

Towards high-accuracy simulations of strongly correlated materials using tensor networks

Thesis by
Matthew J. O'Rourke

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy

The logo for the California Institute of Technology (Caltech), featuring the word "Caltech" in a bold, orange, sans-serif font.

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2023
Defended January 13, 2023

© 2023

Matthew J. O'Rourke
ORCID: 0000-0002-5779-2577

All rights reserved

ACKNOWLEDGEMENTS

I would first like to thank my advisor Garnet Chan. I have known Garnet for many years, dating back to my undergraduate days when we were both at Princeton. His unrelenting and extraordinarily broad passion for science has inspired me since we first met, and will serve as an example for many years to come. Since those early days, his careful mentorship has enabled me to grow tremendously as a scientist, and also more generally as a human being. He has both supported me and challenged me in the proper doses at the proper times. Over the years we have shared many laughs, enjoyed many meals, commiserated over mutual frustrations, and endured the unending name/logo changes of Google's instant messaging service. For all of this I am exceedingly grateful.

I also want to express my gratitude to my thesis committee members: Professors Nick Hutzler, Xie Chen, and Scott Cushing. I had the great pleasure to collaborate with Nick on an interesting project regarding laser-cooled molecules that was initially outside of my comfort zone, but eventually served to broaden my skills and experience. Xie has provided me much help in understanding the subtleties involved in detecting topological order from complicated numerical calculations. And finally, Scott has saved me from a massive headache by not only agreeing to join my committee on short notice, but also serving as the chair!

All members of the Chan group over the past five years have collectively contributed an immense amount to my work and happiness by sustaining an exciting, diverse, and stimulating intellectual environment. I would like to thank in a special way Zhendong Li, with whom I worked closely as a young graduate student. His patience to answer my numerous questions, and to carefully explain the origins of my misunderstandings, helped me to mature as a researcher far more quickly than if I had to work alone. I would also like to thank my close friends and collaborators, Mario Motta and Reza Haghshenas, with whom I have enjoyed many exciting scientific discussions as well as interesting non-scientific experiences.

Furthermore, I would like to acknowledge the large majority of Chan group members who are not American and have travelled from overseas to work in the group. I have tremendously enjoyed learning from you all about the cultures, customs, cuisines, and histories of your home countries. You all have also earned my deep respect for the challenges you have overcome and sacrifices you have made in order to simply

work in the group.

To my close friends outside of the group—Lincoln, Katherine, and Conor—thank you for all the wonderful time we have spent together that has made life during graduate school balanced and fulfilling.

To my parents, thank you for your unwavering commitment to support my progress through life. Anything I accomplish is owed to your humble and gracious sacrifices over the years.

Finally, my deepest appreciation is owed to my fiancée Elena. The existence of this document reflects just a small slice of the myriad ways you manage to support and encourage me.

ABSTRACT

Accurate and verifiable computation of the properties of real materials with strong electron correlation has been a long-standing challenge in the fields of chemistry, physics, and material science. Most existing algorithms suffer from either approximations that are too inaccurate, or fundamental computational complexity that is too high. In studies of simplified models of strongly-correlated materials, tensor network algorithms have demonstrated the potential to overcome these limitations. This thesis describes our research efforts to develop new algorithms for two-dimensional (2D) tensor networks that extend their range of applicability beyond simple models and toward simulations of realistic materials.

We begin by describing three algorithms for projected entangled-pair states (PEPS, a type of 2D tensor network) that address three of their major limitations: numerical stability, long-range interactions, and computational efficiency of operators. We first describe (Ch. 2) a technique for converting a PEPS into a canonical form. By generalizing the QR matrix factorization to entire columns of a PEPS, we approximately generate a PEPS with analogous properties to the well-studied canonical 1D tensor network. This connection enables enhanced numerical stability and ground state optimization protocols. Next, we describe (Ch. 3) a technique to efficiently represent physically realistic long-range interactions between particles in a 2D tensor network operator, a projected entangled-pair operator (PEPO). We express the long-range interaction as a linear combination of correlation functions of an auxiliary system with only nearest-neighbor interactions. This allows us to represent long-range pairwise interactions with linear scaling in the system size. The third algorithm we present (Ch. 4) is a method to rewrite the 2D PEPO in terms of a set of quasi-1D tensor network operators, by exploiting intrinsic redundancies in the PEPO representation. We also report an on-the-fly contraction algorithm using these operators that allows for a significant reduction in computational complexity, enabling larger scale simulations of more complex problems.

We then move on to describe (Ch. 5) an extensive study of a “synthetic 2D material”—a two-dimensional square array of ultracold Rydberg atoms—enabled by some of the new algorithms. We investigate the ground state quantum phases of this system in the bulk and on large finite arrays directly comparable to recent quantum simulation experiments. We find a greatly altered phase diagram compared to earlier numerical and experimental studies, and in particular, we uncover

an unexpected entangled nematic phase that appears in the absence of geometric frustration.

Finally, we finish by describing (Ch. 6) a somewhat unrelated, but topically similar project in which we investigate the feasibility of laser cooling small molecules with two metal atoms to ultracold temperatures. We study in detail the properties of the molecules YbCCCa and YbCCAl for application in precision measurement experiments.

PUBLISHED CONTENT AND CONTRIBUTIONS

1. O'Rourke, M. J. & Chan, G. K. Entanglement in the quantum phases of an unfrustrated Rydberg atom array. *arXiv preprint arXiv:2201.03189* (2022).
M.J.O. designed the project, designed the algorithms, implemented the computer codes, analyzed the results, and wrote the manuscript.
2. O'Rourke, M. J. & Chan, G. K.-L. Simplified and improved approach to tensor network operators in two dimensions. *Phys. Rev. B* **101**, 205142. doi:10.1103/PhysRevB.101.205142 (20 May 2020).
M.J.O. contributed to the conception of the project, designed the algorithms, implemented the computer codes, analyzed the results, prepared the figures, and wrote the manuscript.
3. Haghshenas, R., O'Rourke, M. J. & Chan, G. K.-L. Conversion of projected entangled pair states into a canonical form. *Phys. Rev. B* **100**, 054404. doi:10.1103/PhysRevB.100.054404 (5 Aug. 2019).
M.J.O. contributed to the design of the PEPS canonical form, implemented computer codes and performed simulations for the energy optimizations, contributed to analyzing the results, and contributed to writing the manuscript.
4. O'Rourke, M. J. & Hutzler, N. R. Hypermetallic polar molecules for precision measurements. *Phys. Rev. A* **100**, 022502. doi:10.1103/PhysRevA.100.022502 (2 Aug. 2019).
M.J.O. performed the electronic structure calculations, analyzed the results, prepared the figures, and wrote the manuscript.
5. O'Rourke, M. J., Li, Z. & Chan, G. K.-L. Efficient representation of long-range interactions in tensor network algorithms. *Phys. Rev. B* **98**, 205127. doi:10.1103/PhysRevB.98.205127 (20 Nov. 2018).
M.J.O. contributed to the conception of the project, designed the algorithms, implemented the computer codes, analyzed the results, prepared the figures, and wrote the manuscript.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	v
Published Content and Contributions	vii
Table of Contents	vii
List of Illustrations	x
List of Tables	xiii
Chapter I: Introduction	1
1.1 Context and motivation	1
1.2 Tensor networks	7
1.3 Summary of research	14
Chapter II: Conversion of projected entangled pair states into a canonical form	17
2.1 Abstract	17
2.2 Introduction	17
2.3 PEPS definition and background	19
2.4 PEPS canonical form and column QR ansatz	19
2.5 Summary	31
Chapter III: Efficient representation of long-range interactions in tensor net- work algorithms	33
3.1 Abstract	33
3.2 Introduction	33
3.3 Correlation function valued PEPOs	36
3.4 CF-PEPOs and the auxiliary lattice	38
3.5 Computational cost	40
3.6 Results	42
3.7 Summary	44
Chapter IV: A simplified and improved approach to tensor network operators in two dimensions	45
4.1 Abstract	45
4.2 Introduction	45
4.3 Matrix Product Operators (MPOs)	47
4.4 PEPO expectation value via generalized MPOs	53
4.5 Results	60
4.6 Summary	76
4.7 Short Appendix	77
Chapter V: Entanglement in the quantum phases of an unfrustrated Rydberg atom array	79
5.1 Abstract	79
5.2 Introduction	79
5.3 Numerical strategy and techniques	82

5.4 Bulk phases	83
5.5 Finite phase diagram	87
5.6 Summary	90
Chapter VI: Ultracold hypermetallic polar molecules for precision measurements	92
6.1 Abstract	92
6.2 Introduction	92
6.3 Electronic Structure	95
6.4 Vibrational Structure	102
6.5 Computational Details	107
6.6 Discussion	111
6.7 Summary	113
6.8 Short Appendix	114
Bibliography	117
Appendix A: Supplementary Information: Entanglement in the quantum phases of an unfrustrated Rydberg atom array (Chapter 5)	137
A.1 Numerical methods	137
A.2 Entangled nematic phase	149
A.3 Bulk phase diagram degenerate region	155
A.4 Bulk phase transitions	155
A.5 Finite phase diagram: 15×15 and 16×16	157
A.6 Comparing to experiment: 9×9 and 13×13 lattices	159

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
2.1 Overview of the canonicalization algorithm which proceeds via a column-wise QR procedure.	20
2.2 Detailed tensor network diagrams corresponding to the steps of the QR decomposition algorithm.	22
2.3 Accuracy of the variational ansatz for QR decomposition of a single bulk column.	26
2.4 Convergence of errors in the global norm of the PEPS.	28
2.5 Enhanced numerical stability of the canonical PEPS compared to a standard PEPS.	29
2.6 Ground state optimization for various models using the PEPS canonical form.	30
3.1 Tensor networks diagrams detailing the layout of a CF-PEPO and the construction of its tensors.	36
(a)	36
(b)	36
(c)	36
3.2 Convergence of the fitting procedure for the Coulomb potential using lattice correlation functions as a basis.	41
3.3 Accuracy of energy expectation values using the CF-PEPO.	42
3.4 Ground state optimization of a long-range Heisenberg model using the CF-PEPO and an exact brute force algorithm for comparison.	43
4.1 Tensor network diagram of a generalized-MPO, compared to a typical MPO and PEPO.	48
4.2 Tensor network diagrams specifying the operator-valued elements of a general form of a typical MPO.	52
4.3 Simple example of how a gMPO contraction encodes a Hamiltonian defined on only two rows.	54
4.4 The first full iteration of the boundary gMPO energy contraction algorithm for a 5×5 PEPS.	57

4.5	Relative errors of the gMPO-based expectation value calculation compared to typical PEPO-based calculations and a brute force technique.	65
4.6	Tensor network diagrams specifying the operator-valued gMPO tensor entries for long-range Gaussian interactions.	71
4.7	The complexity of representing a Gaussian interaction potential, and the accuracy of using a basis of lattice Gaussians to fit the Coulomb potential.	74
5.1	Overview of the algorithms and numerical strategy used to resolve the physics of various square lattice Rydberg atom arrays.	81
5.2	Phase diagrams of a “bulk” Rydberg atom array under various physical assumptions.	84
5.3	Detailed physical picture of the bulk striated and nematic phases.	86
5.4	Phase diagram of the 15×15 finite Rydberg atom array and the unique orders that emerge due to the presence of boundaries.	88
5.5	Direct comparison between our numerical results and recent experimental data.	89
6.1	Electron and hole orbitals for the ground state and low lying excited states of the molecule YbCCCa.	100
6.2	Electron and hole orbitals for the ground state and low lying excited states of the molecule YbCCAl.	101
6.3	Potential energy surfaces obtained with DFT and MRCI along the Ca-C bond coordinate in YbCCCa.	109
A.1	Convergence of Γ -point DMRG with respect to both control parameters in the most complex region of the phase diagram.	138
A.2	Entanglement entropy of the ground states along various slices of the phase diagram parameter space.	139
A.3	Examples of some typical ground state optimization trajectories for long-range PEPS using different automatic differentiation schemes.	143
A.4	Phase diagrams of the finite 15×15 and 16×16 Rydberg arrays, detailing convergence.	146
A.5	An example of systematic convergence of the long-range PEPS algorithm on the 15×15 lattice for the frustrated star phase.	147
A.6	Accuracy of 2D DMRG on the 9×9 and 13×13 finite lattices.	148

A.7	Fourth-order energies from Rayleigh-Schrodinger perturbation theory for various low energy Rydberg crystals in the vicinity of the nematic phase.	150
A.8	Example of a perturbative hopping process for a single Rydberg excitation that emerges in the second-order Rayleigh-Schrodinger correction to a classical initial wavefunction.	151
A.9	Structure of the entanglement spectrum for the 1D low-energy projector model of the nematic order.	154
A.10	Details of the bulk phase diagram: the extent of the nematic phase and a region with numerically degenerate solutions.	156
A.11	Rudimentary numerical analysis of the order of the crystallization phase transition for the star and striated phases.	156
A.12	Detailed comparison between numerical and experimental phase diagrams.	158
A.13	Distinguishing the striated and square orders on the 13×13 lattice. .	159

LIST OF TABLES

<i>Number</i>	<i>Page</i>
4.1 The average computational speedups of the boundary gMPO energy expectation value algorithm over PEPO-based calculations for a representative set of 2D Hamiltonians.	61
4.2 The average computational speedups of the boundary gMPO energy expectation value algorithm over “brute force” calculations for a representative set of 2D Hamiltonians.	62
6.1 Bond lengths, bond energies, and molecular frame permanent dipole moments for the molecule YbCCCa.	97
6.2 Excitation energies and transition dipole moments for the lowest-energy electronic excitations of YbCCCa.	99
6.3 Bond lengths, bond energies, and molecular frame permanent dipole moments for the molecule YbCCAl.	99
6.4 Excitation energies and transition dipole moments for the lowest-energy electronic excitations of YbCCAl.	100
6.5 Vibration frequencies for each of the vibrational modes of the triplet ground state of both YbCCCa and YbCCAl.	103
6.6 Franck-Condon factors for the metal-centered electronic transitions of YbCCCa.	104
6.7 Franck-Condon factors for the metal-centered electronic transitions of YbCCAl.	106

Chapter 1

INTRODUCTION

This chapter will begin by giving a short but broad overview of context and motivation for the research described in subsequent chapters of this thesis. It will then proceed by giving a basic introduction to tensor networks, which are the primary technical topic of the work. Finally, the chapter will conclude by summarizing the research contributions presented in this thesis.

1.1 Context and motivation

The properties of molecules, materials, and their chemical reactions are described by the equations of quantum physics. Solving these equations allows researchers to investigate the atomic-level origin of interesting phenomena, as well as to predict and design new systems with desirable functions. Quantum systems are able to host an extraordinarily wide variety of interesting behaviors that arise due to quantum correlations. When a large number of nuclei and electrons are combined together into molecules and materials, the resultant behavior is not simply a “sum” of the behavior of individual particles. Coordinated, collective physics can emerge such as the famous quantum Hall effect [1, 2] and high-temperature superconductivity [3–5]. Judging by the frenzied pace at which new discoveries are currently being made, many collective phenomena remain to be fully understood, with other important ones likely yet to be discovered!

The basics

The many-body Schrödinger equation allows us to study all types of quantum systems and make predictions about them. It is often written in the following remarkably simple notation,

$$i\hbar \frac{d}{dt} |\Psi(t)\rangle = \hat{H} |\Psi(t)\rangle. \quad (1.1)$$

The Hamiltonian of the system, \hat{H} , encodes all the components of its total energy—typically the motion of individual particles and the interactions between each pair of them. This operator changes depending on the system one wants to study, since different molecules or materials will contain different types of atoms and different numbers of electrons. The time-dependent quantum state of the system, or wavefunction, is denoted $|\Psi(t)\rangle$. This object describes all the probabilities of

finding the system in *any* possible configuration of particles' positions and momenta.

One simple property of the above equation is that if \hat{H} itself does not explicitly depend on time (which is often the case), then quantum systems admit a set of stationary states. Such states do not change over time if they are not perturbed, and thus have fixed energies E . These states are fundamental to describing all possible behaviors of the system, and are given by the solutions to the even simpler looking time-independent Schrödinger equation,

$$\hat{H} |\Psi\rangle = E |\Psi\rangle. \quad (1.2)$$

Nevertheless, this simple notation belies an extraordinary complexity. Since $|\Psi\rangle$ encodes the probability of *all* possible configurations of the system, it is a function whose complexity grows exponentially with the number of degrees of freedom. This exponential growth is a classic example of the so-called “curse of dimensionality”, and prohibits even the world’s largest supercomputers from directly solving the equation for more than ~ 15 degrees of freedom.

In order to solve Eq. (1.2) for complex problems of practical interest, which typically contain much more than 15 electrons, a wide variety of numerical algorithms with sub-exponential scaling have been designed to obtain *approximate* solutions. Developing techniques to obtain more accurate approximations while still requiring a reasonable amount of computing resources has been an active field of research across chemistry, physics, and material science since the advent of electronic computers in the 1950s. Broadly speaking, this is the focus of this entire thesis.

Low energy physics

A related concept to the exponential complexity of the quantum state $|\Psi\rangle$ is the fact that, in general, Eq. (1.2) has an exponential number of valid solutions with different energies E_0, E_1, \dots, E_N and corresponding states $|\Psi_0\rangle, |\Psi_1\rangle, \dots, |\Psi_N\rangle$. In this thesis we are interested in the low energy properties of the systems we study. Simply, the lowest energy states describe how the system behaves when it is cooled to low temperatures. At absolute zero, the complexity mentioned above is eliminated because the system has no choice but to assume its lowest energy state, known as the ground state. As the temperature is raised, the additional energy allows the system to access other states with low energies close to the ground state. Even at room temperature, it is typically true that the large majority of solutions $|\Psi_i\rangle$ are too high in energy to contribute meaningfully to the description of the system. In this sense,

molecules and materials in an “ordinary” environment are described by their low energy physics.

It is common practice across the broad field of numerical simulation for quantum systems to focus most on studying the ground state. In part, this is because systems at finite temperature can often be described by the ground state plus some small perturbations coming from low-energy excited states. In part, it is simply because it makes the problem of solving Eq. (1.2) more tractable by eliminating one source of exponential complexity (the number of solutions). The remaining problem of determining the ground state is still exponentially challenging in general, and thus there are a variety of different approaches to approximate the solution.

Ab initio approach

One predominant approach for gleanng useful knowledge about the ground state despite our incomplete ability to solve Eq. (1.2) is the *ab initio* approach. This perspective is more common in the fields of chemistry and material science. The central idea is to make as few approximations as possible to the true Hamiltonian describing all the nuclei and electrons in the molecule or material under investigation. In other words, the goal is to keep almost all atomic-level detail in the description of the problem and then try to solve it as well as possible given the complexity of the system’s description. This results in a (non-relativistic) Hamiltonian of the form

$$\begin{aligned} \hat{H} = & - \sum_A \frac{\hbar^2}{2m_i} \nabla_i^2 - \sum_{i,A} \frac{Z_A e^2}{|\vec{r}_i - \vec{r}_A|} + \sum_{i \neq j} \frac{e^2}{2|\vec{r}_i - \vec{r}_j|} \\ & - \sum_A \frac{\hbar^2}{2M_A} \nabla_A^2 + \sum_{A \neq B} \frac{Z_A Z_B e^2}{2|\vec{r}_A - \vec{r}_B|}, \end{aligned} \quad (1.3)$$

where i and j label electrons, m is the electron mass, A and B label nuclei, M_A and Z_A are the nuclear mass and proton number, \vec{r} is the spatial coordinate of each particle, ∇^2 is the kinetic energy operator, e is the electron charge, and \hbar is the reduced Planck constant. In the field of quantum chemistry, which is primarily concerned with the electronic structure, the Hamiltonian is made slightly simpler by invoking the Born-Oppenheimer approximation [6]. The positions of the nuclei are “frozen” and the solution becomes a function of a fixed set of nuclear coordinates $\{r_A\}$,

$$\hat{H}(\{r_A\}) = - \sum_A \frac{\hbar^2}{2m_i} \nabla_i^2 - \sum_{i,A} \frac{Z_A e^2}{|\vec{r}_i - \vec{r}_A|} + \sum_{i \neq j} \frac{e^2}{2|\vec{r}_i - \vec{r}_j|}. \quad (1.4)$$

These descriptions are often very faithful to the details of the real matter that one

wants to study, and accurate solutions yield results that are in direct agreement with appropriate experimental measurements. However the complexity and generality of this description often makes it difficult to solve with high accuracy. For this reason, a large hierarchy of algorithms has been developed which have a gradual set of trade-offs between accuracy and computational efficiency. For a fixed budget of computational resources, smaller systems with few atoms can be studied with highly accurate methods, while larger systems require methods that are more efficient, but which contain significantly more approximations.

One widely-used set of techniques is density functional theory (DFT) [7, 8]. This framework maps the many-body problem into an auxiliary system of non-interacting electrons, which instead interact with an average potential based on the position of all other electrons. Specifically, this potential is generated by an exchange-correlation functional of the 3-dimensional electron density. By mapping the interacting problem to a non-interacting problem, standard DFT algorithms scale with the number of electrons n_e as $O(n_e^3)$, a dramatic reduction from the general $O(2^{n_e})$ complexity. This efficiency makes DFT very popular for studying large systems. However, this reduction of complexity stems from the approximate nature of the exchange-correlation functional, which accounts for all the many-body effects. Although an extraordinary number of different functionals have been proposed for studying different types of systems, there is no rigorous and systematic way to improve the accuracy of a given functional for *all* systems [9]. This generically prevents DFT algorithms from providing verifiable, high-accuracy quantitative solutions for the ground state.

Another common set of techniques, often referred to as “wavefunction theory”, forms a more rigorous hierarchy. The most efficient, and most approximate, algorithm within this set is Hartree-Fock theory, which represents the wavefunction in a compact form known as a Slater determinant [10]. This representation of the wavefunction contains no quantum correlations between electrons. More accurate techniques are then built on top of this simple wavefunction by systematically adding in an increasing amount of quantum correlations using a variety of strategies. As more correlations are built in, the wavefunction can represent a larger class of possible ground states, but the scaling of the computations also increases. The most accurate method of this hierarchy that is commonly used in realistic calculations is coupled cluster theory [11]. While Hartree-Fock theory scales with the number of electrons n_e as $O(n_e^4)$, common versions of coupled cluster scale as $O(n_e^6) - O(n_e^8)$.

This inherent trade-off between efficiency and accuracy makes solving the *ab initio* problem with high accuracy for large systems a grand challenge of the field. In the case of studying materials this becomes an important issue because many materials contain numerous heavy atoms in a unit cell, each of which has many electrons¹. If the electrons in the ground state are strongly correlated, obtaining accurate solutions becomes extremely challenging. Unfortunately, many of the most interesting collective quantum phenomena are stabilized by the presence of these correlations. This challenge has recently spurred significant renewed research efforts [12–15].

The “big picture” motivation for the research presented in the subsequent chapters of this thesis is to develop new algorithms that are able to overcome this limitation in the study of materials. To do so, we seek to adapt high-accuracy algorithms that are utilized in a different paradigm of material simulation. These algorithms are based on tensor networks, which are systematically improvable like the wavefunction methods mentioned above, but do not suffer from the concomitant penalty of an increased computational complexity². They are typically used in the context of solving simplified models. We will conclude Section 1.1 by briefly discussing the simplified models paradigm of material simulation, and then in Section 1.2 turn to a more thorough introduction of tensor networks, their associated advantages over the typical *ab initio* methods, and their limitations that have so far prevented them from studying realistic descriptions of materials.

Simplified models approach

Another major approach for gleaning useful knowledge about the ground state of a system despite our incomplete ability to solve Eq. (1.2) is based on simplified models. This approach is more common in the field of condensed matter physics, and is primarily used for studying materials. This is slightly different than the *ab initio* approach, which is regularly used for both molecules and materials. The

¹It is important to note that simply having many electrons does not intrinsically preclude a material from being “properly” studied with a purely *ab initio* approach. Many materials (such as band insulators and basic metals) have very simple electronic properties which contain little or no electron correlation. This allows highly approximate algorithms such as DFT to nevertheless obtain an accurate solution.

²Of course, the computational *effort* still increases for an increased accuracy, but the scaling relation does not change.

primary goal of using simplified models to study materials is twofold: (i) to write down a Hamiltonian that can qualitatively reproduce a specific material phenomenon of interest, while (ii) simultaneously being simple enough to solve exactly or to very high accuracy, even for large systems. To this end, the most successful model Hamiltonians include only the smallest number of physical details necessary to generate a ground state that stabilizes a particular phenomenon.

