Fourier transform of fermionic systems and the spectral tensor network

Andrew J. Ferris
Département de Physique, Université de Sherbrooke, Québec, J1K 2R1, Canada
ICFO—Institut de Ciencies Fotoniques, Parc Mediterrani de la Tecnologia, 08860 Barcelona, Spain and
Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Str. 1, 85748 Garching, Germany
(Dated: October 30, 2013)

Leveraging the decomposability of the fast Fourier-transform, I propose a new class of tensor network that is efficiently contractible and able to represent many-body systems with local entanglement that is greater than the area law. Translationally invariant systems of free fermions in arbitrary dimensions as well as 1D systems solved by the Jordan-Wigner transformation are shown to be exactly represented in this class. Further, it is proposed that these tensor networks be used as generic structures to describe more complicated systems, possibly leading to highly-efficient calculations in the Fermi-liquid phase. This class shares some similarities with Evenbly & Vidal's branching MERA, but with some important differences and greatly reduced computational demands.

PACS numbers: 05.10.Cc, 71.10.Fd, 71.18.+y

INTRODUCTION

Solving strongly-correlated many-body systems is one of the primary challenges of modern quantum mechanics that arises in a wide variety of fields and contexts. Because direct approaches are usually intractable, an array of analytical and numerical techniques and approximations have been developed over the last century, each associated with its own region of validity, accuracy and difficulty.

One relatively modern approach is based on tensor network decompositions of many-body wavefunctions [1–4], motivated by our understanding of quantum information in this context. By decomposing the wavefunction into a particular product of tensors, one defines the structure of entanglement that is permitted (for example, the low-energy states of many local Hamiltonians obey an *area law* scaling for entanglement entropy [5–8]). Typically, one applies the variational principle within this class of wavefunctions to determine, for example, the tensor network state closest to the ground state. Increasing the rank of the decomposition increases the amount of allowable entanglement, and potentially leads to a more accurate description of the many-body state.

A wide variety of tensor network geometries have been proposed, beginning with the matrix-product state (MPS) [9–11] that is the foundation of the famous density matrix renormalization group [12] and Wilson's numerical renormalisation group [13]. This one-dimensional structure connects nearest neighbors in a chain, and is extremely accurate for 1D systems that have spatially-local correlations (proven to be exponentially close to the low-energy states of gapped, local 1D Hamiltonians [5]). Real-space renormalization group ideas lead to the development of the tree-tensor network (TTN) [14, 15] and the multi-scale entanglement renormalization ansatz (MERA) [16, 17]. The MERA in particular is able to reproduce larger amounts of entanglement such as that necessary to describe critical, 1D quantum systems — and has been found to correctly reproduce the features of the corresponding universality class at the renormalization fixed-point, as described by conformal field theory [18].

Extensions to higher dimensional systems are also possible. Although the MPS and TTN do not have the correct geometry to represent the area law for large 2D systems [15, 19, 20], the MPS can be generalized to to a two-dimensional tensor network state commonly referred to as a projected entangled-pair state (PEPS) [21–23]. The MERA is also generalized to higher dimensions in a straightforward fashion [24]. All the above techniques have shown success in solving 2D systems (e.g. [25–28]), but are somewhat handicapped by relatively large computational demands for even slightly-entangled systems.

Furthermore, both PEPS and the 2D MERA are limited to obey the area-law for entanglement entropy, where the entanglement entropy of a region is proportional to its boundary area (or perimeter, in 2D). This excludes important classes of physical systems which violate the area law — notably any 2D system exhibiting a 1D Fermi surface (resulting in a logarithmic violation [29]), and exotic phases such as Bose metals [30–32].

This situation left the field of tensor networks in an awkward position, unable to accurately deal with what many outside the field consider to be the simplest models of non-interacting free-fermions, and their perturbations (i.e. Fermi liquids). More recently, a new class of tensor network that exhibits very large amounts of entanglement (up to volumetric) was proposed by Evenbly & Vidal, which they have dubbed the 'branching MERA' [33]. However, its properties and relationship to Fermi liquids has not been determined, as far as the author is aware.

In this work, I propose a new tensor-network geometry that mimics the structure of a spectral transformation and includes the quantum Fourier transform (for fermions) as a trivial example. Up to a reordering of the lattice indices, this 'spectral tensor network' can be seen to be a generalization of TTN, or equally a simplification of branching MERA, and local expectation values can be extracted from a wavefunction of n lattice sites in time $n \log n$. Similarly, it is related to previously suggested circuits for use in linear-optics quantum computing [34] and its fermionic analogue [35]. Here, I will give

exact wavefunctions for translationally-invariant free fermion systems in arbitrary dimensions as well as examples of 1D systems that are exactly solvable via Jordan-Wigner transformations (Heisenberg XX chains, transverse-field Ising model, etc). I will then describe a procedure for approximating the ground state of arbitrary local Hamiltonians in an efficient manner, which may be particularly relevant to simulations of 2D systems in the Fermi liquid phase, before concluding with remarks on future directions.

