Exact NMR simulation of protein-size spin systems using tensor train formalism

D. V. Savostyanov,[1,*] S. V. Dolgov,[2] J. M. Werner,[3] and Ilya Kuprov[1]

[1]*School of Chemistry, University of Southampton, Highfield Campus, Southampton, SO17 1BJ, United Kingdom*
[2]*Max-Planck Institute for Mathematics in the Sciences, Inselstrasse 22, Leipzig 04103, Germany*
[3]*Centre for Biological Sciences, University of Southampton, Highfield Campus, Southampton, SO17 1BJ, United Kingdom*

We introduce a new method, based on alternating optimization, for compact representation of spin Hamiltonians and solution of linear systems of algebraic equations in the tensor train format. We demonstrate the method's utility by simulating, without approximations, a ^{15}N NMR spectrum of ubiquitin—a protein containing several hundred interacting nuclear spins. Existing simulation algorithms for the spin system and the NMR experiment in question either require significant approximations or scale exponentially with the spin system size. We compare the proposed method to the SPINACH package that uses heuristic restricted state space techniques to achieve polynomial complexity scaling. When the spin system topology is close to a linear chain (e.g., for the backbone of a protein), the tensor train representation is more compact and can be computed faster than the sparse representation using restricted state spaces.

I. INTRODUCTION

The amount of patience required to simulate exactly a nuclear magnetic resonance (NMR) spectrum of an N-spin system scales approximately as $\mathcal{O}(2^N)$. That much is rarely available, and considerable thought has consequently been given over the last decade to more efficient methods [1–5], particularly those that promise to achieve that objective in polynomial time. Such algorithms do exist [2,6], but they make significant *a priori* assumptions about the spin system evolution—it is usually assumed that the system stays weakly correlated for the duration of the experiment [6,7].

Outside the NMR community, significant progress was recently made with the development of tensor structured methods [8–12], all of which descend broadly from the density matrix renormalization group (DMRG) [13,14] as well as matrix product state (MPS) [15,16] and matrix product operator (MPO) [17] formalisms. Typical applications of DMRG in condensed matter theory are (one-dimensional) 1D spin chains [8,18–20] with recent extensions to 2D lattices [21–24]. DMRG has also been put to good use in electronic [25–31] and nuclear [32,33] structure theory, but magnetic resonance spectroscopy has so far received little attention; the spin systems encountered in the daily practice of NMR and EPR (proteins, radicals, polynucleotides, polysaccharides) are irregular three-dimensional room-temperature networks with multiple interlocking loops in the spin coupling graph and no identical couplings [34]. When the strict requirement for correct wave function phase during the very long (milliseconds to seconds) dissipative spin system trajectories is added to the list, time-domain DMRG methods are currently struggling.

There are some biologically relevant cases, however, that may still be treated as linear chains; for the purposes of simulating simple backbone NMR experiments, protein side chains may often be ignored. This makes the corresponding spin system a weakly branched linear chain that is amenable to DMRG-type treatment. Simple NMR experiments can also be reformulated as a matrix-inverse-times-vector problem in the frequency domain, for which efficient algorithms in tensor product formats have recently emerged [35–37]. We report in this paper the behavior of the alternating minimal energy (AMEn) algorithm [37,38], applied to the solution of the NMR simulation problem in the frequency domain, as well as to the technical task of adding together, without loss of accuracy, tensor train representations of thousands of spin Hamiltonian terms for a protein.

Having integrated the algorithms described below into SPINACH (a large-scale magnetic resonance simulation library [4]), we are reporting here the exact quantum mechanical simulation of a liquid-state 1D NMR spectrum for a protein backbone spin system with several hundred coupled spins. Beyond the physical assumptions made by chemists at the problem formulation stage and the controllable numerical rounding error of the tensor train format itself [39], there are no approximations.

II. TENSOR PRODUCT FORMATS

Tensor product expressions appear naturally in spin dynamics because the state space of a multi-spin system is a direct product of state spaces of individual spins [40]. A simple example is the nuclear Zeeman interaction Hamiltonian

$$\hat{H}_Z = \sum_{n=1}^{N} \vec{B}_0 \cdot \mathbf{A}^{(n)} \cdot \hat{\vec{S}}^{(n)}, \tag{1}$$

where N is the number of spins, $\vec{B}_0$ is the applied magnetic field, $\mathbf{A}^{(n)}$ are nuclear chemical shielding tensors, and the sum runs over all nuclei. Cartesian components of nuclear spin operators $\hat{\vec{S}}^{(n)} = [\hat{S}_x^{(n)} \quad \hat{S}_y^{(n)} \quad \hat{S}_z^{(n)}]$ have the following tensor product form

$$\hat{S}_{\{x,y,z\}}^{(n)} = \mathbf{1} \otimes \cdots \otimes \mathbf{1} \otimes \hat{\sigma}_{\{x,y,z\}} \otimes \mathbf{1} \otimes \cdots \otimes \mathbf{1}, \tag{2}$$

*d.savostyanov@soton.ac.uk

where **1** denotes an identity matrix of appropriate dimension and Pauli matrices $\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z$ occur at the nth position in the tensor product sequence. This representation is known in numerical linear algebra as the canonical polyadic (CP) format [9]. Although CP representations have been known in magnetic resonance spectroscopy for a long time [41], they suffer in practice from rapid inflation—spin Hamiltonians encountered in NMR and ESR (electron spin resonance) systems can be complicated [40] and, even for simple initial conditions, the number of terms in the canonical decomposition increases rapidly during system evolution. More ominously, the number of CP terms can change dramatically after small perturbations of the Hamiltonian or the system state. A simple example is

$$\hat{S}_z = \sum_{n=1}^{N} \hat{S}_z^{(n)} = \lim_{\varepsilon \to 0} \frac{(\mathbf{1} + \varepsilon \hat{\sigma}_z)^{\otimes N} - \mathbf{1}^{\otimes N}}{\varepsilon}, \quad (3)$$

where $\hat{a}^{\otimes N} = \hat{a} \otimes \cdots \otimes \hat{a}$. The left-hand side of this equation contains N direct product terms, given by Eq. (2), but the expression approximating it on the right-hand side has only two direct product terms, and one could be tempted to use it to reduce storage and CPU time. However, both terms of the approximation grow to infinity when $\varepsilon \to 0$, and the accuracy is lost due to rounding errors. Such instabilities in the CP format make it difficult to use—in finite precision arithmetic the number of terms in the decomposition quickly becomes equal to the dimension of the full state space and any efficiency savings disappear.