For a concrete comparison to Eq. (1.3), a very famous model system for studying Mott insulators [16, 17] is the Hubbard model [18],

$$\hat{H} = -t \sum_{\langle i,j \rangle, \sigma} \left(\hat{a}_{i,\sigma}^\dagger \hat{a}_{j,\sigma} + \hat{a}_{j,\sigma}^\dagger \hat{a}_{i,\sigma} \right) + U \sum_i \hat{n}_{i,\uparrow} \hat{n}_{i,\downarrow}. \quad (1.5)$$

Here, i labels a site on a regular lattice (typically a d -dimensional cubic lattice), $\langle i, j \rangle$ denotes pairs of nearest neighbor sites, $\hat{a}^{(\dagger)}$ is a fermionic annihilation (creation) operator, σ labels fermion spin, and \hat{n}_i is the density operator on site i . Qualitatively, this model describes fermions “hopping” around freely on a cubic lattice, only interacting with each other (with strength U) if two fermions occupy the same site. This physical picture is drastically simpler than the one described in Eq. (1.3). Perhaps remarkably, by tuning the relative strengths of the parameters t and U , the ground state of this simple model displays the essential behavior of a Mott insulator, which emerges in much more complex settings in real materials such as transition metal oxides.

In general, this approach of using simplified models will not yield quantitative results directly comparable to experiments on real materials. The purpose of this sacrifice is to enable the qualitative study of a wide variety of complex collective phenomena. For example, there are an extensive number of studies on topics such as frustrated magnetism, superconductivity, and topological order. However, as discussed earlier, the ground states that support such phenomena contain a significant amount quantum correlations, or to be more precise, quantum entanglement. In order to regularly investigate these types of ground states, high-accuracy algorithms are essential. These requirements have led to the development of a variety of algorithms that have much different advantages and limitations than those described in the *ab initio* section. Some of the most successful algorithms in terms of accuracy, efficiency, and flexibility (to study a variety of different models) are based on tensor networks.

Although tensor network techniques are predominantly used to study models, they

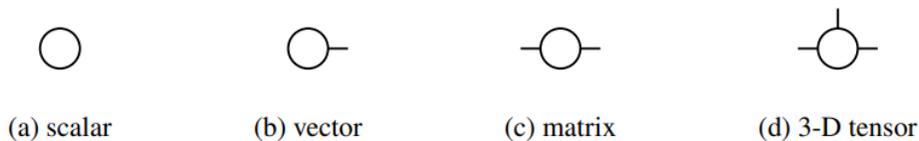
are quite versatile. Much effort has been invested into adapting them to the *ab initio* language for molecules, and it has resulted in great success. They have been used to successfully study some of the most well-known, challenging problems in quantum chemistry, including benzene polymorphs [19], the P-cluster of nitrogenase [20], and the chromium dimer [21]. However, the analogous application of tensor network methods to materials has not yet been realized, due to a lack of appropriate algorithms. In the following section, we will review some basics of tensor networks and describe the origin of these algorithmic limitations. Then, in the final section of this chapter we will provide an overview of the research described in the rest of the thesis, which is aimed at developing new algorithms that will **eventually** enable tensor network simulation of realistic materials. For now, the advances described in this thesis allow for the successful high-accuracy simulation of Hamiltonians with an intermediate level of complexity between simple models and full *ab initio*.

1.2 Tensor networks

This section aims to review some of the important concepts of tensor networks that the following chapters assume as basic knowledge of the reader. The expedited presentation of this extensive topic will likely not contain enough details for complete beginners, whom we refer to Refs. [22, 23] for much more detail.

The basics

The term “tensor network” refers to a broad set of mathematical objects and their associated numerical algorithms. The basic element, a tensor, is defined in this field simply as a multi-dimensional array. Common objects in linear algebra such as matrices, vectors, and scalars are 2, 1, and 0-dimensional tensors, respectively. We will quickly adopt a graphical diagrammatic notation which greatly simplifies the forthcoming algebraic expressions. To represent arbitrary n -dimensional tensors, we will draw them as shapes with n “legs” sticking out. Each leg represents a single axis of the tensor, and the size of this axis is called the bond dimension.



Connecting two legs represents a tensor contraction operation. Some simple instances of this operation include the matrix-vector product and matrix-matrix mul-

tiplication, however it can also be a completely general operation between higher-dimensional tensors.

$$b_i = \sum_j A_{ij}x_j \quad b_{ik} = \sum_j A_{ij}B_{jk} \quad T_{abjk} = \sum_i A_{abi}B_{ijk}$$

(a) matrix-vector mult. (b) matrix-matrix mult. (c) general contraction

A tensor network (TN) can now be simply defined as any diagram that consists of multiple tensors with connected legs, with no completely disconnected sets of tensors. The result of the specified contraction can be arbitrary; it could represent a scalar or an n -dimensional tensor. Although this is very general, there are some specific structures of tensor networks that have important properties for the representation of quantum states, which we will discuss in a moment.

Another basic operation that is very important is matrix decomposition. Specifically, let us consider the singular value decomposition (SVD) of a matrix $T_{ab} = \sum_i U_{ai}\Sigma_{ii}V_{ib}$. This too can be generalized into a tensor form by “vectorizing” multiple indices into the left-hand side and all the other indices into the right-hand side³. For example, the indices of the 4-dimensional tensor T_{wxyz} can be re-grouped as $T_{(wx),(yz)} = T_{ab}$, where $\dim(a) = \dim(w) \times \dim(x)$ and $\dim(b) = \dim(y) \times \dim(z)$. After performing an SVD on the newly defined matrix, the indices a and b can be re-factored like $\sum_i U_{ai}\Sigma_{ii}V_{ib} = \sum_i U_{wxi}\Sigma_{ii}V_{iyz} = T_{wxyz}$.

(a) matrix SVD

(b) tensor SVD

³To avoid any confusion, there is a different operation known as “higher order SVD” that can also be performed on generic tensors. We are not discussing that operation, and it is not used in this thesis.

Matrix product states

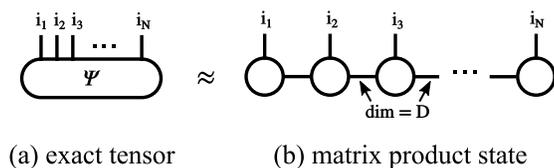
A general many-body quantum state for N degrees of freedom can be written, in second-quantized form, as a superposition of all possible basis states,

$$|\Psi\rangle = \sum_{i_1, \dots, i_N} \Psi_{i_1, \dots, i_N} |i_1, \dots, i_N\rangle, \quad (1.6)$$

where the coefficient for each basis state is contained in the N -dimensional tensor Ψ . If each degree of freedom $|i_k\rangle$ has dimensionality d , then Ψ contains d^N elements. This reflects the exponential scaling of wavefunction that makes simulation of quantum systems challenging, as we discussed more generally in previous sections.

To design an algorithm that can perform efficient computations, one common strategy is to define an *ansatz* for Ψ which contains a much smaller number of parameters. With this more compact form of wavefunction, the objective is then to optimize its parameters to reproduce a target state of interest with as much accuracy as possible given the constraints of the ansatz. When trying to study the ground state of a system, the variational principle of quantum mechanics guarantees that this optimization is specifically a minimization of the energy; the state with minimal energy is the ground state.

A good ansatz for Ψ should be chosen on the basis of physical intuition so that, as the parameter space for coefficients gets restricted, the excluded states are those that are most irrelevant for describing the target state. A very general and powerful ansatz is the matrix product state (MPS) tensor network.



If all the connected legs in the MPS are restricted to a constant bond dimension D , the number of unique parameters in the ansatz is just linear in the number of degrees of freedom, $O(NdD^2)$.

The physical justification for this considerable simplification comes from one-dimensional (1D) simplified models. In this context it has been proven that for the ground state of any Hamiltonian containing only localized interactions, away

from criticality, the entanglement entropy at any bipartition of the system grows proportionally to the size of the boundary of the partition [24]. This important result is known as the “area law of entanglement” for ground states. In one dimension, the boundary of a bipartition is simply a point, which means that the size of the boundary remains constant regardless of the size of either partitioned component of the system. As a consequence we have a remarkable property, that the entanglement between any two partitions is bounded by a constant.

This property is central to the success of tensor networks as a variational ansatz for ground states, since it is built in by construction. Consider a general bipartition of the degrees of freedom in the exact tensor $\Psi_{i_1, \dots, i_k, \dots, i_N} = \Psi_{(i_1, \dots, i_k), (i_{k+1}, \dots, i_N)}$. Performing a tensorized SVD according to this grouping of indices gives us the Schmidt decomposition of the quantum state

$$|\Psi\rangle = \sum_{i_1, \dots, i_N} \sum_{\alpha} U_{(i_1, \dots, i_k), (\alpha)} \Sigma_{\alpha} V_{(\alpha), (i_{k+1}, \dots, i_N)} |i_1, \dots, i_k\rangle \otimes |i_{k+1}, \dots, i_N\rangle, \quad (1.7)$$

where Σ_{α} are the Schmidt values. The von Neumann entanglement entropy between these two parts of the system is given by $S = -\sum_{\alpha} \Sigma_{\alpha}^2 \log \Sigma_{\alpha}^2$. For one-dimensional models, the area law tells us that this quantity is a constant, which means that the number of non-zero entries in Σ_{α} will be constant despite the clear mathematical possibility for it to grow exponentially with the total size of the smaller partition.

If we now look back at the structure of the MPS with constant bond dimension D , we can imagine choosing any point k along the 1D chain and contracting all the tensors together on either side of the partition, leaving only the bond at position k (which connects the partitions) uncontracted. This diagram would be identical to the tensorized SVD of the original exact state, with the Schmidt values on the bond in between the left and right partitions. Above, we just detailed how we only need a constant number of Schmidt values at this partition to accurately represent the 1D ground state, which corresponds exactly to the constant bond dimension D in the MPS. Returning to the original MPS diagram, this analysis can be repeated successfully for any partition point k , revealing that the bond dimension D MPS can represent nearly all 1D ground states with extremely high accuracy.

An analogous 1D tensor network structure exists for efficiently representing operators, such as the Hamiltonian. This is known as a matrix product operator (MPO), and its basic properties are discussed in detail in Chapter 4. These efficient and highly accurate representations of both states and operators can be used to formulate

an efficient variational minimization algorithm over the manifold of all matrix product states, which is known as the density matrix renormalization group (DMRG) [22, 25, 26]. It has been used with extraordinary success to solve for the ground states of nearly any imaginable 1D quantum system [27]. Importantly, it scales linearly with the number of degrees of freedom in the system, so it can be used to study very large systems and even infinitely-sized systems [28] if there is some degree of translational invariance.

Aside from its efficient representability of 1D ground states, which we discussed above, another key property of the MPS that enables the success of DMRG is that the MPS can be transformed into a “canonical form”. Specifically, this form allows for efficient computational operations⁴ and numerically stable optimization. The canonical form makes use of the inherent gauge freedom present in a MPS, which can be revealed by inserting a pair of matrices XX^{-1} on any bond of the MPS. Although these matrices multiply to form an identity, which leaves the full state invariant, the individual matrices can be absorbed into the adjacent tensors with X being contracted to the left and X^{-1} to the right. The result is an MPS with two new tensors

$$\tilde{L}_{\alpha_{k-1}, i_k, \gamma} = \sum_{\alpha_k} L_{\alpha_{k-1}, i_k, \alpha_k} X_{\alpha_k, \gamma} \quad (1.8)$$

$$\tilde{R}_{\gamma, i_{k+1}, \alpha_{k+1}} = \sum_{\alpha_k} X_{\gamma, \alpha_k}^{-1} R_{\alpha_k, i_{k+1}, \alpha_{k+1}}, \quad (1.9)$$

where L and R were the original tensors at positions k and $k + 1$ in the MPS, respectively. We now have a different MPS that represents the exact same quantum state. We can exploit this freedom to change the MPS representation of a given state to one with useful properties. The canonical form refers to the form in which either one or both ends of the MPS $|\Psi_{MPS}\rangle$ reduce to identity matrices when contracted with the conjugate state $\langle\Psi_{MPS}|$.

Importantly, there will always be at least one tensor which does not contract to identity, which is called the “canonical center” and contains all the non-unitary components of the global quantum state. As a result, this can make large portions of the tensor network entirely trivial to contract when evaluating certain important

⁴It is important to note the distinction between efficient representation and efficient computational operations in tensor networks. The representation is a memory requirement, but just because a given tensor network provides a memory-efficient representation of a state does not necessarily guarantee that it also provides efficient computational operations on the state. This will be important in the upcoming discussion of 2D PEPS.



quantities such as expectation values $\langle \Psi_{MPS} | \hat{O} | \Psi_{MPS} \rangle$. In addition to the clear computational efficiency, this also dramatically enhances the numerical stability of optimizations because “local” changes to a single tensor at the canonical center can be chosen in a globally optimal way for the entire state. This is discussed in detail in Chapter 2.

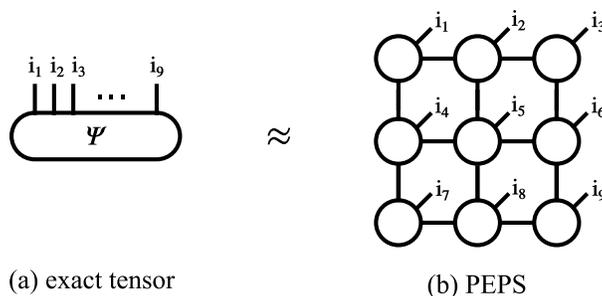
As a final point about MPS and DMRG, we will comment on their utility beyond 1D models. Strictly, outside of ground states in 1D we do not expect the 1D area law to remain true. This means we should no longer expect the required bond dimension for high-accuracy solutions to remain constant as the size of the system grows. However, the area law is only an *asymptotic* statement that describes the expected behavior of entanglement as the system size becomes very large. The ground states of systems that are at some fixed finite size and in an arbitrary geometry will have different, unknown amounts of entanglement between different bipartitions of the degrees of freedom. Whether or not an MPS can accurately represent this type of arbitrary, finite ground state is a purely quantitative question that can only be answered by performing numerical simulations.

The low computational scaling of DMRG with respect to the bond dimension D (typically $O(D^3)$) allows for simulations with large D of the size $O(10000)$ using the most advanced codes. In practice, this has enabled many successful, high-accuracy studies of systems beyond 1D models. One prominent example is molecules in the *ab initio* picture, since the complete description of a single molecule is finite-sized. In particular, DMRG has been most useful in addressing outstanding challenges in molecular simulation related to molecular states containing a large amount of quantum correlations [20, 21, 29]. Unfortunately, extension of molecular *ab initio* DMRG to materials is prohibitively expensive in most cases due to both the problem size and geometry of entanglement. Another important example where DMRG is successful is 2D lattice models with a finite width of ~ 10 lattice sites [30]. There are a large number of studies (including Chapter 5) that have employed DMRG in this

way to gain insight into the rich physics present in many 2D models [27]. However, the scaling of entanglement in 2D ground states requires an exponential growth in the MPS bond dimension as the width of the system increases, so larger systems cannot be studied accurately. To overcome this scaling mismatch, a different tensor network structure is required for the wavefunction ansatz.

Projected entangled-pair states

A projected entangled-pair state is an extension of the ideas of MPS to two (or more) spatial dimensions. In this type of ansatz, the associated tensor network forms a grid of tensors, connected in a square lattice along both the x- and y-axes.



The structure of the PEPS network is designed to satisfy the area law of entanglement for 2D ground states [31, 32]. In this case, the entanglement entropy is still expected to scale proportionally to the boundary between the two partitioned components of the system. While in 1D that scaling relation was a constant, in 2D the length of the boundary grows linearly as the size of both the x- and y-dimensions increase linearly. A constant bond dimension PEPS captures this scaling relation by construction.

Analogous to the MPS, a constant D PEPS can represent 2D ground states with high accuracy while also achieving an exponential reduction of the parameter space compared to the exact wavefunction tensor, from $O(d^N)$ to $O(NdD^4)$. However, the more complex network structure of the PEPS limits the applicability of other analogies to the properties of MPS. Most importantly, the diagrams corresponding to expectation values $\langle \Psi_{PEPS} | \hat{O} | \Psi_{PEPS} \rangle$ and even normalization $\langle \Psi_{PEPS} | \Psi_{PEPS} \rangle$ cannot be contracted exactly in polynomial computing time [33]. A related consequence is that the PEPS lacks a canonical form equivalent to that of the MPS.

Despite these challenges, numerous systematically controllable, approximate contraction algorithms have been developed [34–37]. Additionally, new forms of opti-

mization algorithms that do not rely on the canonical form have been designed⁵ [36, 38–40]. As software libraries have continuously improved, high-accuracy PEPS simulations of the ground state of 2D models on large, finite lattices as well as infinite lattices (with translation invariance) have become common and quite successful [41–44].

At the time that the research in this thesis began (2017), the success of ground state simulations using PEPS was nevertheless strictly limited to the domain of highly localized lattice models. For the purposes of computational efficiency and conceptual simplification, all of the approximate contraction and optimization algorithms that had been well-studied were explicitly designed for Hamiltonians with at most next-nearest neighbor interactions. Directly generalizing these techniques to make them capable of handling more complicated Hamiltonians rendered them inefficient or inaccurate; and the practical use of general projected entangled-pair operators (PEPOs), the 2D generalization of MPOs, was quite scarce in the literature [45, 46]. These algorithmic limitations broadly prohibited the use of PEPS to study the ground states of more complex problems, although the PEPS ansatz itself faces no fundamental barriers to the task of representing these states.

1.3 Summary of research

In order to move beyond the study of highly local model Hamiltonians, and toward the simulation of realistic materials with PEPS, a number of different limitations are addressed in this thesis with the development of new algorithms. One persistent challenge with PEPS is a lack of numerical stability, which fundamentally originates from the inability to perform exact contraction. In the study of local Hamiltonians, a variety of *ad hoc* tricks are often employed to get the ground state optimization procedure to converge stably [36–38]. Since we anticipated the need to develop qualitatively different algorithms for studying more complex Hamiltonians, we wanted to try to address the problems with numerical stability in a more fundamental way that might be applicable to a broader class of algorithms.

In Chapter 2 we describe these efforts. We present a technique for approximately converting a PEPS into a canonical form that is directly analogous to the MPS canonical form. In MPS, a QR matrix factorization is commonly used to shift the canonical center from one site to the next. We generalize the standard QR to the factorization of an entire PEPS column by using a shallow unitary circuit to

⁵The details of some of these complex algorithms are discussed throughout the following chapters of this thesis.

represent Q . Even though this representation only approximates the exact Q , the orthogonality condition is exactly preserved. This allows for exact enforcement of the useful MPS-style canonical property that the left and right sides of the PEPS exactly contract to identity operators when combined with their conjugate state. We are able to show that this form of the PEPS is substantially more numerically stable than a standard PEPS by testing the sensitivity of various contractions to numerical perturbations of the individual tensors. We additionally show that the canonical form of the PEPS can be used to implement ground state optimization for local Hamiltonians based on imaginary time evolution.

In Chapters 3 and 4, a considerable amount of research efforts were invested into the following question: how to best represent and operate with complex Hamiltonians in 2D? Since previously the focus of the literature was overwhelming on local Hamiltonians, this question had been almost entirely unexplored. Importantly, to make significant progress towards the ultimate goal of studying “realistic matter”, we want to efficiently include long-range pairwise interactions between particles, such as the Coulomb potential that appears in the electronic structure Hamiltonian.

In Chapter 3, we detail an approach to efficiently represent long-range interaction potentials in real space using PEPOs. By expressing the interaction approximately as a linear combination of correlation functions from a simple Ising model auxiliary system, we are able to keep the bond dimension of the PEPO constant despite the quadratic (with system size) number of terms in Hamiltonian. This constant bond dimension results in a representation of the long-range Hamiltonian which has a number of unique parameters that scales just linearly with system size.

In Chapter 4, we confront the problem that contracting the expectation value of a PEPO with a PEPS is very computationally expensive. This operation is critical to any possible ground state optimization protocol, but even when using a PEPO with constant bond dimension D_p , the cost of approximate contraction for the expectation value increases by a factor of $O(D_p^7)$ compared to algorithms for a local Hamiltonian. To overcome this problem, we develop an approach to rewrite a PEPO in terms of a set of quasi-1D tensor network operators by exploiting redundancies intrinsic to the PEPO representation. We also develop a new on-the-fly algorithm for contraction of the expectation value using these quasi-1D operators, which dramatically reduces the cost factor to $O(D_p^3)$.

These advances allowed us to perform large-scale ground state simulations of long-range interacting Hamiltonians using PEPS for the first time. In Chapter 5, we

describe our extensive investigation of the ground state quantum phases of a “synthetic 2D material,” namely a 2D square array of ultracold Rydberg atoms. The bosonic degrees of freedom in this system interact with a long-range van der Waals potential. We study the ground state phase diagram of this system in the bulk and on large finite arrays directly comparable to recent quantum simulation experiments. We find that our high-accuracy treatment of the long-range interactions yields a greatly altered phase diagram compared to earlier numerical results with truncated interactions. In particular, we discover an unexpected entangled nematic phase. We also find that a reinterpretation of the previously reported experimental phase diagram is required based on our new results.

Finally, in Chapter 6, we describe a somewhat different research project that is unrelated to tensor networks, but rather related in spirit to the cold atom system from the previous chapter. Here, we investigate the feasibility of laser cooling metal atoms that are not typically amenable to cooling, but are of interest for other applications in precisely controlled quantum science experiments. To do so, we combine such a metal with a different, coolable metal atom in a well-designed small molecule. For some prototypical molecules of this type (YbCCCa and YbCCAl) we find that it may be realistically possible to cool the entire molecule by cycling photons only on Ca or Al. This suggests that the hypermetallic approach may be a new, versatile tool for experimental control of metal species that do not otherwise efficiently cycle photons.

*Chapter 2***CONVERSION OF PROJECTED ENTANGLED PAIR STATES
INTO A CANONICAL FORM**

This chapter is based on the following publication:

1. Haghshenas, R., O'Rourke, M. J. & Chan, G. K.-L. Conversion of projected entangled pair states into a canonical form. *Phys. Rev. B* **100**, 054404. doi:10.1103/PhysRevB.100.054404 (5 Aug. 2019).

2.1 Abstract

We propose an algorithm to convert a projected entangled pair state (PEPS) into a canonical form, analogous to the well-known canonical form of a matrix product state. Our approach is based on a variational gauging ansatz for the QR tensor decomposition of PEPS columns into a matrix product operator and a finite depth circuit of unitaries and isometries. We describe a practical initialization scheme that leads to rapid convergence in the QR optimization. We explore the performance and stability of the variational gauging algorithm in norm calculations for the transverse-field Ising and Heisenberg models on a square lattice. We also demonstrate energy optimization within the PEPS canonical form for the transverse-field Ising and Heisenberg models. We expect this canonical form to open up improved analytical and numerical approaches for PEPS.

2.2 Introduction

Tensor network states (TNS) [25, 26, 47–49] are widely used as variational wave functions to approximate low-energy states of quantum many-body systems [23, 50]. Their power arises from their ability to efficiently capture global behaviors of quantum correlations in the system, as described by entanglement area laws [51, 52]. As a consequence, the global wave function is encoded in local tensors with finite bond dimension. A concrete example is the matrix product state (MPS) [22, 53–55], a class of tensor-network states that capture the area law in 1D, and which underlie the success of the density matrix renormalization group (DMRG) [25, 26].

The local tensors in a TNS are not uniquely defined and contain redundant parameters known as a local gauge. In MPS, such gauges can be fixed by bringing the MPS into a canonical form where all tensors but one are isometric [22]. The canonical form

is simple to compute through QR decompositions, and has many applications, such as in defining optimal local truncations [22, 56], the DMRG algorithm, constructing the tangent space of excitations [57–60], and providing a framework to characterize phases [61, 62].

Projected entangled pair states (PEPS) [23, 63–65] are higher-dimensional generalizations of MPS with analogous area laws. The PEPS has widely been used as a variational ansatz to explore physical properties of quantum many-body systems [41, 42, 66–71]. It has already been observed that partially fixing the gauge of local tensors can dramatically improve the efficiency and stability of PEPS algorithms [36, 38, 72, 73]. However, unlike in MPS, computing a fully canonical form for a PEPS remains a challenge.

Here, we introduce a gauging variational ansatz that efficiently brings a PEPS wave function into a full canonical form in direct analogy with that of an MPS. To do so, we re-express the columns of the PEPS as a QR tensor product, where Q is an isometric column tensor network and R is a matrix product operator (MPO). We show that Q can be compactly parametrized by a finite-depth circuit of block isometries and unitaries that can be determined by variational optimization. After transforming all columns but one (a central column) to be isometric, we obtain the (column) canonical form of the PEPS, where part of the entanglement in the environment is transferred to the central column. We explore the stability and performance of the QR decomposition and PEPS canonical representation in calculating the norm in the 2D transverse-field Ising (ITF) and Heisenberg models on a square lattice. We also analyze the behavior of imaginary-time energy optimization in the canonical PEPS form in the context of the ground-state of the ITF and Heisenberg models.

The paper is organized as follows. The basic concepts of the PEPS ansatz are introduced in Sec. 2.3. We first discuss the canonicalization procedure in the context of MPS in Sec. 2.4 as a basis to describe our approach to addressing this problem for PEPS. We then study the cost, accuracy, and stability of our gauging variational ansatz in calculations on the ITF and Heisenberg models. We also compare the results of direct energy optimization in the canonical form to results from standard PEPS optimization algorithms. Finally, we summarize our findings in Sec. 2.5 and discuss future research directions.

