

X-symbols for non-Abelian symmetries in tensor networks

Andreas Weichselbaum ^{*}

Department of Condensed Matter Physics and Materials Science, Brookhaven National Laboratory, Upton, New York 11973-5000, USA and Physics Department, Arnold Sommerfeld Center for Theoretical Physics, and Center for NanoScience, Ludwig-Maximilians-Universität, Theresienstrasse 37, 80333 Munich, Germany



(Received 30 October 2019; revised manuscript received 22 January 2020; accepted 15 May 2020; published 23 June 2020)

The full exploitation of non-Abelian symmetries in tensor network states (TNSs) derived from a given lattice Hamiltonian is attractive in various aspects. From a theoretical perspective, it can offer deep insights into the entanglement structure and quantum information content of strongly correlated quantum many-body states. From a practical perspective, it allows one to push numerical efficiency by orders of magnitude. Physical expectation values based on TNSs require the full contraction of a given tensor network, with the elementary ingredient being a pairwise contraction. While well established for no or just Abelian symmetries, this can become quickly extremely involved and cumbersome for general non-Abelian symmetries. As shown in this paper, however, the elementary step of a pairwise contraction of tensors of arbitrary rank can be tackled in a transparent and efficient manner by introducing so-called X-symbols. These deal with the pairwise contraction of the generalized underlying Clebsch-Gordan tensors (CGTs). They can be computed deterministically once and for all, and hence they can also be tabulated. Akin to $6j$ -symbols, X-symbols are generally much smaller than their constituting CGTs. In applications, they solely affect the tensors of reduced matrix elements and therefore, once tabulated, allow one to completely sidestep the explicit usage of CGTs, and thus to greatly increase numerical efficiency.

DOI: [10.1103/PhysRevResearch.2.023385](https://doi.org/10.1103/PhysRevResearch.2.023385)

Tensor network states (TNSs) provide a powerful natural framework for the numerical treatment of strongly correlated quantum many-body physics on lattice Hamiltonians [1–5]. Starting in one dimension (1D) with matrix product states (MPSs), the numerical renormalization group (NRG) [6–9] and the density-matrix renormalization group (DMRG) [10–13] represent powerful, nonperturbative, and accurate methods to deal with strongly correlated systems at arbitrary temperature both statically and dynamically. An attractive extension of the one-dimensional MPS structure was provided by the multiscale-entanglement renormalization ansatz (MERA) [14,15], also with an eye on higher dimensions, even if significantly more expensive numerically. The very flexible framework of TNSs for lattice Hamiltonians has already also seen a wide range of applications in two dimensions (2D) via projected entangled pair states (PEPS) [16–20] or higher [21–23], including applications in quantum chemistry [5,24,25]. By providing a natural algebraic structure to study entanglement in strongly correlated systems, this also generated significant interest from a quantum information perspective [26–28] including symmetry protected topological quantum phases [29–33]. While highly efficient in 1D, the numerical cost for dealing with TNSs in 2D or higher, however,

grows exorbitantly, even though at least still polynomially. Therefore the exploitation of symmetries is extremely relevant and important also on practical grounds [34–41]. In particular, this applies for correlated systems in quasi-1D, i.e., for long systems of narrow width, or for tree-tensor network (TTN) states [24,42,43] in the presence of multiple (symmetric) flavors in condensed-matter systems [44–49] or optical lattices [50–53].

Symmetries on *all* indices in a tensor network (TN) state, physical and bond indices alike, are well defined only in 1D or, more generally, in TTN states [24,42,43], i.e., tensor network states without loops along virtual bonds that link tensors. In the presence of loops, significant ambiguities arise. Nevertheless, on practical grounds, one typically sees that enforcing symmetries on all indices also in a general TNS with loops shows clear gains in efficiency [54,55]. An intuitive handwaving argument for this may be provided based on the interpretation of bonds as actual auxiliary state spaces which motivated PEPS to start with [16,18].

In the presence of non-Abelian symmetries [56–61], all tensors can be decomposed into a tensor product structure of reduced matrix element tensors (RMTs) and generalized Clebsch-Gordan coefficient tensors (CGTs) [36,45] as a direct generalization of the Wigner-Eckart theorem. This results in two immediate consequences that can be used to greatly improve numerical efficiency.

(i) By splitting of a CGT factor for each elementary, i.e., simple symmetry present, this allows one to (strongly) reduce the effective dimensionality $D \rightarrow D^*$ of a given index or state space by switching from a state-based description to a description based on entering the RMTs.

*weichselbaum@bnl.gov

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

(ii) The CGTs are purely related to the symmetry of a given problem. Hence much of it can be dealt with once and for all by tabulating the relevant information.

A general bottom-up framework for dealing with general non-Abelian symmetries in TNSs was introduced in Weichselbaum [45] based on a general transparent tensor representation referred to as QSPACE (see Appendix B). The approach taken there was based on the explicit utilization, i.e., generation and subsequent decomposition of CGTs. This is in stark contrast to other approaches based on fusion trees, F moves, etc., which are essentially built on $6j$ -symbols [40,62]. That approach works well for the symmetry $SU(2)$, where $6j$ -symbols are readily available analytically. But it becomes much more difficult for general non-Abelian symmetries such as $SU(N > 2)$, $SO(N > 4)$, or $Sp(2N > 2)$ where due to the presence of outer multiplicity $6j$ -symbols are not at all readily available. In contrast, the QSPACE approach is bottom up, as it solely relies on the bare structure of the Lie algebra [45]. The irreducible representations (ireps), multiplet fusion rules, and corresponding CGTs thus generated were already all tabulated in [45]. However, contractions of CGTs were not tabulated due to the presence of outer multiplicity (OM). The prescription in [45] to deal with OM in the pairwise contraction of tensors was to always reconstruct all CGTs based on their particular instantiation in OM space. However, from a practical point of view, this led to significant computational overhead for larger non-Abelian symmetries, such as $SU(N \gtrsim 4)$. As will be shown in the present paper, however, there also exists a transparent general way to deal with the problem of OM in the pairwise contraction of CGTs based on the introduction of so-called X-symbols (where “X” is simply a reference to generalized tensor multiplication, i.e., contraction). These can be computed once and for all, and thus also be tabulated. X-symbols provide an alternative approach to $6j$ -symbols. Yet they are much more naturally suited to tensor network algorithms, since they strictly deal with the elementary operation of the contraction of a pair of tensors (and hence a pair of CGTs) on an arbitrary subset of shared indices.

Given the many reviews and detailed publications that already exist on TNSs (e.g., see [4,5,9,45,63] and references therein), an elementary understanding of tensor network states is assumed in this paper. With this in mind, the paper is organized as follows. Section I sets the stage with focus on symmetries in TNSs, which strongly builds on Weichselbaum [45]. Section II provides conventions and preliminaries required for the rest of the paper. Section III then introduces X-symbols and discusses their relevance in TNSs, followed by summary and outlook.

I. SYMMETRIES IN TENSOR NETWORK STATES

Tensor network states typically describe lattice Hamiltonians, with whom they share the same lattice structure. Whereas the Hamiltonian may be longer ranged TNSs have nearest-neighbor bonds only in order to minimize index loops. For the same reason, physical sites may also be grouped into supersites [55,64]. Each lattice site n is assigned a tensor A_n that maps the attached (variationally determined) auxiliary

state spaces $|a_x\rangle$ to the physical state space $|\varphi_\sigma\rangle_n$. In a pictorial language, the indices of a tensor are drawn as lines, also referred to as the legs of a tensor. All state spaces are indexed. E.g., the index σ above spans the local state space of a single physical site, whereas the index $x = l, r, \dots$ spans specific named bonds, such as left (l), right (r), etc. [e.g., see (1)].

A. Arrows on all legs

The physical state spaces are orthonormal from the very outset, having ${}_n\langle\varphi_{\sigma'}|\varphi_\sigma\rangle_n = \delta_{\sigma'}^\sigma\delta_n^{n'}$. For generalized tree-tensor networks [42] including matrix product states which, by definition, contain no loops along any path of bond indices, cutting any auxiliary bond separates the TNS into two disconnected blocks. As a direct consequence, all auxiliary or bond state spaces can be made orthonormal. If a given TTN is an exact symmetry eigenstate globally, then all bond indices can be fully symmetrized, i.e., assigned symmetry labels without increasing the bond dimension [43]. In the case of a TTN, these auxiliary bond state spaces can be translated into well-defined orthonormal effective quantum-many-body state spaces that represent entire blocks of the system. Each such block only contains one open bond index, starting from which it necessarily stretches all the way to the open or infinite outer boundary of the physical system, and hence also of the TTN considered.

For a TTN, many-body state spaces are typically generated iteratively by adding one site after another to a block. This way, by construction, any index or leg describes an orthonormal many-body state space that either enters a tensor as part of a tensor product space or leaves a tensor with the interpretation of an combined effective state space. So while several lines may enter a tensor there is always at most one line that can leave a tensor. For finite TTN simulations then, auxiliary indices all flow towards the orthogonality center (OC) [65]; this is the only tensor in an entire TTN that may have no outgoing index. It combines the orthonormal state spaces of various blocks into a normalized global wave function. If the global state is a singlet, this singleton dimension may be skipped, in which case the OC has no outgoing index. If multiple global states are targeted, e.g., if the global symmetry multiplet is not a scalar, then the OC also needs to carry along an outgoing leg, namely, the index that resolves the global state or multiplet.

In a pictorial description, this naturally suggests that each index (leg or line) in a TN is given an arrow. Correspondingly, in mathematical notation, an outgoing (incoming) index to a tensor can be written as a lowered (raised) index, which is equivalent to covariant (contravariant) index notation, respectively. A contracted, i.e., summed over, index then necessarily is outgoing from one tensor and ingoing into another, consistent with Einstein summation convention that an index is summed over if it appears twice, namely, as a raised and lowered index.

By having adopted the convention that the physical state space of site n is denoted by $|\varphi_\sigma\rangle_n$, i.e., with lowered indices, this implies that the index (leg) in a tensor where it enters must be a raised index. Thus combining, for example, the state space of site n with an effective bond state space a_l (“left”),

the fused state space a_r (“right”) is given by

$$|a_r\rangle = \left(\sum_{l\sigma} \right) A_r^{l\sigma} |a_l\rangle \otimes |\varphi_\sigma\rangle \equiv \begin{array}{c} \boxed{A_r^{l\sigma}} \\ \leftarrow |a_l\rangle \quad \leftarrow |\varphi_\sigma\rangle \end{array} \rightarrow a_r \quad (1)$$

where the summation over double indices is implicit, hence the bracket around the sum. TTNs guarantee that orthonormal state spaces, and hence also symmetries, are well defined throughout. This is crucially important for efficient algorithms [4,12], since these orthogonal state spaces ensure an orthonormal environment and hence optimal conditions [66] when truncating bond dimensions to the most entangled and thus most relevant quantum-many-body states constituting a given global state. In this sense, orthogonality and hence arrows are crucially important in any TTN.

For a general TNS in the presence of loops, however, it is no longer possible to associate all indices with well-defined strictly orthonormal state spaces in the unique strong sense this can be achieved with TTN [43]. Instead, one can adopt the much less rigorous PEPS-like interpretation that bonds host actual auxiliary state spaces where one simply imposes by fiat that these are orthonormal and adhering to the symmetries of the overall Hamiltonian. When dynamically truncating, and thus adapting auxiliary state spaces, it is no longer possible then to have a perfectly orthonormal environment. Instead, one needs to resort to optimal conditioning [66]. If one decides to exploit global symmetries, in any case, this necessitates that symmetries must be enforced at every step in a calculation. Here, in particular, symmetries need to be enforced locally with each tensor. As a consequence, all lines in a TNS need to carry arrows. For TNSs with loops, such as a two-dimensional PEPS, certain individual tensors necessarily also need to carry multiple outgoing legs then. Independent of whether or not global symmetries are enforced, however, when describing specific algorithms in tensor networks, arrows on all legs, or equivalently raised and lowered indices, generally represent an important natural concept for the underlying tensor algebra.