Unlike the CP format, which is an open tensor network, closed tensor network formats are stable to small perturbations. The most popular closed tensor network format was repeatedly rediscovered and is currently known under three different names: DMRG in condensed-matter physics [13,14], MPS [15,16,42,43] and MPO [17,44] in computational physics, and tensor train (TT) in numerical linear algebra [39]. A tensor train is defined, using the standard notation of numerical linear algebra [45], as follows:

$$\hat{x} = \tau(\hat{x}^{(1)}, \dots, \hat{x}^{(N)})$$
$$= \sum_{\alpha_1, \dots, \alpha_{N-1}} \hat{x}_{\alpha_1}^{(1)} \otimes \hat{x}_{\alpha_1, \alpha_2}^{(2)} \otimes \cdots \otimes \hat{x}_{\alpha_{N-1}}^{(N)}. \quad (4)$$

The TT representation of the total $\hat{S}_z$ operator in Eq. (3) is similar to the high-dimensional Laplacian [46]:

$$\hat{S}_z = \sum_{\alpha_1=1}^{2} \cdots \sum_{\alpha_{N-1}=1}^{2} \hat{H}_{\alpha_1}^{(1)} \otimes \hat{H}_{\alpha_1,\alpha_2}^{(2)} \otimes \cdots \otimes \hat{H}_{\alpha_{N-1}}^{(N)}, \quad (5)$$

with $\hat{H}^{(1)} = [\hat{\sigma}_z \quad \mathbf{1}], \hat{H}^{(2)} = \dots = \hat{H}^{(N-1)} = \begin{bmatrix} \mathbf{1} & 0 \\ \hat{\sigma}_z & \mathbf{1} \end{bmatrix}$, and $\hat{H}^{(N)} = [\mathbf{1} \quad \hat{\sigma}_z]^{\top}$. The number of terms in each summation (known as bond dimension, or TT rank) is two, and the number of entries of the decomposition is now bounded. The TT representation of $\hat{S}_z$ in Eq. (5) has $4N - 4$ single-spin operators, each of which is either zero, or identity **1**, or the Pauli matrix $\hat{\sigma}_z$. The CP representation of $\hat{S}_z$ in Eq. (3) has N^2 such operators—the tensor train representation is clearly more memory efficient.

Another notable example is the ZZ coupling Hamiltonian that often makes an appearance in models of simple linear spin

chains:

$$\hat{J} = \sum_{m>n} \hat{S}_z^{(n)} \hat{S}_z^{(m)}. \quad (6)$$

As written, this is a CP format with $N(N-1)/2$ terms and $N^2(N-1)/2$ single-spin operators entering direct products. The corresponding TT representation is

$$\hat{J} = \sum_{\alpha_1=1}^{3} \cdots \sum_{\alpha_{N-1}=1}^{3} \hat{J}_{\alpha_1}^{(1)} \otimes \hat{J}_{\alpha_1,\alpha_2}^{(2)} \otimes \cdots \otimes \hat{J}_{\alpha_{N-1}}^{(N)}, \quad (7)$$

with $\hat{J}^{(1)} = [0 \quad \hat{\sigma}_z \quad \mathbf{1}]$, and

$$\hat{J}^{(2)} = \dots = \hat{J}^{(N-1)} = \begin{bmatrix} \mathbf{1} & 0 & 0 \\ \hat{\sigma}_z & \mathbf{1} & 0 \\ 0 & \hat{\sigma}_z & \mathbf{1} \end{bmatrix}, \quad \hat{J}^{(N)} = \begin{bmatrix} \mathbf{1} \\ \hat{\sigma}_z \\ 0 \end{bmatrix}.$$

Here each summation runs over three terms only, and the total number of single-spin operator matrices appearing in $\hat{J}_{\alpha_{n-1},\alpha_n}^{(n)}$ is $9N - 12$, much fewer than the one of the CP format in Eq. (6).

Storage requirements of tensor structured representations (both CP and TT) stand in sharp contrast with the classical approach to magnetic resonance simulations [1,5], where the Hamiltonian is represented as a $2^N \times 2^N$ sparse matrix with all nonzero entries stored in memory. As soon as the matrix is assembled, CPU and memory resources grow exponentially with the number of spins N, making the simulation prohibitively difficult for large systems. Tensor structured methods avoid this problem (it is known colloquially as the curse of dimensionality) by keeping all data in compressed formats of the form given in Eqs. (1) and (4) and manipulating it without ever opening up the Kronecker products.

A very considerable body of literature exists on manipulating expressions directly in tensor product formats [9,11,12]. In particular, a given matrix may be converted into the TT format using sequential singular value decompositions [8,39,43]. Given tensors in the TT format, one can perform linear or bilinear operations (addition, elementwise multiplication, matrix-vector multiplication) [10,39], Fourier transform [47,48], and convolution [49] directly in the TT format, avoiding exponentially large arrays and computational costs.

These developments would have permitted large-scale magnetic resonance simulations entirely in the TT format, were it not for a significant obstacle: the summation operation in tensor train representations is an expensive procedure that carries a significant accuracy penalty due to the need to recompress the representation to keep the bond dimensions low. Spin Hamiltonians of practically interesting biological systems contain many thousands of one- and two-spin terms of the kind shown in Eq. (2). Intermediate expressions in spin dynamics simulations also frequently involve large sums. We demonstrate below that in those circumstances the standard bundle-and-recompress tensor network summation procedure leads either to the bond dimension expansion beyond the limits of modern computing hardware, or to a catastrophic accuracy loss. This problem also occurs with three-dimensional potentials encountered in electronic structure theory [50–52]. Here we propose an alternative algorithm for computing large sums, based on alternating tensor train optimization, and use it to enable NMR simulations on protein-size spin systems.

III. SIMULATION SETTING AND EXPERIMENTAL CONTEXT

Fully ^{13}C- and ^{15}N-labeled protein human ubiquitin (PDB code 1D3Z, Fig. 1) containing over a thousand magnetic nuclei in 76 amino acid residues was chosen for testing purposes with two types of spin subsystem selection: backbone (H, N, C, CA, HA) and extended backbone (H, N, C, CA, CB, HA, HB). Both cases involve a weakly branched continuous chain of spin-spin couplings and are encountered in the simulation of a large class of protein backbone NMR experiments that map out the protein bonding network and thereby assist in molecular structure determination: HNCO [53], HNCOCA [54], HNCA [53], and HSQC [55]. The isotropic NMR Hamiltonian was assembled using chemical shift values from the BMRB database [56] and J couplings from the literature data [57–61]. In the cases where an experimental value of a particular J coupling was not available in the literature, it was estimated based on the known values for structurally similar substances [62–64]; for most NMR simulation purposes and certainly for the purpose of the demonstration of the performance of the tensor train algorithm the accuracy of such coupling estimates (about 20%) is sufficient. The raw data for the magnetic couplings used in this work is available in the example set supplied with the current public version of the SPINACH library [4].