2.3 PEPS definition and background

A PEPS is a TNS defined by a set of local tensors $\{A_i^{s_i}\}$ connected by virtual bonds along the grid of the physical lattice. The bond dimension of the virtual bonds is denoted D , which controls the number of parameters (or, more physically, the amount of entanglement in the wavefunction) and hence the accuracy of the ansatz. The physical indices s_i encode the local physical Hilbert space of dimension d . A PEPS wave function $|\Psi\rangle$ on the $l_x \times l_y = 4 \times 4$ square lattice with open boundary conditions is depicted in Fig. 2.1(a),

$$|\Psi\rangle = \sum_{\{s_i\}} \mathcal{F}(A_1^{s_1}, A_2^{s_2}, \dots, A_{l_x \times l_y}^{s_{l_x \times l_y}}) |s_1, s_2 \dots, s_{l_x \times l_y}\rangle, \quad (2.1)$$

where \mathcal{F} denotes tensor contraction of the virtual bonds. The tensors are all colored differently to indicate that we do not assume translational invariance in the tensor network.

The tensor contraction in Eq. (2.1) is invariant under insertion of a gauge matrix and its inverse G , G^{-1} between two tensors (along with a virtual bond). In an MPS, the canonical form at site i is defined as the choice of gauges such that the environment tensor \mathcal{G}_i (constructed by partial norm-contraction over all sites except i) is the identity tensor

$$\mathcal{G}_i = \mathcal{F}\left(\prod_{j \neq i} E_j\right) = \mathbb{1}, \quad (2.2)$$

with $E_j = \sum_{s_j} A_j^{s_j \dagger} A_j^{s_j}$; $\mathbb{1}$ denotes the tensor $\delta_{i_1 i'_1} \delta_{i_2 i'_2} \dots$ where $i_1 i_2 \dots, i'_1 i'_2 \dots$ index the virtual bonds of $A_i^{s_i \dagger}, A_i^{s_i}$ respectively. By ensuring that the PEPS tensors satisfy Eq. (2.2), it also defines an analogous canonical form for a PEPS, depicted in Fig. 2.1(b). In the case of an MPS, we can convert an arbitrary MPS into canonical form by sequential QR (LQ) decompositions of tensors to the left (right) of site i , $A_j^{s_j} \rightarrow Q_j^{s_j} R_j (A_j^{s_j} \rightarrow L_j Q_j^{s_j})$ where Q_j is orthogonal in the sense $\sum_{s_j} Q_j^{s_j \dagger} Q_j^{s_j} = \mathbb{1}$ (for LQ, $\sum_{s_j} Q_j^{s_j} Q_j^{s_j \dagger} = \mathbb{1}$). For simplicity, we henceforth do not distinguish between QR and LQ, with the choice implicit from the diagrammatic representation. R_j is then absorbed into the adjacent tensor for the subsequent QR decomposition until the full canonical form is reached [22].

2.4 PEPS canonical form and column QR ansatz

To similarly canonicalize a PEPS, we sequentially decompose the PEPS columns, denoted M (composed of tensors $\{m_i\}$), as a QR tensor contraction, where the

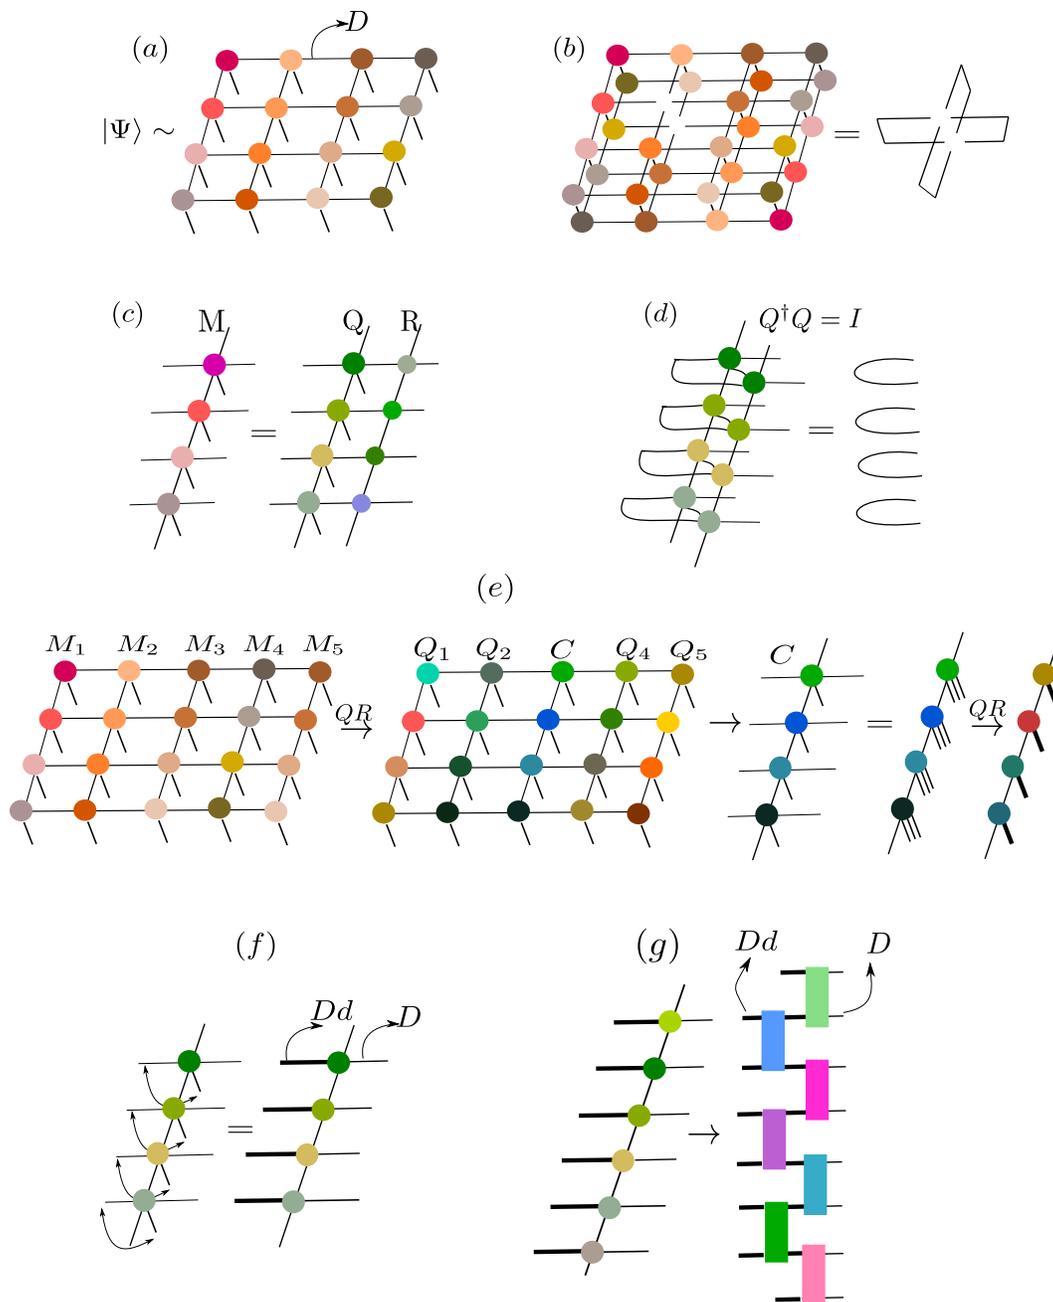


Figure 2.1: (a) Tensor network diagram of the PEPS $|\Psi\rangle$ on a 4×4 square lattice with open boundary conditions. Different colors are used for each tensor to explicitly indicate a non-translationally invariant tensor network. (b) A PEPS canonicalized around a single site, showing that the environment around that site contracts to the identity tensor. (c) A decomposition of a bulk column $M \approx QR$, where (d) the tensor network Q is isometric, i.e. $Q^\dagger Q = \mathbb{1}$. (e) A graphical representation of the steps based on our QR scheme to bring a PEPS into canonical form. Note in the final step, MPS canonicalization is used on the central column C . (f) The tensor network Q is reshaped into an MPO by fusing the virtual bond and physical bond, as shown by the arrows. The thick virtual bonds have bond dimension Dd . (g) The isometric tensor network Q is parameterized by a finite depth circuit of l -site isometries/unitaries $\{u_i\}$.

column tensor network Q is isometric, satisfying $Q^\dagger Q = \mathbb{1}$, see Fig. 2.1(c, d). The gauge column tensor network R (composed of tensors $\{r_i\}$) is an MPO acting on the horizontal virtual bonds. Once all columns (around a central column) are decomposed to be isometric, the central column C can be viewed as an MPS by grouping the horizontal bonds with the physical bonds. This central column can then be canonicalized around a chosen site using the MPS canonicalization algorithm above, to yield a complete PEPS canonicalization (Fig. 2.1(e)). Note that the PEPS canonicalization condition around a site (Fig. 2.1(b)) does not itself specify that columns to the left and right of the central column separately contract to the identity; the conditions we impose are thus sufficient and convenient when canonicalizing a PEPS, but are more constrained than the necessary conditions for Fig. 2.1(b).

To explicitly carry out the QR decomposition, we first rewrite M and thus Q as MPOs by fusing physical bonds with the left virtual bonds (Fig. 2.1(f)); for the equivalent LQ decomposition, the physical bonds should be fused with the right virtual bonds, as in MPS. Then, to explicitly enforce the isometric constraint on Q , we write it as a finite depth- n circuit of block-size l isometries and unitaries $\{u_i\}$, where the isometries appear in the edge layer of the circuit (Fig. 2.1(g): the isometries have thick bonds (dimension Dd) and thin bonds (dimension D) while the unitaries only have thick bonds (dimension Dd)) [74, 75]. The tensors in the first layer are chosen to be unitary and those in the remaining layers are isometries. The layer depth and block size control the distribution of entanglement between Q and R . In practice, to obtain a faithful QR decomposition we have found it sufficient to use $n = 2$ (a single layer of unitaries and isometries), increasing l if necessary. In addition, we set the vertical bond dimension of R equal to that of M .

To determine the tensors in the QR ansatz, we minimize the distance (cost function) $F = \| M - Q(\{u_i\})R(\{r_i\}) \|$ with respect to $\{u_i, r_i\}$ ($\| \cdot \|$ is the Hilbert-Schmidt norm) using standard tensor network techniques [76–78]. We optimize the tensors one at a time and sweep until convergence. The cost function depends quadratically on $\{r_i\}$, i.e. $F = r_i^\dagger N r_i - S r_i + \text{const}$, which is explicitly minimized by solving the linear equation $N r_i' = S$. To update the isometric/unitary tensors $\{u_i\}$, we observe that the cost function only depends linearly on them due to cancellations, i.e. $F = u_i^\dagger Y + \text{const}$, thus the optimal solution is given by $u_i' = -V U^\dagger$, where V , U appear in the singular value decomposition $Y = U s V^\dagger$ [76]. The tensor network diagrams of N , S , and Y appear in Fig. 2.2(a, b, c) respectively.

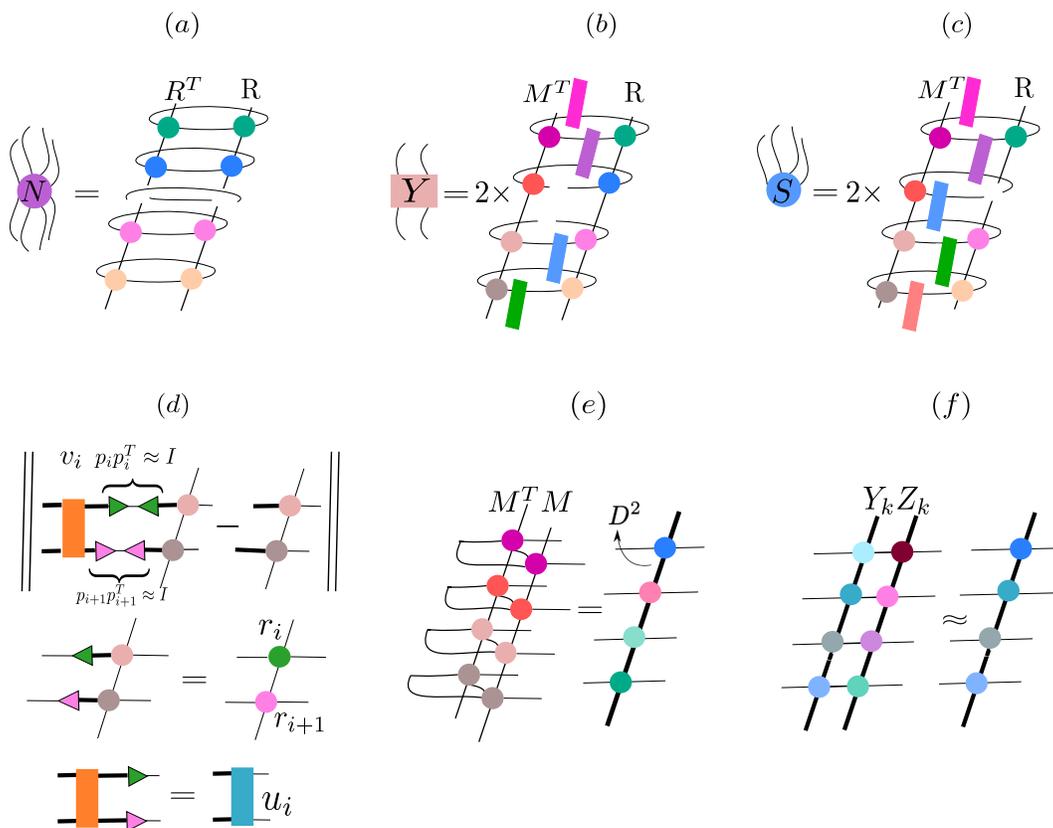


Figure 2.2: (a, b, c) Tensor network representation of the tensors N , S , and Y appearing in the QR optimization procedure; free indices in the first column correspond to the left indices of the tensors in matrix form, free indices in the second column correspond to the right indices of the tensors in matrix form. (d) The local distance (cost function) used to obtain a good initial guess for local tensors u_i , r_i , and r_{i+1} . The cost function is minimized with respect to tensors p_i , p_{i+1} , and v_i , which are used to build u_i , r_i , and r_{i+1} as depicted. (e, f) Graphical representation of $M^\dagger M$ and $Y_k Z_k$ used in the Schulz algorithm.

To accelerate the QR optimization (and to avoid local minima) we start with a good initial guess for $\{u_i, r_i\}$. We have used two techniques. The first uses a local projective truncation on the tensors $\{m_i\}$ to initialize $\{u_i\}$ and $\{r_i\}$. To this end, we contract an approximate resolution of the identity (i.e. $\mathbb{1}_D \approx p_i p_i^\dagger$ where p_i is a local isometry) and a unitary v_i into two adjacent tensors m_i, m_{i+1} and optimize p_i, p_{i+1} , and v_i to minimize the local cost function shown in Fig. 2.2(d). Once we have the optimized tensors p_i, p_{i+1} , and v_i , we construct a guess for r_i, r_{i+1} , and u_i as shown in Fig. 2.2(d). This initialization is purely local but in practice, we find that it performs well.

A second strategy is based on an accurate estimate of the $\{r_i\}$ tensors using a Schulz iteration for the matrix square-root [79]. Note that R is formally the square root of $M^\dagger M$, due to the isometric property of Q . We thus rewrite $M^\dagger M$ as an MPO as in Fig. 2.2(e). Then, starting from $Y_0 = M^\dagger M$ and $Z_0 = \mathbb{1}$, the coupled Schulz iteration, $Y_{k+1} = \frac{1}{2}Y_k(3\mathbb{1} - Z_k Y_k)$, $Z_{k+1} = \frac{1}{2}(3\mathbb{1} - Z_k Y_k)Z_k$, gives $Y_{k \rightarrow \infty} = R$ and $Z_{k \rightarrow \infty} = R^{-1}$. The vertical bond dimension of Y_k and Z_k increases with each MPO multiplication (Fig. 2.2(f)) thus we perform MPO compression after each iteration (viewing the MPO as an MPS). The vertical bond dimension of the final $Y_k (R)$ is compressed back to the original bond dimension of M . Also, since Z_k approximates R^{-1} which may have arbitrarily large norm, we regularize the iteration using $M^\dagger M \rightarrow M^\dagger M + \delta I$, where δ is a small number ($\sim 10^{-6}$). The Schulz iteration converges rapidly (see SM) and we use this accurately estimated R to initialize the optimization of the tensors in Q with respect to the cost function F . Although computing the Schulz iteration is more expensive than the local initialization, we expect it to be better when canonicalizing PEPS with more entanglement.

PEPS canonicalization sweep and truncations

To canonicalize all columns $M^{[1]}M^{[2]} \dots M^{[L]}$ in the PEPS, we sweep over all the columns in a prescribed order (say from left to right) and compute the QR decomposition to each. After column $M^{[1]}$ has been converted to QR form, we then absorb the R gauge into the neighboring $M^{[2]}$ column, creating a combined column $C^{[2]}$ with an increased vertical bond dimension of D^2 . To avoid increasing the vertical bond dimension of subsequent columns, we compress $C^{[2]}$ to a smaller vertical bond dimension $D_c < D^2$. We can perform this column truncation as an MPS truncation with enlarged physical bond dimension $D^2 d$. The role played by D_c is somewhat related to the χ in PEPS contraction algorithms [36, 80], but here D_c is an (auxiliary) bond dimension for a single PEPS layer, rather than for a double

layer. A more relevant comparison is therefore to the χ used in single-layer PEPS algorithms [81, 82], which is argued to be $\propto D$, thus leading us to conjecture that (asymptotically) $D_c \propto D$. Continuing, we perform the QR decomposition on the truncated $C^{[2]}$ column (with vertical bond dimension D_c), absorb R into $M^{[3]}$ and proceed as before over the remaining columns, to finally produce a PEPS in the canonical form as $Q^{[1]}Q^{[2]} \dots Q^{[l_x-1]}C^{[l_x]}$. Note that if MPS truncation is used to compress C , then the final central column will be in canonical form around a single site.

Canonicalization redistributes entanglement in the PEPS, and thus the canonicalized PEPS has different bond dimensions than the original PEPS. If we use the $n = 2, l = 2$ ansatz for Q , then when viewed as an MPO the Q columns have a vertical bond dimension of $\mathcal{O}(D^2)$, while the central column C has a vertical bond dimension of D_c . The formally large bond dimension of Q is primarily an artifact of expressing the isometric constraint in terms of gates. Thus it is computationally most efficient to use the structure of Q (i.e. viewing the column of isometries and column of unitaries separately) in the tensor network contractions.

Given a canonical PEPS, the canonicalization sweep can be used to convert between canonical forms (where we move the central column) which is important in algorithms such as energy optimization. The only difference then from the canonicalization sweep discussed above is that the Q columns have a larger vertical bond dimension than M . Thus, when absorbing R into a neighboring Q column (in the $n = 2, l = 2$ ansatz), we create a central column C of vertical bond dimension D^2D_c , which we subsequently compress to D_c .

From the above, we see that in computing the canonical form, and in moving the central column, there are two potential sources of error that must be controlled. The first is the *QR approximation error*, controlled by the finite-depth/block-size (n, l) of $\{u_i\}$ and the vertical MPO bond dimension D of R (which we fix). The second is the *absorption error*, which arises from the truncation of the central column's vertical bond dimension to D_c .

Cost of conversion to canonical form

We now discuss the leading costs of the computational steps in the conversion to the canonical form, assuming the $n = 2, l = 2$ ansatz for Q and setting the vertical bond dimension of R to always be equal to that of the column that is being decomposed.

QR optimization. For a column of vertical bond dimension D , the cost to determine

the isometries/unitaries is $O(l_y D^6)$ (non-iterative) and $O(l_y D^4)$, using an iterative algorithm to solve the linear equation $F = u_i^\dagger Y + \text{const}$. The cost to determine $\{r_i\}$ tensors is $O(l_y D^{12})$ (non-iterative) and $O(l_y D^8)$, using a minimization algorithm to solve $F = r_i^\dagger N r_i - S r_i + \text{const}$.

Absorption step. As discussed above, when converting a standard PEPS with columns $M^{[1]} M^{[2]} \dots M^{[l_x]}$ that have vertical bond dimension D into canonical form, the typical absorption step during the sweep creates a central column C with an enlarged vertical bond dimension DD_c (because R has dimension D_c). Compressing this down to a vertical bond dimension D_c , using sequential SVD on the MPS bonds, costs $O(l_y D^5 D_c^3)$. Alternatively, if we consider direct minimization $\| |\phi\rangle - |\psi\rangle \|^2$ (the cost typically reported in boundary PEPS algorithms) where $|\phi\rangle$ and $|\psi\rangle$ are MPSs with physical bond dimension dD^2 and virtual bond dimensions DD_c and D_c respectively, the cost is $O(l_y D^4 D_c^4)$.

Alternatively when carrying out the absorption step for a PEPS already in canonical form, e.g. when moving C from l_x to l_{x-1} in the PEPS with columns $Q^{[1]} Q^{[2]} \dots C^{[l_x]}$, then the absorption step involves compressing $C^{[l_x-1]}$ from vertical bond dimension $D^2 D_c \rightarrow D_c$. Using sequential SVD on the MPS bonds, the cost is $O(l_y D^8 D_c^3)$, while direct minimization gives a cost of $O(l_y D^6 D_c^3) + O(l_y D^4 D_c^4)$.

However, in this case, since $C^{[l_x-1]} = Q^{[l_x-1]} R$, we can use the ansatz structure of Q to reduce the cost of the truncation, by absorbing and truncating first the column of isometries, then the column of unitaries. (In both these truncations, the surrounding columns are canonical, and thus each can be performed as an MPS truncation.) With this technique, the cost of the absorption step is reduced to $O(l_y D^4 D_c^4)$ by using direct minimization. The cost is the same for truncating the column of isometries and for the column of unitaries. Using sequential SVD truncation, the leading term is $O(l_y D^6 D_c^3)$.

Accuracy of QR ansatz

To assess the accuracy of the QR ansatz, we first study its performance for a single PEPS column. As our initial state, we use the (approximate) ground-state of the spin- $\frac{1}{2}$ ITF and Heisenberg models on the square lattice. The ITF model and the

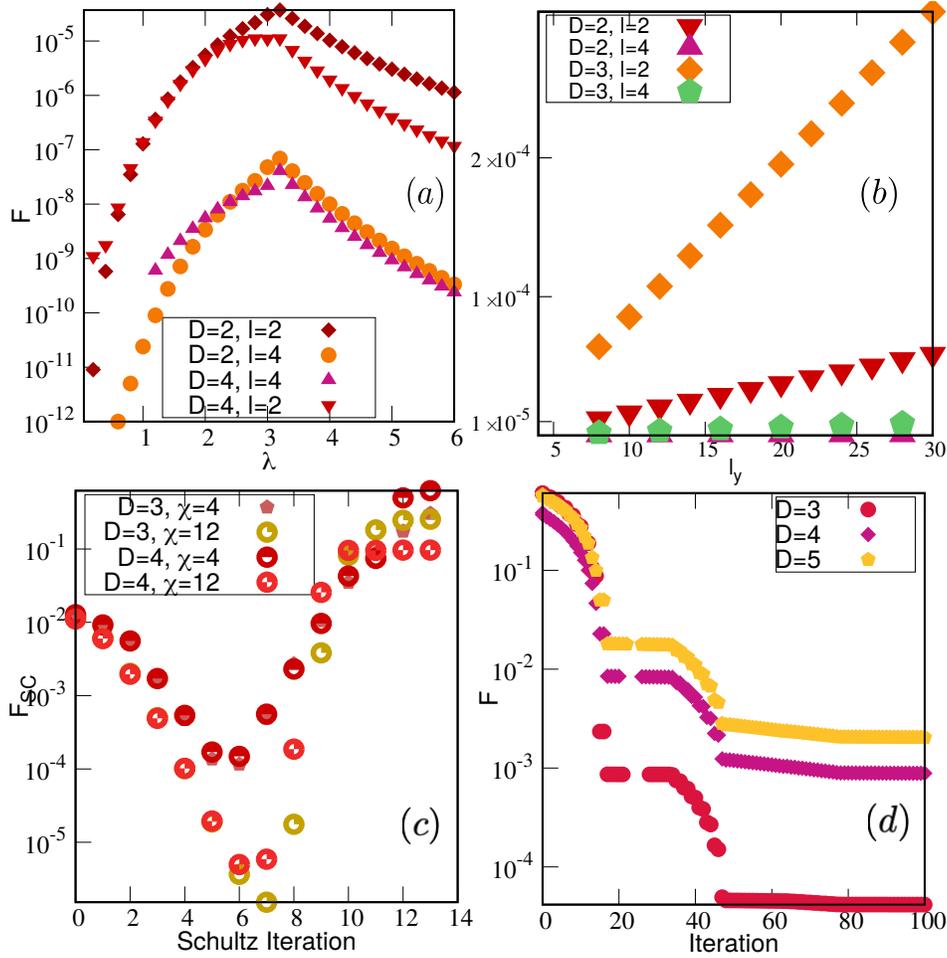


Figure 2.3: Accuracy of the variational ansatz for QR decomposition of a single bulk column. (a) The relative distance F as a function of transverse Ising field λ , for a single bulk column with $l_y = 16$, and for given values of bond dimension D and block size l used in Q . The distance F rapidly decreases with increasing l , as more entanglement is included in the Q circuit. (b) The relative distance as a function of PEPS column length l_y for different block sizes l . The error of the variational ansatz increases linearly with the length l_y . (c) The distance F_{SC} as a function of Schulz iteration for a single bulk column with $l_y = 16$ at magnetic field $\lambda = 2.0$. (d) The relative distance F of the Heisenberg model versus iteration number of the QR optimization for a single bulk column with $l_y = 10$ and with different bond dimensions D .