B. General tensor decomposition

In the presence of (non-Abelian) symmetries, the tensor coefficients A (without hat) of any tensor operator \hat{A} (with hat) can be decomposed into a sum over tensor products of reduced matrix elements $\|A\|$ times generalized Clebsch-Gordan coefficients C [36,45]:

$$A = \bigoplus_q \|A\|_q \otimes C_q. \quad (2)$$

The tensor A can have arbitrary rank r which is defined here as the number of legs (indices) attached. The sum over q indicates the sum over symmetry sectors. For each independent elementary (i.e., simple) symmetry considered, another CGT can be split off [45], collectively written as C_q above. To be precise, in this paper, a CGT always refers to a specific elementary (semi-) simple symmetry such as $SU(N)$, $Sp(2n)$, or $SO(n)$ (see also Appendix A), and can always also be chosen real. Yet in the same way that an algorithm needs to deal with RMTs and CGTs separately, also CGTs for different symmetries can be dealt with completely separately and in parallel. Hence, while for the sake of simplicity of the

argument only a single elementary non-Abelian symmetry is considered, all arguments in this paper straightforwardly also generalize to the presence of multiple symmetries.

The decomposition in Eq. (2) holds for general symmetries, non-Abelian and Abelian symmetries alike. For Abelian symmetries, however, all multiplets contain only a single state. Hence the tensorial structure of the CGTs reduces to plain number, namely 1 (permissible from a symmetry point of view), or 0 (not permissible). For non-Abelian symmetries, on the other hand, one still also needs to systematically account for outer multiplicity as explained in detail below. The overall tensorial structure such as rank or directions of incoming and outgoing indices, i.e., the arrow configuration, is exactly inherited by all terms, RMTs as well as CGTs. Therefore e.g., in Eq. (2) the RMT $\|A\|$ and the CGT C have the same label structure, in that *both* have subscript q which is then summed over via the direct sum.

In order to avoid excessive proliferation of indices, for the sake of readability, simplified shortcut notations are adopted as explained in the following. In particular, this concerns the at times somewhat loose distinction between raised or lowered indices if the distinction is not explicitly important in a specific context. In Eq. (2), for example, the combined symmetry labels for all legs have simply been written as subscript q , even though legs may have mixed raised and lowered indices. The symmetry labels (simply also referred to as “ q labels”) for all $r = l + l'$ legs of a given CGT of rank r are given by the set

$$q \equiv (q^1, \dots, q^l, q_{l'}, \dots, q_{l'}) \equiv (q_1, \dots, q_l | q_{l'}, \dots, q_{l'}). \quad (3)$$

Here each q^i (or $q_{i'}$) is a tuple of multiplet labels of fixed length that specifies the symmetry sector on a given leg i for all symmetries considered. While the order of legs within the group of incoming (or outgoing) legs is important, incoming and outgoing indices can be arbitrarily interspersed. This means that the position of raised indices relative to lowered indices is irrelevant. Therefore, e.g., all incoming indices can be listed first, as shown in Eq. (3). While the first decomposition of q in Eq. (3) explicitly specifies raised and lowered indices, the equivalent last decomposition splits the group of incoming indices (l legs) via the bar “|” from the group of outgoing indices (l' legs). It uses all lower indices, if relevant, otherwise. This is useful, for example, when discussing standard rank-3 CGTs $q \equiv (q_1 q_2 | q_3)$ which fuse q_1 and q_2 into the combined total q_3 , also referred to as CG3 in this paper. E.g., consider $(q\bar{q}|0)$, which fuses multiplet(s) q with their dual(s) into the scalar representation, simply denoted as 0. This has no more raised or lowered indices whatsoever, which thus requires the last notation in Eq. (3).


Depending on the context, the combined label q as a whole as in Eq. (2) may also be written as a superscript, instead, with no specific meaning of the location, unless explicitly stated otherwise. Only when q appears paired up as a raised and lowered index, e.g., $C_q D^q$, the legs q are considered traced over with arrows reverted in D relative to C (see conjugate tensors below). Decomposed q labels with the same CGT always adhere to the interpretation of raised and lowered indices. The CGTs corresponding to the CG3s above then, for example, correspond to $C_q \equiv C_{q_3}^{q_1 q_2}$ or $C_0^{q\bar{q}}$. The same index convention as in Eq. (3) also holds for the RMTs, i.e., using

the notation $n \equiv (n^1 \dots n^1 n_{l'}, \dots, n_{l'})$ for the index into the tensor of reduced matrix elements. Here n_i indexes individual symmetry multiplets in a given symmetry sector q_i in the state space decomposition of leg i as in $|qn; q_z\rangle_i$ with q_z the internal states of a single multiplet in q [45]. By convention, in the notation of q in Eq. (3) or also n , incoming indices are listed first. Note, however, that this is reverse to the bra-ket notation, e.g., as in Eq. (1), which places incoming states as kets, and hence needs to be read right to left.

C. Outer multiplicity

In the presence of non-Abelian symmetries, tensors as in Eq. (2) are typically faced with the problem of OM in their CGTs. This means that for exactly the same symmetry labels q multiple ($m_q > 1$) orthogonal CGTs can arise [see Eq. (7) later]. Hence the CGTs acquire an additional multiplicity index $\mu = 1, \dots, m_q$, i.e., $C_q \rightarrow C_{q,\mu}$. If a given CGT has no outer multiplicity, then $m_q = 1$, which represents a singleton dimension and thus may be safely skipped. If a given combination of symmetry labels is not permitted from a symmetry point of view, then $m_q = 0$. CGTs that are always free from outer multiplicity are CGTs of rank $r = 0$ (scalars) or $r = 2$, as these are just proportional to the identity matrix. Moreover, for rank 3, CGTs that contain the defining irep or its dual at least on one of their legs, i.e., *primitive* CGTs, are also always OM free. Furthermore, CG3s ($q_1 q_2 | q_3$) where $q_3 = q_1 + q_2$ carries the maximum weight states of q_1 and q_2 combined [45] are also OM free since the maximum weight state is unique.

While there is no OM in the well-known SU(2) at the level of CG3s, OM routinely also occurs in SU(2) for CGTs of rank $r > 3$ (see Appendix A). A simple example for the emergence of OM in SU(2) for a rank-4 CGT is shown in (4a):



which represents the contraction of two multiplicity-free CGTs C and C' of rank 3 over the shared intermediate multiplet $q_i = S_i$. However, typically there will exist multiple choices for the intermediate spin multiplet S_i . Therefore when “zooming out,” i.e., contracting the intermediate index, the resulting rank-4 tensor (brown circle) acquires outer multiplicity $\mu = 1, \dots, m_q$. Here m_q is the number of permitted intermediate S_i for fixed open outer multiplets $S_1^{(i)}$ and $S_2^{(i)}$. In the present case, the presence of several permitted internal multiplets S_i directly translates into OM indexed by μ . Recoupling of the internal (contracted) structure results in an orthogonal rotation in OM space. For overall consistency, the OM index [green line in (4a)] also carries an arrow. Since this comes “out” of a given CGT decomposition for $C * C'$, arbitrarily but fixed, the index μ is chosen as an outgoing index in (4a) [and of C in (4c)]. The above point of view can be used to systematically determine the level of OM for arbitrary symmetries and CGTs, as discussed later in Sec. III E.

The presence of OM introduces an additional vector space for a given CGT, referred to as outer multiplicity space. The CGTs in Eq. (2) then become a linear superposition in OM space, described by the normalized coefficients $w^{\mu\mu'}$:

$$A = \bigoplus_q [\|A\|_{q\mu} \otimes (w^{\mu\mu'} C_{q\mu'})], \quad (5)$$

using $\|w^\mu\| = 1$, with implicit regular summation over the multiplicity indices μ and μ' (Einstein summation convention), in contrast to the *direct* sum over symmetry sectors q . Note, in particular, that the (block) summation over q adds to the overall dimension of the tensor A , whereas the multiplicity indices no longer do [OM leads to additional multiplets, which must have already been allocated at the level of the RMTs $\|A\|_q$ (e.g., see Fig. 2 in [45]). All OM related indices are denoted in green color consistent with the graphical notation as in Eq. (4). Effectively, the role of the coefficient matrix w is such that it “ties” together [67] the otherwise plain tensor product between RMTs $\|A\|_q$ and CGTs C_q , as graphically depicted in Eq. (4c).

For a given symmetry sector q in Eq. (5), $\|A\|_{q\mu}$ are the reduced matrix elements that come with the CGT component $C_q^\mu \equiv \sum_{\mu'} w^{\mu\mu'} C_{q\mu'}$. In other words, for different orthogonal outer multiplicity components, there can be a completely different set of reduced matrix elements. For this reason, the matrix w in Eq. (5) can have at most as many rows (indexed by μ) as there are columns (indexed by $\mu' \leq m_q$). With all CGTs real, the matrix w can always be chosen also real and such that $ww^T = 1$. Fewer rows than m_q are allowed in w (the tensor A then simply does not span the full OM space), such that $w^T w \neq 1$. However, if there had been more than m_q rows, QR or singular value (SV) decomposition [68] can be employed to decompose $w = z\tilde{w}$. Contracting z onto the RMT $\|A\|_{q\mu}$ in Eq. (5), w can be replaced by the orthogonal \tilde{w} now with at most m_q rows.

The CGTs C_q can be computed once and for all and then stored (tabulated) in a central, dynamically generated database. For the storage of any tensor A then only a *reference* to the CGTs is required in terms of metadata. Associating w with C_q , i.e., explicitly using $C_q^\mu \equiv \sum_{\mu'} w^{\mu\mu'} C_{q\mu'}$ and storing this with $\|A\|_{q\mu}$ separately for each μ , this results in a nonunique listing of symmetry sectors q in the listing of nonzero blocks of A in Eq. (5). This was the procedure in [45]. However, instead, one can explicitly add an OM index μ onto the RMT to the already existing indices n [67], and contract w onto the RMT, instead, $\|A\|_q^{\mu'} \equiv \sum_{\mu} \|A\|_{q\mu} w^{\mu\mu'}$. Then the tensor A can be written as a *unique* listing over symmetry sectors q (see Appendix B). While the matrix w may be eliminated this way, for practical purposes, it may be explicitly kept with the tensor decomposition in Eq. (5) nevertheless. For example, this then allows one to simply absorb operations in OM space such as orthogonal transformations into w , while the likely much larger CGTs and RMTs can remain unaltered.


II. CONVENTIONS AND PRELIMINARIES

A. Conjugate tensors

Conjugate (or daggered) tensors arise whenever computing expectation values $\langle \psi | \dots | \psi \rangle \sim \text{tr}(A^\dagger \dots A)$ or whenever expressing terms in a Hamiltonian that arise out of a scalar

product, such as spin-spin interactions, $\hat{S}_i \cdot \hat{S}_j \sim S_i^\dagger \cdot S_j$ in between some sites i and j . In general, the *conjugate tensor* A^\dagger of any given tensor A can be defined by the following prescription: take the complex conjugate of all of its matrix elements, i.e., A^* (relevant for RMTs only, since CGTs are real), then revert all arrows, and take a mirror image of the tensor when depicted graphically [45]. The latter is simply the generalization of the transposition part in A^\dagger , e.g., when an operator sandwiched in between bra-ket states acts to the left (bra) rather than to the right (ket). In terms of index order convention, taking a mirror image implies with respect to the original tensor that if indices are read clockwise in A in a pictorial description starting from some arbitrary but fixed leg, they need to be read counter clockwise in A^\dagger . As pointed out with Eq. (3), though, incoming and outgoing indices can be dealt with separately.

Taking the mirror image, actually, is a rather common graphical procedure in TNSs, even if one typically does not further dwell on this. For example, when computing expectation values $\langle \psi | \dots | \psi \rangle$ for a given MPS, this involves the contraction of an A -tensor as in (1) with itself, as shown in (6a):



Here A^\dagger (empty square at the bottom) is the mirror image of A (filled square at the top) with the arrows on all lines reverted, while also taking complex conjugated matrix elements. The mirror plane is indicated by the horizontal yellow marker line. Similarly, the above procedure of conjugating tensors also shows up systematically when dealing with fermionic PEPS [19] where in a pictorial description $\langle \psi | \dots | \psi \rangle$, the bra state is a full mirror image of the ket state.