NMR experiments were performed at 25° C on a Varian Inova 600 MHz (14.1 Tesla) spectrometer equipped with a Z-gradient triple-resonance cryogenic probe using a 0.5 mM sample of uniformly ^{13}C- and ^{15}N-labeled human ubiquitin in 10% D_2O. ^{15}N spectra were collected as 2D ^{1}H-^{15}N HSQC [65] spectra incorporating gradient-enhanced coherence selection [66] and water flip back. The spectra were recorded with acquisition times of 150 ms (t_1, ^{15}N) and 500 ms (t_2, ^{1}H). During the ^{15}N evolution period, $^1J_{HN}$ and $^{1,2}J_{NC}$ couplings were either allowed to evolve, or decoupled by insertion of a rectangular ^{15}N or a shaped 200μs ^{13}C inversion pulse using the central lobe of the sinc function. During ^{1}H acquisition ^{15}N nuclei were either evolved or decoupled using 40 ppm broadband WURST sequence [67].
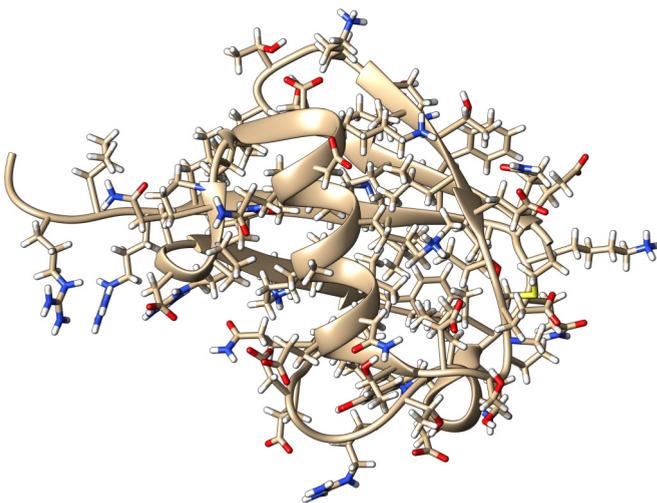
FIG. 1. (Color online) Human ubiquitin protein (PDB code 1D3Z): 76 amino acids, 563 magnetic nuclei in the extended backbone (H, N, C, CA, CB, HA, HB)

The liquid state NMR Hamiltonian of ^{13}C-, ^{15}N-labeled ubiquitin is:

$$\hat{H}(t) = \sum_k \omega_k \hat{S}_z^{(k)} + 2\pi \sum_{l>m} J_{strong}^{(l,m)} \hat{\vec{S}}^{(l)} \cdot \hat{\vec{S}}^{(m)}$$
$$+ 2\pi \sum_{p>q} J_{weak}^{(p,q)} \hat{S}_z^{(p)} \hat{S}_z^{(q)}$$
$$+ \omega_x(t) \sum_r \hat{S}_x^{(r)} + \omega_y(t) \sum_s \hat{S}_y^{(s)}, \quad (8)$$

where canonical NMR spectroscopy notation is used [40], k index runs over all nuclei, l and m indices run over pairs of nuclei that belong to the same isotope, p and q run over pairs of nuclei that belong to different isotopes, r and s run over the nuclei influenced by radio frequency pulses, $\omega_x(t)$ and $\omega_y(t)$ are time profiles of those pulses, ω_k are offset frequencies arising from the chemical shielding of the corresponding nuclei [68], $J_{strong}^{(l,m)}$ are strong NMR J couplings [69], $J_{weak}^{(p,q)}$ are weak NMR J couplings [40], and spin operators $\hat{S}_x^{(n)}, \hat{S}_y^{(n)}, \hat{S}_z^{(n)}$ are defined by Eq. (2). In the case of extended ubiquitin backbone, the Hamiltonian in Eq. (8) contains 563 shielding terms, 1840 coupling terms, and 1126 radio frequency terms. All calculations reported below were performed by extending the functionality of SPINACH library [4] to the tensor train formalism and interfacing it to TT-Toolbox [70] where appropriate.

Due to the abundance of complicated multipulse NMR experiments with time-dependent Hamiltonians [40], magnetic resonance simulations are generally carried out in the time domain. They always require long-term evolution trajectories with accurate phases (at least 100 ms, much longer than the reciprocal Hamiltonian norm) for the density operator $\hat{\rho}(t)$ under the Liouville–von Neumann equation:

$$\frac{d}{dt}\hat{\rho}(t) = -i[\hat{H}(t), \hat{\rho}(t)] + \hat{\hat{R}}(\hat{\rho}(t) - \hat{\rho}_{eq}),$$
$$O(t) = \langle \hat{O} \mid \hat{\rho}(t) \rangle = \text{Tr}[\hat{O}^\dagger \hat{\rho}(t)],$$
$$\hat{\rho}_{eq} = \frac{\exp(-\hat{H}/k_B T)}{\text{Tr}\exp(-\hat{H}/k_B T)}, \quad (9)$$

where $\hat{\hat{R}}$ is the relaxation superoperator ($T_{1,2}$ model with literature values for relaxation times [34] was used in the present work), $\hat{\rho}_{eq}$ is the thermal equilibrium state, and $\hat{O}$ is the observable operator, usually a sum of $\hat{S}_x$ or $\hat{S}_y$ operators on the spins of interest. In very simple cases where the Hamiltonian is not time dependent, the general solution to Eq. (9) can be written as:

$$O(t) = \langle \hat{O} \mid \exp[-i(\hat{\hat{H}} + i\hat{\hat{R}})t] \mid \hat{\rho}_0 \rangle, \quad (10)$$

where $\hat{\hat{H}}$ is the Hamiltonian commutation superoperator.