Heisenberg model are respectively defined by

$$H_{\text{ITF}} = - \sum_{\langle ij \rangle} \sigma_z^i \sigma_z^j - \lambda \sum_i \sigma_x^i,$$

$$H_{\text{Heisenberg}} = \sum_{\langle i,j \rangle} \mathbf{S}_i \cdot \mathbf{S}_j,$$

where $\mathbf{S}_i \equiv (\sigma_x, \sigma_y, \sigma_z)$ and σ_α are the Pauli matrices. The ITF model has a critical point at $\lambda_c \approx 3.05$. Our initial PEPS is constructed from the bulk tensors of an infinite PEPS ground-state [34] (optimized with a full-update scheme [38, 80] and a 2×2 unit cell) that is repeated periodically across the finite PEPS lattice.

We measure the accuracy of the QR ansatz by the value of its optimization cost function F . Here, the parameter controlling the accuracy is the block size l of the isometric/unitary circuit (the number of layers is kept as $n = 2$, and the vertical bond dimension of R is kept as D). In Fig. 2.3(a), we show the plot of the distance F versus ITF magnetic field λ . As expected, when the system is close to criticality, the accuracy is reduced as the ground state becomes more entangled. Increasing l increases the disentangling effect of the unitaries, and the accuracy increases rapidly (we conjecture exponentially with l), especially far from criticality.

Next, we investigate the QR accuracy as a function of system size l_y for the ITF model at field strength $\lambda = 3.5$. As shown in Fig. 2.3(b), the relative error in F increases linearly with system size, although the slope shows a rapid decay with the isometry/unitary block size l . Thus, the variational gauging ansatz introduces a constant error per lattice site, consistent with a fidelity that goes like $e^{-\epsilon l_y} \sim 1 - \epsilon l_y$.

We can also study the accuracy of the Schulz iteration. We show this by evaluating the distance $F_{SC} = \| M^\dagger M - R^2 \|$ as a function of the Schulz iteration. The MPO compression is controlled by a truncated bond dimension χ (in the final iteration in the canonicalization algorithm, this is always set to D). In Fig. 2.3(c), we show how the accuracy of the Schulz iteration depends on χ for different initial bond dimensions D for the ITF model at magnetic field $\lambda = 2.0$. The regularization parameter is always set to $\delta \sim 10^{-6}$.

Finally, we give additional results for the QR optimization in the Heisenberg model. We plot the relative distance F versus iteration number for a single bulk column with $l_y = 10$ in Fig. 2.3(d), using block size $l = 2$. We see that the relative error of the norm contraction Δ is similar to that of the ITF model at the critical point ($\sim 10^{-2}$) for a lattice of size 10×10 .

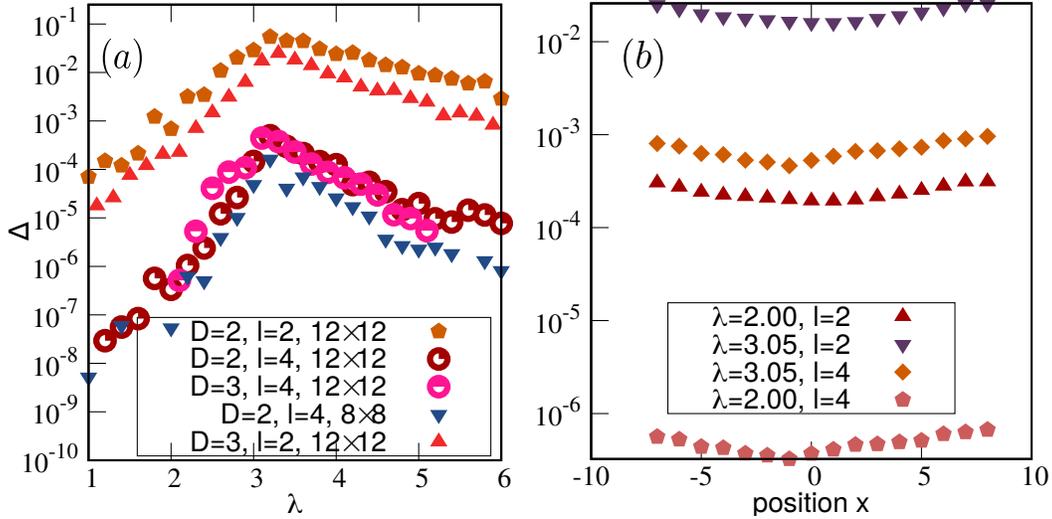


Figure 2.4: (a) The relative error of the norm contraction $\Delta = \frac{N-N_b}{N_b}$ as a function of λ . N_b is a reference norm obtained by boundary method contraction. For the norm computed via the canonical form, we use $D_c = 3D$. (b) Δ as a function of column position x in the PEPS for different values of λ . The bond dimensions are $D = 3$ and $D_c = 12$, and the lattice size is 16×16 .

Accuracy of PEPS canonical form

We next investigate the accuracy and stability of the full PEPS canonical form constructed from a sweep of the QR approximation and absorption steps across the columns. We estimate the faithfulness of the canonical form from the norm contraction $\mathcal{N} = \langle \Psi | \Psi \rangle$. We compute the norm in the canonical form using only the central column C since all other columns contract exactly to the identity. The relative error of norm contraction is then defined as $\Delta = \frac{N-N_b}{N_b}$, where the reference value N_b is obtained using an accurate boundary contraction of the original (uncanonicalized) PEPS keeping a large boundary auxiliary bond dimension [35, 50, 64]. In Figs. 2.4(a, b), we show a plot of the relative error Δ as a function of ITF magnetic field λ (using the same approximate ground-state as above) and central column position. Similarly to the single-column results above, the accuracy of the full canonical form depends on the correlation length of the model, and the canonicalization error decreases rapidly (exponentially) as the block size l is increased.

Norm stability of the canonical form

An important property to assess is the numerical stability of the PEPS in the canonical form. We can study this by examining the stability of the total norm N with

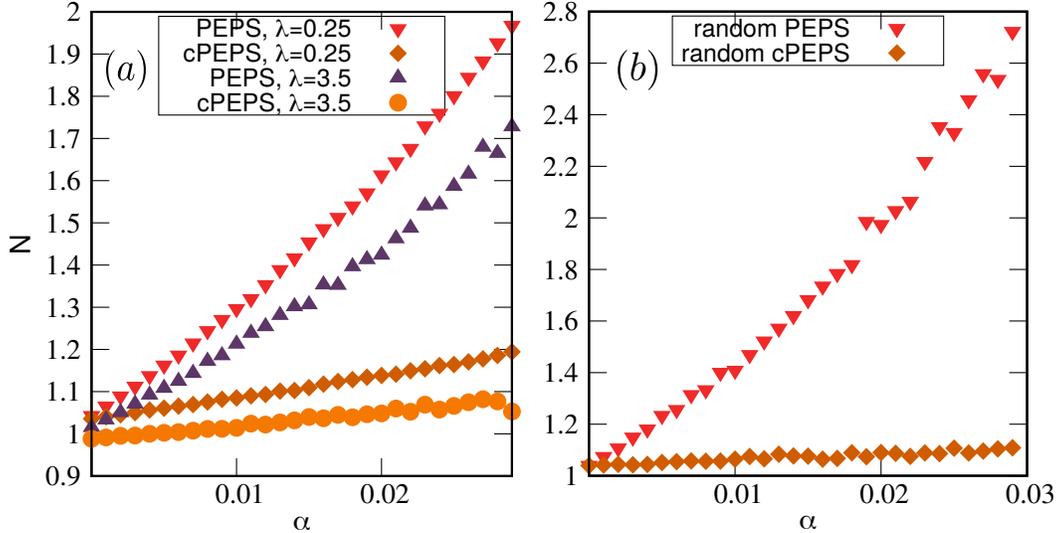


Figure 2.5: The total norm of a 8×8 PEPS as a function of perturbation strength α for (a) ITF ground-state PEPS at different coupling parameters λ and for (b) a uniform random PEPS. cPEPS denotes canonical PEPS derived from the standard PEPS. The bond dimensions are chosen to be $D = 3$ and $D = 2$ for the ITF PEPS ground state and random PEPS respectively.

respect to small perturbations of the tensors. A numerically stable form is one where small perturbations to the tensors result in small perturbations in the norm. To study this property, we compare the stability of the norm in a standard PEPS versus in its canonical form. We add small perturbations to each tensor $A_{s_i}^i \rightarrow A_{s_i}^i + \alpha P$ in a canonical PEPS and a normal PEPS (the unperturbed canonical PEPS is obtained by first converting the normal PEPS to canonical form). The normalized tensors P ($|P| = 1$) are chosen to form a uniform random distribution in the interval $[0, 1]$. In Fig. 2.5(a, b), we plot the norm N as a function of perturbation strength α for the ITF ground-state PEPS and a random PEPS. We observe, in both cases, that the canonical PEPS (cPEPS) remains more stable, as the norm changes less with the perturbation as compared to the standard PEPS. This is a result of the fact that the Q tensors are isometric and adding small perturbations to them only affects the isometric property on the unit scale.

Energy optimization in canonical form

A natural application of the PEPS canonical form is to ground-state energy optimization, which mimics the use of the MPS canonical form in energy optimization. To show this, we perform imaginary time evolution on the ITF model, which we

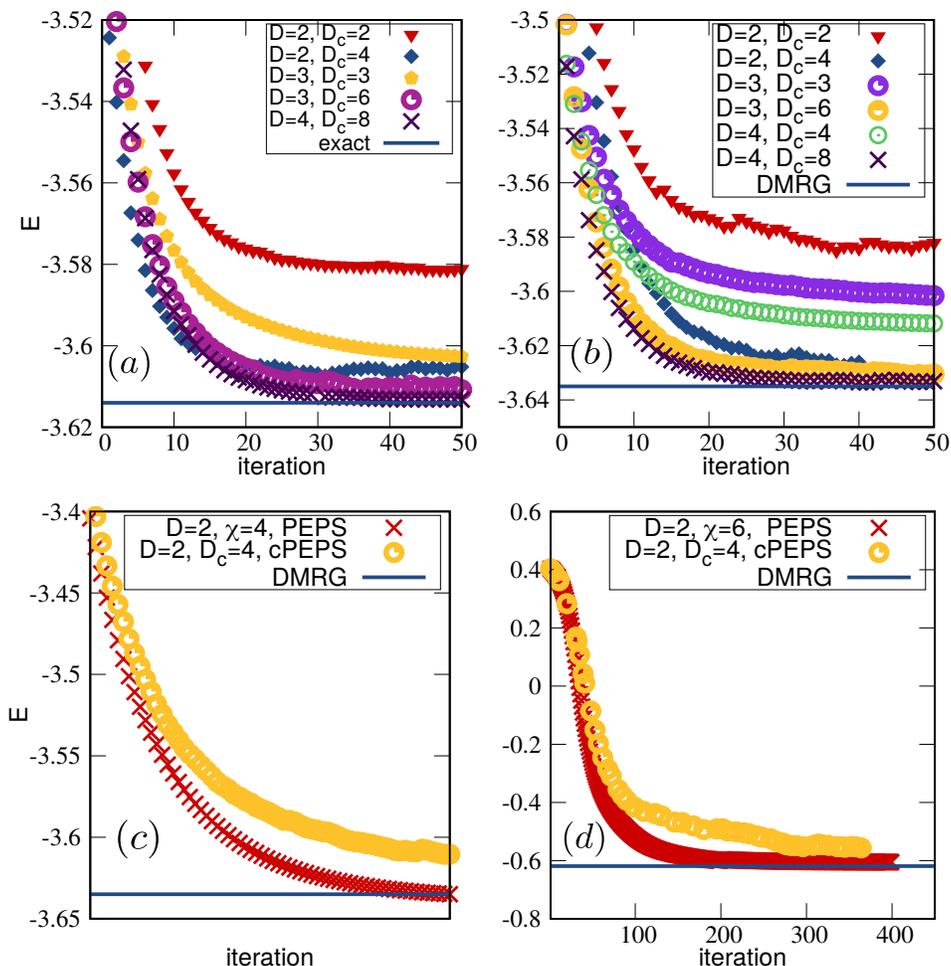


Figure 2.6: Imaginary time energy optimization based on the PEPS canonical form. The canonical PEPS energy of the ITF model as a function of imaginary time sweep for (a) 4×4 and (b) 8×8 square lattices with field $\lambda = 3.5$. A direct variational energy comparison between canonicalized PEPS (cPEPS) and standard uncanonicalized PEPS [35] for the (c) ITF model at $\lambda = 3.5$ and (d) Heisenberg model, on a 8×8 square lattice. Both algorithms start from the same random initial PEPS. The horizontal DMRG line is the result of a converged 2D DMRG calculation and can be taken to be numerically exact. The parameter χ stands for boundary bond dimension in environment calculations in the standard PEPS optimization. cPEPS denotes a canonical PEPS.

carry out with a sequence of gates $e^{-\tau h}$ on the horizontal and vertical bonds [36, 50]. Evolution on a column of vertical bonds is conveniently carried out on the central column C of a canonical PEPS with bond dimension D_c , where it reduces to an MPS imaginary time evolution followed by an MPS truncation with an enlarged physical bond dimension $D^2 d$. Evolution on a column of horizontal bonds can be carried out using a two-column canonical PEPS (analogous to the two-site MPS canonical form) where there are two central columns, and columns to the left and right of these two are isometric tensors Q , thus reducing the optimization problem to one of a PEPS with only two columns. In this case, rather than canonicalizing the remaining environment around the bond in the two-column PEPS, we contract it exactly, which is straightforward. Since there are only two columns, these can be reduced to an MPS with enlarged vertical bond dimension $D^2 D_c$ and physical bond dimension $D^2 d^2$. In Fig. 2.6(a, b) we show the energy as a function of the number of full imaginary time sweeps for the ITF model at field strength $\lambda = 3.5$ compared to a near-exact DMRG result. Note that both D (which controls the variational space of the standard PEPS) and D_c (vertical bond dimension of the central column, which controls the accuracy of the absorption step in the canonicalization sweep) affect the final converged energy; in this setting, increasing D_c has a larger effect than increasing D . The relative error of the energy per site reached for the largest bond dimension $D = 4$, $D_c = 8$ for both 4×4 and 8×8 lattice sizes is on the order of 10^{-4} .

In Fig. 2.6(c, d), we benchmark the variational energy for the canonical PEPS and the standard PEPS for the ITF model at $\lambda = 3.5$ and the Heisenberg model. For the standard PEPS we used the optimization algorithm from Ref. [36] to obtain the ground state. Both algorithms are initialized by the same random PEPS. The maximum relative errors of the QR and absorption steps using the two layer, $l_y = 2$ QR ansatz are on the order of 10^{-4} and 10^{-3} respectively in the ITF and Heisenberg simulation. This relative error bounds the ultimate accuracy in the energy that the canonical PEPS simulation can achieve. Thus we observe that standard PEPS optimization leads to lower energies than the canonical PEPS energy in both cases as expected.

2.5 Summary

In summary, we have described a procedure to convert a PEPS into a canonical form analogous to that of an MPS where all columns but one are isometric, by sequentially decomposing columns through a variational QR ansatz. We find that

the canonicalization is stable and can be carried out with a small and controllable error. Canonicalization redistributes entanglement in the PEPS, resulting in a central column with increased bond dimension. Our procedure introduces the possibility to formulate canonical PEPS algorithms which make explicit use of an isometric environment, which we demonstrated in an imaginary time optimization of the ground-state energy. The canonical form is clearly numerically more stable, as we show in calculations of the stability of the norm with respect to perturbations of the tensors. However, a faithful comparison of the cost of algorithms using the canonical PEPS and standard PEPS requires considerably more analysis. This is because the canonicalization leads to a non-homogeneous PEPS with different bond dimensions on the vertical and horizontal bonds, quite different from the standard PEPS scenario, and in addition, the canonical PEPS introduces additional numerical parameters (to control the accuracy of the QR ansatz and R absorption) which must be converged. To eliminate the additional errors of the QR decomposition, it may be more expedient to directly optimize the underlying network of isometries and unitaries suggested by the QR ansatz, without explicitly converting into the standard PEPS column-row form. Future investigations will focus on more detailed analysis of these and other algorithms as well as the general representational power of canonicalized PEPS.

Note: At the time this project was completed, a similar work was also disseminated [83]. In that work, the authors similarly pursue a full canonicalization of a PEPS, but use a different set of isometric conditions on the Q tensors which are more constrained than the ones that we use. Further work is needed to understand the relationship between these techniques.

EFFICIENT REPRESENTATION OF LONG-RANGE INTERACTIONS IN TENSOR NETWORK ALGORITHMS

This chapter is based on the following publication:

1. O'Rourke, M. J., Li, Z. & Chan, G. K.-L. Efficient representation of long-range interactions in tensor network algorithms. *Phys. Rev. B* **98**, 205127. doi:10.1103/PhysRevB.98.205127 (20 Nov. 2018).

3.1 Abstract

We describe a practical and efficient approach to represent physically realistic long-range interactions in two-dimensional tensor network algorithms via projected entangled-pair operators (PEPOs). We express the long-range interaction as a linear combination of correlation functions of an auxiliary system with only nearest-neighbor interactions. To obtain a smooth and radially isotropic interaction across all length scales, we map the physical lattice to an auxiliary lattice of expanded size. Our construction yields a long-range PEPO as a sum of ancillary PEPOs, each of small, constant bond dimension. This representation enables efficient numerical simulations with long-range interactions using projected entangled pair states.

3.2 Introduction

The accurate description of strongly correlated quantum many-body systems is a major challenge in contemporary physics. Nonetheless, some of the most intriguing macroscopic quantum phenomena, such as high-temperature superconductivity and the fractional quantum Hall effect, arise from strong quantum correlations. In recent years, tensor network states (TNS) [25, 26, 47–49, 84], including matrix product states (MPS) [22, 53–55] and projected entangled-pair states (PEPS) [23, 63–65], have emerged as promising classes of variational states to numerically approximate the low energy physics of correlated quantum systems with area or near-area law physics. Their power stems from systematically improvable accuracy through increasing the tensor bond dimension D [85], and the $O(A)$ linear complexity of the associated algorithms with respect to the system size A (under assumption of contractibility of the underlying tensor network, as is common in many physical applications, using approximate contraction methods [63, 64, 86–90].)

One promising application of TNS is to accurate calculations of electronic structure of realistic materials. While the electronic structure Hamiltonian can be represented in multiple ways [91–94], the simplest—and the one of interest in this work—is a real-space grid formulation [95–99],

$$\begin{aligned}\hat{H} &= -t \sum_{\langle i,j \rangle} (a_{i\sigma}^\dagger a_{j\sigma} + h.c.) + \sum_i v_i^{ne} n_i + \hat{V}^{ee}, \\ \hat{V}^{ee} &= \sum_i v_{ii}^{ee} n_{i\alpha} n_{i\beta} + \sum_{i<j} v_{ij}^{ee} n_i n_j,\end{aligned}\tag{3.1}$$

where i, j label lattice sites, $\sigma \in \{\alpha, \beta\}$ labels spin, t is the kinetic energy matrix element, and a^\dagger , a , and n are fermion creation, annihilation, and number operators, respectively. As the spacing between grid points (h) goes to zero, the parameters scale as $t \propto h^{-2}$ and $v_{ij}^{ee} \propto h^{-1}$; these become exact representations of $-\frac{1}{2}\nabla^2$ and the continuum Coulomb potential $1/r_{ij}$ with $r_{ij} \triangleq |\mathbf{r}_i - \mathbf{r}_j|$ [96, 98]. This simple form of the electronic structure Hamiltonian is especially suited to TNS algorithms as the Coulomb interaction is a pairwise operator as opposed to a general quartic operator when using a non-local basis, and Eq. (3.1) can be viewed as an extended Hubbard model with long-range terms. Ground states of such grid Hamiltonians have been computed in 1D using MPS and the density matrix renormalization group (DMRG), yielding near exact electronic structure benchmarks for small lattice spacings [25, 26, 95]. In principle, this success in 1D should be extensible to 2D and 3D by using PEPS instead of MPS, and would then provide a route to simulate arbitrarily complex electronic structure problems with arbitrarily improvable accuracy.

However, current state-of-the-art PEPS applications to physical problems have not yet advanced beyond local lattice models in 2D [41, 88, 100–103]. There are two principal complications. The first is that long-range interactions can in principle lead to increased entanglement, and even volume-law entanglement, that would be difficult or impossible to capture with a PEPS with a finite bond dimension. Fortunately, in applications of the density matrix renormalization group using the Coulomb interaction (for example, to electronic structure) it is seen that the increase in entanglement is modest and volume law entanglement is not observed [93, 95, 97, 104–106]. The second complication is simply the increased cost of all operations when long-range interactions are considered, even for a fixed bond dimension. To see the basic challenge, consider the evaluation of the energy expectation value: for a Hamiltonian with localized interactions, the number of terms in a standard term-by-term calculation scales linearly with the size of the system, $O(A)$. However, for a Hamiltonian with long-range interactions, the number of terms scales like

$O(A^2)$, which is prohibitively expensive in two (or higher) dimensions, as we take the continuum limit. Alternatively, one might try to use an exact tensor network operator, or projected entangled pair operator (PEPO), to represent the long-range interaction [107], avoiding the explicit term-by-term evaluation. However, the exact PEPO representation for arbitrary long-range interactions in 2D has a bond dimension that scales as $O(A^{1/4})$, causing the overall cost to compute expectation values to scale as $O(A^2)$ [45].

In 1D, the increased computational cost of long-range interactions can be eliminated if they are smooth and decaying. In this case one can approximate the exact matrix product operator (MPO) by a compressed MPO of constant bond dimension D that generates a sum of exponential interactions, and smoothly decaying interactions can be approximated well by such sums [46, 107, 108]. Exponential interactions in MPOs arise naturally from the matrix product structure, which also gives rise to the exponential decay of two-point correlation functions in MPS. Extending the correlation function analogy to 2D leads to an efficient representation of long range interactions in 2D when their form exactly coincides with the correlation function of a 2D lattice model. This was demonstrated in Ref. [107], which constructed a compact pair interaction PEPO whose interaction potential was given by the critical 2D Ising correlation function.

Building on these ideas, in this work we describe how general long-range interactions in two dimensions, including the Coulomb interaction, can be efficiently encoded as a sum of low rank correlation function valued PEPOs. Although superficially similar to the problem of approximating a smooth interaction in 1D by a sum of exponentials, additional complications arise in two dimensions because physical interactions possess different analytic properties from two-point correlation functions on the same lattice. For example, the Coulomb interaction is radially isotropic at all distances, while the two-point lattice correlation functions are isotropic only at large distances due to the lattice discretization. We show how to overcome these and other difficulties by introducing an expanded auxiliary lattice, and demonstrate the effectiveness of the representation in a ground-state finite PEPS simulation of a 2D spin model with Coulombic Heisenberg interactions. Although we specifically treat only the Coulomb interaction and two dimensions in our numerical examples, our arguments naturally extend to representing smooth and radially isotropic interactions in any dimension.