Another example concerns the representation of an operator \hat{S} that acts on a local state space of a site as depicted in (6b), e.g., which may be inserted into the vertical line labeled by σ in (6a). In general, the operator $\hat{S} \equiv \{\hat{S}^\alpha\}$ consists of a set of operators indexed by α which are irreducible with respect to the symmetries under consideration. This makes it an irreducible operator (irop), which thus naturally acquires the third index α [45]. Nonscalar irops are also non-Hermitian, examples being fermionic hopping, or spin raising and lowering operators. Here (6c) summarizes what one means by taking the Hermitian conjugate of an operator: matrix elements are complex conjugated ($S \rightarrow S^*$) and top and bottom index are exchanged, resulting in tangled yet noncrossing lines. Here also all arrows are reverted, and a mirror image is taken with respect to the vertical yellow marker line which flips the irop index α to the left. Untangling the lines by rotating the tensor clockwise by 180° leaves the tensor unaltered, otherwise, and also does not introduce any crossing of lines [19]. This leads to the final S^\dagger in (6c) which looks much the same as (6b), but now with reverted index α . This is required, e.g., when constructing scalar products of operators as in $\hat{S}_i^\dagger \cdot \hat{S}_j$ where the dot product simply translates into a contracted directed line indexed by α .

This paper also adopts the graphical convention that tensor conjugation switches from a filled tensor (box, circle, etc.) to an empty outlined object, or vice versa, as already seen in Eq. (6). Similarly, in (4a) the CGT C with two incoming legs was depicted by a filled circle, in contrast to the conjugate tensor C' with two outgoing legs (empty circle). In a mathematical notation, finally, tensor conjugation can be denoted by raising or lowering its indices, e.g., $(C_{q\mu})^\dagger \equiv C^{q\mu}$.

B. CGT normalization convention in OM space

When considering a CGT C_q for an arbitrary but fixed set of symmetry labels q , a convenient normalization convention in OM space adopted throughout is

$$\text{Tr}(C_{q\mu} C_{q\mu'}^\dagger) \equiv \text{Tr}(C_{q\mu} C^{q\mu'}) = \delta_{\mu'}^{\mu}, \quad (7)$$

where $\text{Tr}(C_{q\mu} C_{q\mu'}^\dagger)$ stands for the full contraction (tensor trace) of C_q with the conjugate of itself, only keeping OM indices open and having fixed q . By pairing with its conjugate tensor, this also reverts the arrow of the OM index, hence raising the index μ' in Eq. (7). The CGT normalization in Eq. (7) simply generates an orthonormal basis in the OM “vector space,” which is convenient when performing explicit decompositions or projections in OM space. Equation (7) fixes the CGT C_q up to an orthogonal transformation in OM space (see matrix w in Sec. IC). As the CGT will be generated dynamically depending on the calculation, this makes it history dependent, adding OM components as they occur [45]. This fixes C_q up to an overall sign convention. Including μ as the last index in C_q and assuming column-major index order, the standard sign convention of Clebsch-Gordan coefficients is followed, namely, that the first nonzero matrix element of the full CGT C_q is chosen positive. Therefore the component for $\mu = 1$ fixes the sign of the entire C_q for $\mu \leq m_q$.

For rank-3 CGTs, the normalization in Eq. (7) is closely related to the normalization of Wigner- $3j$ symbols [69]. As such, it is different from the normalization of the standard Clebsch-Gordan coefficients, e.g., for $SU(2)$ ($S_1 S_2 | S_3$) that fuse (S_1, S_2) into an orthonormal basis in S_3 (and not in OM). The latter would result in a norm in Eq. (7) that is equal to $|S_3|$, i.e., the dimension of multiplet S_3 . Consequently, this would make the normalization dependent on the direction of arrows in a given CGT. This is rather inconvenient on general grounds, and specifically so in the TNS context, since the presence or not of OM *does not depend on the direction of arrows or on the specific order* of symmetry labels. This simply follows from the discussion below that raising or lowering of an index on a CGT is equivalent to applying a unitary matrix on that index while also switching to the dual representation on that leg (see Sec. IID). Conversely, a permutation of indices on a CGT can only induce orthogonal rotations in OM space (see Sec. IIIC) which clearly also leaves the OM dimension invariant.

Therefore based on the normalization in Eq. (7), the CG3s $q = (q_1, q_2 | q_3) \equiv (q^1 q^2 q^3)$ have the same normalization for any permutation of the symmetry labels together with arbitrary raising or lowering of indices. Raising and lowering of labels, however, changes the interpretation of which multiplets are fused together. By raising all indices in a CGT, for $q = (q^1 q^2 q^3)$, this is equivalent to fusing all three multiplets

of a CG3 into a total scalar representation denoted by $q_{\text{tot}} = 0$, i.e., $q = (q_1, q_2, q_3|0)$, and subsequently skipping the trailing singleton dimension. The resulting rank-3 CGT has all-incoming legs, i.e., $q = (q^1 q^2 q^3)$, which is exactly equivalent to the Wigner-3j (here “3q”) symbol up to an overall sign.

C. Fully contracted CGT networks

When a tensor network is fully contracted, all indices are paired up and summed over. This holds for both the RMTs as well as the CGTs. Here, for the sake of the argument, however, the focus is on an isolated fully contracted TNS solely comprised of CGTs. The OM for each participating CGT is assumed fixed to some arbitrary but fixed linear superpositions in OM space [e.g., see Eq. (5)], such that there is no open index left in the TNS, and the full contraction yields some number x . By construction, this number can be nonzero only if the CGT network is permitted from a symmetry point of view. Now if one opens up a single leg in this otherwise fully contraction CGT network, say with symmetry label q_i , the result is proportional to a rank-2 CGT $C_{q_i}^{q_i}$, i.e., with one incoming and one outgoing index. The only such CGT that exists is the identity matrix up to normalization. Its graphical representation is a single directed line. Therefore it follows that a CGT network that is fully contracted up to a single opened up index i necessarily is proportional to the identity matrix $1^{|q_i|}$ of dimension $|q_i|$, i.e., the size of multiplet q_i .

As a specific example, consider the tensor network in Eq. (7) of just one CGT fully contracted with itself for arbitrary but fixed μ and μ' . Then, opening up one bond index i with symmetry label q_i , one obtains

$$\text{Tr}_{|q_i|}(C_{q_i\mu} C_{q_i\mu'}^\dagger) = \frac{\delta_{\mu\mu'}}{|q_i|} 1^{|q_i|}. \tag{8}$$

The normalization is determined by the requirement that the final trace over i , when performed, results back in Eq. (7).

D. Reverting arrows and 1j-symbols

Reverting the arrow on a given bond in a TNS changes its interpretation, as well as the interpretation of the associated tensors. As this will be useful also later in the CGT context, consider first the elementary process where the orthogonality center in an MPS is iteratively propagated from site $n \rightarrow n+1$ with associated tensors \tilde{A}_n and A_{n+1} [e.g., see Eq. (1)], contracted on their shared auxiliary bond [4]:

$$\tilde{A}_n A_{n+1} = (A_n \underbrace{\tilde{X}_n}_{\equiv \tilde{A}_{n+1}}) A_{n+1} \equiv \underbrace{\tilde{A}_n (\tilde{X}_n^{-1} \tilde{X}_n)}_{\equiv A_n} \underbrace{A_{n+1}}_{\equiv \tilde{A}_{n+1}}, \tag{9}$$


where the tilde indicates the tensor that carries the OC. Before the iteration step, the OC is located on site n , and the bond in between sites $(n, n + 1)$ describes an orthonormal effective many-body state space for the entire right block of sites $n' > n$. After the iteration step, conversely, the arrow on the bond changed its direction, and now describes the orthonormal effective many-body state space for the entire left block of sites $n' \leq n$. On a procedural level, one starts with the tensor \tilde{A}_n in Eq. (9) that carries the OC and performs QR or SV decomposition on it [4,45]. This yields $\tilde{A}_n = A_n \tilde{X}_n$, where A_n is a new isometry, and the OC is now shifted onto the tensor

\tilde{X}_n located on the bond in between sites $(n, n + 1)$. When contracted onto A_{n+1} , this makes the former isometry A_{n+1} the new OC \tilde{A}_{n+1} . Formally, in the last equality of Eq. (9), the direction of a bond is flipped by using the Gauge freedom inherent to TNS [4]. This consists of inserting the identity $1 = \tilde{X}_n^{-1} \tilde{X}_n$ and then associating the tensors \tilde{X}_n^{-1} and \tilde{X}_n with the left and right A-tensors, respectively. However, as seen from the left of Eq. (9), it is not necessary to actually compute the inverse of \tilde{X}_n here as this may be ill conditioned. Overall, one can exactly flip the orientation of a leg in a TNS without changing the global physical state. One only changes the local perspective and interpretation by splitting off the tensor \tilde{X}_n from \tilde{A}_n and contracting it onto the neighboring tensor A_{n+1} .

Now the discussion of flipping the direction of an arrow on a leg of a given elementary CGT C_q with fixed symmetry labels q follows much of the same spirit. In contrast to RMTs, however, the procedure is naturally much more constrained for CGTs. For example, a CGT is already always in a canonical form (see higher-order SV decomposition [21]). Even more, with Eq. (8), it has a constant singular value spectrum with respect to any of its bonds, since by symmetry all states in a given multiplet are necessarily equally important. As will be shown below then, the matrix \tilde{X} in Eq. (9) to flip the direction of an arrow must be unitary for CGTs, while in the same process the irep label q also needs to be switched to its dual \bar{q} .

In order to show this, it is sufficient to narrow the discussion further down to a single directed line with symmetry label q , assuming a single elementary non-Abelian symmetry for the sake of the argument without restricting the case. This line may represent, e.g., an auxiliary bond in a TNS, and may be associated with the CGT $C_q^q \propto 1^{|q|}$. In general, now irep q has a unique dual representation \bar{q} . This dual shares the same multiplet dimension, i.e., $|q| = |\bar{q}|$, and it is the only irep that, when fused together with q , permits the scalar representation as an outcome, i.e., having $(q\bar{q}|0)$ exactly once. The corresponding CGT $C_0^{q\bar{q}}$ therefore never has OM, but is unique up to an overall sign convention. By definition, the scalar representation, always denoted by the label 0 here, is fully symmetric under symmetry operations. Hence its multiplet only consists of a single state, i.e., $|0\rangle = 1$. For example, SU(2) is self-dual, i.e., $\bar{q} = q$ for all q , such that the product space of spin $q = S$ with itself (and itself only) always also yields a singlet with $S_{\text{tot}} = 0$, having $(SS|0)$. The same argument can be further extended also to U(1) Abelian symmetries such as charge (N) or spin (S_z). There the dual is simply given by $\bar{q} = -q$, since $q + (-q) = 0$.

By making use of the dual representation, this allows one to define the unitary matrix $U_{(0)}^{q\bar{q}} \equiv \sqrt{|q|} C_0^{q\bar{q}}$ of dimension $|q|$ for any representation q as depicted in (10a):



The singleton dimension in the scalar representation, indicated by the dashed line to the right, will be frequently skipped in U to emphasize that U is a matrix, indeed. Hence the subscript 0 has been written in brackets. Now $\text{Tr}(U^\dagger U)$ [Eq. (10b)] is proportional to a rank-3 CGT fully contracted with its conjugate. So when opening up a single index, here the one for q , this must be proportional to the identity matrix [see Eq. (8)], and with the normalization as chosen in Eq. (10a) $U^{q\bar{q}}$ must therefore be unitary (strictly speaking, orthogonal, since CGTs are always real).