FOURIER TRANSFORMATIONS AND THE SPECTRAL TENSOR NETWORK

We now discuss the unitary circuit required to transform a quantum lattice system between real- and momentum-space, and write it as an efficiently-contractible tensor network. The discrete Fourier transform has long been known to be decomposable into a series of sparse operations [36], in a technique generally referred to as the fast Fourier transform (FFT). Here we will leverage this decomposition to describe the unitary circuit for the quantum Fourier transform of fermionic lattice models (similar to [35]), and show that local expectation values of wavefunctions based on this circuit can be efficiently determined.

We begin with a chain of $n = 2^k$ sites possibly containing a spinless fermion, that begin in the product state $|\Psi_0\rangle = |\psi_1\rangle|\psi_2\rangle \dots |\psi_n\rangle$. The Fourier transform can be decomposed in a series of k steps, with two-site gates acting on pairs of sites separated by distance 2^{k-1}, 2^{k-2}, ..., 2, 1, followed by a (bit-reversal) reordering. The two-site gates, we will show, consist of the fermionic two-site Fourier transform $\hat{F}_2$

$$\hat{F}_2 = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 2^{-1/2} & 2^{-1/2} & 0 \\ 0 & 2^{-1/2} & -2^{-1/2} & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix}, \tag{1}$$

written above in the number basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$, combined with a the one-body phase-gate, $\hat{\omega}_b^a = \hat{Z}^{2a/b}$ ($\hat{Z} = \hat{\mathbb{1}} - 2\hat{n}$ is the Pauli matrix, $\hat{n}$ being the occupation operator), where the phases are commonly called the 'twiddle factors'. Two example circuits for 8 sites are given in Fig. 1, displaying the freedom in choosing the order of the layers. The structure of the circuit is identical to the Cooley-Tukey fast Fourier transform [36] and the determination of the twiddle factors is well studied in classical computer science. It is easy to verify that the circuit produces the correct transformation because it is (a) linear in the field operators, and (b) performs the correct transformation on the single-particle subspace, in analogy with the classical FFT.

This circuit corresponding to the inverse-Fourier transform can be applied to an (unentangled) product-state wavefunction in momentum space to yield a (highly-entangled) wavefunction in real space, as depicted in Fig. 2. I call this tensor network structure a 'spectral tensor network', and the wavefunction is a spectral tensor network state. In general, the connectivity of the diagram is that of a hypercube with side length 2 (even in higher spatial dimensions). Ignoring completely the permutation of sites at the bottom of Fig. 2, the state resembles a branching MERA with half of the gates removed (making it significantly cheaper to contract), or similarly a generalized TTN. However, the site reordering turns the traditional concept of real-space renormalization on its head, as the tensor network couples sites with *maximal* separation first (an intuitive explanation is more likely to come from the decomposition of the Abelian group corresponding to translation invariance).

More generally, a (linear) circuit with the same structure can be applied to bosonic, fermionic and possibly anyonic states. For bosons, $\hat{F}_2$ is the beam-splitter operation

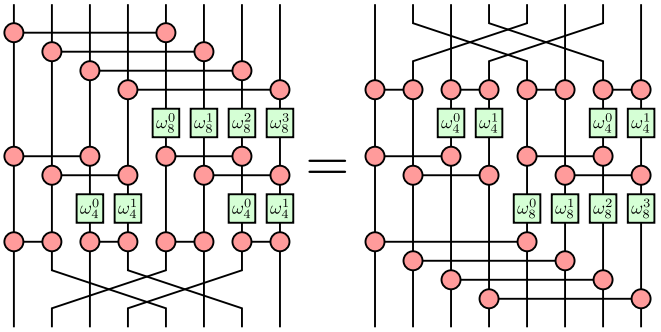

FIG. 1: (Color online) Unitary circuit for performing the quantum Fourier transform on 8 fermionic sites, $\hat{F}_8$. The red gates are the two-site Fourier transform, $\hat{F}_2$. The layers can be permuted, resulting in different 'twiddle factors' $\hat{\omega}_b^a$ and reorderings. The reordering of the sites at the top or bottom of the circuit is given by the bit-reversal operation.