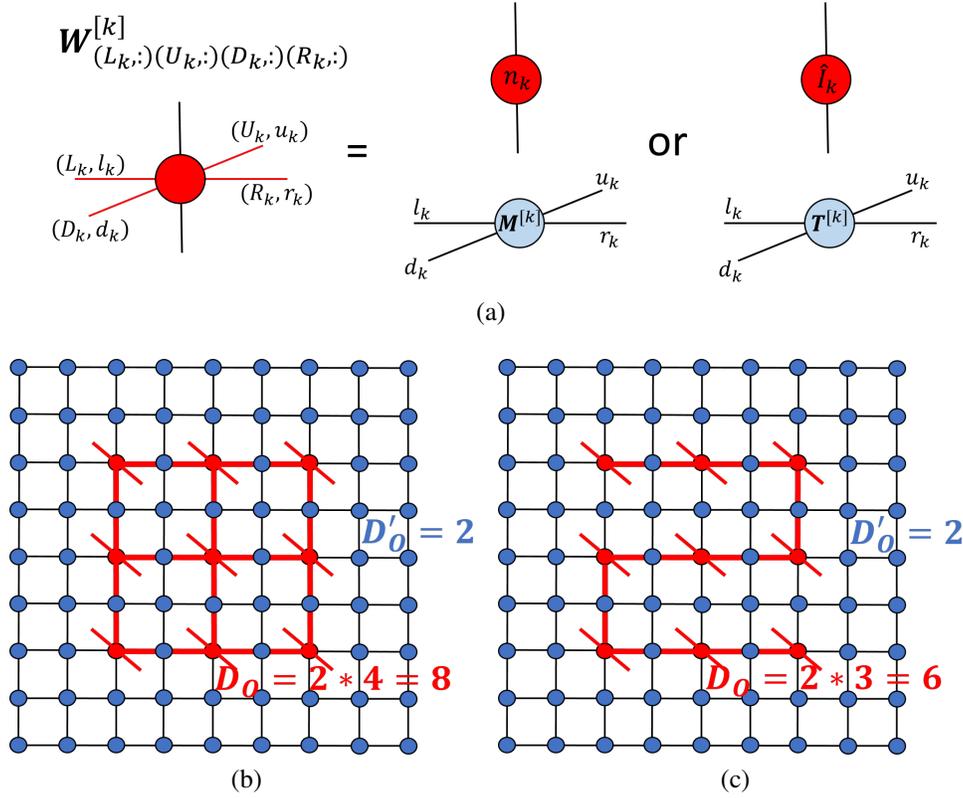


Figure 3.1: (a) The construction of the nonzero parts of the CF-PEPO tensor $\mathbf{W}^{[k]}$ via the coupling of the finite state machine (FSM) tensor (red) with the Ising correlation function tensors (blue). Note that here the physical indices of $\mathbf{W}^{[k]}$ are explicitly shown, whereas they are suppressed in Eq. (3.3). (b)-(c) Two possible constructions of the long-range PEPO for a 3x3 physical system with 1 fictitious Ising site (blue) in between adjacent physical sites (red) and a 2 site buffer to help mitigate boundary effects in the encoding of the potential. Black bonds are $D'_O = 2$ and red bonds are $D_O = 8$ (b) and 6 (c).

3.3 Correlation function valued PEPOs

We first define correlation function valued PEPOs (CF-PEPOs), which are central to this work. As motivation, we recall the construction of MPOs for smooth interactions approximated by sums of exponentials. This is usually done in the language of finite state machines (FSM), where the MPO is viewed as an operator valued MPS, and the incoming and outgoing bonds of each MPO tensor are interpreted as machine states [46, 108]. An FSM can encode an exponentially decaying interaction strength $e^{-\lambda r_{ij}}$ via a single non-zero element in each MPO tensor with value $e^{-\lambda}$, which gets multiplied along the lattice as long as the FSM stays in a specified state. The pairwise operator $\sum_{i < j} e^{-\lambda r_{ij}} n_i n_j$ can then be represented by an MPO with bond dimension

3, with the two additional states in the FSM acting to combine the exponential scalar values with the operators $n_i n_j$. The construction can be extended to the general 1D interaction $\sum_{i<j} V(r_{ij}) n_i n_j \approx \sum_{i<j} \sum_{t=1}^{N_t} c_t e^{-\lambda_t r_{ij}} n_i n_j$ by introducing additional states for each of the N_t exponential decays, for a total MPO bond dimension of $N_t + 2$ (or alternatively, N_t MPOs of bond dimension 3). However, while this representation is natural in 1D, its direct extension to 2D is not. This is because multiplying the element $e^{-\lambda}$ along any single FSM path between two sites i and j creates an exponentially decaying strength as a function of the Manhattan distance $|x|+|y|$, not the desired Euclidean distance $(x^2+y^2)^{1/2}$, as the elements are multiplied out along the grid lines [45].

A different starting point, that is more natural in higher dimensions, is to consider scalar interaction strengths generated by the two-point correlation function $\langle o(\mathbf{r}_i) o(\mathbf{r}_j) \rangle_\beta$ of a classical model at inverse temperature β . We term the PEPO for the operator $\sum_{i<j} \langle o(\mathbf{r}_i) o(\mathbf{r}_j) \rangle_\beta n_i n_j$, a correlation function valued PEPO (CF-PEPO). Using a classical model with local interactions yields a CF-PEPO with low bond dimension, as noted in Ref. [107]. As a concrete example, consider the spin-spin correlation function $\langle \sigma_i \sigma_j \rangle$ of the 2D Ising model, which has the Hamiltonian $H = -\sum_{\langle m,n \rangle} \sigma_m \sigma_n$, $\sigma \in \{+1, -1\}$. For two given points on the lattice i and j , this correlation can be exactly represented by the Ising PEPS with $D = 2$ [65, 109], viz.,

$$\langle \sigma_i \sigma_j \rangle_\beta = \frac{1}{Z} \text{Tr} \left(\prod_{k \neq i,j} T_{l_k u_k d_k r_k}^{[k]} M_{l_i u_i d_i r_i}^{[i]} M_{l_j u_j d_j r_j}^{[j]} \right). \quad (3.2)$$

Here $Z = \text{Tr} \prod_k \mathbf{T}^{[k]}$ is the partition function and the tensors \mathbf{T} and \mathbf{M} are the local tensors of the PEPS off and on the correlation function sites, respectively. These tensors are obtained from the eigenvalue decomposition $X = U \lambda U^T$ of the familiar 2×2 Ising model transfer matrix $X_{ij} = \exp((-1)^{\delta_{ij+1}} J \beta)$, which encodes the local terms of the partition function for a pair of nearest neighbor spins [110]. In tensor network language, these X matrices would be placed on each bond of the square lattice. In order to create a local tensor network description of the system, we define the ‘‘square root’’ of this transfer operator as $P = U \sqrt{\lambda} U^T$, and define the local tensors as $T_{ludr} = \sum_a P_{la} P_{ua} P_{ad} P_{ar}$ and $M_{ludr} = \sum_{ab} P_{la} P_{ua} P_{da} \sigma_{ab}^z P_{br}$, where σ^z is the standard Pauli matrix.

To obtain the Ising CF-PEPO, we combine the tensors $\mathbf{T}^{[k]}$, $\mathbf{M}^{[k]}$ of the Ising PEPS at each site with (translationally invariant) tensors $\mathbf{Y}^{[k]}$ of a PEPO for the interaction $\sum_{i<j} n_i n_j$. As demonstrated in a general fashion in [45] based on work in [46], the

$\mathbf{Y}^{[k]}$ tensors can be obtained by a FSM construction in 2D, where each element of the tensor $Y_{L,U,D,R}$ at a given site corresponds to a specific local state of the FSM and returns a specific local operator $\{0, \hat{I}, n\}$. The Ising CF-PEPO tensors are then formed by a selective direct product between $\mathbf{Y}^{[k]}$, $\mathbf{T}^{[k]}$, and $\mathbf{M}^{[k]}$,

$$\begin{aligned} \sum_{i < j} \langle \sigma_i \sigma_j \rangle_{\beta} n_i n_j &= \text{Tr} \left(\prod_k W_{(L_k, l_k)(U_k, u_k)(D_k, d_k)(R_k, r_k)}^{[k]} \right), \\ \mathbf{W}_{(L_k, :)(U_k, :)(D_k, :)(R_k, :)}^{[k]} &= Y_{L_k, U_k, D_k, R_k}^{[k]} \otimes \mathbf{T}^{[k]} \text{ if } Y = \hat{I}_k, \\ \mathbf{W}_{(L_k, :)(U_k, :)(D_k, :)(R_k, :)}^{[k]} &= Y_{L_k, U_k, D_k, R_k}^{[k]} \otimes \mathbf{M}^{[k]} \text{ if } Y = n_k, \\ \mathbf{W}_{(L_k, :)(U_k, :)(D_k, :)(R_k, :)}^{[k]} &= 0 \text{ if } Y = 0. \end{aligned} \quad (3.3)$$

Here $\mathbf{W}^{[k]}$ (Fig. 3.1(a)) is the operator valued tensor in the Ising CF-PEPO and (L_k, l_k) is a composite index of the bond L_k for the 2D FSM and the bond l_k of the Ising PEPS. Note that the selective direct product can be formed unambiguously due to the 1 : 1 correspondence between possible states of $\mathbf{Y}^{[k]}$ and the Ising PEPS tensors $\mathbf{M}^{[k]}$ and $\mathbf{T}^{[k]}$.

Since the FSM tensors $\mathbf{Y}^{[k]}$ only need to encode the two operators $n_i n_j$ and contain no information about the distance between them, there is some flexibility in the possible topologies of the FSM (see Fig. 3.1). The snake geometry in (c) has a significantly reduced computational complexity compared to the original FSM from [45] shown in (b), and it also imposes an ordering that allows for a simple way to include fermionic statistics (via Jordan-Wigner strings) at the operator level, eliminating the need for swap gates in fermionic PEPS [80]. The full specifications for constructing the tensors $\mathbf{Y}^{[k]}$ according to both FSM geometries are given in Appendix A. As an important note, both of these constructions are compatible with existing iPEPS [34] algorithms.

3.4 CF-PEPOs and the auxiliary lattice

Using the above arguments, we might now consider approximating the form of a physical, smooth, and isotropic interaction $V(r_{ij})$ by a sum of N_t lattice correlation functions at different temperatures, $V(r_{ij}) \approx V_{\text{fit}}(r_{ij}) = \sum_{t=1}^{N_t} c_t f_{\beta_t}(r_{ij})$ [$f_{\beta_t}(r_{ij}) \triangleq \langle o(\mathbf{r}_i) o(\mathbf{r}_j) \rangle_{\beta_t}$], giving the interaction operator as a sum of CF-PEPOs. In Fig. 3.2(a) we show the maximal absolute error in a direct fit of $1/r_{ij}$ using Ising correlation functions on an $L \times L$ lattice. For large r_{ij} , the maximal error (at a given radius) can be seen to converge rapidly, with a fitted convergence rate of $\sim O(r_{ij}^{-2.7})$ (Fig. 3.2(a)), showing we can easily capture the long distance behavior of the Coulomb potential

that is sampled at large system sizes. However, for small r_{ij} , the maximal errors are much larger, and the expansion does *not* converge even with very many terms, as seen in Fig. 3.2(b). This is because the lattice discretization of the correlation functions prevents radial isotropy in the basis $\{f_{\beta_t}\}$ at short lattice distances. In addition, for finite lattices, boundary effects also cause errors in the isotropy and translational invariance.

The short distance anisotropy error can be remedied by representing the isotropic physical interaction by correlation functions generated on an *expanded auxiliary lattice* with additional “fictitious” sites. The physical distance r_{ij} (on the original lattice) maps to the expanded distance $R_{ij} = (N_f + 1)r_{ij}$ on the auxiliary lattice (N_f denotes the number of fictitious sites added to the sides of one unit square on the original lattice). This gives us a rescaled potential that is easier to fit at small r_{ij} ,

$$\tilde{V}_{\text{fit}}^{[N_f]}(r_{ij}) \triangleq (N_f + 1)V_{\text{fit}}(R_{ij}) = (N_f + 1) \sum_{t=1}^{N_t} c_t f_{\beta_t}(R_{ij}), \quad (3.4)$$

where the specific rescaling in Eq. (3.4) has been shown for the Coulomb potential. Choosing a sufficiently large expansion factor N_f ensures that the fitting basis becomes isotropic up to an error ϵ , and the radial fit can then be performed to increasing accuracy with increasing N_t up to a similar ϵ . Further, choosing a suitably large side length of the auxiliary lattice buffering the physical region also removes the boundary effects in a finite lattice simulation.

In Figs. 3.2(b)-(c) we show the behavior of the maximal error in fitting $\tilde{V}_{\text{fit}}^{[N_f]}(r_{ij})$ to $1/r_{ij}$ for several values of r_{ij} , as a function of both the number of fictitious sites N_f and fitting terms N_t . They demonstrate that for $N_f = 10$ and a modest $N_t = 8$, we are able to obtain a maximum error of 10^{-3} with $\tilde{V}_{\text{fit}}^{[10]}(r_{ij})$. In Fig. 3.2(c), note that the $r_{ij} = 1$ curve (i.e. the maximal error curve) converges as $\sim (N_f + 1)O(N_f^{-2.7}) \propto N_f^{-1.7}$ due to the rescaling factor in Eq. (3.4). Thus by further increasing N_f the error can be continually decreased.

Up to this point in this section, we have implicitly considered r_{ij} only on the unit lattice, i.e., $r_{ij} \triangleq |(x, y)_i - (x, y)_j|; x, y \in \mathbb{Z}$, which is to say that the lattice spacing $h = 1$. In addition to the above discussion of increasing N_f to reduce the fitting error for a fixed spacing $h = 1$, an alternative (but equivalent) viewpoint is that N_f can be increased to maintain a given maximal error in the potential as $h \rightarrow 0$. Precisely, the maximal error in the new potential will occur at the new shortest physical distance, $V(h) = h^{-1}\tilde{V}_{\text{fit}}^{[N_f]}(1)$. The error at this point $\epsilon(V(h))$ scales as $\epsilon \propto h^{-1}N_f^{-1.7}$, which

reveals that N_f must increase as $N_f \propto h^{-1/1.7} = h^{-0.59}$ in order to maintain the level of error originally incurred at the point $\tilde{V}_{\text{fit}}^{[N_f]}(1)$ (for $h = 1$).

In summary, the full CF-PEPO is obtained by coupling the FSM of the operators (either in the snake form, or the full 2D FSM) to the Ising CF-PEPS on an expanded lattice as specified by Eq. (3.4), and as shown in Fig. 3.1(b)-(c). The total error of the fit is controlled by the expansion parameter N_f and the number of terms N_t . For the Coulomb interaction and a desired accuracy, N_t is only weakly dependent on the physical lattice discretization and system size. This is similar to what is observed in MPO fits in one dimension [46, 95, 107, 108] as well as analytical work on exponential fits of the Coulomb operator in 2D [111].

3.5 Computational cost

We now consider the evaluation of a finite PEPS expectation value for a PEPS of bond dimension D_S and an Ising CF-PEPO of bond dimension D_O . To define the computational cost, we must choose an approximate contraction scheme. Here we use a simple generalization of the “optimized” contraction scheme proposed in Ref. [82] to include a PEPO. Using the full 2D FSM (Fig. 3.1(b)), the CF-PEPO has bond dimension $D_O = 8$ for the bonds emanating from the physical sites and $D'_O = 2$ for bonds that only connect fictitious sites, and the leading contraction cost can be derived to be $N_t [O(A\chi^3 D_O^3) + O(AN_f \chi^3 D'_O{}^2 D_O) + O(AN_f^2 \chi^3 D'_O{}^3) + O(A\chi^3 D_S^3) + O(AN_f \chi^3 D'_O{}^2 D_S)]$, where χ is the maximum bond dimension appearing in the approximate contraction scheme and can be taken as $\chi \sim D_S^2 D_O$. For the snake FSM construction (Fig. 3.1(c)) $D_O = 6$ instead of 8, and the physical PEPO tensors only have two large bond dimensions instead of four. This reduces the overall scaling to $N_t [O(A\chi^3 D'_O{}^2 D_O) + O(AN_f \chi^3 D'_O{}^2 D_O) + O(AN_f^2 \chi^3 D'_O{}^3) + O(A\chi^3 D_S^3) + O(AN_f \chi^3 D'_O{}^2 D_S)]$.

In both cases, the cost is linear in the system area A as we originally desired. However, it is instructive to compare these costs to an implementation without a PEPO. In a naive implementation of the exact term-by-term contraction of each $n_i n_j$ operator in the Coulomb potential, a single term would involve a contraction of cost $O(A\chi^3 D_S^3)$ with $\chi \sim D_S^2$, and there would be $O(A^2)$ such terms, giving an $O(A^3)$ cost. Assuming a reasonably large value for D_S , this cost can be compared to the analogous term in the (snake) PEPO contraction cost, which gives an approximate crossover when $A^2 \sim N_t D_O^3$, which for $N_t = 10$, $D_O = 6$, corresponds to $A \sim 50$. In a more sophisticated exact implementation, we could rewrite $\sum_{ij} V_{ij} n_i n_j$ as $\sum_i n_i \hat{O}_i$,

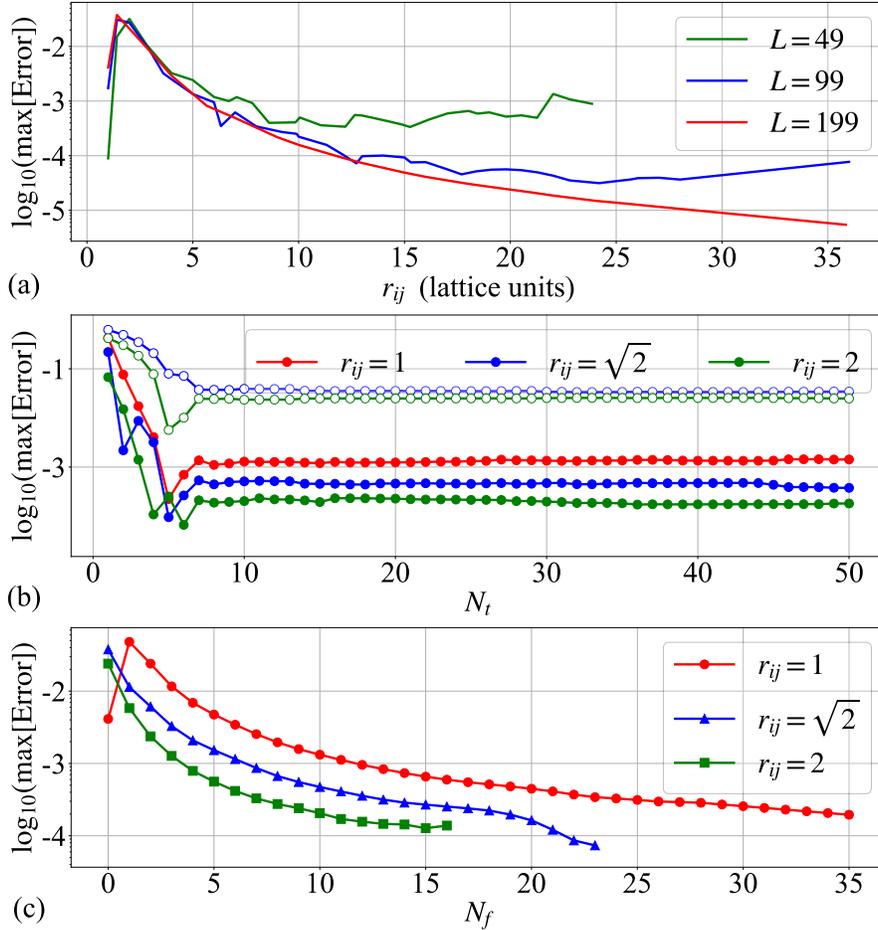


Figure 3.2: Convergence properties of Coulomb fitting. For all plots $r_{ij} = 0$ is the central point on the lattice. (a) The upper envelope of $|V_{\text{fit}}(r_{ij}) - 1/r_{ij}|$ obtained with $N_t = 12$, $r_{ij} = R_{ij}$, a least squares weight function of $r_{ij}^{1.5}$, and Ising model lattices with different side lengths L . The fits were performed on a disc with radius equal to the maximum r_{ij} displayed for a given curve. (b) and (c): The maximum fitting error $|\tilde{V}_{\text{fit}}^{[N_f]} - 1/r_{ij}|$ at selected values of r_{ij} as functions of N_t (b) and N_f (c). In (b), the open circles correspond to $N_f = 0$ and the closed circles to $N_f = 10$. In (c), $N_t = 12$. The fits in (b) and (c) were performed on discs of radius $r_{ij} = 36$ with $L = 199$ and a weight function of $r_{ij}^{1.5}$.

with $\hat{O}_i = \sum_j V_{ij} n_j$. Each \hat{O}_i can be represented as a snake-like MPO with bond dimension $D = 3$, and the cost of contracting a single \hat{O}_i expectation value is then $O(A\chi^3 D_S^3)$ with $\chi \sim DD_S^2$, with $O(A)$ such terms. The crossover with our (snake) PEPO representation then occurs when $A \sim 8N_t$, which for $N_t = 10$ corresponds to $A < 100$. Thus in either comparison, a crossover between our PEPO representation

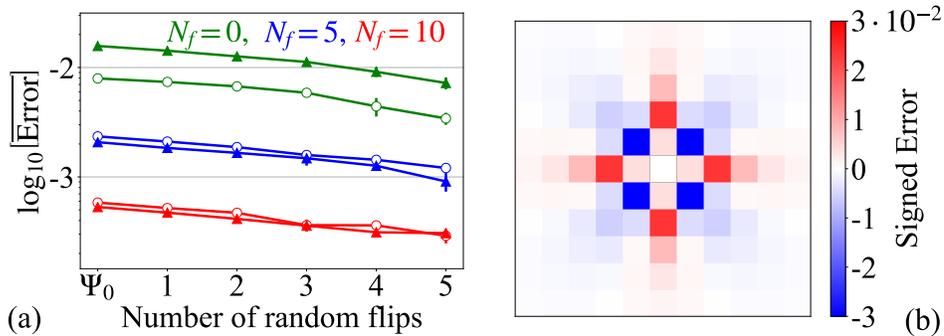


Figure 3.3: (a) Average accuracy of energy per site expectation values for 6×6 FM and AFM trial PEPS with $D_S = 1$. The solid triangular markers show FM states while the open circles show AFM states. Ψ_0 is a true FM or AFM state, while the “ x flip” regions are Ψ_0 perturbed by x random spin flips. The average error is taken over 5 PEPS for each x and each N_f . (b) The signed error $1/r_{ij} - \tilde{V}_{\text{fit}}^{[0]}(r_{ij})$, where $r_{ij} = 0$ is the white square in the center, each adjacent square is $r_{ij} = 1$, etc. For (a)-(b) the fitted potentials are obtained from Eq. (3.4) with $N_t = 12$.

and other implementations of the long-range operator is achievable already at modest lattice sizes.

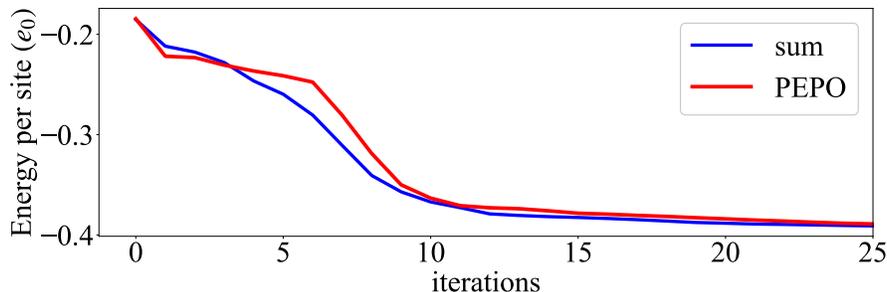
3.6 Results

To numerically test our PEPO’s faithful discretized representation of long range interactions, we have explicitly constructed a long-range $S=1/2$ Heisenberg Hamiltonian on 4×4 , 6×6 , and 8×8 square lattices,

$$\hat{H} = \sum_{i<j} \frac{\vec{S}_i \cdot \vec{S}_j}{r_{ij}}, \quad (3.5)$$

in which every pair of spins has an interaction strength of Coulomb form. To represent this operator, we first used the fitting scheme described in Eq. (3.4) with $N_t = 12$. Figure 3.3(a) shows the accuracy of the energy per site expectation value (e_0) for 6×6 trial ferromagnetic (FM) and anti-ferromagnetic (AFM) PEPS with $D_S = 1$. The FM and AFM states show similar levels of error for a given value of N_f , indicating that the fitted operator can obtain similar levels of error even for states which have different structures of the *signed* error.

We next performed a simple gradient-based variational optimization for the ground state PEPS with $D_S = 1, 2$ [39, 40]. Note that our goal here is not to demonstrate fully converged physics with respect to the PEPS bond dimension, which will be discussed in future studies, but rather to show that our PEPO leads to a stable



	χ	sum e_0	PEPO e_0	$\langle \psi_0^{[P]} \psi_0^{[S]} \rangle$
$4 \times 4, D_S = 1$	40	-0.184314	-0.184425	0.999244
$4 \times 4, D_S = 2$	100	-0.408209	-0.408492	0.999070
$4 \times 4, \text{exact}$	–	0.424577	–	–
$8 \times 8, D_S = 1$	40	-0.193983	-0.193861	0.994549
$8 \times 8, D_S = 2$	120	-0.414653	-0.414422	0.989271
$8 \times 8, \text{exact}$	–	-0.431648	–	–

Figure 3.4: *Top*: The trajectories over the first 25 iterations of the energy optimization for the 4×4 $D_S = 2$ system using the PEPO and the explicit sum over all $O(A^2)$ terms in (3.5). The long tails of the trajectories are excluded for clarity. *Bottom*: Ground state energies per site e_0 for the Hamiltonian (3.5) with various system sizes and bond dimensions. The fifth column is the overlap of the normalized ground states obtained with the two different methods. In all cases $N_f = 4$ and $N_t = 12$. The “exact” rows are the results of converged DMRG calculations.

optimization procedure. Here we refined the fit for each lattice size to ensure that the maximum PEPO fitting error was limited to $\sim 4.5 \cdot 10^{-4}$ with only $N_f = 4$, $N_t = 12$. Fig. 3.4 shows the initial convergence behavior of the energy optimization using the PEPO compared to the same optimization using the more expensive sum over terms formalism. We observe that the trajectories are similar and the use of the PEPO does not change the stability of the gradient optimization, although it does require a larger value of χ . The small- D_S converged energies and normalized wavefunction overlaps are given in Fig. 3.4. In all cases, the CF-PEPO nicely reproduces the explicit sum-over-terms algorithm, as the maximum fitting error is faithfully reflected in the accuracy of e_0 . It is also interesting to see that the error of the ground-state energy using $D_S = 2$ is $\sim 3\%$ for both the 4×4 and 8×8 lattice, suggesting that the entanglement does not grow significantly with system size despite the long-range interaction, which is a similar observation to other simulations of physical Coulombic systems.