This now permits us to insert a simple identity $U^\dagger U = 1$ on a given leg in a TNS, as depicted in Eq. (10c), which provides a simple recipe to revert the arrow on any contracted line with symmetry label q that links CGTs, much in the spirit of Eq. (9) earlier. By simply contracting U and U^T onto the two neighboring tensors linked via the contraction, as indicated by the broad arrows to the left and to the right in light color in (10d), the arrow on the black line in the center now points in the opposite direction as compared to (10c). At the same time, the symmetry label on the bond flipped to its dual, i.e., $q \rightarrow \bar{q}$. When not mentioned explicitly, the latter will be implicitly assumed whenever individual indices are lowered or raised. Since the unitary matrix $U^{q\bar{q}}$ effectively only refers to a single symmetry label q , the unitary $U^{q\bar{q}}$ is referred to as a $1j$ -symbol with reference to the literature on $SU(2)$ [70,71].

One needs to be careful, however, when skipping the singleton dimension [dashed line in (10c)], which is reduced to the little residual “stems” to the top in (10d): these stems are important to keep track of the order of $U^{q\bar{q}}$ versus $U^{\bar{q}q}$, e.g., for self-dual q . They must point in the same direction (here up), to ensure that the conjugate of precisely the *same* object $U^{q\bar{q}}$ is inserted together with $U^{\bar{q}q}$ in order to guarantee an identity matrix overall. Specifically, the dashed line in (10c) must not cross the solid horizontal line. Otherwise sign errors can arise, since $[U^{q\bar{q}}]^T = \pm U^{\bar{q}q}$ where the sign depends on q . For example, for $SU(2)$ a sign arises for all half-integer spins.

In summary, $1j$ -symbols can be utilized to revert arrows on lines in a TNS or, equivalently, to raise or lower indices in mathematical notation (in this sense, the $1j$ -symbol acts like a metric tensor within the tensor algebra of a given multiplet). In principle, therefore it also suffices to tabulate the CGTs with all-upper indices only (all incoming), since indices can be simply lowered by applying $1j$'s.

1. Generation of $1j$ -symbols

A $1j$ -symbol can be computed, obviously, via a full irep decomposition of $(q\bar{q}|*)$. Starting from maximum weight states, however, the largest ireps are always generated first, with the $1j$ -symbol the very last CGT to be generated. This is not practical for very large multiplets, bearing in mind that $|q| = |\bar{q}|$. Also, in a TNS one typically fuses a given large effective state space (bond index) with new local state spaces of small dimension. That is, large (effective) multiplets get fused routinely with (much) smaller ones. But for the most part one can avoid fusing two large multiplets. This is specifically important for large symmetries such as $SU(N \gtrsim 4)$ (see Appendix A).

Therefore an alternative route to computing $1j$ -symbols is desirable. Note that $1j$ -symbols are only square matrices of

dimension $|q|$, which is in stark contrast to fully decompose a $|q|^2$ dimensional vector space into irreducible representations. Moreover, $1j$ -symbols are (close to) antidiagonal, i.e., very sparse like CGTs in general. In the absence of inner multiplicity (IM) [45] such as for $SU(2)$, they are strictly antidiagonal with alternating entries ± 1 .

The nontrivial part for $1j$ -symbols arises from the antidiagonal block structure in the presence of IM, which requires consistent conventions on how to decompose IM spaces [45]. The $1j$ -symbol derives from a CG3. When the underlying symmetry already permits OM for CG3s, this also implies the presence of IM, and hence block structure in $1j$ -symbols. The $1j$ -symbol itself, however, is unique otherwise up to a global sign convention which is simply inherited here from the overall sign convention on CGTs, as discussed with Eq. (7).

The approach taken in QSPACE (see Appendix B) then to compute $1j$ -symbols is based on the fact that the scalar multiplet $q = 0$ is destroyed by every one of the $\alpha = 1, \dots, \tau_{\text{sym}}$ generalized raising and lowering operators $S_\alpha^{(\dagger)}$, with τ_{sym} the rank of a given symmetry (see Appendix A; note that τ_{sym} needs to be differentiated from the rank r of a tensor which is just the number of its legs). Having explicit access to the sparse generators $S_\alpha^{(\dagger)}$ in the representation q of some given symmetry of rank τ_{sym} , one can resort to a variational Krylov based minimization, and compute the ground state of the sparse pseudo-Hamiltonian (cost function):

$$\mathcal{H}_q^{1j} = \sum_{\alpha \leq \tau_{\text{sym}}} (S_\alpha S_\alpha^\dagger + S_\alpha^\dagger S_\alpha)_q. \quad (11)$$

where $S_\alpha^{(\text{tot})} \equiv S_\alpha^q \otimes 1^{\bar{q}} + 1^q \otimes S_\alpha^{\bar{q}}$ are the total generators in the combined state space. This is in general a well-conditioned problem with a unique ground state (the $1j$ -symbol) at “energy” zero and with a “gap” of order 1. By construction, this ground state must be simultaneously maximum *and* minimum weight state, hence representing a scalar multiplet. The Krylov based minimization then allows one to directly converge the $1j$ -symbols via iterative means down to one’s numerical floating-point precision. The sparse nature of the the $1j$ -symbols strongly limits the variational parameter space, and hence leads to fast convergence.

2. $1j$ -symbols via contractions


A useful application of $1j$ -symbols arises when computing a ground state of a system which itself is in a global singlet symmetry sector, i.e., the scalar representation $q_{\text{tot}} = 0$. By skipping this global singleton dimension, this requires access to simple $1j$ -symbols during setup. However, when sweeping through the TNS, the OC repeatedly gets located on an auxiliary bond deep inside the TNS where a wide range of multiplets can be explored. Here the shifting of the OC can also be achieved by contraction, e.g., by projecting onto identity A -tensors (which correspond to simple A -tensors as in Eq. (1) yet initially without truncation [45]). The tensor \tilde{X} that carries the OC onto a bond then has rank 2 with all indices incoming. Therefore, up to normalization, all of its CGTs necessarily must correspond to $1j$ -symbols. In this sense, $1j$ -symbols can also be generated via contractions.

E. Determination of OM dimension

Outer multiplicity of a given CGT is independent of the direction of its legs. This is apparent from the above explicit construction of reverting arrows which solely corresponds to applying a specific unitary on a given leg. Therefore OM is an intrinsic quantity of a CGT. Consequently, the OM index in a pictorial description needs to be attached to the tensor itself (and not to any of its legs), as already shown with the tensor C in Eq. (4c).

The OM dimension for rank-3 CGTs can be determined via standard fusion rules of a pair of irreducible representations. If such a tensor product arises from building a quantum many-body state starting from the vacuum state and iterative fusion of local state spaces, they need to be computed in full via a standard decomposition of a pair of ireps [45]. These CG3s also build the elementary basis and starting point for subsequent contractions of CGTs.

Now if one encounters a CGT of rank $r > 3$, its full OM can be determined iteratively in a constructive way from smaller rank tensors assuming that their OM is known. For example, for a rank-4 CGT, the example in (4a) for $SU(2)$ can be generalized as follows: since arrow directions can be altered at will without affecting the OM dimension (bearing in mind to also switch to dual representations), one can build the sequential MPS-like structure in (11a),



by taking the fused multiplet q_i out of a CGT C_1 into a tensor product space with irep q_3 , and then picking the desired irep q_4 . For fixed q 's then, given that there is no loop in (12a), the combined OM $\mu_{\text{tot}} = 1, \dots, m_{\text{tot}}$ of the CGT described by the brown box is simply the product of the multiplicities of C_1 and C_2 . Bearing in mind that the intermediate contracted multiplet q_i can vary, the total OM of the CGT with $q = (q_1 q_2 q_3 q_4)$ is given by

$$m_{\text{tot}} = \sum_{q_i} m_1(q_i) \times m_2(q_i). \tag{13}$$


This can be shown by building an OM basis for the CGTs derived from (12a) for all $\mu_1 \leq m_1, \mu_2 \leq m_2$, and for all permitted q_i . However, these are already all orthogonal to each other, as seen from computing their overlap as in (12b). Starting with the orthogonality of CG3s, the contraction in the left blue box in (12b) is proportional to the identity matrix, i.e., reduces to a simple line with weight $\propto \delta_{\mu_1 \mu_1'} \delta_{q_i q_i'}$. When repeated iteratively with the next (here last) pair of CGTs C_2 and C_2' , this directly leads to Eq. (13). As emphasized there, the multiplicities m_1 and m_2 clearly depend on the choice of q_i . The above procedure can be extended towards any sequence of CGTs also of higher rank, that are contracted in a linear sequence without loops. As a result this demonstrates that the full outer multiplicity space grows rapidly (exponentially) with increasing rank of a tensor.

As a general strategy then to avoid proliferation of OM spaces, this suggests (i) to reduce the rank of a tensor by fusing indices as far as possible (and practicable) in a TN algorithm. Moreover, the actual level of OM generated also depends on the specific TN calculation performed. For the largest CGTs encountered, typically a far smaller OM space is explicitly generated by contractions than theoretically possible. Hence (ii) one can refrain from insisting on building the full OM space in any circumstances encountered. Rather, one can build the OM space on demand (see Appendix B). If a new OM component is encountered via contractions, it can be added once and for all to one's database. On the downside, a buildup of the CGT database this way becomes dependent on the history of calculations. So one must be extremely careful to ensure consistency across independent calculations or threads that simultaneously access the same central database. This can be achieved by coordinating updates, e.g., via locking mechanisms, in order to avoid race conditions resulting in inconsistent histories.

III. CONTRACTIONS


A. Pairwise contractions and X-symbols

Contractions in any TN state are always tackled by elementary pairwise contractions, in practice, in complete analogy to evaluating the product of multiple matrices. Hence the elementary step for contracting a TN state is the contraction of two tensors. To be specific, consider some rank-4 tensor with arbitrary but fixed index order 1, ..., 4, as shown in (14a):




In the presence of symmetries, it is redrawn schematically in the spirit of (4a) and (4c) in (14b), depicting the decomposition of symmetry sectors into the tensor product of RMTs and CGTs. For simplicity only a single CGT is shown, while in the presence of multiple symmetries each has its own CGT. Also, with reference to Eq. (5), there is a sum over symmetry sectors while in the pictorial representation in (14b) the focus is on one arbitrary but fixed set of symmetry labels q with the matrix w that links the CGT to the corresponding RMT. To further simplify the following discussion, (14b) is redrawn in (14c) with indices and arrows removed, while bearing in mind that, of course, arrows and index order stay intact. Also the matrix $w_q^{\mu\mu'}$ can be fully merged with (i.e., contracted onto) $\|A\|_q$. Hence the matrix w is not explicitly needed for the sake of the argument here, and thus is skipped. The shading of the CGT at the top of (14c), finally, indicates that the CGT itself does not need to be explicitly stored with the tensor A itself, but that a reference to a central database suffices. This way the multiplicity index with the RMT automatically represents an open index, as suggested by the grouping in Eq. (5). Depending on the contraction, however, the resulting CGTs may need to be updated centrally, e.g., if a new OM component is encountered.

Now consider a pair of tensors, A and B , as in (14) contracted on a shared set of indices (state spaces), as shown in (15a):



This is a generalized matrix multiplication, and thus is symbolically written as $A * B$. In the presence of non-Abelian symmetries, they contain references to CGTs [Eq. (15b)]. Via the tensor product structure in Eq. (5), the tensorial structure is exactly the same for both RMTs as well as CGTs. Therefore, when performing a contraction of two tensors on a specified set of indices in a TNS, precisely the same contraction needs to be performed on the level of the RMTs, $||A|| * ||B||$, as well as on the level of CGTs, $C * D$. These are separate from each other, and hence can be dealt with completely independently. The contraction of the RMTs always needs to be performed explicitly, as this is part of the physical problem under investigation. The contraction of CGTs, however, is purely related to symmetries, and hence can be computed once and for all and tabulated. Now consider the contraction of some arbitrary but fixed pair of CGTs, $C * D$ as in (15b), shown in (16a):



The contraction of two CGTs $C * D$ necessarily yields another CGT labeled E in (16b), with its own orthonormal OM space. Assuming its OM space is complete, then the contraction of $C * D$ can be projected by inserting the identity $E^\dagger E = 1$ in (16b). When fully contracting the conjugate E^\dagger onto $C * D$, this results in what is referred to as an *X-symbol*:

$$X_\kappa^{\mu\nu} \equiv \text{Tr}\{(\text{Tr}_{i_C, i_D} [C^\mu * D^\nu]) * E_\kappa^\dagger\}, \quad (17)$$

where C is contracted on the subset of legs i_C with D on legs i_D , the result of which is fully contracted with E^\dagger , while keeping the OM indices μ, ν , and κ (green lines) open.