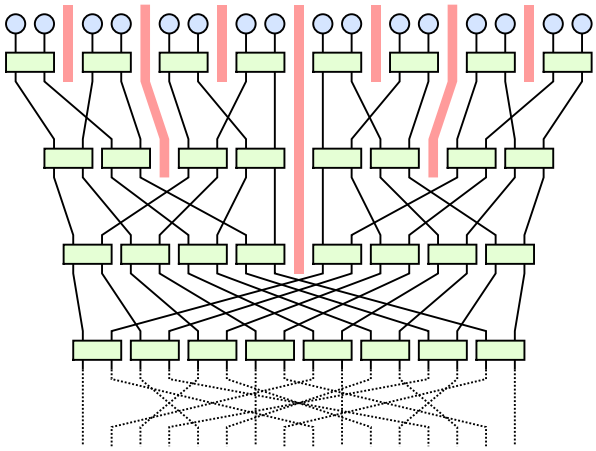

FIG. 2: (Color online) The (inverse) Fourier transform on 16 sites applied to a product momentum-space state is redrawn as a spectral tensor network. At the bottom, the sites are reordered according to a bit-reversal permutation (dotted lines). Working upwards from there, two-body unitaries disentangle the state into two separable systems (emphasized by the wide red barriers). The process is repeated successively until a product state is reached.

that maps two input annihilation operators $(\hat{a}_1, \hat{a}_2)$ to $(\hat{a}_1 + \hat{a}_2, \hat{a}_1 - \hat{a}_2)/\sqrt{2}$, while the phase gate $\hat{Z}^{a/b}$ transforms $\hat{a}_1$ to $(-1)^{a/b}\hat{a}_1$. Unlike fermions, the bond dimension is not fixed at two and in general can grow as large as the total number of particles in the system (and thus becoming numerically inefficient — though this has already been suggested as a powerful tool in linear-optics quantum computation [34]). For fermions, the gates perform a similar linear transformation on the fermionic mode operators as for bosons, but notice that $\hat{F}_2|11\rangle = -|11\rangle$ because of the anti-commutation relations of the creation/annihilation operators. One must also respect the correct statistics of the degrees of freedom in the tensor network diagram by performing the appropriate permutation operation when two wires cross (specifically, introducing a phase of -1 whenever an odd number of fermions cross [37–40]). It would be interesting to determine if the structure is applicable to anyonic models, with more complicated gates and permutation operations (like bosons, the bond dimension may not be bounded for large systems). Note that this transformation is *not* related to the quantum Fourier transformation on qubits used in the quantum computing community.

For systems of higher dimensions, the Fourier transform along each dimension can be applied sequentially along parallel wires. Furthermore, the individual layers of the transforms can be reordered, for instance to recover an alternating structure with gates along the x, y, etc dimensions at each length scale, as shown in Fig. 3.

Together, this gives us a prescription for describing translationally-invariant free-particle systems in arbitrary dimensions, because these Hamiltonians are exactly diagonalized by the Fourier transform. In particular, we have a tensor-network wavefunction for free fermions in arbitrary dimensions. As is well known, such wavefunctions can exhibit quasi-long range correlations (polynomially decaying) and logarithmic violations to the area law. Nonetheless, we now show that extracting local expectation values from these wavefunctions can be achieved by an efficient tensor network contraction scheme, scaling (essentially) linearly with n.

First, we take advantage of the unitary nature of the tensors to vastly simplify the calculation of expectation values of local operators. In Fig. 4 we see the tensor network corresponding to the expectational value of a single-site operator. In general, the number of tensors reduces from $\mathcal{O}(n \log n)$ to $\mathcal{O}(n)$ and by contracting the remaining tree structure from the outside-in, the value can be calculated with computation cost $\mathcal{O}(\chi^5 n)$ (where χ is the dimension of the local Hilbert space). Full details of the contraction are found in the Appendix. Further, a set of all one-site expectation values (or indeed, every one-site reduced density matrix) can be calculated with cost $\mathcal{O}(\chi^5 n \log n)$ by recycling previous calculations. In general, arbitrary two-site expectation values can be calculated with cost $\mathcal{O}(\chi^8 n)$ (see Appendix), and using similar recycling techniques all two-body correlations in a translationally invariant system can be recovered with cost $\mathcal{O}(\chi^8 n \log n)$. For the systems we study here, both ground states and (grand-canonical) thermal states can be studied because the transformation exactly diagonalizes the Hamiltonian.