Direct time-domain evaluation of this equation in tensor train format, either using explicit operator exponentiation or Krylov-type propagation techniques, does not appear to be possible—in all cases described by Eq. (8) the ranks in the tensor train expansion quickly grow beyond the capacity of modern computers. Increasing the singular value cutoff threshold at the representation compression stage leads to catastrophic loss of accuracy. Fortunately, there are simple

cases (most notably pulse-acquire 1D NMR spectroscopy) where amplitudes at only a few specific frequencies are actually required for the Fourier transform of Eq. (10), meaning that the problem can be reformulated in the frequency domain:

$$O(\omega) = -i\langle \hat{O} | (\hat{\hat{H}} + i\hat{\hat{R}} + \omega \mathbf{1})^{-1} | \hat{\rho}_0 \rangle. \qquad (11)$$

That is, to compute the observable at the point ω in the frequency domain, we need to solve a linear system $(\hat{\hat{H}} + i\hat{\hat{R}} + \omega \mathbf{1})|\hat{x}\rangle = |\hat{\rho}_0\rangle$. The problem formulation in Eq. (11) sacrifices a great deal of generality compared to Eq. (9) (simulation of arbitrary NMR pulse sequences is no longer possible), but it does serve as a stepping stone and enables the demonstration calculation presented below.

IV. TENSOR TRAIN ALGORITHM FOR THE SUMMATION AND SOLUTION OF LINEAR SYSTEMS

The DMRG algorithm was initially proposed [13,14] to find the ground state of a Hermitian matrix A by the minimization of the Rayleigh quotient $Q(x) = x^*Ax/x^*x$. The dynamical DMRG algorithm [35] was then developed to find the solution of a linear system $Ax = b$ with a Hermitian positive definite matrix A by the minimization of the energy function $J(x) = x^*Ax - 2\Re(x^*b)$. Apart from the change of the minimization target function, the two algorithms are similar.

In DMRG formalism the solution is sought in the form of a tensor train introduced in Eq. (4), but the minimization over all cores $x^{(n)}$ simultaneously is a complicated nonlinear problem. To make the procedure feasible, it is replaced by a sequence of optimizations carried over one core at a time:

$$x_\star^{(n)} = \arg\min_{x^{(n)}} J[\tau(x^{(1)}, \dots, x^{(n)}, \dots, x^{(N)})]. \qquad (12)$$

The TT format is linear in all cores $x^{(n)}$. This fact may be expressed as $x = X_{\neq n} x^{(n)}$, where the frame matrix $X_{\neq n}$ maps the parameters of the TT core $x^{(n)}$ to the vector x. The linearity allows to rewrite Eq. (12) as $x_\star^{(n)} = \arg\min J_n(x^{(n)}) = A_n^{-1} b_n$, where J_n is the energy function for the local problem $A_n x^{(n)} = b_n$ with $A_n = X_{\neq n}^* A X_{\neq n}$ and $b_n = X_{\neq n}^* b$. Using the nonuniqueness of the tensor train representation (4), one can always construct the representation with the unitary frame matrix $X_{\neq n}$, that guarantees the stability of the local problem. Such a choice is known as gauge condition in the MPS literature, and canonical form in the DMRG literature. After the solution $x_\star^{(n)}$ is computed, we substitute $x^{(n)} := x_\star^{(n)}$ in the tensor train, and continue for $n = 1, \dots, N$, and then back and forth along the chain.

The convergence of the above described one-site DMRG procedure depends on the initial guess and in particular on the initial choice of the TT ranks because they remain the same during the sequence of updates defined by Eq. (12). This is a severe restriction and additional measures are therefore taken to adapt the TT ranks during the computations. One way to do that is to replace the optimization over single cores by the optimization over pairs of neighboring cores, and then to adapt the TT rank between them. Another possibility is to expand the search space by adding auxiliary directions. The first method of the latter type is the corrected one-site DMRG algorithm [71], which targets in addition to x a surrogate of the next Krylov vector Ax.

For the solution of linear systems, the alternating minimal energy (AMEn) algorithm was recently proposed [37], which also uses an additional direction to adapt tensor train ranks. The local optimization step in AMEn is carried over one site only. To adapt TT ranks and improve convergence, TT blocks are expanded by auxiliary information, $x^{(n)} := [x_\star^{(n)} \quad r^{(n)}]$. The enrichment $r^{(n)}$ introduces new directions in the subspace spanned by $X_{\neq n+1}$. A good choice of the enrichment is the component $r^{(n)}$ of the TT representation (exact or approximate) $\tilde{r} = \tau(r^{(1)}, \dots, r^{(N)})$ of the residual $r = b - Ax$. AMEn algorithm is as fast as one-site methods, but as rank adaptive as the two-site DMRG algorithm, and demonstrates comparable or better convergence rates. For the solution of a linear system $Ax = b$ with a Hermitian positive definite matrix, it has a proven global bound on the geometrical convergence rate. Unlike the corrected one-site DMRG method [71], the AMEn algorithm is stable to perturbations and free from tuning parameters and heuristics [38]. The rank adaptation strategy in the enrichment phase of AMEn is determined by a single relative accuracy parameter.

In this work we use the AMEn algorithm for two purposes. First, we apply it to a system with a trivial matrix $A = \mathbf{1}$, but a complicated right-hand side b, which is a sum of many elementary tensors like the one in Eq. (2). This allows us to compress a Hamiltonian returned by the SPINACH package from the CP format given by Eq. (8) into the TT format Eq. (4). The Hamiltonian is stretched into a vector, and the target functional $J(x) = \|x - b\|^2$ is a Frobenius-norm distance between a given Hamiltonian b and Hamiltonian x sought in the tensor train format. The one-site optimization in Eq. (12) is effectively the solution of the overdetermined linear system $X_{\neq n} x^{(n)} = b$ using the least squares method. For the unitary frame matrix we have $x_\star^{(n)} = X_{\neq n}^* b$, and therefore the local optimization step is obtained by contracting the frame matrix with the given Hamiltonian b. The enrichment step uses a low-rank approximation of the error $r = b - x$, which is obtained by one-site DMRG optimization.