The X-symbol is derived from a contraction (hence “X”) of a pair of CGTs of arbitrary rank each and with fixed symmetry labels. It has all TNS-related indices fully contracted [yellow lines in (16b)]. Consequently, X-symbols represent a mapping in OM spaces with the three open indices $(\mu, \nu) \mapsto \kappa$. With E also a CGT fixed by symmetry, it can be referenced in the final object, again suggested by the shading in (16c). The overall result out of (16c) can now be inserted back into (15c).

The pictorial representation in (15c) then exemplifies the central result of this paper: in the process of a pairwise contraction of tensors, it suffices to contract the centrally stored X-symbol onto the OM indices of the corresponding pair of RMTs, thus merging (μ, ν) and decomposing (κ) the


OM spaces. If the w matrices had not been contracted onto the RMTs as in (14c), they can also be contracted here onto the X-symbol, instead. Importantly, in the present context, the summation over the OM indices μ and ν in Eq. (5) has turned into a regular contraction involving RMTs and X-symbols only. With the index κ left open now, it is E here that can be simply referenced as indicated by the shading in (15c). Nothing else remains to be done on the level of CGTs themselves. Therefore if all X-symbols are available and up to date in the database, the CGTs themselves can be completely sidestepped. The X-symbols fully take care of the symmetry related multiplicity spaces in an efficient and general manner.

As apparent from the definition in Eq. (17), each X-symbol needs to remember where it came from via metadata. This includes references to the three participating CGTs C, D , and E (which also specifies all their symmetry labels, order, and direction of legs), and what indices have been contracted [i_C and i_D in Eq. (17)]. Moreover, depending on the context, one may have to contract the conjugate of the input tensors C or D , instead [indicated then by i_C^* or i_D^* in Eq. (17)]. Hence the X-symbol also stores conjugation flags for all three CGTs. If the OM space is built successively via contractions as they occur, the X-symbol further needs to remember identifiers as to the state of CGTs such as a high-resolution time stamp of their last modification time when the X-symbol was computed. Then if any of the CGTs gets updated later along the course of a calculation, these serve as flags as to whether or not also the X-symbol needs to be updated when the same contraction is encountered again at a later stage.

Note that if C and D already have complete OM individually, this does *not at all* imply that also the $C * D$ will exhaust the OM space of the resulting CGT E [see Eq. (13)]. When (re)computing an X-symbol, if E is already present, e.g., from other earlier contractions, $C * D$ needs to be projected onto it. If the OM space of the current E was already complete, the projection can fully represent the result. If the OM space of E was incomplete, then new OM components out of $C * D$ may arise, which need to be extracted and orthonormalized via Schmidt decomposition (performed twice for numerical stability). As this extends the OM space of the CGT E , finally, it needs to be updated centrally.

B. Relation to 6j-symbols

Consider the fully contracted TN of four tensors with a total of six contracted lines in (18a):



The part left of the yellow vertical marker, for example, may represent the matrix elements of an operator S^\dagger acting on the local site incorporated by tensor A in a matrix product state [4,9,45] [see also Eq. (6)]. For the sake of argument, the symmetry sectors are arbitrary but fixed for all legs as indicated. In the presence of non-abelian symmetries then, the

contraction can be written in terms of RMTs and X-symbols as in (18b). Now the contraction in (18a) can be performed pairwise. The order of contractions is somewhat arbitrary, where the one chosen in (18b) is $[A^\dagger * (S^\dagger * A)] * Y$. Each of the pairwise contractions of CGTs can make use of X-symbols. The first contraction $S^\dagger * A$, say, makes use of the X-symbol x_1 . The resulting CGT contracted with A^\dagger makes use of the X-symbol x_2 , the result of which when contracted with Y makes use of x_3 . For the pairwise contraction of full tensors which include RMTs and CGTs, the X-symbols need to be contracted onto the RMTs as discussed with (15c). Now the sequence of X-symbols depicted in (18b), in green lines, can be isolated, thus keeping OM indices open. With (16c) then, the resulting MPS-like sequence of X-symbols in (18c) exactly corresponds to the contraction in (18a) but now purely in terms of CGTs. Moreover, with the TN in (18a) fully contracted, the CGT describing the overall result is a rank-zero CGT, which itself clearly has no outer multiplicity. Consequently, the sequence in (18c) stops with a singleton dimension indicated by a dashed line with label “0.”

In terms of CGTs then, the contraction of the TN left of the yellow marker in (18a) line results in a rank-3 CGT $(q_2|q_2q_4) \equiv (q_4q_2|q_2)^\dagger$. The result of the contraction $A^\dagger * S^\dagger * A$ therefore resembles the conjugate of a CG3. In order to determine the precise decomposition, it can be projected onto the CGT $(q_4q_2|q_2)$, marked as Y in (18a) and (18b). The TN in (18a) contains six contracted indices, and therefore six symmetry sectors $q = \{q_1, q_2, q_2', q_3, q_3', q_4\}$. They connect four CG3s, resulting in a fully contracted TN. Overall, therefore (18a) represents a “6q” symbol or, in the SU(2) context, the well-known $6j$ symbol. Contrary to the case of SU(2), however, which has no OM at the rank-3 level such that $6j$ -symbols are plain numbers, for general symmetries, each of the four participating CGTs in (18a) can carry outer multiplicity. Therefore the resulting $6q$ symbol has four open OM indices for a general non-Abelian symmetry, i.e., represents a rank-4 tensor purely in terms of OM indices.

It follows therefore from the above constructive approach that any fully contracted TN built from CGTs can be decomposed into a linear sequence of contractions based on X-symbols. In this sense, X-symbols are equally general as $3n-j$ symbols, in that any $3n-j$ symbol can be computed from them. Yet X-symbols are much more naturally suited to tensor network algorithms, in that they provide a general prescription for the very elementary operation of a pairwise contraction of two tensors of arbitrary rank each on an arbitrary subset of indices.

C. Permutations and sorted CGTs

The indices of any tensor in a TNS need to be chosen in some arbitrary but fixed order. The precise choice of order within the TNS is typically a matter of convention, but of no further concern otherwise. Permutations, if performed, correspond to resorting of matrix elements according to the new index order. The same also holds for CGTs. Therefore, by convention, it suffices to only tabulate *sorted* CGTs which have their symmetry labels (q labels) sorted within the set of incoming or outgoing indices, e.g., in a lexicographical style. Any reference to a specific CGT then includes a reference

to a sorted CGT together with a permutation p describing the actual index order. The adopted sign convention is that the stored sorted CGT starts with a positive coefficient. The permuted references adhere to this original sorted tensor up to the permutation only, i.e., there is no further sign adaptation after permutation. This approach of making use of sorted CGTs, i.e., with sorted q labels, allows one to significantly reduce redundancy of the CGTs that need to be stored or, subsequently, also contracted.

The above is a well-defined prescription for CGTs that have different q labels on all of their legs. However, subtleties arise if symmetry labels are degenerate, i.e., when precisely the same representation occurs on more than one leg within the group of either incoming or outgoing indices. Then the prescription back and forth to sorted q labels is not unique, i.e., there can be a different permutation p' that also leads to the same sorted q labels as the default permutation p used in one’s algorithm. In this case the permutation p that transforms p into p' only operates within degenerate q -label subspaces. It generates a nontrivial orthogonal rotation U_p in OM space that needs to be included in the permutation. This U_p can be explicitly computed by fully contracting C_q with the conjugate of itself permuted by p , the result of which yields an X-symbol, say X_p . As such, the matrix U_p has all the properties of an X-symbol, and hence can also be stored as such. Since X_p represents the full contraction of two CGTs, the resulting rank-zero tensor cannot have OM. Therefore similar to the discussion with (18c), X_p has a singleton trailing dimension, which can be skipped. This way, the X-symbol X_p reduces to the matrix U_p above. It can be absorbed into the matrix w as in (14b) if w is kept track of, or directly contracted onto the RMT for the situation in (14c).

D. Generating standard rank-3 CGTs

An elementary starting point for TN calculations are the standard CG3s that fuse the typically small representations of a physical site. As a TN grows, however, the multiplets on the auxiliary bond indices can quickly explore a far larger set of representations. This implies for contractions that often one does not require the full tensor-product decomposition for every single rank-3 CGT encountered, but only very specific combinations. Since the full tensor-product decomposition of the fusion of two large multiplets can become prohibitively expensive for large symmetries (see Appendix A), the question arises as to what extent specific (standard) rank-3 CGTs can be obtained by other means.


Being interested in some specific CGT $(q_1q_2|q_3)$, in a TNS setting, the typical situation is such that at least one of the legs belongs to either the local state space of a physical site or to the ireps according to which irreducible operators transform. All of these are typically multiplets of small dimension. Therefore in TN simulations, in practice, one of the ireps q_i ($i = 1, 2, 3$) in a rank-3 CGT can be considered small. Frequently, it may even refer to the defining representation or its dual, i.e., a primitive CGT [72,73]. One simple strategy to compute such rank-3 CGTs with their full OM is to exploit the freedom that arrows can be reverted at will (while also switching to dual ireps). Hence, for example, the multiplets q_i can be sorted according to their dimension $|q_i|$. Then the full

tensor-product decomposition may be performed by taking the tensor product of the *smallest* two irreps, and subsequently reverting arrows as needed making use of $1j$ -symbols.

Rank-3 CGTs also appear routinely as the result out of contractions. The remainder of this section therefore is dedicated to the question as to what extent this can be used more systematically to generate new (larger) CG3s with full OM. As this only concerns symmetries, RMTs are fully ignored in the following discussion.

1. Triangular CGT networks


A minimal CGT network required to compute a CG3 from the contraction of other CG3s consists of the three CG3s in the triangular configuration as shown in (19a):



Once the closed lines along the triangle in (19a) are contracted, one obtains another CGT with three open legs. In the present arrow configuration, the result is proportional to a CG3 $(q_1 q_2 | q_3)$ labeled E in (19b). The proportionality factor is obtained by projecting E onto the result out of (19a), as depicted in (19c). This situation then is completely analogous to computing $6j$ -symbols as discussed with (18), with the set of “ $6j$ ”s $\{q_1, q_2, q_3, q'_1, q'_2, q'_3\}$. In general, E itself can have outer multiplicity, in which case the proportionality factor in (19a) becomes a matrix factor. Specifically, the result out of contraction (19a) needs to be decomposed into the OM space of E already present. If E has not yet been computed, the OM space arising out of (19a) needs to be orthonormalized (e.g., via QR decomposition), thus defining the new E . If E already existed yet had not been obtained from a full tensor-product decomposition, the contraction in (19a) may yield new OM components, and thus expands the OM space in E .

2. Iterative schemes for CG3s with full OM

New standard rank-3 CGTs, i.e., with two incoming and one outgoing index, may be obtained systematically via a recursive scheme based on contractions of smaller CG3s for the sake of numerical efficiency [72–75]. To start with, consider the CG3 $(q_1 q_2 | q_3)$ in (20a):



Any non-Abelian multiplet relevant in TN simulations, by construction, is generated by building a quantum many-body

Hilbert space iteratively by adding one particle after another in the defining irrep (if all possible multiplets can be reached this way, then the representation is said to be faithful (Burnside theorem [72]). The number of particles relevant in a given multiplet for $SU(N)$ then relates to the number of boxes. In this spirit, one can trivially add a disconnected line in the defining irrep [labeled “1,” e.g., one box in a $SU(N)$ Young tableau] along indices 1 and 2 as indicated by the blue line in (20b). In order to have consistent directions with q_1 and q_2 , however, the arrow on the blue line is reverted with respect to, say, leg 2, which introduces the $1j$ -symbol U_1 while also switching to the dual $\bar{1}$ on leg 2. This way the particle entering on leg 1 is annihilated on leg 2, and never affects leg 3.