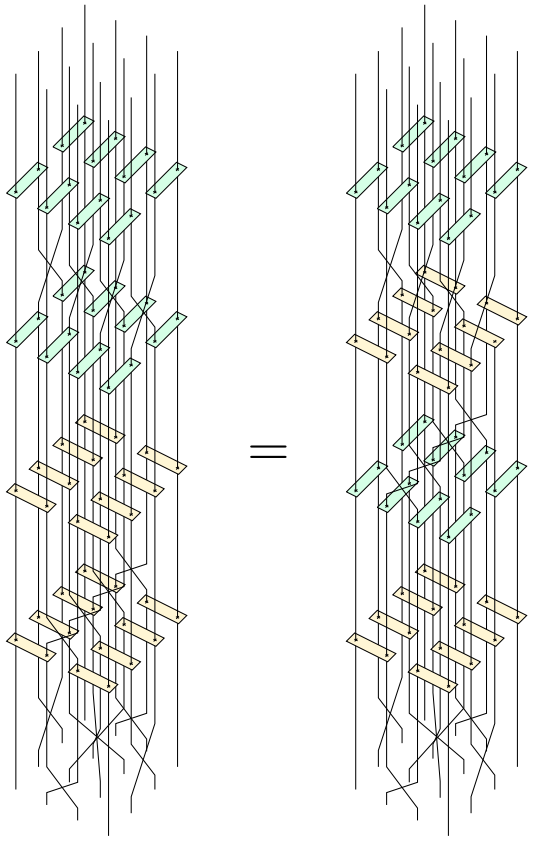

FIG. 3: (Color online) Unitary circuit for performing the Fourier transform on a 2D square lattice system of 4×4 sites. There is significant freedom in choosing the order of the transformation layers.

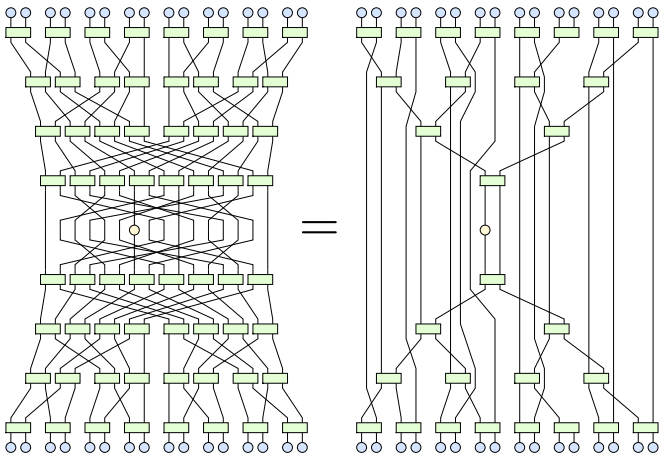

FIG. 4: (Color online) The tensor network corresponding to the expectation value of a operator with support on a single site out of 16. Most of the tensors cancel due to unitarity, leaving a network that is contractible with cost $\mathcal{O}(\chi^5 n)$, where χ is the bond-dimension of the wires and n is the number of sites.

As an example, we examine free-fermions in 1D and 2D

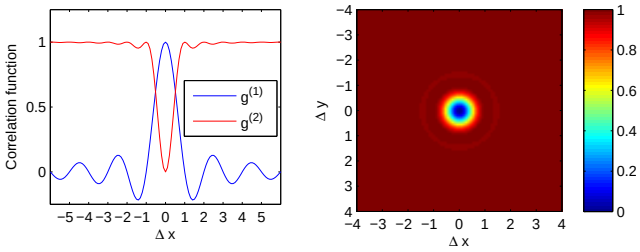

FIG. 5: (Color online) (a) The first-order and second-order normalized correlation functions $g^{(1)}$ and $g^{(2)}$ for the ground state of a 1D system of free fermions (103 fermions on 1024 sites). (b) Second-order normalised correlation function $g^{(2)}$ for a 2D system of free fermions (2093 fermions on a 512×512 lattice). In both cases the distance (Δx and Δy) is normalized to the mean particle spacing.

with nearest-neighbour tunneling

$$\hat{H} = -\sum_{<i,j>} \hat{c}_i^\dagger \hat{c}_j + \text{h.c.} \tag{2}$$

where $<i,j>$ corresponds to neighboring sites on the lattice with $i < j$ (or similar in higher dimensions). We create the ground state wavefunction for a low filling factor and calculate the normalized first- and second-order correlation functions:

$$g^{(1)}(\Delta x) = \frac{\langle \hat{c}_0^\dagger \hat{c}_{\Delta x}\rangle}{\langle \hat{n}_0\rangle}, \quad g^{(2)}(\Delta x) = \frac{\langle \hat{n}_0 \hat{n}_{\Delta x}\rangle}{\langle \hat{n}_0\rangle^2}. \tag{3}$$

The numerical results for 1D and 2D are shown in Fig. 5. These results should be considered exact, to numerical precision. The 2D simulation describes a highly-entangled system (in real-space) of more than quarter of a million sites, yet the calculation was easily completed on a laptop computer using MATLAB (of course, in this case the results could be obtained analytically or by alternative numerical techniques).