After the Hamiltonian is compressed into the tensor train format, we compute 1D NMR spectra by solving the linear system in Eq. (11). Since the matrix $\hat{\hat{A}} = \hat{\hat{H}} + i\hat{\hat{R}} + \omega \mathbf{1}$ is not expected to be Hermitian positive definite, we consider instead an equivalent symmetrized problem $(\hat{\hat{A}}^* \hat{\hat{A}})\hat{\rho}(\omega) = \hat{\hat{A}}^* \hat{\rho}_0$. For demonstration purposes, we chose a simple nonselective damping relaxation model $\hat{\hat{R}} = -\mu \mathbf{1}$, and the same $\hat{S}_+ = \hat{S}_x + i\hat{S}_y$ operator for the initial and the detection state, where $\hat{S}_{\{x,y\}} = \sum_n \hat{S}_{\{x,y\}}^{(n)}$ are the total spin operators of all ^{15}N nuclei in the system. This avoids explicit radio frequency pulses in the system and makes the Hamiltonian in Eq. (8) time independent and real valued, $\hat{H}(t) = \hat{H} = \hat{H}^* = \hat{H}^T$, those properties are also inherited by the commutation superoperator $\hat{\hat{H}}$. Since the detection state $\hat{O}$ is also real valued, the NMR spectrum $O(\omega)$ in Eq. (11) can be computed from $\Im\hat{\rho}(\omega)$, that we obtain as follows:

$$\Im\hat{\rho}(\omega) = \mu(\hat{\hat{H}}^*\hat{\hat{H}} + 2\omega\hat{\hat{H}} + (\omega^2 + \mu^2)\mathbf{1})^{-1}\hat{\rho}_0. \qquad (13)$$

This equation is solved by the AMEn algorithm at each point ω in the user-specified frequency interval.

V. RESULTS

As discussed above, a major problem in the application of tensor train methods to magnetic resonance simulation of large systems is the calculation of lengthy sums involved in the construction of spin Hamiltonians and density matrices, and their compression into the TT format. Figure 2 illustrates the performance of our proposed solution to this problem in the case of minimal (H, N, C, CA, HA) and extended (H, N, C, CA, HA, CB, HB) ubiquitin backbone spin systems. Storage requirements for the TT format in Eq. (4) depend on all TT ranks (bond dimensions) $k_1, \ldots, k_{N-1}$, and are characterized by the effective TT rank k, defined by $Nk^2 = \sum_{n=1}^{N} k_{n-1}k_n$. It is clear from the left panels of Fig. 2 that the primary obstacle —rapid growth in the tensor train rank—has been removed by the AMEn method: the effective ranks stay below 50 for the extended backbone and below 40 for the minimal backbone, well within the capability of modern desktop workstations. Since k^2 is smaller than the number of terms in the CP representation, the TT format with Nk^2 operators provides more compact storage than the CP format.

The alternative to AMEn is binary summation, which adds up Hamiltonian terms pairwise and recompresses the representation after each addition. As demonstrated in Fig. 2, binary summation drives tensor train ranks up to several hundred and thereby makes the solution of the linear system in Eq. (13) exceedingly difficult. It is clear from the right panels of Fig. 2 that the CPU time requirements of AMEn summation compared to binary summation are essentially the same, making AMEn procedure clearly superior for all practical purposes. The resulting representation of the ubiquitin backbone spin Hamiltonian matrix is, up to the
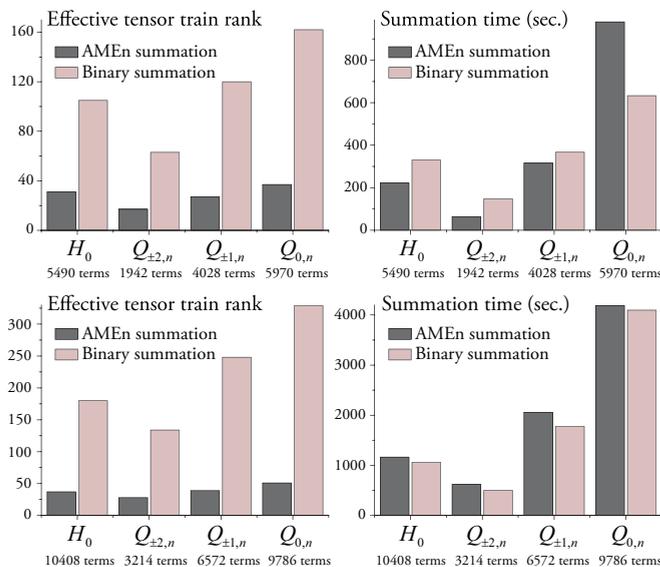
FIG. 2. (Color online) Performance comparison for binary and AMEn tensor train summation with relative accuracy parameter $\varepsilon = 10^{-12}$ during the construction of the NMR spin Hamiltonian. Top: human ubiquitin backbone (H, N, C, CA, HA); bottom: human ubiquitin extended backbone (H, N, C, CA, HA, CB, HB). Here H_0 refers to the isotropic part of the Hamiltonian and $Q_{k,n}$ to the irreducible spherical components of the anisotropic part.
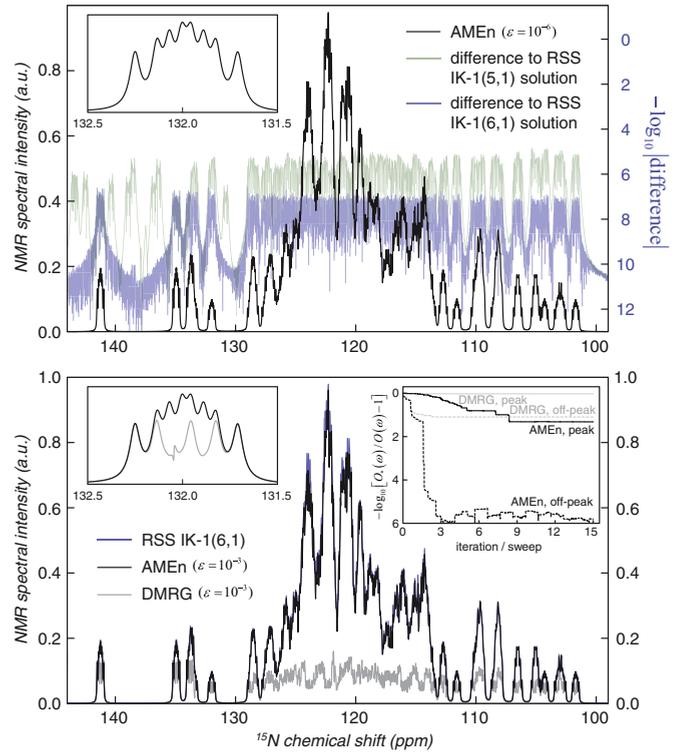
FIG. 3. (Color online) Amino group region of the pulse-acquire ^{15}N NMR spectrum of human ubiquitin. Top: spectrum computed by AMEn algorithm [37] with accuracy parameter $\varepsilon = 10^{-6}$ is compared to the results obtained by the restricted state space (RSS) approximation [4] with basis containing local spin correlations of orders up to 5 and 6. Bottom: accurate RSS computation is used as a reference $O_\star(\omega)$ and compared to the spectra $O(\omega)$ computed by AMEn and DMRG [35,36], both using the accuracy parameter $\varepsilon = 10^{-3}$. Right subgraph: convergence of AMEn and DMRG methods at two points of the frequency domain (dashed lines: an off-peak point at 100 ppm, solid lines: a peak at 122 ppm).

rounding error of the complex double precision arithmetic, exact. In magnetic resonance spectroscopy this is an unprecedented development—ubiquitin NMR simulation is currently just about feasible [72], with significant approximations and colossal computational resources. Tensor train representation is therefore a large step forward, even though Eq. (11) is not in general applicable to arbitrary NMR experiments.