3.7 Summary

In summary, we have detailed the efficient construction of a PEPO capable of encoding long-range interactions in 2D TNS that maintains the strengths of tensor network algorithms: systematically improvable accuracy and linear computational complexity in the system size. Despite an increased cost prefactor compared to local simulations, this approach allows for the possibility of practically including long-range interactions in numerical studies of physically realistic systems that have an entanglement structure consistent with PEPS. The crossover between our approach and other more naive implementations of long-range interactions can be achieved at modest system sizes. In the context of *ab initio* electronic structure calculations, while there remain many issues to explore, in particular associated with the continuum limit of relevance to such applications, this advance presents a first step towards these calculations using higher dimensional tensor networks.

*Chapter 4***A SIMPLIFIED AND IMPROVED APPROACH TO TENSOR NETWORK OPERATORS IN TWO DIMENSIONS**

This chapter is based on the following publication:

1. O'Rourke, M. J. & Chan, G. K.-L. Simplified and improved approach to tensor network operators in two dimensions. *Phys. Rev. B* **101**, 205142. doi:10.1103/PhysRevB.101.205142 (20 May 2020).

4.1 Abstract

Matrix product states (MPS) and matrix product operators (MPOs) are one-dimensional tensor networks that underlie the modern density matrix renormalization group (DMRG) algorithm. The use of MPOs accounts for the high level of generality and wide range of applicability of DMRG. However, current algorithms for two dimensional (2D) tensor network states, known as projected entangled-pair states (PEPS), rarely employ the associated 2D tensor network operators, projected entangled-pair operators (PEPOs), due to their computational cost and conceptual complexity. To lower these two barriers, we describe how to reformulate a PEPO into a set of tensor network operators that resemble MPOs by considering the different sets of local operators that are generated from sequential bipartitions of the 2D system. The expectation value of a PEPO can then be evaluated on-the-fly using only the action of MPOs and generalized MPOs at each step of the approximate contraction of the 2D tensor network. This technique allows for the simpler construction and more efficient energy evaluation of 2D Hamiltonians that contain finite-range interactions, and provides an improved strategy to encode long-range interactions that is orders of magnitude more accurate and efficient than existing schemes.

4.2 Introduction

The density matrix renormalization group (DMRG) algorithm [25, 26] is a popular and successful [56] technique for finding the variational ground state of the Schrödinger equation in one spatial dimension (1D). In its modern form, the variational wave function and the Hamiltonian are represented as 1D tensor networks (TNs), namely matrix product states (MPS) [22, 53–55] and matrix product operators (MPOs) [94, 107, 112–114]. The widespread use of MPOs has allowed for

the development of very general, efficient implementations of the algorithm [115], permitting the study of large classes of complex problems in a relatively black-box manner.

However, the two-dimensional (2D) generalization of MPS, known as projected entangled-pair states (PEPS) [23, 63–65], and their associated ground state algorithms [34, 36, 39, 40, 87, 114] have not yet come close to the same level of generality or range of applicability. One significant reason for this is that projected entangled-pair operators (PEPOs) [45, 46, 70, 107], the 2D generalization of MPOs, have been scarcely used in the tensor network literature to date. The ground state optimization algorithms employed by most authors instead utilize a significantly less general representation of the Hamiltonian that is restricted to relatively local interactions [34, 36, 39, 40, 87, 114, 116]. We conjecture that this under-utilization of PEPOs in favor of simpler operator representations can be attributed to two facts. Firstly, the construction of a PEPO for an arbitrary 2D Hamiltonian is more conceptually complicated than the construction of the MPO for the analogous Hamiltonian in 1D, which itself is still more complicated than building the local operators currently used in 2D simulations. Secondly, when compared to the local operators currently used in 2D, the use of PEPOs in a ground state optimization significantly increases the computational cost of the approximate contraction algorithms for 2D tensor networks in both the finite [35] and infinite (iPEPS) [34, 37, 87] cases.

In this article we describe how to overcome both the computational and conceptual complexity of using general tensor network operator representations of the Hamiltonian in 2D algorithms. To do so, we first briefly summarize the MPO formalism and review some well-known examples that are central ideas in this work (Section 4.3). We then introduce a new type of tensor network operator known as a generalized MPO (gMPO), which is closely related to the traditional MPO (Section 4.4). Next we show how to reformulate the calculation of the expectation value of a general PEPO into a series of operations involving only MPOs and gMPOs, which we call the boundary gMPO method (Section 4.4). Since the language of MPOs is much better known than that of PEPOs, this reformulation serves to simplify the construction of general 2D Hamiltonians for most readers. In Section 4.5 we demonstrate this simplicity by reporting the explicit forms of the gMPOs for various representative types of 2D Hamiltonians. We also show that the new scheme sacrifices no accuracy compared to the explicit usage of a PEPO, while providing large speedups in computational time. In addition, a new scheme for efficiently constructing and

evaluating Hamiltonians with long-range interactions is shown to be many orders of magnitude more accurate and efficient than existing PEPO-based approaches [70, 117, 118].

4.3 Matrix Product Operators (MPOs)

Since many detailed and comprehensive presentations of MPOs already exist [22, 45, 46, 94, 107], this section will simply contain a brief overview in order to establish notation, as well as some simple examples which we will call upon in later sections.

Overview

Consider a 1D system which has been discretized into L localized sites, each with a local Hilbert space \mathcal{H}_i of dimension d_i . A general operator \hat{O} acting on such a system can be written as,

$$\hat{O} = \sum_{\{\hat{o}_i\}} O^{o_1 o_2 \dots o_L} \hat{o}_1 \hat{o}_2 \dots \hat{o}_L, \quad (4.1)$$

where $\{\hat{o}_i\}$ is the set of local operators acting on \mathcal{H}_i and O is a rank- L tensor with indices o_i whose dimensions are equal to the cardinality of their respective set $\{\hat{o}_i\}$. O contains the weights associated with all possible configurations of the local operators \hat{o}_i .

By fixing the indices, a specific element $O^{o_1 o_2 \dots o_L}$ of the tensor O can then be decomposed into a product of matrices $W[i]$,

$$O^{o_1 o_2 \dots o_L} = \sum_{\{\alpha\}} W_{\alpha_1}^{o_1} [1] W_{\alpha_1 \alpha_2}^{o_2} [2] \dots W_{\alpha_{L-1}}^{o_L} [L], \quad (4.2)$$

where α indexes the so-called ‘‘virtual’’ or ‘‘auxiliary’’ indices which are introduced to perform the matrix multiplication. In Eq. (4.2) the o_i are simply labels, intended to indicate that each matrix $W[i]$ is chosen specifically so that their product reproduces the element $O^{o_1 o_2 \dots o_L}$. However, if the labels are all reinterpreted as their corresponding indices from Eq. (4.1), then we see that the full tensor O can be reconstructed as the contraction over rank-3 tensors $W[i]$.

Similar to MPS, this decomposition of a rank- L tensor into L rank-3 tensors is motivated by the fact that most operators of interest do not contain general L -body interactions, but instead are usually limited to few-body terms. This means that, while in general this decomposition could be exponentially expensive, often the O tensor is quite sparse and such a transformation can be a highly efficient way to represent the full tensor.

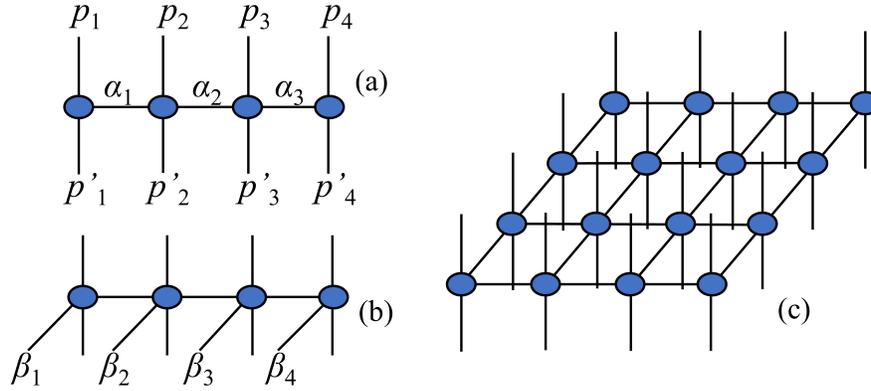


Figure 4.1: Tensor network diagrams of (a) an MPO, (b) a gMPO, and (c) a PEPO.

It is common and frequently useful to associate the operators \hat{o}_i with their corresponding coefficient tensor $W[i]$ according to,

$$\hat{W}_{\alpha_{i-1}\alpha_i}[i] = \sum_{o_i} W_{\alpha_{i-1}\alpha_i}^{o_i}[i] \hat{o}_i. \quad (4.3)$$

This yields matrices $\hat{W}[i]$ in which every element is a $d_i \times d_i$ local operator acting on \mathcal{H}_i . The full operator \hat{O} is thus reconstructed via simple matrix multiplication,

$$\hat{O} = \sum_{\{\alpha\}} \hat{W}_{\alpha_1}[1] \hat{W}_{\alpha_1\alpha_2}[2] \dots \hat{W}_{\alpha_{L-1}}[L], \quad (4.4)$$

and the set of matrices $\{\hat{W}[i]\}$ are referred to as the MPO representation of \hat{O} . This form of an MPO is commonly used throughout the literature, and will be heavily utilized in the remainder of this work.

We will now relate the MPO form in Eq. (4.4) to the common diagrammatic representation, as seen in Fig. 4.1. Since every element of $\hat{W}_{\alpha_{i-1}\alpha_i}[i]$ is itself a $d_i \times d_i$ matrix, each individual numerical element can be exposed by introducing two new indices p_i and p'_i , each of dimension d_i . By fixing each of α_{i-1} , α_i , p_i , and p'_i , the expression $(\hat{W}_{\alpha_{i-1}\alpha_i}[i])_{p_i p'_i}$ yields a single number. More commonly written as $\hat{W}_{\alpha_{i-1}\alpha_i}^{p_i p'_i}[i]$, the correspondence to the rank-4 tensors shown in MPO diagrams becomes apparent. The new indices p_i and p'_i are the so-called “physical” indices, which map the action of the local operators onto the corresponding site tensors of an MPS.

Examples

Frequently the operator that one wants to encode as an MPO is a Hamiltonian \hat{H} , so that the DMRG algorithm can be used to find its ground state in the form of an MPS.

Here we will explicitly write out the well-known matrices $\hat{W}[i]$ which make up the MPO representations of several common Hamiltonians consisting of 1- and 2-body terms. There are multiple techniques that can be used to derive these matrices, each with their own conventions and notation, but in this work we will remain agnostic to these different languages in an attempt to make the presentation in the following sections as conceptually simple and widely accessible as possible. To do so, we will simply refer back to these explicit examples. In lieu of derivations we will point to helpful references for readers who do not already have a preferred technique for understanding the form of MPO matrices.

Nearest-neighbor interactions

Consider a system of L sites, which are indexed by $i \in \{1, 2, \dots, L\}$, and a Hamiltonian consisting of local terms and nearest-neighbor interactions of the form $\hat{H} = \sum_{i=1}^L \hat{C}_i + \sum_{i=1}^{L-1} \hat{A}_i \hat{B}_{i+1}$. In the MPO literature this Hamiltonian is usually written with $\hat{B} = \hat{A}$ so that the interaction is symmetric and \hat{H} is Hermitian, however in this paper we will always keep the operators distinct for purposes of notational clarity, even though this means that some Hamiltonians under consideration will be non-Hermitian when $\hat{B} \neq \hat{A}$. The MPO matrices for this Hamiltonian, denoted \hat{W}_{NN} , are given by,

$$\begin{aligned} \hat{W}_{NN}[1] &= \begin{pmatrix} \hat{C} & \hat{A} & \hat{I} \end{pmatrix}, \quad \hat{W}_{NN}[L] = \begin{pmatrix} \hat{I} & \hat{B} & \hat{C} \end{pmatrix}^T, \\ \hat{W}_{NN}[i] &= \begin{pmatrix} \hat{I} & \hat{O} & \hat{O} \\ \hat{B} & \hat{O} & \hat{O} \\ \hat{C} & \hat{A} & \hat{I} \end{pmatrix}, \end{aligned} \quad (4.5)$$

where \hat{I} is the identity operator and \hat{O} is the zero operator.

If instead the interaction is symmetric so that $\hat{H} = \sum_{i=1}^L \hat{C}_i + \sum_{i=1}^{L-1} (\hat{A}_i \hat{B}_{i+1} + \hat{B}_i \hat{A}_{i+1})$, then the MPO matrices (\hat{W}_{NN-sym}) are given by,

$$\begin{aligned} \hat{W}_{NN-sym}[1] &= \begin{pmatrix} \hat{C} & \hat{A} & \hat{B} & \hat{I} \end{pmatrix}, \\ \hat{W}_{NN-sym}[L] &= \begin{pmatrix} \hat{I} & \hat{B} & \hat{A} & \hat{C} \end{pmatrix}^T, \\ \hat{W}_{NN-sym}[i] &= \begin{pmatrix} \hat{I} & \hat{O} & \hat{O} & \hat{O} \\ \hat{B} & \hat{O} & \hat{O} & \hat{O} \\ \hat{A} & \hat{O} & \hat{O} & \hat{O} \\ \hat{C} & \hat{A} & \hat{B} & \hat{I} \end{pmatrix}. \end{aligned} \quad (4.6)$$

In general, for an exact MPO representation of a Hamiltonian \hat{H} , the required bond dimension of the MPO matrices is $D = 2 + b \cdot r$, where r is the maximum distance over which interactions occur and b is the number of unique operators that act “first” in the interactions. This is reflected in Eq. (4.5) where $r = 1$ and $b = 1$, and in Eq. (4.6) where $r = 1$ and $b = 2$. To understand these patterns, as well as the form of the MPO matrices in this section, we recommend Ref. [46].

Exponentially decaying interactions

One important exception to the above result is the MPO representation of a Hamiltonian which has long-range interactions that decay exponentially, such as $\hat{H} = \sum_i \hat{C}_i + \sum_{i < j} e^{-\lambda(j-i)} \hat{A}_i \hat{B}_j$. Here we have introduced a second index j which runs from $i + 1$ to L . Despite the fact that $r = L$ in this case, the Hamiltonian has an exact, compact representation with $D = 3$ MPO matrices (\hat{W}_{exp}) of the form,

$$\begin{aligned} \hat{W}_{exp}[1] &= \begin{pmatrix} \hat{C} & \hat{A} & \hat{I} \end{pmatrix}, \quad \hat{W}_{exp}[L] = \begin{pmatrix} \hat{I} & e^{-\lambda} \hat{B} & \hat{C} \end{pmatrix}^T, \\ \hat{W}_{exp}[i] &= \begin{pmatrix} \hat{I} & \hat{O} & \hat{O} \\ e^{-\lambda} \hat{B} & e^{-\lambda} \hat{I} & \hat{O} \\ \hat{C} & \hat{A} & \hat{I} \end{pmatrix}. \end{aligned} \quad (4.7)$$

Refs. [45, 107, 108, 117] provide insight into why this is possible for the unique case of exponential interactions.

A special case of this representation, which will prove useful in later sections, is when $\lambda = 0$. The Hamiltonian then has long-range interactions between every pair of sites but the strength of the interactions are all the same, $\hat{H} = \sum_i \hat{C}_i + \sum_{i < j} \hat{A}_i \hat{B}_j$. We will denote this special case with its own MPO notation: $\hat{W}_{uniform}$.

Much like before, if the interactions are symmetric so that $\hat{H} = \sum_i \hat{C}_i + \sum_{i \neq j} e^{-\lambda|j-i|} \hat{A}_i \hat{B}_j$ (where now both $i, j \in \{1, \dots, L\}$), the MPO matrices become,

$$\begin{aligned} \hat{W}_{exp-sym}[1] &= \begin{pmatrix} \hat{C} & \hat{A} & \hat{B} & \hat{I} \end{pmatrix}, \\ \hat{W}_{exp-sym}[L] &= \begin{pmatrix} \hat{I} & e^{-\lambda} \hat{B} & e^{-\lambda} \hat{A} & \hat{C} \end{pmatrix}^T, \\ \hat{W}_{exp-sym}[i] &= \begin{pmatrix} \hat{I} & \hat{O} & \hat{O} & \hat{O} \\ e^{-\lambda} \hat{B} & e^{-\lambda} \hat{I} & \hat{O} & \hat{O} \\ e^{-\lambda} \hat{A} & \hat{O} & e^{-\lambda} \hat{I} & \hat{O} \\ \hat{C} & \hat{A} & \hat{B} & \hat{I} \end{pmatrix}. \end{aligned} \quad (4.8)$$

Again, we will give the special case of $\lambda = 0$ its own notation, $\hat{W}_{uniform-sym}$, which will prove useful in the coming sections.

General two-body long-range interactions

As mentioned previously, exact MPO representations of Hamiltonians with general long-range interaction coefficients $\hat{H}_{gen} = \sum_i \hat{C}_i + \sum_{i<j} V_{ij} \hat{A}_i \hat{B}_j$ require a bond dimension which is proportional to L [45]. However, if V_{ij} is a smoothly decaying function of the distance between two sites, $V_{ij} = f(j - i)$, then highly accurate *approximate* MPO representations of \hat{H}_{gen} can often be found which have finite, constant bond dimensions. The traditional technique is to fit $f(j - i)$ by a sum of exponentials [107, 108],

$$f(j - i) = \sum_{k=1}^K a_k e^{-\lambda_k(j-i)}. \quad (4.9)$$

This yields an MPO representation of \hat{H}_{gen} with bond dimension $K + 2$, where the MPO matrices take the form,

$$\begin{aligned} \hat{W}_{K-exp}[1] &= \left(\hat{C} \quad a_1 \hat{A} \quad a_2 \hat{A} \quad \cdots \quad a_K \hat{A} \quad \hat{I} \right), \\ \hat{W}_{K-exp}[L] &= \left(\hat{I} \quad e^{-\lambda_1} \hat{B} \quad e^{-\lambda_2} \hat{B} \quad \cdots \quad e^{-\lambda_K} \hat{B} \quad \hat{C} \right)^T, \\ \hat{W}_{K-exp}[i] &= \begin{pmatrix} \hat{I} & \hat{0} & \hat{0} & \cdots & \hat{0} & \hat{0} \\ e^{-\lambda_1} \hat{B} & e^{-\lambda_1} \hat{I} & \hat{0} & \cdots & \hat{0} & \hat{0} \\ e^{-\lambda_2} \hat{B} & \hat{0} & e^{-\lambda_2} \hat{I} & \cdots & \hat{0} & \hat{0} \\ \vdots & \vdots & \vdots & \ddots & \hat{0} & \hat{0} \\ e^{-\lambda_K} \hat{B} & \hat{0} & \hat{0} & \cdots & e^{-\lambda_K} \hat{I} & \hat{0} \\ \hat{C} & a_1 \hat{A} & a_2 \hat{A} & \cdots & a_K \hat{A} & \hat{I} \end{pmatrix}. \end{aligned} \quad (4.10)$$

The accuracy of the representation $\{\hat{W}_{K-exp}\}$ is determined by the quality of the fit in Eq. (4.9).

Although this is often a reasonably accurate approach, several more sophisticated techniques have been developed in recent years which are based on the singular value decomposition (SVD) of blocks of V_{ij} [94, 97]. These methods also work most effectively when V_{ij} is a smooth function of the distance, but they are able to fit more general functions f that may be challenging to represent directly with exponentials like those in Eq. (4.9) [94]. They also can be a bit more efficient, producing a higher accuracy representation of \hat{H}_{gen} with a smaller bond dimension than Eq. (4.10) [97].

In this work, we utilize the technique described in Ref. [97]. The basic idea is that

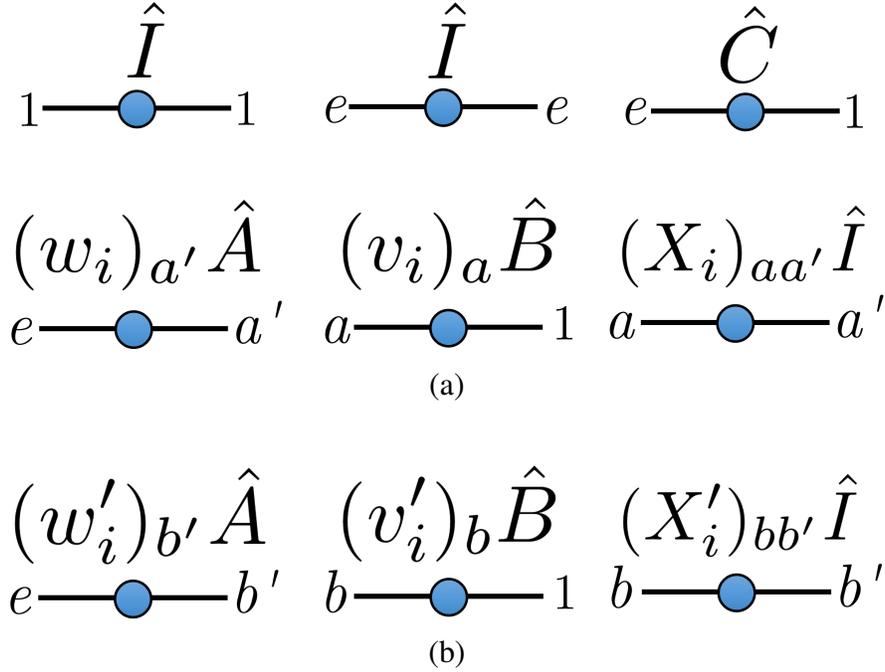


Figure 4.2: A set of tensor network diagrams that represent the elements of the operator-valued MPO matrix $\hat{W}_{gen}[i]$ (a), along with the additional elements needed for $\hat{W}_{gen-sym}[i]$ (b). Here we assume $\hat{W}_{gen}[i]$ is a $(2 + l_i) \times (2 + r_i)$ matrix and $\hat{W}_{gen-sym}[i]$ is a $(2 + 2l_i) \times (2 + 2r_i)$ matrix. We use the symbol “1” to denote the first value of a given index, e to denote the final value of a given index, a to denote the set of values ranging from 2 to $l_i + 1$, b to denote the set of values ranging from $l_i + 2$ to $2l_i + 1$, a' to denote the set of values ranging from 2 to $r_i + 1$, and b' to denote the set of values ranging from $r_i + 2$ to $2r_i + 1$. This index labelling corresponds directly to the expressions in Equations (4.11) and (4.12).

the MPO matrices \hat{W}_{gen} for the general Hamiltonian \hat{H}_{gen} can be written as,

$$\hat{W}_{gen}[i] = \begin{pmatrix} \hat{I} & \hat{0} & \hat{0} \\ (v_i)_a \hat{B} & (X_i)_{aa'} \hat{I} & \hat{0} \\ \hat{C} & (w_i)_{a'} \hat{A} & \hat{I} \end{pmatrix}, \quad (4.11)$$

where \vec{v}_i is a column vector of coefficients that has length l_i and is indexed by a , X_i is an $l_i \times r_i$ matrix of coefficients indexed by a and a' , and \vec{w}_i is a row vector of coefficients that has length r_i and is indexed by a' , yielding a $(2 + l_i) \times (2 + r_i)$ MPO matrix. We write the indexed elements of \vec{v}_i , \vec{w}_i , and X_i in Eq. (4.11) to remind the reader of the shape of these quantities. For clarity, tensor network diagrams for this matrix are given in Fig. 4.2(a). If the coefficients contained in $\hat{W}_{gen}[i]$ can be, to a good approximation, related to the coefficients contained in $\hat{W}_{gen}[i + 1]$ by a linear transformation, then the MPO matrices for each site can be successively generated

by finding the correct linear transformation on the coefficients contained in the MPO matrix on the previous site. These linear transformations can be found by taking SVDs of certain blocks of the upper triangle of V_{ij} . It is observed in [97] that if V_{ij} is a smooth function of the distance $|j - i|$, the transformations are often compact (i.e. their dimensions do not scale with L) and highly accurate because sub-blocks of the upper triangle of V_{ij} are low-rank. These ideas are developed in full detail in the supplementary information of Ref. [97]¹.