JORDAN-WIGNER TRANSFORMATION FOR SPIN CHAINS

An additional result is a compact tensor-network representation of strongly-correlated 1D systems that can be analytically solved using a Jordan-Wigner transformation. A straight-forward example is the fully-anisotropic Heisenberg spin-1/2 chain (XX model). It is slightly more complicated to deal with a Bogoliubov transformation between opposite momentum modes, such as used to solve the quantum transverse-field Ising chain on n sites with Hamiltonian

$$\hat{H} = \sum_i \hat{X}_i \hat{X}_{i+1} + h\hat{Z}_i. \tag{4}$$

This system corresponds to a free fermion system by the transformation $\hat{c}_i = \frac{1}{2}(\hat{X}_i + i\hat{Y}_i)\prod_{j<i}\hat{Z}_j$ where $\hat{c}_i$ are fermion annihilation operators, resulting in

$$\hat{H} = 2\sum_i (\hat{c}_i^\dagger - \hat{c}_i)(\hat{c}_{i+1}^\dagger + \hat{c}_{i+1}) + h\hat{c}_i^\dagger \hat{c}_i - 2n. \tag{5}$$

This Hamiltonian is quadratic in the fermionic operators and translationally invariant, and is traditionally solved by a Fourier transform followed by a Bogoliubov transformation coupling $\hat{\tilde{c}}_k$ with $\hat{\tilde{c}}_{-k}^\dagger$, where $\hat{\tilde{c}}_k = \sum_j \exp(2\pi ijk)\hat{c}_j/\sqrt{n}$. Applying this final transformation to the 'top' of the tensor network diagram makes it slightly more complicated to contract, but it is possible with cost scaling as $\mathcal{O}(n \log n)$ using a 'folding' scheme (see the Appendix for details and numerical results).

SIMULATION OF ARBITRARY FERMIONIC SYSTEMS

The quantum circuit structure used above to perform the Fourier transformation can be utilized to construct a wide array of strongly-correlated many-body wavefunctions. The spectral tensor network that defines this class was depicted in Fig. 2.

We now suggest how to apply this structure to generic systems. As a first step, we consider more complicated bilinear fermionic systems than the spinless case considered above. A simple example is the solution of bilinear fermionic systems which is invariant under translations of l sites. By combining l sites together into a Hilbert space of $\chi = 2^l$, we end with fully-translationally-invariant Hamiltonian that is exactly diagonalized by the Fourier transform (followed by some local operation). This is exactly equivalent to allowing each site to contain more varieties of fermions (with different spin or local orbital configuration, for example) and generally results in systems with multiple bands. The fermionic (fast) Fourier transform can be implemented in its naïve form, individually transforming each 'type' of fermion. This partially diagonalizes the Hamiltonian, separating out the momentum states. Full diagonalization is achieved by mixing components of the *same* momentum — and thus every eigenstate is explicitly given by a spectral tensor network. Anomalous terms can be dealt with a simple Bogoliubov transformation between opposite momenta as discussed in the Appendix.

We further note that a spectral tensor network can be converted into another spectral tensor network with one less layer by combining pairs of neighboring physical sites in the lattice. In this case, the dimension of the local Hilbert space, or more generically the bond dimension of the tensor network χ, is squared in this transformation. This provides a route for increasing the fraction of Hilbert space covered by this decomposition, at the price of increased computation cost (which we will show is a polynomial function of χ).

More generic systems are free to be studied by employing the variational method on the manifold of states defined by the tensor network diagram (in addition to the unitary constraints and any symmetry constraints of the model, e.g. the parity of the number of fermions). For instance, simple energy minimization may yield states very close to highly-entangled ground states, particularly in the Fermi-liquid phase where the entanglement structure should be related to that of free fermions. This is an area of intense interest to the author and

is the subject of further work. Furthermore, it would be interesting to see if the spectral tensor network can accurately represent systems with different boundary conditions or impurities, or be applied effectively to spin systems.

CONCLUSION

The spectral tensor network has been introduced here as a way of efficiently representing the quantum fast Fourier transform and certain classes of quantum states. The Fourier transform can be used to construct free-fermion states with a logarithmic violation of the area law, and as far as the author is aware is the first efficiently contractible tensor network that has been shown to *exactly* capture this phase in two- and higher-dimensions.

The tensor network is shown to be a convenient form for capturing analytic solutions to 1D spin problems that are solved by the Jordan-Wigner transformation. Further, we suggest that this tensor network may be used as a variational ansatz for low-energy states of many-body systems. This may be especially effective in the case of Fermi-liquid phases, where from a *local* perspective the system is highly entangled and is somewhat difficult to treat with existing tensor network techniques, while the Fourier transform is able to convert the system into a relatively disentangled momentum-space basis. Work in this direction is ongoing.

The spectral tensor network is somewhat related to the branching MERA construction of Evenbly and Vidal, with two major differences. The spectral tensor network contains only half the tensors of the branching MERA and is more efficient numerically. A more striking difference is the nonlocal reordering of lattice sites placed at the bottom of the unitary circuit. In complete opposition to the spirit of real-space renormalization, the unitaries at the bottom of the circuit deal with the *longest* length-scales of the system, with the length scales *decreasing* as we ascend the diagram.