The blue line can then be fused with both q_1 and q_2 [(20c), giving rise also to larger multiplets q'_1 and q'_2 (indicated by thicker lines). The contracted result in (20c) is again a CG3, but now in $(q'_1, q'_2 | q_3)$. The $1j$ -symbol may be simply contracted onto any of the connected CG3s (empty circles), thus resulting in a triangular configuration similar to the one already encountered in (20a).

The procedure above suggests that a CG3 in some larger irreps q'_1 and q'_2 can be computed via a contraction of three smaller CG3s (circles) in a triangular TN that has smaller irreps (thinner lines) on the contracted lines. This also includes primitive CG3s, i.e., CG3s that contain either the defining irrep $q_i = 1$ or its dual $\bar{1}$ on one of their legs [72,73]. Such CG3s [blue circles in (20)] are always OM free. Starting in (20a) from a valid CG3, followed by plain operations of adding and fusing a line in (20b), the result in (20c) is expected to be nonzero for any valid primitive CG3s $(q'_1, 1 | q_1)$ and $(q'_2, \bar{1} | q_2)$, as long as they also result in a valid overall CG3 $(q'_1 q'_2 | q_3)$.

The simple argument with (20a)–(20c) can be reformulated into a more general strategy on how to compute larger CG3s recursively from smaller ones. Suppose one is interested in computing the CG3 with larger irreps $(q_1 q_2 | q_3)$, as depicted by thicker lines in (20d), meaning that they are not the defining irrep or its dual. Then one may split off the defining irrep from q_1 , such that $(q_1 | 1, q'_1)$ with $q'_1 = q_1 - 1 < q_1$. Here q'_1 is “smaller” as compared to q_1 in the sense that q'_1 has one “fewer particle in the defining representation,” e.g., for $SU(N)$ it has one box fewer in the Young tableau (note q'_1 does not necessarily have to be an irrep that is also smaller in dimension, since a Young tableau with fewer boxes may represent a multiplet of larger dimension). Similarly, the dual of the defining representation, $\bar{1}$, is split off from q_2 , such that $(q_2 | \bar{1}, q'_2)$ with $q'_2 = q_2 - \bar{1} < q_2$ in the same sense as for q'_1 above, except that an “antiparticle” (hole) was removed from q_2 where, e.g., for $SU(N)$ the irrep $\bar{1}$ consists of $N - 1$ boxes in a Young tableau, such that $N - 1$ boxes are split off from the Young tableau for q_2 . Then merging the $1j$ -symbol as in (20c) (blue dot) with the lower CGT, this is equivalent to having a simple line in irrep 1 directed downward in (20e). Depending on the external irreps, one may also have to consider the reversed process in (20f) where one particle moves upward. For general non-Abelian symmetry, typically also multiple intermediate multiplets q'_1 and q'_2 can occur with the same number of “particles” or “boxes,” all of which need to be included, as this gives rise to OM in the targeted $(q_1 q_2 | q_3)$. Besides, the CG3 $(q'_1 q'_2 | q_3)$ itself may already have

OM which, however, is assumed to be complete from earlier iterations, given that this CG3 is “smaller” with respect to two of its legs in the above sense.

Aiming at reducing the size of the constituting CG3s as in (20e) fails for a maximum weight CG3, having $q_3 = q_1 + q_2$. More generally, it fails for ireps q_3 which have all particles or boxes present from $q_1 + q_2$, still. For these cases, q_3 can no longer be reached by the smaller ireps $q'_1 = q_1 - 1$ and $q'_2 = q_2 - \bar{1}$. However, in the present case this problem is simply taken care of by remembering that directions on CGTs can be easily altered using $1j$ -symbols. Therefore adopting a sorted order $q_1 \geq q_2 \geq q_3$ (e.g., sorted lexicographically in terms of q labels) this excludes $(q_1, q_2|q_3)$ from ever becoming a maximum weight CG3 for nontrivial $q_i > 0$. As an aside, note that for $SU(N)$ the multiplets q_3 out of $(q_1, q_2|q_3)$ all have the same number of boxes in their Young tableau up to modulo N , i.e., up to having full columns of $N = 1 + (N - 1)$ boxes removed, the latter corresponding to $1 + \bar{1}$. In (20e), by comparison, the maximum weight state for $q'_1 = q_1 - 1$ and $q'_2 = q_2 - \bar{1}$ also has N fewer boxes, indeed, effectively having removed one column in the Young tableau.

Consider still the specific example of $SU(2)$ where, for convenience and also consistency with general $SU(N)$, symmetry labels are taken as the integers $q_i \equiv 2S_i \geq 0$ for spin label S_i [45], as this also directly specifies the number of boxes in the corresponding Yang tableau. When computing the CG3 $(q_1 q_2 | q_3)$ as in (20e) then, with $1 = \bar{1}$ for $SU(2)$, two fewer boxes in a Young tableau can reach q_3 , i.e., $\delta q = 2$, or equivalently $\delta S = 1$. Since the maximum weight CG3 was excluded by reference to sorted CG3s above, this is no issue. This shows that the decomposition in (20e), indeed, works well for $SU(2)$, in that two primitive CG3s (blue circles) are contracted with a CG3 that strictly contains smaller ireps on its legs. The present prescription is thus analogous, e.g., to the three-term recursive schemes for $SU(2)$ introduced in [74,75] which also make use of primitive CG3s.

The strategy above to compute CG3s for specific q labels via a recursive approach is general for non-Abelian symmetries from its outline. What is explicitly demonstrated by construction in (20a)–(20c) is that (20e) permits us to split off and “route” a particle from leg 1 to 2 or vice versa (20f), while it never reaches or participates in the multiplet on leg 3 [see (20b)]. This approach above appears to work well empirically to also generate the full OM space of the larger CG3 $(q_1 q_2 | q_3)$ for arbitrary non-Abelian symmetries. At the present stage this remains a conjecture, though, and a rigorous mathematical proof is left for the future.

In any case, explicit generation of all encountered CGTs becomes prohibitive with increasing symmetry rank τ , since typical multiplets grow exponentially in τ , in practice like $\lesssim 10^\tau$ (see Appendix A). However these tensors are generated, even the explicit evaluation of the primitive CG3s based on tensor product decomposition fusing the defining irep only will quickly hit a hard wall for $\tau \gtrsim 5$ [e.g., with $SU(N \gtrsim 6)$ already giving rise to a full tensor-product decomposition of state spaces frequently exceeding several millions in dimension, even if the problem is sparse at a density of $\gtrsim 10^{-3}$]. On the other hand, given that the X-symbols introduced in this paper represent fully contracted CGT networks up to

OM indices, this also makes them susceptible to recursive buildup schemes. This is an attractive route, since X-symbols are generally *much* smaller in dimension than the contracted CGTs they represent. For the purposes of this paper, however, this is left as an outlook.

E. Summary

This paper introduces X-symbols for the efficient treatment of pairwise contractions in tensor networks in the presence of general non-Abelian symmetries. Once computed from CGTs and tabulated, they permit us to completely sidestep the explicit usage of CGTs at a latter stage, as they are contracted onto the multiplicity indices of the involved RMTs. X-symbols represent a general framework that is also trivially applicable to Abelian symmetries. As such, they provide a coherent concept for any type of symmetry setting. Much of this paper is a summary of significant extensions that have been implemented and already tested thoroughly in the QSPACE v3 tensor library (see Appendix B) with strong applications, e.g., in [55,76–81]. In this sense, the present paper provides a concise, polished, and proven version of the underlying concepts.

ACKNOWLEDGMENTS

This paper has greatly benefited from a few intensive discussions on non-Abelian symmetries in general while visiting Anthony Kennedy (University of Edinburgh, Scotland). I also acknowledge general discussions and much appreciated support from the group of Jan von Delft in Munich specifically also including Seung-Sup Lee, Katharina Stadler, Benedikt Bruognolo, Bin-Bin Chen, Jheng-Wei Li, and Arne Alex, as well as Wei Li (University of Beihang, China), Wang Yilin (Brookhaven National Laboratory), Thomas Quella (University of Melbourne, Australia), László Borda, and Gergely Zaránd (Budapest University, Hungary). This work was supported by the Heisenberg fellowship (Deutsche Forschungsgemeinschaft Grant No. WE4819/2-1), as well as by the US Department of Energy, Office of Basic Energy Sciences, under Contract No. DE-SC0012704 without temporal overlap.

APPENDIX A: SYMMETRY RANK AND TYPICAL MULTIPLY DIMENSION

The rank τ of a symmetry is an intrinsic property of a given simple non-Abelian symmetry that needs to be differentiated from the rank r of tensors. The Lie algebra of a simple non-Abelian symmetry possesses at most τ simultaneously commuting generators, which thus defines its rank. This *Cartan subalgebra* S_α^z with $\alpha = 1, \dots, \tau$ can be simultaneously diagonalized [56–61], and hence gives rise to an τ -dimensional label, or weight space [45]. The Cartan subalgebra is complemented by generalized raising and lowering operators $S_\alpha^{(\dagger)}$ with $\alpha = 1, \dots, \tau$ which correspond to the simple roots of the Lie algebra.

In practice, then the typical multiplet dimension encountered grows exponentially with the rank τ as in $|q| \sim 10^\tau$ [45]. The essential reason for this is that multiplets explore an τ dimensional volume of label space derived from S_α^z . This makes large symmetries such as $SU(N)$ with $\tau_{SU(N)} = N - 1$

for $N \gtrsim 4$ computationally challenging. In contrast, the well-known $SU(2)$ is a comparatively very simple yet elementary non-Abelian symmetry, in that it corresponds to a symmetry of rank $\tau_{SU(2)} = 1$. Therefore a single label suffices (i.e., the spin S). For the same reason, there is no IM for $SU(2)$, and also no OM for standard Clebsch-Gordan coefficient tensors (rank-3 CGTs).

APPENDIX B: QSPACE v3

The concepts for non-Abelian symmetries presented in this paper have been implemented over the past few years in the tensor library QSPACE [45]. They were thoroughly tested in a range of papers [7,9,45,55,76–81] with many more applications via student and other research projects. QSPACE was introduced in 2006 [as version 1 (v1)] on the level of arbitrary combinations of Abelian symmetries. This is effectively the present state of the ITensor Library [82]. QSPACE v1 presented a convenient framework, with Weichselbaum and von Delft [7] its first notable application. QSPACE v2 [45] was also able to handle non-Abelian symmetries on a generic level, by introducing an additional tensor layer for generalized CGTs. It built a database for irreducible representations (R-store), as well as CG3s (C-store), i.e., with two arrows in and one out. The C-store also stored a listing of the branching rules out of each tensor product decomposition.

However, in QSPACE v2 the OM resolved CGTs were explicitly attached to each tensor in full together with the RMTs. Individual entries (fixed symmetry sectors) contained linear superpositions in outer multiplicity, i.e., CGTs of type $C_q^\mu \equiv \sum_{\mu'} w^{\mu\mu'} C_{q\mu'}$ [see discussion after Eq. (5)]. This turned out cumbersome since CGTs were re-contracted in every QSPACE tensor contraction. Specifically for larger symmetries, e.g., for $SU(N \gtrsim 4)$, this quickly shifted the dominant numerical cost from the actual physical calculation with respect to the RMTs $\|A\|_q$ to the treatment of the CGTs C_q .

Therefore for QSPACE v3 (developed and thoroughly tested since 2015), not just CG3s but all CGTs of arbitrary rank are computed on demand once and for all, properly orthonormalized (see Sec. II B), and stored in the C-store. The QSPACE tensor no longer carries the full CGTs but only a reference. The X-symbols for pairwise contractions are also computed on demand once and for all, and stored in an additional centralized database (X-store).

