the representation grid, N, on which the DOS is evaluated. The approach can be applied to a wide class of rank-structured symmetric spectral problems.

In Theorems 3.2 and 3.3, we prove linear scaling of the rank-structured trace calculation algorithm in the matrix size. This result is confirmed by numerical experiments performed for the DOS of the TDA matrix associated with some molecular systems as shown in Figure 3.3.

In Theorem 4.1 we justify the low rank QTT approximation of the DOS in the case of Gaussian regularization. The efficiency of low-rank QTT approximation to DOS is illustrated

25

numerically on the example of discrete Laplacian as well as for the BSE/TDA spectral problem for several moderate size molecules. Numerical tests demonstrate the logarithmic number of functional calls of the QTT-cross approximation scheme in the case of large grid size N, applied to the discretized DOS as depicted in Figure 4.6.

It is worth noting that the presented approach serves to recover DOS on the whole spectral interval which is demonstrated in a number of numerical tests. However, the algorithms are applicable to any fixed subinterval of interest in the whole spectrum, which will correspondingly reduce the overall computational time.

The presented methods introduce a new efficient tool for numerical approximation of the DOS function for large rank-structured matrices arising in various applications in condensed matter physics, computational quantum chemistry as well as in large-scale problems of numerical linear algebra and scientific computing.

References

[1] Z. Bai and R.-C. Li. Minimization principle for linear response eigenvalue problem, I: Theory. *SIAM J. Matrix Anal. Appl.*, 33(4):10751100, 2012.

[2] Z. Bai and R.-C. Li. Minimization principle for linear response eigenvalue problem, II: Computation. *SIAM J. Matrix Anal. Appl.*, 34(2):392–416, 2013.

[3] P. Benner, S. Dolgov, V. Khoromskaia, and B. N. Khoromskij. Fast iterative solution of the Bethe-Salpeter eigenvalue problem using low-rank and QTT tensor approximation. *J. Comp. Phys.*, (334):221–239, 2017.

[4] P. Benner and H. Faßbender. An implicitly restarted symplectic Lanczos method for the Hamiltonian eigenvalue problem. *Linear Algebra Appl.*, 263:75–111, 1997.

[5] P. Benner, H. Faßbender, and C. Yang. Some remarks on the complex J-symmetric eigenproblem. *Linear Algebra and its Applications*, 544:407 – 442, 2018.

[6] P. Benner, V. Khoromskaia, and B. N. Khoromskij. A reduced basis approach for calculation of the Bethe-Salpeter excitation energies using low-rank tensor factorizations. *Mol. Physics*, 114(7-8):1148–1161, 2016.

[7] P. Benner, V. Mehrmann, and H. Xu. A numerically stable, structure preserving method for computing the eigenvalues of real Hamiltonian or symplectic pencils. *Numerische Mathematik*, 78(3):329–358, 1998.

[8] F. Bruneval, T. Rangel, S. M. Hamed, M. Shao, C. Yang, and J. B. Neaton. MOLGW 1: Many-body perturbation theory software for atoms, molecules, and clusters. *Comp. Phys. Comm.*, 208:149–161, 2016.

[9] A. Bunse-Gerstner and H. Faßbender. Breaking Van Loan's curse: A quest for structure-preserving algorithms for dense structured eigenvalue problems. In P. Benner, M. Bollhöfer, D. Kressner, C. Mehl, and T. Stykel, editors, *Numerical Algebra, Matrix Theory, Differential-Algebraic Equations and Control Theory*, pages 3–23. Springer International Publishing, 2015.

[10] J. Deslippe, G. Samsonidze, D. A. Strubbe, M. Jain, M. L. Cohen, and S. Louie. BerkeleyGW: A massively parallel computer package for the calculation of the quasi-particle and optical properties of materials and nanostructures. *Comp. Phys. Comm.*, 183:1269–1289, 2012.

[11] S. V. Dolgov, B. N. Khoromskij, and I. V. Oseledets. Fast solution of multi-dimensional parabolic problems in the tensor train/quantized tensor train–format with initial application to the Fokker-Planck equation. *SIAM J. Sci. Comp.*, 34(6):A3016–A3038, 2012.

[12] D. A. Drabold and O. F. Sankey. Maximum entropy approach for linear scaling in the electronic structure problem. *Phys. Rev. Lett.*, 70:3631–3634, 1993.

[13] F. Ducastelle and F. Cyrot-Lackmann. Moments developments and their application to the electronic charge distribution of d bands. *J. Phys. Chem. Solids*, 31:1295–1306, 1970.

[14] H. Faßbender and D. Kressner. Structured eigenvalue problem. *GAMM Mitteilungen*, 29(2):297–318, 2006.

[15] G. H. Golub and C. F. Van Loan. *Matrix Computations, 4th ed.* Johns Hopkins University Press, Baltimore, 2013.

[16] R. Haydock, V. Heine, and M. J. Kelly. Electronic structure based on the local atomic environment for tight-binding bands. *J. Phys. C Solid State Phys.*, 5:2845–2858, 1972.

[17] L. Hedin. New method for calculating the one-particle Green's function with application to the electron-gas problem. *Phys. Rev.*, 139:A796, 1965.

[18] A. Hermann, W. G. Schmidt, and P. Schwerdfeger. Resolving the optical spectrum of water: Coordination and electrostatic effects. *Phys. Review Lett.*, 100:207403, 2008.

[19] V. Kazeev and C. Schwab. Quantized tensor-structured finite elements for second-order elliptic PDEs in two dimensions. *Numer. Math.*, 138:133–190, 2018.

[20] V. Khoromskaia and B. N. Khoromskij. Grid-based lattice summation of electrostatic potentials by assembled rank-structured tensor approximation. *Comp. Phys. Comm.*, 185(12):3162–3174, 2014.