Acknowledgement

I would like to thank Guifré Vidal and Luca Tagliacozzo for stimulating discussions. This work was supported by TOQATA (Spanish grant PHY008-00784), the EU IP SIQS, the MPQ-ICFO collaboration, and NSERC and FQRNT through the network INTRIQ.

[1] F. Verstraete, V. Murg, and J. Cirac, Adv. Phys. **57**, 143 (2008).
[2] J. I. Cirac and F. Verstraete, J. Phys. A: Math. Theor. **42**, 504004 (2009).
[3] R. Orus, arXiv:1306.2164 (2013).
[4] J. Eisert, arXiv:1308.3318 (2013).
[5] M. B. Hastings, J. Stat. Mech. p. P08024 (2007).
[6] M. Wolf, F. Verstraete, M. Hastings, and J. Cirac, Phys. Rev. Lett. **100**, 070502 (2008).
[7] L. Masanes, Phys. Rev. A **80**, 052104 (2009).
[8] J. Eisert, M. Cramer, and M. B. Plenio, Rev. Mod. Phys. **82**, 277 (2010).
[9] M. Fannes, B. Nachtergaele, and R. F. Werner, Commun. Math. Phys. **144**, 443 (1992).
[10] S. Östlund and S. Rommer, Phys. Rev. Lett. **75**, 3537 (1995).
[11] G. Vidal, Phys. Rev. Lett. **91**, 147902 (2003).
[12] S. R. White, Phys. Rev. Lett **69**, 2863 (1992).
[13] K. Wilson, Rev. Mod. Phys. **47**, 773 (1975).
[14] Y. Shi, L. Duan, and G. Vidal, Phys. Rev. A **74**, 022320 (2006).
[15] L. Tagliacozzo, G. Evenbly, and G. Vidal, Phys. Rev. B **80**, 235127 (2009).
[16] G. Vidal, Phys. Rev. Lett. **99**, 220405 (2007).
[17] G. Vidal, Phys. Rev. Lett. **101**, 110501 (2008).
[18] R. N. C. Pfeifer, G. Evenbly, and G. Vidal, Phys. Rev. A **79**, 040301 (2009).
[19] E. M. Stoudenmire and S. R. White, Annual Review of Condensed Matter Physics **3**, 111 (2012).
[20] A. J. Ferris, Phys. Rev. B **87**, 125139 (2013).
[21] F. Verstraete and J. I. Cirac, eprint arXiv:cond-mat/0407066 (2004).
[22] J. Jordan, R. Orus, G. Vidal, F. Verstraete, and J. I. Cirac, Phys. Rev. Lett. **101**, 250602 (2008).
[23] Z.-C. Gu, M. Levin, and X.-G. Wen, Phys. Rev. B **78**, 205116 (2008).
[24] G. Evenbly and G. Vidal, Phys. Rev. Lett. **102**, 180406 (2009).
[25] V. Murg, F. Verstraete, and J. I. Cirac, Phys. Rev. A **75**, 033605 (2007).
[26] P. Corboz, A. M. Laeuchli, K. Penc, M. Troyer, and F. Mila, Phys. Rev. Lett. **107**, 215301 (2011).
[27] S. Yan, D. A. Huse, and S. R. White, Science **332**, 1173 (2011).
[28] K. Harada, Phys. Rev. B **86**, 184421 (2012).
[29] D. Gioev and I. Klich, Phys. Rev. Lett. 96 **96**, 100503 (2006).
[30] D. N. Sheng, O. I. Motrunich, and M. P. A. Fisher, Phys. Rev. B **79**, 205112 (2013).
[31] Y. Zhang, T. Grover, and A. Vishwanath, Phys. Rev. Lett. **107**, 067202 (2011).
[32] H.-H. Lai, K. Yang, and N. E. Bonesteel, arXiv:1306.2698 (2013).
[33] G. Evenbly and G. Vidal, arXiv:1210.1895 (2012).
[34] R. Barak and Y. Ben-Aryeh, J. Opt. Soc. Am. B **24**, 231 (2007).
[35] F. Verstraete, J. I. Cirac, and J. I. Latorre, Phys. Rev. A **79**, 032316 (2009).
[36] J. W. Cooley and J. W. Tukey, Math. Comp **19**, 297 (1965).
[37] T. Barthel, C. Pineda, and J. Eisert, Phys. Rev. A **80**, 042333 (2009).
[38] P. Corboz, G. Evenbly, F. Verstraete, and G. Vidal, Phys. Rev. A 81 **81**, 010303 (2010).
[39] P. Corboz, R. Orus, B. Bauer, and G. Vidal, Phys. Rev. B **81**, 165104 (2010).
[40] C. V. Kraus, N. Schuch, F. Verstraete, and J. I. Cirac, (preprint) arXiv:0904.4667 (2009).
[41] A. J. Ferris and G. Vidal, Phys. Rev. B **85**, 165146 (2012).
[42] G. Evenbly and G. Vidal, Phys. Rev. B **79**, 144108 (2009).
[43] J. Haegeman, J. I. Cirac, T. J. Osborne, I. Pizorn, H. Verschelde, and F. Verstraete, Phys. Rev. Lett. **107**, 070601 (2011).