The form of this MPO matrix can be viewed as a direct generalization of $\hat{W}_{uniform}$. The ‘‘coefficients’’ in adjacent $\hat{W}_{uniform}$ matrices can be related to each other via the simplest possible linear transformation ($X_{1 \times 1} = 1$, $\vec{v} = \vec{w} = 1$) because all the interactions are of identical strength and thus all sub-blocks of V_{ij} are rank 1. However, when the interaction coefficients vary with distance and the sub-blocks of the upper triangle of V_{ij} are rank- l , the single \hat{I} in the center of $\hat{W}_{uniform}$ gets generalized to an $l \times l$ block $X_{l \times l} \hat{I}$ in \hat{W}_{gen} . By extension, \vec{v} and \vec{w} undergo the same generalization. The MPOs \hat{W}_{exp} and \hat{W}_{K-exp} are special, simple cases of this generalization.

As a final note, if the interactions in the general Hamiltonian \hat{H}_{gen} become symmetric so that $\hat{H} = \sum_i \hat{C}_i + \sum_{i \neq j} V_{ij} \hat{A}_i \hat{B}_j = \sum_i \hat{C}_i + \sum_{i < j} V_{ij} \hat{A}_i \hat{B}_j + \sum_{j < i} V_{ij} \hat{B}_j \hat{A}_i$, then the general MPO matrices become,

$$\hat{W}_{gen-sym}[i] = \begin{pmatrix} \hat{I} & \hat{0} & \hat{0} & \hat{0} \\ (v_i)_a \hat{B} & (X_i)_{aa'} \hat{I} & \hat{0} & \hat{0} \\ (v'_i)_b \hat{A} & \hat{0} & (X'_i)_{bb'} \hat{I} & \hat{0} \\ \hat{C} & (w_i)_{a'} \hat{A} & (w'_i)_{b'} \hat{B} & \hat{I} \end{pmatrix}. \quad (4.12)$$

Tensor network diagrams representing this matrix are given in Fig. 4.2. Here we have introduced the additional indices b and b' to index the new vectors \vec{v}'_i , \vec{w}'_i and the new matrix X'_i , as described in Fig. 4.2. If we have the additional property that interaction coefficients themselves are symmetric, $V_{ij} = V_{ji}$, then the above expression can be simplified according to $\vec{v}'_i = \vec{v}_i$, $\vec{w}'_i = \vec{w}_i$, $X'_i = X_i$.

4.4 PEPO expectation value via generalized MPOs

Generalized MPOs (gMPOs)

In order to relate the contraction of PEPOs to the well-known 1D MPOs described in Section 4.3, we must first introduce the notion of a generalized MPO (gMPO).

¹It should be noted that in the referenced article there is a typo in the explicit expressions for the compressed MPO matrices. The local operators associated with the X block of each matrix should be \hat{I} , not \hat{n} , as in Eq. (4.11).

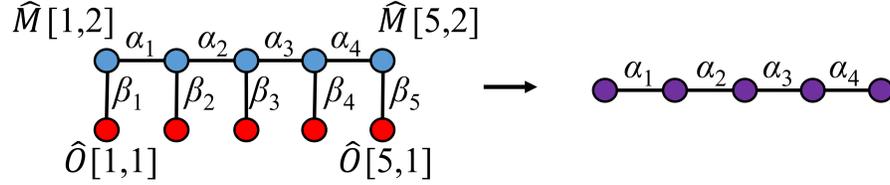


Figure 4.3: A gMPO-based representation of the two-row example Hamiltonian in Section 4.4 for a 2×5 system. *Left*: The gMPO tensors (blue) appear on the sites in row 2, while the complementary operator vectors (red) appear on the sites in row 1. Physical indices are suppressed for simplicity. *Right*: The resulting tensor network along row 2 (an MPO) after contractions over the β indices have been performed.

In a gMPO, the operator-valued MPO matrices $\hat{W}[i]$ are elevated to rank-3 tensors, which will be indicated by the addition of a virtual index $\beta_i \in \{1, 2, \dots, g\}$. The new operator-valued, rank-3 gMPO tensors will be denoted by $\hat{M}_{\beta_i}[i]$. Exposing all the indices explicitly, this gives a rank-5 tensor $M_{\alpha_{i-1}\alpha_i\beta_i}^{p_i p'_i}[i]$, which is shown in diagrammatic form in Fig. 4.1.

The basic notion of a gMPO is that for each value of β_i , a different MPO *matrix* $\hat{W}[i]$ can be encoded in the gMPO tensor. In the simplest case β_i only takes a single value ($g = 1$) and thus every gMPO tensor can only represent a single MPO matrix, reducing the gMPO back to a regular MPO. If instead β_i takes two values ($g = 2$), then every tensor can represent two different MPO matrices, and the gMPO can encode 2^L different 1D MPOs. In practice, however, the β_i are not “free” indices but are instead summed over in the final network just like the α indices in a regular MPO (see Eq. (4.4)). The proper notion of a gMPO is therefore as a tensor network operator that can represent a sum of many regular 1D MPOs after the β_i are appropriately summed over. This formulation is useful because it provides a flexible framework in which operators in regular MPOs can be coupled with other operators that act “outside” of the 1D domain of the regular MPO. In general it allows for the complete coupling of two distinct MPOs into one, however in this work we only utilize a simpler special case in which specific local operators are coupled together. Much like how a local operator on site i can be coupled to a local operator on site $i + 1$ by summing over the index α_i in a regular MPO, we use the gMPO formalism to couple a local operator that acts “below” site i to the local operators on site i by performing an appropriate sum over β_i .

For clarity, let us consider a simple example. Given a 2D system of size $2 \times L$ consisting of two rows with L sites each, we can label each site by (i, y) , where

$i \in \{1, \dots, L\}$ as usual and $y \in \{1, 2\}$, as depicted in Fig. 4.3. Consider the Hamiltonian $\hat{H} = \hat{H}_1 + \hat{H}_2 = \sum_i (\hat{A}_{i,1} \hat{B}_{i,2} + \hat{A}_{i,2} \hat{B}_{i+1,2})$, where there are nearest-neighbor interactions between row 1 and row 2 (\hat{H}_1), as well as nearest-neighbor interactions within row 2 (\hat{H}_2). This Hamiltonian can be represented by a simple gMPO (\hat{M}) acting on row 2 along with the complementary operators (\hat{O}) that act locally on the sites in row 1.

Since there are no interactions between sites in row 1, the operators $\{\hat{O}[i, 1]\}$ that are applied in this row take the form of vectors, like those at the ends of a regular MPO, but applied along the β_i index instead of α (see Fig. 4.3),

$$O_{\beta_i}^{p_{i,1} p'_{i,1}}[i, 1] \rightarrow \hat{O}_{\beta_i}[i, 1] = \begin{pmatrix} \hat{I}_{i,1} & \hat{A}_{i,1} \end{pmatrix}. \quad (4.13)$$

To couple these operators with the local operators in row 2, as well as to encode the nearest-neighbor interactions within row 2, gMPO tensors can be used in row 2. They take the form,

$$\begin{aligned} \hat{M}_1[i, 2] &= \hat{W}_{NN}[i], \\ \hat{M}_2[1, 2] &= \begin{pmatrix} \hat{B}_{1,2} & \hat{O} & \hat{O} \end{pmatrix}, \quad \hat{M}_2[L, 2] = \begin{pmatrix} \hat{O} & \hat{O} & \hat{B}_{L,2} \end{pmatrix}^T, \\ \hat{M}_2[1 < i < L, 2] &= \begin{pmatrix} \hat{O} & \hat{O} & \hat{O} \\ \hat{O} & \hat{O} & \hat{O} \\ \hat{B}_{i,2} & \hat{O} & \hat{O} \end{pmatrix}, \end{aligned} \quad (4.14)$$

where \hat{W}_{NN} is from Section 4.3 (with $\hat{C} = \hat{O}$). The reason why the matrix $\hat{M}_2[i, 2]$ takes this form can be understood by explicitly considering what happens during the contraction over β_i for a given column i .

$$\begin{aligned} \sum_{\beta_i} \hat{O}_{\beta_i}[i, 1] \hat{M}_{\beta_i}[i, 2] &= \\ \hat{I}_{i,1} \cdot \begin{pmatrix} \hat{I}_{i,2} & \hat{O} & \hat{O} \\ \hat{B}_{i,2} & \hat{O} & \hat{O} \\ \hat{O} & \hat{A}_{i,2} & \hat{I}_{i,2} \end{pmatrix} + \hat{A}_{i,1} \cdot \begin{pmatrix} \hat{O} & \hat{O} & \hat{O} \\ \hat{O} & \hat{O} & \hat{O} \\ \hat{B}_{i,2} & \hat{O} & \hat{O} \end{pmatrix} \\ &= \begin{pmatrix} \hat{I} & \hat{O} & \hat{O} \\ \hat{B}_{i,2} & \hat{O} & \hat{O} \\ \hat{A}_{i,1} \hat{B}_{i,2} & \hat{A}_{i,2} & \hat{I} \end{pmatrix}. \end{aligned} \quad (4.15)$$

The resulting tensor network operator now looks like a regular MPO along row 2 (see Fig. 4.3), and the form of its matrices looks very similar to \hat{W}_{NN} (Eq. (4.5)),

which encodes non-symmetric nearest neighbor interactions. The only difference is that in the place of \hat{C} , the 1-body on-site term in Section 4.3, there is now the inter-row interaction term \hat{H}_1 for column i . Thus, if these MPO matrices are now all contracted together along the α indices in row 2, we will exactly recover all the terms in our original two row Hamiltonian.

The function of $\hat{M}_2[i, 2]$ is thus evident: it couples the inter-row interactions into an intra-row MPO matrix in a consistent manner with the structure of the intra-row MPO. Without \hat{M}_2 , the action of $\hat{A}_{i,1}$ could not be selectively coupled into specific matrix elements of \hat{M}_1 . Thus, the form of \hat{M}_2 can be simply determined based on an understanding of the structure of the “in-row” MPO matrix \hat{M}_1 ; namely, to which matrix elements the “external” operators should couple. Although this formalism may appear unnecessarily general in the context of this simple example, its full utility will become apparent in the subsequent sections as more complicated Hamiltonians are considered.

Evaluation of PEPO expectation values using gMPOs

To this point, the Hamiltonians under consideration have acted on lattices that are either strictly or quasi- one dimensional. In this section we will present an algorithm that utilizes the gMPO formalism to evaluate the expectation value of fully 2D Hamiltonians with the same level of generality as PEPOs, but with simpler and more familiar concepts. This presentation will focus on the case of a finite $L_x \times L_y$ rectangular lattice, but prospects for its extension to the infinite case will be discussed in Section 4.6. The concepts for this technique begin with consideration of the three subsets of local operators that are distinguished by a bipartitioning of the system. Namely, given the full system Hamiltonian \hat{H} represented by a localized structure such as a PEPO and a horizontal bipartition of it (as depicted in Fig. 4.4(a)), all the local operators in \hat{H} can be grouped into three mutually exclusive groups: (i) those for which there are interactions between sites that are all below the line (\hat{H}_{bot}), (ii) all above the line (\hat{H}_{top}), or (iii) those for which interactions occur across the line (\hat{O}_{int}). This decomposition,

$$\hat{H} = \hat{H}_{\text{bot}} + \sum_{ij} h_{ij} \hat{O}_i \hat{O}_j + \hat{H}_{\text{top}}, \quad (4.16)$$

where i indexes sites below the partition, j indexes sites above the partition, and h_{ij} contains the coefficients for the interactions that get “cut”, is a familiar concept in 1D for the analysis of MPOs and is the basis of an efficient implementation of the

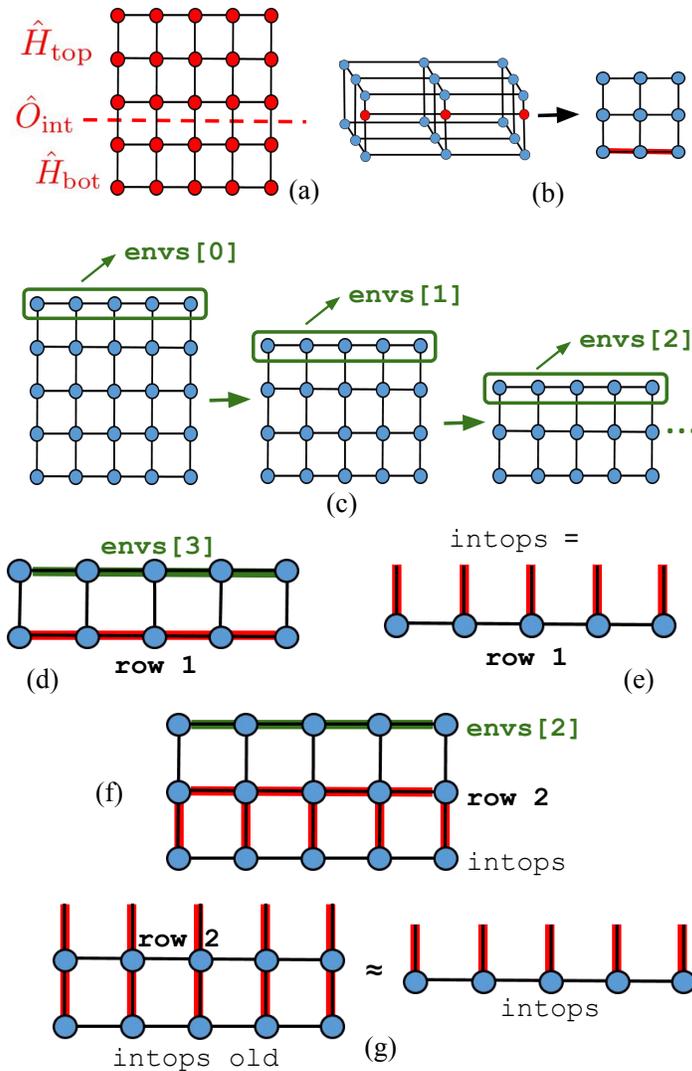


Figure 4.4: The first full iteration of the boundary gMPO algorithm for a 5×5 PEPS. (a) A 5×5 PEPO (with physical indices that are suppressed) that is bipartitioned by a cut between rows 2 and 3. (b) A useful diagrammatic definition: when a flat square lattice TN diagram is drawn with some red bonds and some black bonds, the red corresponds to where a tensor network operator has been sandwiched between the bra and ket. Black bonds contain just bra and ket virtual indices. Figures (c)-(g) are the diagrams that directly correspond to the algorithm steps 1-5, respectively (see Section 4.4). Green bonds are used to denote the pre-computed environments from step 1.

DMRG algorithm [94]. In 2D, it allows for the evaluation of $\langle \psi | \hat{H} | \psi \rangle$ on-the-fly using gMPOs.

To see how, first consider the contraction of the finite, 2-layer, 2D tensor network

corresponding to $\langle \psi | \psi \rangle$ for some PEPS $|\psi\rangle$ using the “boundary MPS” method [35]. Starting from the bottom, the first point of reference is row 1 and as the contraction progresses, it shifts upward to row 2, then row 3, etc. During this process the Hamiltonian can be successively partitioned along with the reference row of the norm contraction, so that the first line lies between row 1 and row 2, then the next is between row 2 and row 3, etc. Using this idea, the total energy $\langle \psi | \hat{H} | \psi \rangle$ can be accumulated as follows (shown graphically in Fig. 4.4):

1. Pre-compute all the partial contractions of $\langle \psi | \psi \rangle$ using the boundary method, starting from the top with row L_y and working downward. They should be stored as $\{\text{envs}[\emptyset], \dots, \text{envs}[L_y - 2]\}$ (Fig. 4.4(c)).
2. Construct an MPO which contains all the 1-body terms in \hat{H} that act locally in row 1 as well as all the interactions between sites in row 1. In other words, this should be the MPO representation of \hat{H}_{bot} when the partition is between row 1 and row 2. Apply this MPO between the bra and ket tensors of row 1, and evaluate $E_{\text{bot}} = \langle \psi | \hat{H}_{\text{bot}} | \psi \rangle$ by contracting this partial TN with $\text{envs}[L_y - 2]$ (Fig. 4.4(d)).
3. Construct complementary operator vectors which contain the local operators \hat{O}_{int} that act in row 1 but have interactions with sites above row 1 (as in Section 4.4). Apply these vectors between the corresponding row 1 ket and bra tensors along the vertical bonds. This partial TN will be called *intops* (Fig. 4.4(e)).
4. Shift the partition line up by 1 row (in general, now in between rows y and $y + 1$). Construct a gMPO to be applied in row y that encodes all the terms in the new \hat{H}_{bot} that have not already been evaluated. Apply the gMPO between the row y bra and ket tensors, and contract this TN with *intops* (below) and $\text{envs}[L_y - y - 1]$ (above). Add the resulting scalar to E_{bot} to obtain a new E_{bot} , which now accounts for all the terms in $\langle \psi | \hat{H}_{\text{bot}} | \psi \rangle$ given the new partition position. For clarity, the case immediately following step 3 would be when $y = 2$. To accumulate the proper terms, this gMPO should include interactions within row 2, as well as all the interactions between sites in row 2 and sites in the rows beneath it, which is just row 1 for now ($y = 2$ case shown in Fig. 4.4(f)).

5. Construct an updated (approximate) `intops`. This step can be understood as iteratively building up MPOs along the vertical bonds. First a complementary operator *matrix* (which is just an MPO matrix) is constructed for each column, which relates the \hat{O}_{int} in a given column of row $y - 1$ to the \hat{O}_{int} in the same column of row y . This is exactly like how a regular MPO matrix relates the operators on site $x - 1$ to the operators on site x . Then these complementary operator matrices are applied between each of the bra and ket tensors of row y along the vertical indices. This row can then be contracted with the old `intops` and its horizontal bond dimension can be compressed according to the boundary method contraction routine. This yields a new approximate `intops` that contains the action of all the local operators \hat{O}_{int} that lie below the partition when it is between rows y and $y + 1$ ($y = 2$ case shown in Fig. 4.4(g))
6. Iterate steps 4 and 5 until the top of the PEPS is reached. When the final gMPO is applied to row L_y and contracted with `intops`, the expectation values of all the terms in \hat{H} will have been tallied in the running total E_{bot} .

Given a Hamiltonian with general interactions of the form $\hat{A}_i \hat{B}_j$, where $i < j$, the big picture of this algorithm (which we will call the “boundary gMPO” method for future reference) can be succinctly summarized as follows: To compute $\langle \psi | \hat{H} | \psi \rangle$, we think about classifying terms in \hat{H} into 3 non-mutually exclusive groups according to the bipartition of a PEPO between rows y and $y + 1$. Group (1) contains terms where \hat{A} and \hat{B} are both below the partition. Group (2) contains terms where \hat{A} is below the partition but \hat{B} is somewhere above it. Group (3) contains terms where \hat{A} and \hat{B} are both below the previous partition (when it was between rows $y - 1$ and y). At each iteration of the algorithm, we first compute $\langle \psi | \hat{H} | \psi \rangle$ for the set of terms in the difference (1) - (3) by contracting a gMPO with `intops`, and then we construct a new `intops` for the next iteration that accounts for all the terms in (2) by slightly modifying the previous `intops`.

This can be viewed as a “decomposed” contraction of the expectation value of a PEPO. As the partition is iteratively shifted upwards, MPOs are sequentially constructed and applied tensor-by-tensor along the vertical bonds and gMPOs are applied along the horizontal bonds in order to “extract” the expectation values of the terms in \hat{H}_{bot} , as it is defined based on the current progress of the contraction. When explicitly contracting the expectation value of a PEPO, the boundary tensors

accumulate the identical terms but they are not fully evaluated until the entire contraction is complete. By extracting the “completed” terms along the way, the boundary gMPO method allows for the energy evaluation of the same set of general 2D Hamiltonians that can be represented by PEPOs while only invoking MPOs and gMPOs. Since the ideas for constructing MPOs, and thus also gMPOs, are more familiar and well-established in the literature than PEPOs, we expect that this will be a useful conceptual simplification.

Additionally, this formulation leads to a reduction in computational cost because `intops` can always be constructed with operator virtual indices pointing only in the vertical direction ². When compared to the contraction of a PEPO, the cost of boundary absorption and compression (the time-dominant step; step 5 and Fig. 4.4(g) above) is reduced because the boundary tensors no longer contain any operator virtual indices along the horizontal bonds. This decreases the cost of boundary absorption by a factor of D_{op}^4 and compression by a factor of D_{op}^6 (where D_{op} is the virtual bond dimension of the PEPO/vertical bond dimension of `intops` operators) ³.

In the context of a variational [39, 40] ground state optimization of a PEPS with respect to the Hamiltonian \hat{H} , this algorithm fits very nicely within the framework of the newly-developed differentiable programming techniques for tensor networks [119]. Since the expectation values of different sets of operators are evaluated during different iterations, each iteration of steps 4 and 5 can be differentiated separately. This allows for the gradient of the energy to also be computed on-the-fly as the energy itself is being computed, leading to a highly efficient computational formulation.

4.5 Results

In this section, we will present the explicit constructions of the MPOs and gMPOs needed to implement the boundary gMPO algorithm described in Section 4.4 for various types of 2D Hamiltonians. From the set of Hamiltonians that we explicitly describe, we expect that the construction of most other Hamiltonians of potential interest will be conceptually straightforward. We will also demonstrate the speed and

² Although this is always possible, it is not required. It may be the case that for some Hamiltonians not explicitly considered in this work, allowing horizontal operator virtual indices in `intops` results in a more efficient representation.

³ These factors are determined under the assumption that the bond dimension χ of the boundary MPS during boundary method contraction [35] must be proportional to D_{op} for accurate results when using a full PEPO [70].

	\hat{H}_{NN}	\hat{H}_D	\hat{H}_{LRNC}	\hat{H}_{LRAC}
$D = 2, \chi = 5$	2.52	3.70	3.66	18.6
$D = 2, \chi = 10$	6.94	11.9	11.6	18.4
$D = 2, \chi = 20$	13.2	27.9	27.2	19.1
$D = 3, \chi = 15$	20.5	39.2	37.5	1.63
$D = 3, \chi = 30$	25.3	52.6	51.7	1.36
$D = 3, \chi = 40$	24.0	50.9	51.2	1.41
$D = 4, \chi = 15$	19.0	33.5	34.2	
$D = 4, \chi = 30$	27.8	59.3	60.3	
$D = 4, \chi = 40$	32.7	62.8	62.5	

Table 4.1: The average computational speedups of the boundary gMPO algorithm over PEPO-based expectation value calculations for a representative set of 2D Hamiltonians (Eqns. (4.17), (4.19), (4.22), and (4.25)). The gMPO-based scheme is generically and significantly faster than the PEPOs for all the Hamiltonians except the one with long-range interactions mediated by a distance-dependent potential (LRAC). The reported numbers are averages taken over multiple calculations for each of multiple different trial wavefunctions: PEPS ground states for the 8×8 AFM Heisenberg model and FM transverse field Ising model ($h = 3.5$). D denotes the bond dimension of the trial PEPSs and χ denotes the maximum boundary bond dimension used during contraction [35]. Both algorithms were implemented in a straightforward manner in order to compare their runtimes as fairly as possible. This data should be used in conjunction with Fig. 4.5.

accuracy of the new algorithm, and compare it to the performance of expectation value computations using explicit PEPOs as well as “brute force” application of all the Hamiltonian terms separately (this technique is analogous to the current technique used in 2D simulations, as mentioned in Section 4.2).

In our brute force implementations we do not utilize any caching strategies for contraction intermediates that are recyclable between the evaluation of multiple different Hamiltonian terms. This would lead to a faster routine, and might allow for a more direct comparison to the boundary gMPO algorithm since it inherently utilizes a (quite limited) caching strategy. However, while the implementation of the `envs` intermediates in the boundary gMPO method is very straightforward, proper caching for the brute force technique is more complicated, especially for Hamiltonians which include long-range interactions. To keep the results for all Hamiltonians comparable, we thus always refrain from caching in the brute force method.