After the Hamiltonian is compressed, we compute ^{15}N pulse-acquire NMR spectra using Eq. (13) with the AMEn algorithm and compare it to the simulation produced by the restricted state space (RSS) approximation [4], which is currently the only other method that is capable of handling NMR systems of this size. As demonstrated in Fig. 3 (top), When the basis set used by RSS is increased, its result converges to the one produced by AMEn, and the relative deviation between two methods falls below 10^{-6} across the frequency interval.

It is instructive to compare the results of AMEn simulations with those produced by the dynamical DMRG [35] technique. As shown in Fig. 3 (bottom), the NMR spectrum computed by AMEn matches the reference spectrum returned by RSS with only minor deviations, while the accuracy of the result

computed by the dynamical DMRG algorithm at the same relative accuracy parameter is unacceptable. DMRG does of course produce the right answer if a much tighter accuracy parameter is specified, but the simulation time goes up by several orders of magnitude. AMEn does therefore appear to have a better accuracy-to-effort ratio. This is also confirmed by the convergence graph of AMEn and DMRG, given in the same figure, where the relative deviation between the computed and the reference values $O(\omega)$ is shown during the iterations (sweeps) for both DMRG and AMEn. Note also that the inexact values of the spectrum, computed by AMEn and DMRG are always below the reference values; this was first noted by Jeckelmann [35]. The comparison in Fig. 3 is made using $\varepsilon = 10^{-3}$ to visually emphasize the observed difference between the two methods; the same conclusion also holds for more accurate calculations using $\varepsilon = 10^{-8}$.

Due to the intrinsically low sensitivity of liquid-state ^{15}N protein NMR spectroscopy, it is not possible to record the experimental equivalent of Fig. 3 directly with a sufficient signal-to-noise ratio; we have therefore taken a somewhat longer route to the experimental validation of the tensor train simulation. Figure 4 shows experimental proton-detected ^{1}H-^{15}N HSQC spectra of ubiquitin, compared to the simulations obtained at the basis set limit of the RSS formalism [6]. Perfect agreement is apparent in both cases. This provides experimental evidence to the accuracy of the restricted state space method. The tensor train results in Fig. 3 can now be justified by comparison to the RSS results—it is clear that the TT formalism performs as intended.

VI. DISCUSSION

The successful 1D NMR simulation notwithstanding, very significant obstacles remain on the path to practical applications of the tensor train formalism to NMR spectroscopy. The following issues should be addressed in future work to fully uncover the potential of the DMRG/MPS/TT formalism for spin dynamics simulations:

(i) The requirement for the spin system to be a chain or a tree should be lifted. Biological magnetic resonance spin systems are irregular polycyclic interaction networks with multiple interlocking loops in the coupling graph, particularly in solid-state NMR, where internuclear dipolar couplings form very dense meshes. A generalization of tensor train algorithms to general contraction networks that fully mimic the molecular structure is therefore required.

(ii) Rank explosion problem for time-domain simulations should be solved. It is clear from the success of the restricted state space approximation [4,7] that the order of spin correlation in many evolving magnetic resonance spin systems either is or may safely be assumed to be quite low. This suggests data sparsity and separability, and indicates that some kind of low-rank decomposition is possible. One likely direction is through the enforcement of symmetries and conservation laws within the tensor train format itself during time evolution.

(iii) Our experience indicates that tensor train objects are very far from being drop-in replacements for their matrix counterparts in standard simulation algorithms and software—it does actually appear that nearly everything in the very considerable body of magnetic resonance simulation methods
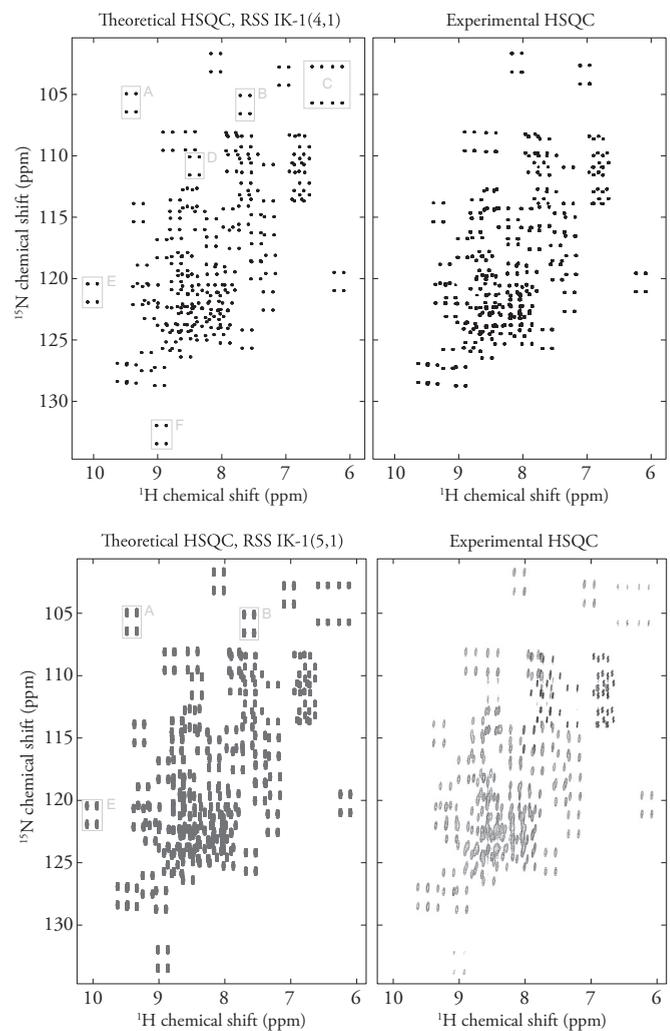
FIG. 4. Theoretical (left) and experimental (right) ^{1}H-^{15}N HSQC spectra of ^{15}N-,^{13}C-labeled human ubiquitin. Top: proton decoupling switched off in the indirect dimension and nitrogen decoupling switched off in the direct dimension to demonstrate accurate quantum mechanical treatment of spin-spin coupling by the simulation. Bottom: additionally, carbon decoupling is switched off in the indirect dimension. Signal groups marked A–F in the theoretical spectrum are not visible in the experimental data due to partial deuteration and slow conformational exchange of the corresponding amino acid residues.

needs to be adapted to the realities of DMRG. Current implementation of tensor product methods still requires a number of tuning parameters (approximation accuracies, TT ranks of the enrichment, etc.). Broad adoption of tensor network algorithms would require basic linear algebra operations to be handled transparently and seamlessly by the existing simulation software packages, in the same way as sparse matrices currently are.