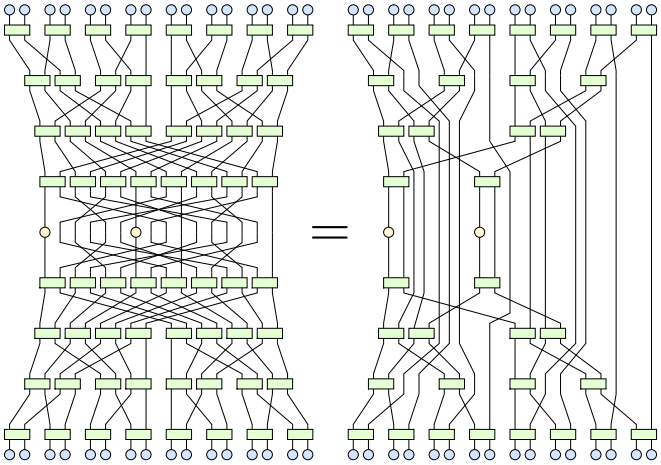

FIG. 6: (Color online) The tensor network contraction for computing the expectation of a two-body operator. Many of the tensors have cancelled due to unitarity, and the remaining can be contracted with cost $\mathcal{O}(n\chi^8)$.

APPENDIX

Here we detail the tensor contractions necessary for performing calculations with the spectral tensor network, including derivations the computational cost of measuring one- and two-site local variables and performing ground state optimizations in arbitrary dimensions. We also consider the extended case corresponding to a Fourier transform combined with a Bogoliubov transform that couples modes with opposite momentum, which should be contracted in a slightly different order to remain efficient.

Tensor contractions for one- and two-site operators, and optimization

The expectation value of few-site observables can also be calculated efficiently with cost $\mathcal{O}(n)$, although with a cost which grows exponentially with the size of the operator in question. In Fig. 6, the necessary contraction for a widely-separated, two-site operator is shown.

This diagram (and that of Fig. 4 in the main paper) can be contracted using simple steps starting from the 'outside' of the diagram and working towards the physical indices, layer-by-layer. These steps each cost $\mathcal{O}(\chi^5)$ for single-site operators and up to $\mathcal{O}(\chi^8)$ for two-site operators, and are shown in Fig. 7. For a given observable, n such steps are necessary, making the total cost $\mathcal{O}(n\chi^5)$ and $\mathcal{O}(n\chi^8)$ for one- and two-site observables, respectively.

Fortunately, when computing a range of observables many computations can be reused to reduce the total cost. For instance, in a translationally-invariant system, the entire set of two-body observables $\langle \hat{A}_0 \hat{B}_i \rangle$ can be calculated with cost $\mathcal{O}(n \log n \chi^8)$ (as was done to generate Fig. 5 in the main paper). The same scaling applies to the expectation value of

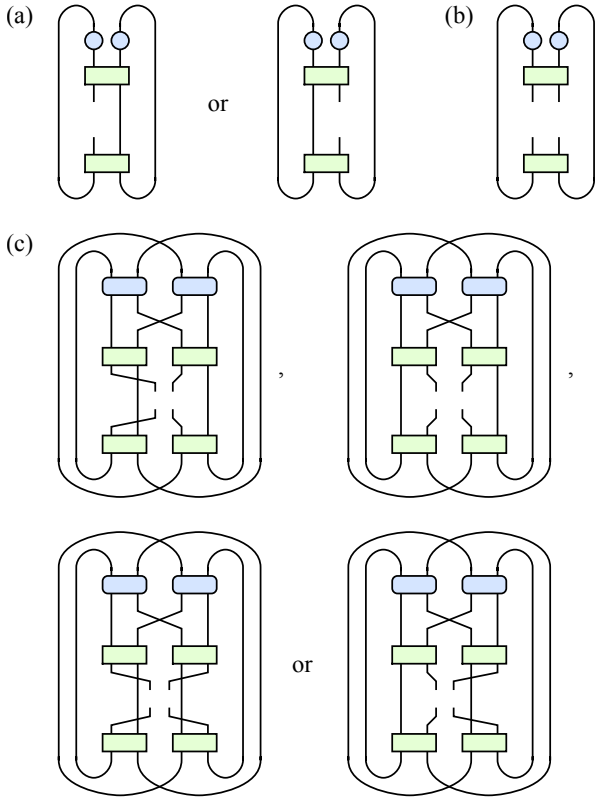

FIG. 7: (Color online) The scaling operators necessary to contract expectation values of one- and two-body operators. The open wires are arranged to make the necessary fermionic wire-crossing operations more obvious. (a) Two single-site density operators (blue) are transformed into a new single-site operator, with open indices either in the left or right 'direction' (cost $\mathcal{O}(\chi^5)$. (b) Two single-site density operators are transformed (unitarily) into a new two-site operator. (c) Two two-site density operators are transformed into a single two-site operator, with four possible descending directions (cost $\mathcal{O}(\chi^5)$.

local, two-body Hamiltonian (with no assumption on translational invariance).