In QSPACE, databases are generally built on demand, except for the very elementary initialization when a symmetry is used for the very first time. Preemptive calculation of all possible objects, e.g., up to some prespecified multiplet dimension, quickly proliferates to hundreds of thousands of entries, even though redundancy in storage has been minimized to a large degree by using sorted CGTs, i.e., sorted with respect to their q labels, etc. In the sense that it is impossible to build a complete database for non-Abelian symmetries that permit an infinite number of ireps, it is mandatory eventually in any case, to build entries on demand. While it would not really matter for $SU(2)$ or $SU(3)$, since all objects are (re-)computed quickly in these cases, for larger non-Abelian symmetries, the concepts of *on demand* and *once and for all* become crucially

important for numerical performance. For example, note that, starting with $SU(N \gtrsim 4)$, the typical size of individual multiplets in CGTs quickly surpasses the number of multiplets in RMTs used in a calculation. For example, in DMRG simulations one barely exceeds an effective dimension $D^* \lesssim 8000$ within the space of RMTs. However, while typical multiplets for $SU(N)$ already reach dimensions up to $\lesssim 10^{N-1} = 1000$ for $SU(4)$, the largest generated multiplets there already reach $\lesssim 10^N = 10\,000$ (see Appendix A).

Now building one's database on demand quickly makes it dependent on the history of a calculation. Even if CGTs are computed with full OM, there can be an arbitrary orthogonal rotation in OM space. In the absence of OM, this reduces to a simple sign convention. However, for example, in $SU(3)$, the CG3 for $(qq|q)$ with $q = (mn)$ (using Dynkin labels [45,56]) with $n \geq 0$ has OM $m = n + 1$, and hence can already be made arbitrarily large. Using symmetries under exchange of the three legs is of limited use here to fully fix rotations in OM space. Given that larger-rank CGTs are computed on demand via contractions, their specific rotation in OM space necessarily will depend on the specific tensors contracted first, and hence on the history of the calculation. There exists no (at least to the author) known convention that naturally and fully fixes the basis of an OM decomposition in all circumstances for CGTs of arbitrary rank.

Standard CG3s can be generated in two ways: (i) via explicit tensor product decomposition ("standard Clebsch-Gordan coefficients") or (ii) via contractions including multiple rank-3 CGTs (see Sec. III D) or a pair of tensors with rank $r > 3$. Option (i) is the clearly preferred option, since it generates the full OM at once and in a deterministic and thus well-defined manner. However, this can become prohibitive for large symmetries, e.g., when considering the fusion of two large multiplets on the bonds of a tensor network state, while actually the full tensor-product decomposition is not required (e.g., since the large bond multiplets only fuse into the symmetry label of the global wave function, which often is a singlet). The full tensor product of two large multiplets may also be far too large to actually occur as a multiplet at the level of a physical site in the lattice Hamiltonian of interest. In this sense, there is (necessarily) a cutoff in dimension: if a tensor product decomposition is required, e.g., when building a product state space in a TNS setting, it needs to be performed in any case. But if one needs to compute a contraction that results in a rank-3 CGT, one may opt to perform a full tensor-product decomposition first, and then project the result of the contraction onto it. Or, alternatively, one may be satisfied just with the result of the contraction itself. If the resulting CGT already exists, the result is projected onto it, thus possibly extending the existing OM space. Either way may result in a different basis in OM space. It is crucially important then that one strictly ensures consistency across one's calculations.

Now given a history dependent database that is accessed and maintained centrally, this implies when running multiple jobs at the same time or when parallelizing within a single job that threads need to be coordinated. That is, threads may have to wait, if another thread is currently in the process of updating the same object in the C-store (via contraction or

tensor-product decomposition) or the X-store (if a contraction between a new pair of CGTs needs to be performed, or if a new OM component was encountered such that the derived X-symbol needs to be updated). This coordination can be enforced on the level of the database (e.g., locks on affected objects) and also in memory in between different threads (thread locks).

In summary, QSPACE v3 consists of three databases (referred to as “RCX store” as a whole):

(1) The R-store for irep representations generated in a calculation. These include a full basis decomposition in terms of weight labels, a sparse representation of the diagonal Cartan subalgebra, as well as of the simple roots of the Lie algebra (the generalized raising and lowering operators) under consideration (see Appendix A).

(2) The C-store for storage of all CGTs of arbitrary rank $r \geq 2$. The CGTs are stored in sparse format with typical average sparsity $\gtrsim 10^{-3}$. The sparse format necessarily requires a framework for sparse tensors of arbitrary rank which has been coded from scratch into QSPACE v2. The C-store also stores the fusion rules out of full tensor product decompositions, as well as all $1j$ -symbols, which simply represent a special case of CGTs, namely, $(q\bar{q}|0)$.

(3) The X-store for the X-symbols that derive from any encountered pairwise CGT contraction that are *not* trivially zero due to non-permissible combinations of symmetry labels (e.g., when a CGT contraction results in a nondiagonal rank-2 CGT or a rank-3 CGT that was not listed in an earlier full tensor product decomposition).

The data in the R- and C-store are computed in better than double precision (roughly quad), since the entire RCX store is built iteratively along a TN calculation starting from the very elementary defining representation (and its dual, for convenience). This guards against accumulated error and ensures that all entries are numerically exact in double precision. It is also important for sparse storage in order to reliably distinguish actual possibly small CGT coefficients from numerical noise. The X-symbols are computed from CGTs in the C-store, but can eventually be cast into plain double precision as they are contracted onto RMTs anyway.

The C-store for larger symmetries is extremely heterogeneous, as it contains tensors that represent scalars, all the way up to individual CGTs that [e.g., for rank-4 CGTs in

SU(4) quickly] require 1 TB of space or larger. The X-store contains by far the most of the entries. Many contractions are known to be trivially zero since the symmetry labels of the resulting CGT are not permitted from a symmetry point of view, and hence can be excluded from the X-store. Still, e.g., by not insisting that OM spaces are complete, the X-store also contains many X-symbols that are actually zero, meaning that the pairwise contractions of two CGTs results in a CGT of finite dimensions, yet with (Frobenius) norm resembling numerical noise.

When running multiple jobs, it is convenient to maintain a central global RCX store that has strictly read-only access (except for times when it is updated manually) which contains the bulk of all symmetry related data. In addition, a differential store that is local to each job allows each job to compute and store whatever is needed in addition. Since the latter is decoupled from other running jobs, at least at this level interference between different simulations leading to possible inconsistencies is avoided. Once an RCX store is complete for a given calculation (e.g., if the same calculation is run a second time), only metadata are read from the R and C-store (such as branching rules in tensor products or irep dimensions). For contractions, only the required X-symbols in the X-store need to be read once, e.g., at runtime into a program-internal hash table, with no need to explicitly load the full CGTs from the C-store. The X-symbols are typically much smaller than the involved CGTs, much like $6j$ -symbols.

Since any tensor is stored as a QSPACE with the data comprised as the tensor product in Eq. (2), in principle, it has access to *all* matrix elements in the full state space. This makes QSPACE tensors versatile [45] in that all of the elementary tensor operations are allowed that one is used to when performing calculations *without* symmetries, as long as they do not explicitly break a symmetry. For example, it is very difficult (since inconsistent) to represent a finite magnetic field BS_z if the calculation was initialized with SU(2) spin symmetry. A representation of S_z would require us to break up CGTs into specific components which when preserving symmetries, however, are considered inseparable units. In the presence of spontaneous symmetry breaking of an otherwise symmetric Hamiltonian, on the other hand, symmetries can be turned on and off at will. In particular, non-Abelian symmetries can also be reduced to their Abelian center. This is a valuable approach to shed light on physical scenarios where spontaneous symmetry breaking is weak or debated.

-
- [1] J. Jordan, R. Orús, G. Vidal, F. Verstraete, and J. I. Cirac, Classical Simulation of Infinite-Size Quantum Lattice Systems in two Spatial Dimensions, *Phys. Rev. Lett.* **101**, 250602 (2008).
- [2] H. C. Jiang, Z. Y. Weng, and T. Xiang, Accurate Determination of Tensor Network State of Quantum Lattice Models in two Dimensions, *Phys. Rev. Lett.* **101**, 090603 (2008).
- [3] R. Orús and G. Vidal, Simulation of two-dimensional quantum systems on an infinite lattice revisited: Corner transfer matrix for tensor contraction, *Phys. Rev. B* **80**, 094403 (2009).
- [4] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys. (NY)* **326**, 96 (2011).
- [5] R. Orús, Tensor networks for complex quantum systems, *Nat. Rev. Phys.* **1**, 538 (2019).
- [6] K. G. Wilson, The renormalization group: Critical phenomena and the Kondo problem, *Rev. Mod. Phys.* **47**, 773 (1975).
- [7] A. Weichselbaum and J. von Delft, Sum-Rule Conserving Spectral Functions from the Numerical Renormalization Group, *Phys. Rev. Lett.* **99**, 076402 (2007).
- [8] R. Bulla, T. A. Costi, and T. Pruschke, Numerical renormalization group method for quantum impurity systems, *Rev. Mod. Phys.* **80**, 395 (2008).
- [9] A. Weichselbaum, Tensor networks and the numerical renormalization group, *Phys. Rev. B* **86**, 245124 (2012).