(iv) Transparent and clear tensor train approximation accuracy criteria, rank control and *a priori* error bounds should be developed in order to estimate the influence of the representation compression errors on the accuracy of the final result. This problem is particularly acute for the state vector phase in time-domain simulations: magnetic

resonance experiments rely critically on the phase being correctly predicted.

All of that having been said, we are very optimistic about the future of DMRG/MPS/TT methods, having also found them useful in Fokker-Planck type formalisms related to NMR and EPR spectroscopy [73]. Their primary strength is the lack of heuristic assumptions and the controllable nature of the representation accuracy. An experimental implementation of tensor train magnetic resonance simulation paths, via an interface to the TT-Toolbox [70], is available in version 1.3.1980 of our SPINACH library [4].

VII. CONCLUSIONS AND FUTURE WORK

Even with their well-documented limitations (the requirement for the spin system to be close to a chain, difficulty with long-range time-domain simulations, code implementation challenges, etc.), the ability of tensor network formalisms to simulate simple liquid state NMR spectra of large spin systems essentially without approximations is impressive. They cannot yet match the highly optimized dedicated methods developed by the magnetic resonance community [72], but if some of the limitations are lifted by the subsequent research, DMRG methods would have the potential to become a very useful formalism in NMR research.

Having solved in this paper the last purely technical problem on the way to the broad adoption of tensor train formalism in magnetic resonance spectroscopy, we are quite optimistic about its potential. In particular, the following avenues appear promising:

(1) Generalizing AMEn method to arbitrary tensor networks, e.g., tree tensor networks [74–77], that closely match the coupling topology of the spin system.

(2) Development of reliable tensor train methods for solving linear systems of algebraic equations with indefinite matrices, and time evolution problems.

(3) Development of tensor product methods that reduce memory requirements and accelerate convergence by enforcing conservation laws [8,78,79] and matrix symmetries [80].

Elsewhere in magnetic resonance, benefits to electron spin resonance spectroscopy, with its star-shaped spin interactions graphs, are likely to be harder to achieve, but may still be obtained by exploiting the direct product structure of combined spin and spatial dynamics appearing in Fokker-Planck-type problems [73].

[1] M. Bak, J. T. Rasmussen, and N. C. Nielsen, J. Magn. Reson. **147**, 296 (2000).

[2] M. C. Butler, J. N. Dumez, and L. Emsley, Chem. Phys. Lett. **477**, 377 (2009).

[3] R. S. Dumont, S. Jain, and A. Bain, J. Chem. Phys. **106**, 5928 (1997).

[4] H. J. Hogben, M. Krzystyniak, G. T. P. Charnock, P. J. Hore, and I. Kuprov, J. Magn. Reson. **208**, 179 (2011).

[5] M. Veshtort and R. G. Griffin, J. Magn. Reson. **178**, 248 (2006).

[6] A. Karabanov, I. Kuprov, G. T. P. Charnock, A. van der Drift, L. J. Edwards, and W. Köckenberger, J. Chem. Phys. **135**, 084106 (2011).

[7] I. Kuprov, N. Wagner-Rundell, and P. J. Hore, J. Magn. Reson. **189**, 241 (2007).

[8] U. Schollwöck, Rev. Mod. Phys. **77**, 259 (2005).

[9] T. G. Kolda and B. W. Bader, SIAM Review **51**, 455 (2009).

[10] U. Schollwöck, Ann. Phys. (NY) **326**, 96 (2011).

[11] B. N. Khoromskij, Chemometr. Intell. Lab. Syst. **110**, 1 (2012).

[12] W. Hackbusch, *Tensor spaces and numerical tensor calculus* (Springer-Verlag, Berlin, 2012).

[13] S. R. White, Phys. Rev. Lett. **69**, 2863 (1992).

[14] S. R. White, Phys. Rev. B **48**, 10345 (1993).

[15] M. Fannes, B. Nachtergaele, and R. Werner, Comm. Math. Phys. **144**, 443 (1992).

[16] A. Klümper, A. Schadschneider, and J. Zittartz, Europhys. Lett. **24**, 293 (1993).

[17] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac, Phys. Rev. Lett. **93**, 207204 (2004).

[18] R. Bursill, T. Xiang, and G. Gehring, J. Phys.: Condens. Matter **8**, L583 (1996).

[19] N. Shibata, J. Phys. Soc. Jpn. **66**, 2221 (1997).

[20] X. Wang and T. Xiang, Phys. Rev. B **56**, 5061 (1997).

[21] S. Liang and H. Pang, Phys. Rev. B **49**, 9214 (1994).

[22] I. P. Mcculloch, Philos. Mag. B **81**, 1603 (2001).

[23] T. Xiang, J. Lou, and Z. Su, Phys. Rev. B **64**, 104414 (2001).

[24] E. M. Stoudenmire and S. R. White, Annu. Rev. Condens. Matter Phys. **3**, 111 (2012).

[25] S. R. White and R. L. Martin, J. Chem. Phys. **110**, 4127 (1999).

[26] G. K.-L. Chan and M. Head-Gordon, J. Chem. Phys. **116**, 4462 (2002).

[27] D. Zgid and M. Nooijen, J. Chem. Phys. **128**, 014107 (2008).

[28] Y. Kurashige and T. Yanai, J. Chem. Phys. **130**, 234114 (2009).

[29] S. Wouters, P. A. Limacher, D. Van Neck, and P. W. Ayers, J. Chem. Phys. **136**, 134110 (2012).

[30] S. Sharma and G. K.-L. Chan, J. Chem. Phys. **136**, 124121 (2012).

[31] S. Wouters, W. Poelmans, P. W. Ayers, and D. Van Neck, Comput. Phys. Commun. **185**, 1501 (2014).

[32] S. Pittel and N. Sandulescu, Phys. Rev. C **73**, 014301 (2006).