[21] V. Khoromskaia and B. N. Khoromskij. *Tensor Numerical Methods in Quantum Chemistry.* De Gruyter, Berlin, 2018.

[22] V. Khoromskaia, B. N. Khoromskij, and R. Schneider. Tensor-structured calculation of two-electron integrals in a general basis. *SIAM J. Sci. Comp.*, 35(2):A987–A1010, 2013.

[23] B. Khoromskij and A. Veit. Efficient computation of highly oscillatory integrals by using QTT tensor approximation. *Comp. Meth. Appl. Math.*, 16(1):145–159, 2016.

[24] B. N. Khoromskij. Structured $(r_1, \ldots r_d)$- decomposition of function related tensors in $\mathbb{R}^d$. *Comp. Meth. Appl. Math.*, 6:194–220, 2006.

[25] B. N. Khoromskij. $O(d \log N)$-quantics approximation of N-d tensors in high-dimensional numerical modeling. *J. Constr. Approx.*, 34(2):257–289, 2011.

[26] B. N. Khoromskij. *Tensor Numerical Methods in Scientific Computing.* De Gruyter, Berlin, 2018.

[27] T. Kolda and B. W. Bader. Tensor decompositions and applications. *SIAM Review*, 51(3):455–500, 2009.

[28] L. Lin, Y. Saad, and C. Yang. Approximating spectral densities of large matrices. *SIAM Review*, 58(1):34–5, 2016.

[29] E. Napoli, E. Polizzi, and Y. Y. Saad. Efficient estimation of eigenvalue counts in an interval. *Numer. Lin. Algebra Appl.*, 23(4):674 – 692, 2016.

[30] G. Onida, L. Reining, and A. Rubio. Electronic excitations: density-functional versus many-body Green's-function approaches. *Rev. of Modern Physics*, 74, 2002.

[31] I. V. Oseledets. Tensor-train decomposition. *SIAM J. Sci. Comp.*, 33(5):2295–2317, 2011.

[32] I. V. Oseledets and E. E. Tyrtyshnikov. Breaking the curse of dimensionality, or how to use SVD in many dimensions. *SIAM J. Sci. Comp.*, 31(5):3744–3759, 2009.

[33] I. V. Oseledets and E. E. Tyrtyshnikov. TT-cross approximation for multidimensional arrays. *Linear Algebra Appl.*, 432(1):70–88, 2010.

[34] E. Rebolini, J. Toulouse, and A. Savin. Electronic excitation energies of molecular systems from the Bethe-Salpeter equation: Example of H_2 molecule. *In: Concepts and Methods in Modern Theoretical Chemistry (S. Ghosh and P. Chattaraj eds), vol 1: Electronic Structure and Reactivity*, page 367, 2013.

[35] L. Reining, V. Olevano, A. Rubio, and G. Onida. Excitonic effects in solids described by time-dependent density functional theory. *Phys. Rev. Lett.*, 88:066404, 2002.

[36] D. Rocca, R. Gebauer, Y. Saad, and S. Baroni. Turbo charging time-dependent density-functional theory with Lanczos chains. *J. Chem. Phys.*, 128:154104, 2008.

[37] D. Rocca, D. Lu, and G. Galli. *Ab Initio* calculations of optical absorption spectra: Solution of the Bethe-Salpeter euqation within density matrix perturbation theory. *J. Chem. Phys.*, 133:164109 1–10, 2010.

[38] E. E. Salpeter and H. A. Bethe. A relativistic equation for bound-state problems. *Phys. Review*, 82(2):309–310, 1951.

[39] D. Savostyanov and I. V. Oseledets. Fast adaptive interpolation of multi-dimensional arrays in tensor train format. *IEEE Publisher: The 2011 International Workshop on Multidimensional (nD) Systems*, pages 1–8, 2011.

[40] U. Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Ann.Phys.*, 51(326):96–192, 2011.

[41] M. Shao, F. da Jornada, L. Lin, C. Yang, J. Deslippe, and S. Louie. A structure preserving lanczos algorithm for computing the optical absorption spectrum. *SIAM J Matrix Anal Appl.*, 39(2):683–711, 2018.

[42] M. Shao, F. H. da Jornada, C. Yang, J. Deslippe, and S. Louie. Structure preserving parallel algorithms for solving the Bethe-Salpeter eigenvalue problem. *Linear Algebra and its Applications*, 488:148–167, 2016.

[43] R. E. Stratmann, G. E. Scuseria, and M. J. Frisch. An efficient implementation of time-dependent density-functional theory for the calculation of excitation energies of large molecules. *J. Chem. Phys.*, 109:8218, 1998.

[44] L. N. Trefethen and M. Embree. *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators.* Princeton University Press, Princeton and Oxford, 2005.

[45] I. Turek. A maximum-entropy approach to the density of states within the recursion method. *J. Phys. C*, 21:3251–3260, 1988.

[46] J. L. M. Van Dorsselaer and M. E. Hoschstenbach. Computing probabilistic bounds for extreme eigenvalues of symmetric matrices with the Lanczos method. *SIAM J. Matrix Anal. Appl.*, 22:837–852, 2000.

[47] G. Vidal. Efficient classical simulation of slightly entangled quantum computations. *Phys. Rev. Lett.*, 91(14), 2003.

[48] L.-W. Wang. Calculating the density of states and optical-absorption spectra of large quantum systems by the plane-wave moments method. *Phys. Rev. B*, 49:10154–10158, 1994.

[49] J. C. Wheeler and C. Blumstein. Modified moments for harmonic solids. *Phys. Rev. B*, 6:4380–4382, 1972.

[50] S. R. White. Density-matrix algorithms for quantum renormalization groups. *Phys. Rev. B*, 48(14):10345–10356, 1993.