- [10] S. R. White, Density Matrix Formulation for Quantum Renormalization Groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [11] S. R. White and A. E. Feiguin, Real-Time Evolution Using the Density Matrix Renormalization Group, *Phys. Rev. Lett.* **93**, 076401 (2004).
- [12] U. Schollwöck, The density-matrix renormalization group, *Rev. Mod. Phys.* **77**, 259 (2005).
- [13] E. M. Stoudenmire and S. R. White, Studying two-dimensional systems with the density matrix renormalization group, *Annu. Rev. Condens. Matter Phys.* **3**, 111 (2012).
- [14] G. Vidal, Entanglement Renormalization, *Phys. Rev. Lett.* **99**, 220405 (2007).
- [15] G. Evenbly and G. Vidal, Tensor Network Renormalization, *Phys. Rev. Lett.* **115**, 180405 (2015).
- [16] F. Verstraete and J. I. Cirac, Renormalization algorithms for quantum-many body systems in two and higher dimensions [arXiv:cond-mat/0407066v1](https://arxiv.org/abs/cond-mat/0407066v1) (2004).
- [17] C. V. Kraus, N. Schuch, F. Verstraete, and J. I. Cirac, Fermionic projected entangled pair states, *Phys. Rev. A* **81**, 052338 (2010).
- [18] J. I. Cirac and F. Verstraete, Renormalization and tensor product states in spin chains and lattices, *J. Phys. A: Math. Theor.* **42**, 504004 (2009).
- [19] P. Corboz, R. Orús, B. Bauer, and G. Vidal, Simulation of strongly correlated fermions in two spatial dimensions with fermionic projected entangled-pair states, *Phys. Rev. B* **81**, 165104 (2010).
- [20] H. N. Phien, J. A. Bengua, H. D. Tuan, P. Corboz, and R. Orús, Infinite projected entangled pair states algorithm improved: Fast full update and gauge fixing, *Phys. Rev. B* **92**, 035142 (2015).
- [21] Z. Y. Xie, J. Chen, M. P. Qin, J. W. Zhu, L. P. Yang, and T. Xiang, Coarse-graining renormalization by higher-order singular value decomposition, *Phys. Rev. B* **86**, 045139 (2012).
- [22] A. García-Sáez and J. I. Latorre, Renormalization group contraction of tensor networks in three dimensions, *Phys. Rev. B* **87**, 085130 (2013).
- [23] S.-J. Ran, A. Piga, C. Peng, G. Su, and M. Lewenstein, Few-body systems capture many-body physics: Tensor network approach, *Phys. Rev. B* **96**, 155120 (2017).
- [24] N. Nakatani and G. K.-L. Chan, Efficient tree tensor network states (TTNS) for quantum chemistry: Generalizations of the density matrix renormalization group algorithm, *J. Chem. Phys.* **138**, 134113 (2013).
- [25] C. Krumnow, L. Veis, Ö. Legeza, and J. Eisert, Fermionic Orbital Optimization in Tensor Network States, *Phys. Rev. Lett.* **117**, 210402 (2016).
- [26] S. Anders, M. B. Plenio, W. Dür, F. Verstraete, and H.-J. Briegel, Ground-State Approximation for Strongly Interacting Spin Systems in Arbitrary Spatial Dimension, *Phys. Rev. Lett.* **97**, 107206 (2006).
- [27] N. Schuch, I. Cirac, and D. Pérez-García, PEPS as ground states: Degeneracy and topology, *Ann. Phys. (NY)* **325**, 2153 (2010).
- [28] J. Eisert, M. Cramer, and M. B. Plenio, Colloquium: Area laws for the entanglement entropy, *Rev. Mod. Phys.* **82**, 277 (2010).
- [29] M. A. Levin and X.-G. Wen, String-net condensation: A physical mechanism for topological phases, *Phys. Rev. B* **71**, 045110 (2005).
- [30] Z.-C. Gu and X.-G. Wen, Tensor-entanglement-filtering renormalization approach and symmetry-protected topological order, *Phys. Rev. B* **80**, 155131 (2009).
- [31] X.-L. Qi and S.-C. Zhang, Topological insulators and superconductors, *Rev. Mod. Phys.* **83**, 1057 (2011).
- [32] M. Hermele and V. Gurarie, Topological liquids and valence cluster states in two-dimensional $SU(N)$ magnets, *Phys. Rev. B* **84**, 174441 (2011).
- [33] S. Capponi, P. Lecheminant, and K. Totsuka, Phases of one-dimensional $SU(N)$ cold atomic fermi gases: From molecular Luttinger liquids to topological phases, *Ann. Phys. (NY)* **367**, 50 (2016).
- [34] I. P. McCulloch and M. Gulcsi, The non-Abelian density matrix renormalization group algorithm, *Europhys. Lett.* **57**, 852 (2002).
- [35] A. I. Tóth, C. P. Moca, Ö. Legeza, and G. Zaránd, Density matrix numerical renormalization group for non-Abelian symmetries, *Phys. Rev. B* **78**, 245109 (2008).
- [36] S. Singh, R. N. C. Pfeifer, and G. Vidal, Tensor network decompositions in the presence of a global symmetry, *Phys. Rev. A* **82**, 050301(R) (2010).
- [37] G. Alvarez, Implementation of the $SU(2)$ Hamiltonian symmetry for the DMRG algorithm, *Comp. Phys. Commun.* **183**, 2226 (2012).
- [38] S. Singh and G. Vidal, Tensor network states and algorithms in the presence of a global $SU(2)$ symmetry, *Phys. Rev. B* **86**, 195114 (2012).
- [39] Cătălin Pașcu Moca, A. Alex, J. von Delft, and G. Zaránd, $SU(3)$ Anderson impurity model: A numerical renormalization group approach exploiting non-Abelian symmetries, *Phys. Rev. B* **86**, 195128 (2012).
- [40] P. Scholl, S. Singh, M. Rizzi, and R. Orus, A programming guide for tensor networks with global $SU(2)$ symmetry, [arXiv:1809.08180](https://arxiv.org/abs/1809.08180) (2018).
- [41] P. Nataf and F. Mila, Density matrix renormalization group simulations of $SU(n)$ Heisenberg chains using standard young tableaux: Fundamental representation and comparison with a finite-size Bethe ansatz, *Phys. Rev. B* **97**, 134420 (2018).
- [42] V. Murg, F. Verstraete, Ö. Legeza, and R. M. Noack, Simulating strongly correlated quantum systems with tree tensor networks, *Phys. Rev. B* **82**, 205105 (2010).
- [43] S. Singh and G. Vidal, Global symmetries in tensor network states: Symmetric tensors versus minimal bond dimension, *Phys. Rev. B* **88**, 115147 (2013).
- [44] T. A. Costi, L. Bergqvist, A. Weichselbaum, J. von Delft, T. Micklitz, A. Rosch, P. Mavropoulos, P. H. Dederichs, F. Mallet, L. Saminadayar, and C. Bäuerle, Kondo Decoherence: Finding the Right Spin Model for Iron Impurities in Gold and Silver, *Phys. Rev. Lett.* **102**, 056802 (2009).
- [45] A. Weichselbaum, Non-Abelian symmetries in tensor networks: A quantum symmetry space approach, *Ann. Phys. (NY)* **327**, 2972 (2012).
- [46] M. Hanl, A. Weichselbaum, T. A. Costi, F. Mallet, L. Saminadayar, C. Bäuerle, and J. von Delft, Iron impurities in gold and silver: Comparison of transport measurements to numerical renormalization group calculations exploiting non-Abelian symmetries, *Phys. Rev. B* **88**, 075146 (2013).
- [47] K. M. Stadler, Z. P. Yin, J. von Delft, G. Kotliar, and A. Weichselbaum, Dynamical Mean-Field Theory Plus Numerical Renormalization Group Study of Spin-Orbital Separation in a Three-Band Hund Metal, *Phys. Rev. Lett.* **115**, 136401 (2015).

- [48] P. Nataf, M. Lajkó, P. Corboz, A. M. Läuchli, K. Penc, and F. Mila, Plaquette order in the SU(6) Heisenberg model on the honeycomb lattice, *Phys. Rev. B* **93**, 201113 (2016).
- [49] T. A. Toth, A. M. Laeuchli, F. Mila, and K. Penc, Three-Sublattice Ordering of the SU(3) Heisenberg Model of Three-Flavor Fermions on the Square and Cubic Lattices, *Phys. Rev. Lett.* **105**, 265301 (2010).
- [50] M. Greiner, O. Mandel, T. Esslinger, T. W. Hansch, and I. Bloch, Quantum phase transition from a superfluid to a mott insulator in a gas of ultracold atoms, *Nature (London)* **415**, 39 (2002).
- [51] I. Bloch, J. Dalibard, and W. Zwerger, Many-body physics with ultracold gases, *Rev. Mod. Phys.* **80**, 885 (2008).
- [52] M. A. Cazalilla and A. M. Rey, Ultracold Fermi gases with emergent SU(N) symmetry, *Rep. Prog. Phys.* **77**, 124401 (2014).
- [53] F. Scazza, C. Hofrichter, M. Höfer, P. C. De Groot, I. Bloch, and S. Fölling, Observation of two-orbital spin-exchange interactions with ultracold SU(N)-symmetric fermions, *Nat. Phys.* **10**, 779 (2014).
- [54] B. Bauer, P. Corboz, R. Orús, and M. Troyer, Implementing global Abelian symmetries in projected entangled-pair state algorithms, *Phys. Rev. B* **83**, 125106 (2011).
- [55] T. Liu, W. Li, A. Weichselbaum, J. von Delft, and G. Su, Simplex valence-bond crystal in the spin-1 Kagome Heisenberg antiferromagnet, *Phys. Rev. B* **91**, 060403(R) (2015).
- [56] E. B. Dynkin, The structure of semi-simple algebras, *Usp. Mat. Nauk* **2**, 59 (1947).
- [57] R. N. Cahn, *Semi-Simple Lie Algebras and their Representations* (Benjamin, New York, 1984).
- [58] M. A. A. van Leeuwen, A. M. Cohen, B. Lissner, and C. Arjeh, LiE, A Package for Lie Group Computations, Computer Algebra Nederland, Université de Poitiers, Amsterdam, 1992.
- [59] C. Pope, Geometry and group theory, Online lecture notes, Texas A&M University, 2006.
- [60] R. Gilmore, *Lie Groups, Lie Algebras, and Some of their Applications* (Dover, New York, 2006).
- [61] J. E. Humphreys, *Introduction to Lie Algebras and Representation Theory* (Springer, New York, 2012), Vol. 9.
- [62] S. Singh, Tensor network states and algorithms in the presence of Abelian and non-Abelian symmetries, Ph.D. thesis, The University of Queensland, 2012, [arXiv:1203.2222v2](https://arxiv.org/abs/1203.2222v2).
- [63] R. Orús, Advances on tensor network theory: Symmetries, fermions, entanglement, and holography, *Eur. Phys. J. B* **87**, 280 (2014).
- [64] H. J. Liao, Z. Y. Xie, J. Chen, Z. Y. Liu, H. D. Xie, R. Z. Huang, B. Normand, and T. Xiang, Gapless Spin-Liquid Ground State in the $S = 1/2$ Kagome Antiferromagnet, *Phys. Rev. Lett.* **118**, 137202 (2017).
- [65] E. M. Stoudenmire and S. R. White, Minimally entangled typical thermal state algorithms, *New J. Phys.* **12**, 055026 (2010).
- [66] M. Lubasch, J. I. Cirac, and M.-C. Bañuls, Algorithms for finite projected entangled pair states, *Phys. Rev. B* **90**, 064425 (2014).
- [67] M. A. Werner, Cătălin Pașcu Moca, Ö. Legeza, M. Kormos, and G. Zaránd, Spin fluctuations after quantum quenches in the $s = 1$ Haldane chain: Numerical validation of the semiclassical theory, *Phys. Rev. B* **100**, 035401 (2019).
- [68] E. Anderson, Z. Bai, C. Bischof, L. S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney *et al.*, *LAPACK Users' Guide* (SIAM, 1999).
- [69] A. Messiah, *Quantum Mechanics*, English translation from French (Wiley, New York, 1966).
- [70] J.-R. Derome and W. T. Sharp, Racah algebra for an arbitrary group, *J. Math. Phys.* **6**, 1584 (1965).
- [71] P. H. Butler, Coupling coefficients and tensor operators for chains of groups, *Trans. R. Soc. London A* **277**, 545 (1975).
- [72] P. H. Butler and B. G. Wybourne, Calculation of j and jm symbols for arbitrary compact groups. I. Methodology, *Int. J. Quantum Chem.* **10**, 581 (1976).
- [73] B. G. Searle and P. H. Butler, Calculation of primitive 6- j symbols, *J. Phys. A: Math. Gen.* **21**, 3041 (1988).
- [74] K. Schulten and R. G. Gordon, Exact recursive evaluation of 3- j - and 6- j -coefficients for quantum-mechanical coupling of angular momenta, *J. Math. Phys.* **16**, 1961 (1975).
- [75] J. H. Luscombe and M. Luban, Simplified recursive algorithm for Wigner 3- j and 6- j -symbols, *Phys. Rev. E* **57**, 7274 (1998).
- [76] S.-S. B. Lee, J. von Delft, and A. Weichselbaum, Doublon-Holon Origin of the Subpeaks at the Hubbard Band Edges, *Phys. Rev. Lett.* **119**, 236402 (2017).
- [77] B.-B. Chen, L. Chen, Z. Chen, W. Li, and A. Weichselbaum, Exponential Thermal Tensor Network Approach for Quantum Lattice Models, *Phys. Rev. X* **8**, 031082 (2018).
- [78] A. Weichselbaum, S. Capponi, P. Lecheminant, A. M. Tsvelik, and A. M. Läuchli, Unified phase diagram of antiferromagnetic SU(n) spin ladders, *Phys. Rev. B* **98**, 085104 (2018).
- [79] K. M. Stadler, G. Kotliar, A. Weichselbaum, and J. von Delft, Hundness versus Mottness in a three-band Hubbard-Hund model: On the origin of strong correlations in Hund metals, *Ann. Phys. (NY)* **405**, 365 (2019).
- [80] E. Walter, K. M. Stadler, S. S. B. Lee, Y. Wang, G. Kotliar, A. Weichselbaum, and J. von Delft, Uncovering non-Fermi-liquid behavior in Hund metals: Conformal field theory analysis of a SU(2) x SU(3) spin-orbital Kondo model, [arXiv:1908.04362](https://arxiv.org/abs/1908.04362) (2019).
- [81] Y. Wang, E. Walter, S.-S. B. Lee, K. M. Stadler, J. von Delft, A. Weichselbaum, and G. Kotliar, Global Phase Diagram of a Spin-Orbital Kondo Impurity Model and the Suppression of Fermi-Liquid Scale, *Phys. Rev. Lett.* **124**, 136406 (2020).
- [82] ITensor Library (version 2.0.11), <http://itensor.org>.