With a simple modification to the above procedure, one can additionally recover the derivative of the expectation value with respect to the tensors themselves (the so-called 'environment' of each tensor) with the same $\mathcal{O}(\chi^8 n \log n)$ scaling. This allows us to follow a derivative-based optimization procedure for determining the state with minimum energy, such as steepest-descent within the unitary subspace [41], singular-value decomposition updates [42] or more complicated approaches [43].

Bogoliubov transformations

In the case of a quadratic Hamiltonian with anomalous pair terms, a Bogoliubov transformation coupling modes with opposite momentum will diagonalize a translationally-invariant system. This coupling can be added explicitly to the spectral tensor network at the 'top' (momentum space), as shown in

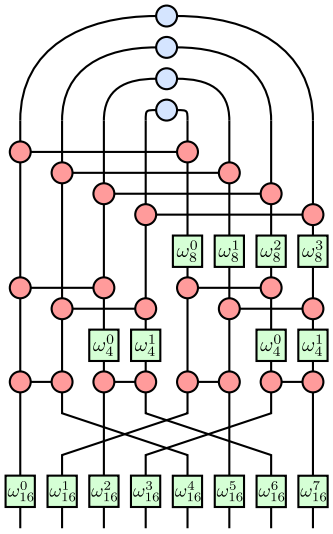

FIG. 8: (Color online) A spectral tensor network state with additional links between $\pm k$ momentum modes and 'half-integer' momenta. The eigenstates of translationally-invariant, quadratic Hamiltonians with anomalous terms can be represented with this tensor network.

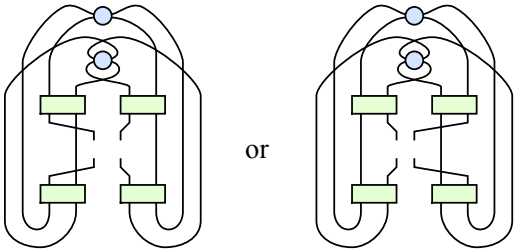

FIG. 9: (Color online) The scaling operators necessary to contract expectation values of one-site operators with the Bogoliubov transformation, each costing $\mathcal{O}(\chi^8)$.

Fig. 8.

However, contracting the tensor network naïvely in the previous order will result in growing bond-dimensions in the intermediates steps. To perform the contraction efficiently, we simply perform the layer-by-layer 'renormalization' on the coupled pairs instead of the individual sites. At each step, we combine two pairs of sites with a unitary transform and partial trace, resulting in a density matrix for a single pair of sites. We can see the steps required for calculating single-site observables in Fig. 9, costing $\mathcal{O}(\chi^8)$ each. Like we saw previously, $\mathcal{O}(n)$ such steps are required for a given observable (and just $\mathcal{O}(n \log n)$ steps for a full set on every site by reusing calculations). A similar procedure applies to the case of observables over two or more sites.

We have implemented and tested this procedure for the transverse-field Ising model. In Fig. 10 we present results for the magnetic susceptibility in the field direction, based on a simple numerical derivative evaluating

$$\chi = -\frac{d\langle \hat{Z} \rangle}{dh}. \tag{6}$$

We observe a sharp peak at the critical point at $h = 1$, as expected.

The scheme presented here might be improved on in several ways. For instance, it has not yet been investigated if the spectral tensor network can represent the Bogoliubov transformation natively, that is, without the extra tensors at the top linking $\pm k$ modes, which is the subject of future work.

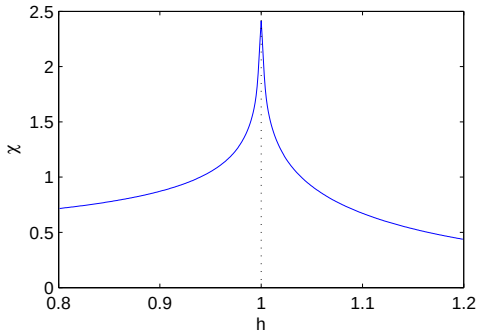

FIG. 10: (Color online) Numerical results for magnetic susceptibility the transverse-field Ising model on a lattice of 1024 sites around the critical point at $h = 1$.