[33] J. Rotureau, N. Michel, W. Nazarewicz, M. Płoszajczak, and J. Dukelsky, Phys. Rev. Lett. **97**, 110603 (2006).

[34] J. Cavanagh, W. J. Fairbrother, A. G. Palmer, III, M. Rance, and N. J. Skelton, *Protein NMR Spectroscopy: Principles and Practice*, 2nd ed. (Elsevier Academic Press, Amsterdam, 2007).

[35] E. Jeckelmann, Phys Rev B **66**, 045114 (2002).

[36] S. V. Dolgov and I. V. Oseledets, SIAM J. Sci. Comput. **34**, A2718 (2012).

[37] S. V. Dolgov and D. V. Savostyanov, arXiv:1301.6068 [SIAM J. Sci. Comp. (to be published)]; arXiv:1304.1222 [SIAM J. Sci. Comp. (to be published)].

[38] S. V. Dolgov and D. V. Savostyanov, arXiv:1312.6542.

[39] I. V. Oseledets, SIAM J. Sci. Comput. **33**, 2295 (2011).

[40] R. R. Ernst, G. Bodenhausen, and A. Wokaun, *Principles of Nuclear Magnetic Resonance in One and Two Dimensions*, Vol. 14 (Clarendon Press, Oxford, 1987).

[41] R. Brüschweiler and R. R. Ernst, J. Magn. Reson. **124**, 122 (1997).

[42] F. Verstraete and J. I. Cirac, Phys. Rev. B **73**, 094423 (2006).

[43] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Quantum Info. Comput. **7**, 401 (2007).

[44] G. M. Crosswhite, A. C. Doherty, and G. Vidal, Phys. Rev. B **78**, 035116 (2008).

[45] S. V. Dolgov, B. N. Khoromskij, I. V. Oseledets, and D. V. Savostyanov, Computer Phys. Comm. **185**, 1207 (2014).

[46] V. A. Kazeev and B. N. Khoromskij, SIAM J. Matrix Anal. Appl. **33**, 742 (2012).

[47] D. E. Browne, New J. Phys. **9**, 146 (2007).

[48] S. V. Dolgov, B. N. Khoromskij, and D. V. Savostyanov, J. Fourier Anal. Appl. **18**, 915 (2012).

[49] V. Kazeev, B. Khoromskij, and E. Tyrtyshnikov, SIAM J. Sci. Comp. **35**, A1511 (2013).

[50] B. N. Khoromskij and V. Khoromskaia, SIAM J. Sci. Comput. **31**, 3002 (2009).

[51] D. V. Savostyanov, Numer. Math. Theor. Meth. Appl. **2**, 439 (2009).

[52] S. A. Goreinov, I. V. Oseledets, and D. V. Savostyanov, SIAM J. Sci. Comput. **34**, A1 (2012).

[53] L. E. Kay, M. Ikura, R. Tschudin, and A. Bax, J. Magn. Reson. (1969) **89**, 496 (1990).

[54] A. Bax and M. Ikura, J. Biomolecular NMR **1**, 99 (1991).

[55] J. Schleucher, M. Schwendinger, M. Sattler, P. Schmidt, O. Schedletzky, S. Glaser, O. W. Sørensen, and C. Griesinger, J. Biomolecular NMR **4**, 301 (1994).

[56] E. L. Ulrich, H. Akutsu, J. F. Doreleijers, Y. Harano, Y. E. Ioannidis, J. Lin, M. Livny, S. Mading, D. Maziuk, Z. Miller, E. Nakatani, C. F. Schulte, D. E. Tolmie, R. K. Wenger, H. Yao, and J. L. Markley, Nucleic Acids Res. **36**, D402 (2008).

[57] D. A. Case, C. Scheurer, and R. Brüschweiler, J. Am. Chem. Soc. **122**, 10390 (2000).

[58] F. Löhr, J. M. Schmidt, and H. Rüterjans, J. Am. Chem. Soc. **121**, 11821 (1999).

[59] C. Pérez, F. Löhr, H. Rüterjans, and J. M. Schmidt, J. Am. Chem. Soc. **123**, 7081 (2001).

[60] B. Vögeli, J. Ying, A. Grishaev, and A. Bax, J. Am. Chem. Soc. **129**, 9377 (2007).

[61] A. C. Wang and A. Bax, J. Am. Chem. Soc. **118**, 2483 (1996).

[62] M. Barfield, R. J. Spear, and S. Sternhell, Chem. Rev. **76**, 593 (1976).

[63] P. E. Hansen, Prog. NMR Spectros. **14**, 175 (1981).

[64] L. B. Krivdin and E. W. Della, Prog. NMR Spectros. **23**, 301 (1991).

[65] G. Bodenhausen and D. J. Ruben, Chem. Phys. Lett. **69**, 185 (1980).

[66] L. Kay, P. Keifer, and T. Saarinen, J. Am. Chem. Soc. **114**, 10663 (1992).

[67] Ē. Kupce and R. Freeman, J Magn. Reson. A **117**, 246 (1995).

[68] N. F. Ramsey, Phys. Rev. **78**, 699 (1950).

[69] N. F. Ramsey, Phys. Rev. **91**, 303 (1953).

[70] I. Oseledets *et al.*, https://github.com/oseledets/TT-Toolbox.

[71] S. R. White, Phys. Rev. B **72**, 180403 (2005).

[72] L. J. Edwards, D. Savostyanov, Z. Welderufael, D. Lee, and I. Kuprov, J. Magn. Reson. **243**, 107 (2014).

[73] L. E. Edwards, D. V. Savostyanov, A. A. Nevzorov, M. Concistrè, G. Pileo, and I. Kuprov, J. Magn. Reson. **235**, 121 (2013).

[74] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Phys. Rev. A **74**, 022320 (2006).

[75] V. Murg, F. Verstraete, O. Legeza, and R. M. Noack, Phys. Rev. B **82**, 205105 (2010).

[76] I. Pižorn, F. Verstraete, and R. M. Konik, Phys. Rev. B **88**, 195102 (2013).

[77] W. Hackbusch and S. Kühn, J. Fourier Anal. Appl. **15**, 706 (2009).

[78] S. Ramasesha, S. K. Pati, H. R. Krishnamurthy, Z. Shuai, and J. L. Brédas, Phys. Rev. B **54**, 7598 (1996).

[79] I. P. McCulloch and M. Gulácsi, Europhys. Lett. **57**, 852 (2002).

[80] T. K. Huckle, K. Waldherr, and T. Schulte-Herbrüggen, Linear and Multilinear Algebra **61**, 91 (2013).