In all cases we will consider a finite two-dimensional system on a rectangular lattice of $L_x \times L_y$ sites labelled (x, y) , where x indexes the sites in a row and y indexes the

	\hat{H}_{NN}	\hat{H}_D	\hat{H}_{LRNC}	\hat{H}_{LRAC}
$D = 2, \chi = 5$	60.8	118.4	1066	34.2
$D = 2, \chi = 10$	59.9	107.6	975.7	33.0
$D = 2, \chi = 20$	49.4	84.6	782.2	24.2
$D = 3, \chi = 15$	43.1	72.8	672.3	16.9
$D = 3, \chi = 30$	36.6	64.6	609.6	16.1
$D = 3, \chi = 40$	37.3	65.4	628.9	16.8
$D = 4, \chi = 15$	39.9	61.8	592.6	
$D = 4, \chi = 30$	35.3	63.4	569.8	
$D = 4, \chi = 40$	37.5	68.7	623.8	
$16 \times 16, D = 2, \chi = 20$				296.1

Table 4.2: The average computational speedups of the boundary gMPO algorithm over “brute force” expectation value calculations for a representative set of 2D Hamiltonians (Eqns. (4.17), (4.19), (4.22), and (4.25)). The gMPO-based scheme is significantly faster for all the Hamiltonians under consideration, especially those which contain long-range interactions (LRNC and LRAC). In the brute force technique, the Hamiltonian is evaluated term-by-term by explicitly applying each pair of local operators. Both algorithms were implemented in a straightforward manner in order to compare their runtimes as fairly as possible. The reported numbers are averages taken in an identical manner to Table 4.1, and the parameters D and χ are also indentially defined. This data should be used in conjunction with Fig. 4.5.

sites in a column. By the conventions of the previous sections, $(1, 1)$ corresponds to the bottom left corner and (L_x, L_y) to the top right corner. When a sum is taken over all the sites in the lattice using a single index, such as $\sum_{i=1}^{L_x \times L_y}$, the order in which the sites are indexed is such that site $(x + 1, y)$ always has a larger label number than site (x, y) , and site $(x, y + 1)$ also has a larger label number than (x, y) . This convention will be important when restrictions are placed on the sums, such as the condition $i < j$.

Local Hamiltonians

Nearest-neighbor interactions

Consider a Hamiltonian with local 1-body terms and non-symmetric nearest-neighbor interactions of the form,

$$\hat{H}_{NN} = \sum_{i=1}^{L_x \times L_y} \hat{C}_i + \sum_{\langle ij \rangle, i < j} \hat{A}_i \hat{B}_j, \quad (4.17)$$

where both i and j index through all $L_x \times L_y$ sites. The MPO in step 2 of the boundary gMPO algorithm is given by \hat{W}_{NN} from Eq. (4.5), Section 4.3. The vertical MPOs

that are applied tensor-by-tensor as the algorithm progresses in order to produce `intops` are given by,

$$\begin{aligned}\hat{O}_{\beta_1}[x, 1] &= \begin{pmatrix} \hat{I} & \hat{A} \end{pmatrix}, \\ \hat{O}_{\beta_{y-1}\beta_y}[x, L_y > y > 1] &= \begin{pmatrix} \hat{I} & \hat{A} \\ \hat{O} & \hat{O} \end{pmatrix}.\end{aligned}\quad (4.18)$$

Note that here we use the index label β_y to denote its position (y) along the vertical bonds within column. This is a slight abuse of notation when compared to Section 4.4, where the subscript on β was used to denote its position (x) within in a single row. A fully consistent notation would require an x and y subscript on every β , but for all Hamiltonians under consideration in Section 4.5 the vertical MPO matrices will be the same for every x , so we always suppress the x label (and sometimes also the y label when the context is unambiguous) on β for simplicity.

Also note that these MPO matrices (Eq. (4.18)) only need to be of dimension 2×2 because each time a new `intops` is created, the expectation values of interaction terms with the row above are immediately extracted by contracting it with an appropriate gMPO. Unlike a typical MPO, we therefore never need to “complete” an interaction with a \hat{B} operator in these matrices because that is taken care of in the gMPO. This eliminates the need for the third row and column to account for \hat{B} . In the current case of nearest neighbor interactions, the bottom row of $\hat{O}_{\beta_{y-1}\beta_y}$ is all \hat{O} s because the action of \hat{A} in row y does not need to be stored once the point of reference is shifted up to row $y + 1$.

The gMPO tensors, used in step 4 to extract the expectation values of terms in \hat{H}_{bot} , were given as the example in Eq. (4.14). To make the notation consistent with a fully 2D Hamiltonian, the coordinates of the tensors in that expression should be transformed according to: $\hat{M}_1[i, 2] \rightarrow \hat{M}_1[x, y > 1]$; $\hat{M}_2[1, 2] \rightarrow \hat{M}_2[1, y > 1]$; $\hat{M}_2[L, 2] \rightarrow \hat{M}_2[L_x, y > 1]$; $\hat{M}_2[i, 2] \rightarrow \hat{M}_2[L_x > x > 1, y > 1]$. Additionally, for generality we do not have $\hat{C} = \hat{O}$ in our current example. In essence, the evaluation of this Hamiltonian’s expectation value amounts to performing the same calculation as the one outlined in the example of Section 4.4 for every row in the system.

The accuracy of the boundary gMPO algorithm using these tensors to evaluate the expectation value of the given Hamiltonian (with $\hat{C} = \hat{O}$, $\hat{A} = \hat{B} = \sigma_z$) with respect to various trial PEPS is shown in Fig. 4.5. It is almost identically accurate to the brute force scheme and its accuracy is also very similar to the PEPO-based implementation in most cases, with the outliers showing an improved accuracy

for the gMPOs. Despite the similar accuracies, using the gMPOs allows for a computational speedup of up to $\sim 30\times$ over the PEPOs and $\sim 40\times$ over the brute force implementation, as seen in Tables 4.1-4.2.

Diagonal-neighbor interactions

Now consider a Hamiltonian that has local 1-body terms as well as both nearest-neighbor and diagonal-neighbor interactions, with strengths J_1 and J_2 respectively,

$$\hat{H}_D = \sum_{i=1}^{L_x \times L_y} \hat{C}_i + J_1 \sum_{\langle ij \rangle, i < j} \hat{A}_i \hat{B}_j + J_2 \sum_{\langle\langle ij \rangle\rangle, i < j} \hat{A}_i \hat{B}_j. \quad (4.19)$$

Although we are again considering the non-symmetric Hamiltonian construction (denoted by $i < j$) for simplicity, if the interaction operators are chosen to be symmetric (i.e. $\hat{B} = \hat{A}$) then the given Hamiltonian differs from the truly symmetric one \hat{H}_{sym} (i.e. $i < j \rightarrow i \neq j$) by a factor of 2 in the interaction coefficients, $\hat{H}_{\text{sym}}(J_1, J_2) = \hat{H}_D(2J_1, 2J_2)$. If the representation of \hat{H}_{sym} is needed when $\hat{B} \neq \hat{A}$, it can be determined by using the results in this section and following the examples in Section 4.3.

The MPO in step 2 is again given by \hat{W}_{NN} (Eq. (4.5), Section 4.3), and the vertical MPO matrices are still given by Eq. (4.18). The gMPO tensors are,

$$\begin{aligned} \hat{M}_1[1, y > 1] &= \begin{pmatrix} \hat{C} & J_1 \hat{A} & \hat{B} & \hat{I} \end{pmatrix}, \\ \hat{M}_1[L_x, y > 1] &= \begin{pmatrix} \hat{I} & \hat{B} & \hat{O} & \hat{C} \end{pmatrix}^T, \\ \hat{M}_1[L_x > x > 1, y > 1] &= \begin{pmatrix} \hat{I} & \hat{O} & \hat{O} & \hat{O} \\ \hat{B} & \hat{O} & \hat{O} & \hat{O} \\ \hat{O} & \hat{O} & \hat{O} & \hat{O} \\ \hat{C} & J_1 \hat{A} & \hat{B} & \hat{I} \end{pmatrix}, \\ \hat{M}_2[1, y > 1] &= \begin{pmatrix} J_1 \hat{B} & J_2 \hat{I} & \hat{O} & \hat{O} \end{pmatrix}, \\ \hat{M}_2[L_x, y > 1] &= \begin{pmatrix} \hat{O} & \hat{O} & J_2 \hat{I} & J_1 \hat{B} \end{pmatrix}^T, \\ \hat{M}_2[L_x > x > 1, y > 1] &= \begin{pmatrix} \hat{O} & \hat{O} & \hat{O} & \hat{O} \\ \hat{O} & \hat{O} & \hat{O} & \hat{O} \\ J_2 \hat{I} & \hat{O} & \hat{O} & \hat{O} \\ J_1 \hat{B} & J_2 \hat{I} & \hat{O} & \hat{O} \end{pmatrix}. \end{aligned} \quad (4.20)$$

These matrices can be understood by noting the similarity between \hat{M}_1 and $\hat{W}_{NN\text{-sym}}$ from Eq. (4.6). The only difference is that in \hat{M}_1 the entry for \hat{A} in the first column

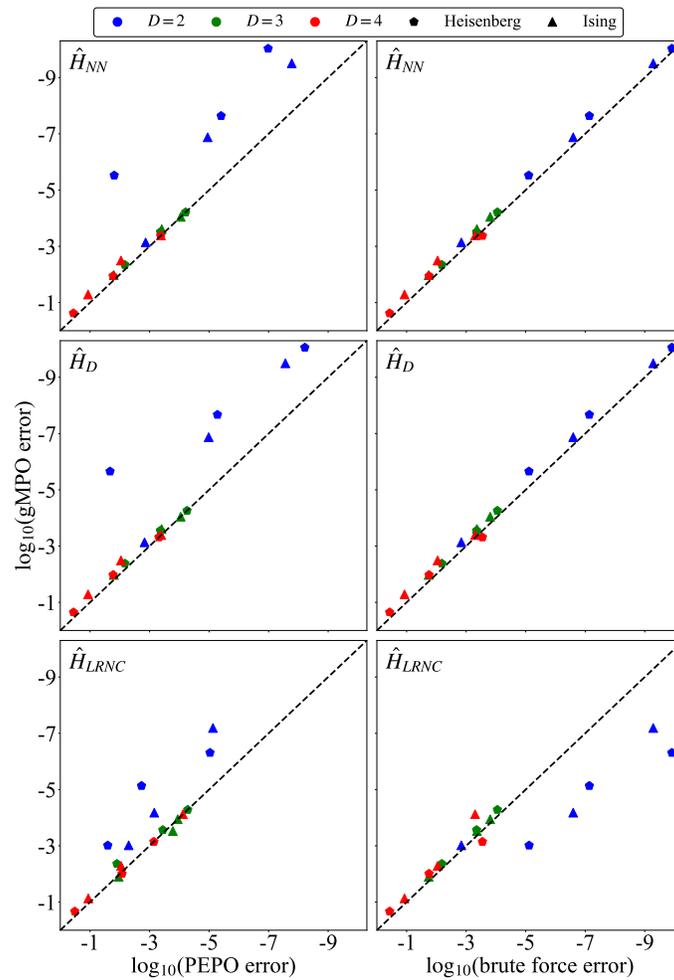


Figure 4.5: The relative error of the gMPO-based expectation values compared to the relative errors obtained using both PEPO-based calculations and the “brute force” technique of evaluating each term in the Hamiltonian separately. For the majority of cases tested, all three techniques exhibit the same level of accuracy. None of these Hamiltonians (Eqns. (4.17), (4.19), (4.22)) contain long-range distance-dependent potentials. The expectation values are calculated with respect to various 8×8 trial PEPS of bond dimensions $D = 2, 3, 4$. A single point compares the relative error of gMPOs with either PEPOs or brute force, with each technique using the same trial state and χ value (the boundary bond dimension during contraction [35]). All errors are measured with respect to a brute force evaluation that is highly converged in χ . The displayed points are for selected values of χ less than the converged value, in order to compare the levels of accuracy that can be obtained with a given computational effort. For a full picture of the computational effort, this data should be used in conjunction with the speedups reported in Tables 4.1-4.2, which also include the χ values for each of the points here.

is made to be \hat{O} (and interaction coefficients are included). This is done to prevent symmetric nearest-neighbor interactions of the form $\hat{B}_{x-1,y}\hat{A}_{x,y}$ from being included along the gMPO row (in these coordinates the gMPO is being applied to row y). However, since the sites are ordered in such a way that the (non-symmetric) diagonal-neighbor interactions occur between site $(x, y - 1)$ and sites $(x - 1, y)$, $(x + 1, y)$, we still want to include the action of \hat{B} “on the left” on site $(x - 1, y)$. This is exactly what the form of \hat{W}_{NN-sym} is designed to do.

If the \hat{M}_1 matrices were the only ones included in the gMPO, then this action of $\hat{B}_{x-1,y}$ “on the left” would never be utilized due to the \hat{O} in place of $\hat{A}_{x,y}$ in the first column. However, \hat{M}_2 couples the action of $\hat{A}_{x,y-1}$ (from `intops`) into the two typical locations of \hat{A} in \hat{W}_{NN-sym} (and also multiplies by J_2). This allows the “on the left” action of $\hat{B}_{x-1,y}$ to interact with the action of $J_2\hat{A}_{x,y-1}$, which is exactly the diagonal interaction that we want to include. \hat{M}_2 also couples $J_2\hat{A}_{x,y-1}$ into the same position as $J_1\hat{A}_{x,y}$ in \hat{M}_1 , which allows for the nearest-neighbor horizontal interaction and diagonal-neighbor interaction “to the right” to be accounted for simultaneously. Specifically, after the β_{y-1} indices have been appropriately contracted over, the subsequent contraction over an α index will yield a term like $(J_2\hat{A}_{x,y-1} + J_1\hat{A}_{x,y})\hat{B}_{x+1,y}$. For clarity, in the spirit of the example in Eq. (4.15), a typical contraction over the β_{y-1} index (with $\hat{C} = \hat{O}$) would look like,

$$\begin{aligned}
& \sum_{\beta_{y-1}} \hat{O}_{\beta_{y-1}}[x, y - 1] \hat{M}_{\beta_{y-1}}[x, y] = \\
& \hat{I} \cdot \begin{pmatrix} \hat{I}_{x,y} & \hat{O} & \hat{O} & \hat{O} \\ \hat{B}_{x,y} & \hat{O} & \hat{O} & \hat{O} \\ \hat{O} & \hat{O} & \hat{O} & \hat{O} \\ \hat{O} & J_1\hat{A}_{x,y} & \hat{B}_{x,y} & \hat{I}_{x,y} \end{pmatrix} + \\
& \hat{A}_{x,y-1} \cdot \begin{pmatrix} \hat{O} & \hat{O} & \hat{O} & \hat{O} \\ \hat{O} & \hat{O} & \hat{O} & \hat{O} \\ J_2\hat{I}_{x,y} & \hat{O} & \hat{O} & \hat{O} \\ J_1\hat{B}_{x,y} & J_2\hat{I}_{x,y} & \hat{O} & \hat{O} \end{pmatrix} \\
& = \begin{pmatrix} \hat{I} & \hat{O} & \hat{O} & \hat{O} \\ \hat{B}_{x,y} & \hat{O} & \hat{O} & \hat{O} \\ J_2\hat{A}_{x,y-1} & \hat{O} & \hat{O} & \hat{O} \\ J_1\hat{A}_{x,y-1}\hat{B}_{x,y} & J_1\hat{A}_{x,y} + J_2\hat{A}_{x,y-1} & \hat{B}_{x,y} & \hat{I} \end{pmatrix}. \tag{4.21}
\end{aligned}$$

The form of this gMPO, which is the simplest case where \hat{A} can interact with a \hat{B}

from a different row *and* column, is the basis for generating all the more complicated finite-range 2D Hamiltonians with interactions between more distant neighbors. In essence, the form of \hat{M}_1 has to be adapted to the desired pattern of operators within the gMPO row, and then $\hat{M}_2, \hat{M}_3, \dots$, etc. take the forms which properly couple the operators from the vertical MPOs (*intops*) into \hat{M}_1 . For a general construction of this form that includes all non-symmetric interactions between neighbors up to range R , see Appendix A.

The speed and accuracy of the boundary gMPOs using these tensors (with $\hat{C} = \hat{O}$, $\hat{A} = \hat{B} = \sigma_z$, $J_2 = J_1/2$) is compared to a PEPO-based implementation and a brute force implementation in Fig. 4.5 and Tables 4.1-4.2. The gMPOs produce accuracies which are nearly identical to the brute force scheme, but with a computational effort that is $\sim 60 - 70\times$ less. When compared to PEPOs, a speedup of up to $\sim 50\times$ is observed and in most cases the gMPOs and PEPOs also produce the same level of accuracy. In cases where they differ, the gMPOs are observed to be more accurate.

Long-range Hamiltonians with no coefficients

We will now consider a Hamiltonian which has local 1-body terms and non-symmetric pairwise interactions of equal strength between every site on the lattice. This can be viewed as the 2D version of the Hamiltonian represented by $\hat{W}_{uniform}$ (see Section 4.3). We have,

$$\hat{H}_{LRNC} = \sum_i \hat{C}_i + \sum_{i<j} \hat{A}_i \hat{B}_j. \quad (4.22)$$

The MPO used in step 2 of the boundary gMPO algorithm is given by $\hat{W}_{uniform}$. The vertical MPOs used for the construction of *intops* are given by,

$$\begin{aligned} \hat{O}_{\beta_1}[x, 1] &= \begin{pmatrix} \hat{I} & \hat{A} \end{pmatrix}, \\ \hat{O}_{\beta_{y-1}\beta_y}[x, L_y > y > 1] &= \begin{pmatrix} \hat{I} & \hat{A} \\ \hat{O} & \hat{I} \end{pmatrix}. \end{aligned} \quad (4.23)$$

These MPO matrices differ from those in Eq. (4.18) because they “remember” the action of all the local operators in a given column x . In Eq. (4.18), the contractions over β_{y-1} that are performed in step 5 result in operator vectors of the form $(\hat{I}, \hat{A}_{x,y})$. This was sufficient because the previous Hamiltonians under consideration were local, so the action of the $\hat{A}_{x,y-1}, \hat{A}_{x,y-2}, \dots$, etc. operators had already been completely accounted for by the time the reference row was shifted up by one. However, in our current Hamiltonian the interactions are long-ranged, so the action

of all the local operators in a given column must be accounted for in a single `intops` tensor. This is achieved by the MPO matrices in Eq. (4.23), for which a contraction over $\beta_1, \beta_2, \dots, \beta_{y-1}$ yields operator vectors of the form $(I, \hat{A}_{x,1} + \hat{A}_{x,2} + \dots + \hat{A}_{x,y})$.

The corresponding gMPO tensors are given by,

$$\begin{aligned}
\hat{M}_1[1, y > 1] &= \begin{pmatrix} \hat{C} & \hat{A} & \hat{B} & \hat{I} \end{pmatrix}, \\
\hat{M}_1[L_x, y > 1] &= \begin{pmatrix} \hat{I} & \hat{B} & \hat{O} & \hat{C} \end{pmatrix}^T, \\
\hat{M}_1[L_x > x > 1, y > 1] &= \begin{pmatrix} \hat{I} & \hat{O} & \hat{O} & \hat{O} \\ \hat{B} & \hat{I} & \hat{O} & \hat{O} \\ \hat{O} & \hat{O} & \hat{I} & \hat{O} \\ \hat{C} & \hat{A} & \hat{B} & \hat{I} \end{pmatrix}, \\
\hat{M}_2[1, y > 1] &= \begin{pmatrix} \hat{B} & \hat{I} & \hat{O} & \hat{O} \end{pmatrix}, \\
\hat{M}_2[L_x, y > 1] &= \begin{pmatrix} \hat{O} & \hat{O} & \hat{I} & \hat{B} \end{pmatrix}^T, \\
\hat{M}_2[L_x > x > 1, y > 1] &= \begin{pmatrix} \hat{O} & \hat{O} & \hat{O} & \hat{O} \\ \hat{O} & \hat{O} & \hat{O} & \hat{O} \\ \hat{I} & \hat{O} & \hat{O} & \hat{O} \\ \hat{B} & \hat{I} & \hat{O} & \hat{O} \end{pmatrix}. \tag{4.24}
\end{aligned}$$

Note that this result is nearly identical to the gMPO tensors in the previous section for diagonal interactions (Eq. (4.20)). The only difference is the replacement of two \hat{O} s with \hat{I} s in \hat{M}_1 (and the removal of the interaction coefficients). The reason for this similarity can be understood in two distinct ways. Firstly, the addition of these identities can be viewed as an elevation of the symmetric nearest-neighbor interactions in \hat{W}_{NN-sym} to symmetric interactions of arbitrary range, which captures all the new terms in \hat{H}_{LRNC} . Secondly, we can see a direct analogy between the relations of the current \hat{M}_1 to $\hat{W}_{uniform-sym}$ (Section 4.3) and the previous \hat{M}_1 (Eq. (4.20)) to \hat{W}_{NN-sym} . In other words, in the previous section we argued that because \hat{M}_1 only differed from \hat{W}_{NN-sym} by a single element, it was clear that it would encode the symmetric nearest-neighbor action of \hat{B} about site (x, y) that was necessary to generate the diagonal interactions. Now in the current case, we replace the modified \hat{W}_{NN-sym} with an identically modified $\hat{W}_{uniform-sym}$ to obtain the symmetric action of \hat{B} on *all* sites to the left and right of (x, y) . This is precisely the pattern of operators that needs to be encoded in order to generate all the terms in \hat{H}_{bot} .

The performance of the boundary gMPOs using these tensors (with $\hat{C} = \hat{O}$, $\hat{A} = \hat{B} = \sigma_z$) is compared to a PEPO-based implementation and a brute force implementation in Fig. 4.5 and Tables 4.1-4.2. In this case, due to the long-range nature of the interactions, the scaling of our brute force evaluation is $O(N^3)$. While this can be slightly reduced with appropriate caching of contraction intermediates, the gMPO- and PEPO-based techniques only scale as $\sim O(N)$ (where N is the total number of sites in the system). Thus in addition to the $\sim 60\times$ speedup over the PEPOs, the gMPOs attain large speedups of $\sim 600\times$ over the brute force algorithm for the $N = 64$ cases that we consider. For larger systems, this speedup will grow rapidly. Given this poor scaling and the fact that the gMPOs can reproduce the accuracy of the brute force calculations in all of the most challenging test cases, it is clear that the brute force technique is not a viable approach to study systems with non-local interactions. Of the two viable strategies, gMPOs show very similar accuracy to PEPOs across most of the test cases, as in the previous sections.

Long-range isotropic Hamiltonians with approximate coefficients

In the previous section, we demonstrated an exact and compact representation of a long-range interacting 2D Hamiltonian when the interactions coefficients were all the same (this can also be done with a PEPO [70]). Despite this, it is a challenging problem to efficiently ⁴ represent a 2D Hamiltonian which has long-range interaction coefficients that depend on the distance between sites, even in an approximate manner [45, 70, 117, 118]. Various solutions to this problem have been proposed recently [70, 117, 118], but they all require the explicit use of PEPOs, making their computational cost high.

The introduction of the gMPO formalism allows for a new, simpler approach to be derived, which we will show to be many orders of magnitude more accurate and efficient than the PEPO-based approaches. We will consider a restricted case of the general long-range interacting Hamiltonian on the 2D lattice,

$$\hat{H}_{LRAC} = \sum_i \hat{C}_i + \sum_{i<j} V_{ij} \hat{A}_i \hat{B}_j, \quad (4.25)$$

where V is a translation invariant, decaying function of the Euclidean distance between sites i and j (i.e. it is isotropic).

The crux of the long-range interaction problem on the 2D lattice is that functions of the Euclidean distance $f(\sqrt{x^2 + y^2})$, which are necessary for physical

⁴Here we define “efficient” to mean that the computational cost to evaluate the expectation value of the Hamiltonian scales linearly with the number of sites in the system.

potentials V , are difficult to represent efficiently within a tensor network structure [70, 117]. Although 1D functions of x and y can be independently constructed with ease (see Section 4.3), the known possibilities for combining them within a 2D tensor network ansatz yield functions of the Manhattan distance $f(|x| + |y|)$ or product functions $f(x)g(y)$, but not the desired radially symmetric ones $f(r) = f(\sqrt{x^2 + y^2})$. However, the Gaussian function $e^{-\lambda r^2}$ has the unique property that $f(x)f(y) = e^{-\lambda(x^2+y^2)} = f(r^2)$. This connection allows for a radially symmetric Gaussian function in 2D to be created from the product of two 1D Gaussians $f(x) = e^{-\lambda x^2}$ and $f(y) = e^{-\lambda y^2}$.

This observation can be directly exploited by the gMPO-based algorithm. If the vertical MPOs encode the interactions $\hat{H}_1 = \sum_{i=1}^{L_y} \sum_{j>i} e^{-\lambda(j-i)^2} (\hat{A}_i \hat{I}_j + \hat{A}_i \hat{B}_j)$ and the gMPOs encode horizontal interactions of the form $\hat{H}_2 = \sum_{i=1}^{L_x} \sum_{j>i} e^{-\lambda(j-i)^2} (\hat{I}_i \hat{B}_j + \hat{B}_i \hat{I}_j + \hat{A}_i \hat{B}_j)$, then they can be combined as a product (as in Sections 4.5, 4.5) to make complete interactions of the form $e^{-\lambda(a^2+b^2)} \hat{A}_{x,y} \hat{B}_{x+a,y\pm b} + e^{-\lambda a^2} \hat{A}_{x,y} \hat{B}_{x+a,y} + e^{-\lambda b^2} \hat{A}_{x,y} \hat{B}_{x,y+b}$. The two-dimensional, radially symmetric Gaussians can then be used as a basis to fit the desired long-range potential,

$$V(x, y) \approx \sum_{k=1}^K c_k e^{-\lambda_k(x^2+y^2)}, \quad (4.26)$$

which is a well-studied problem with highly accurate, compact solutions when V smoothly decays with distance [111, 120, 121]. The expectation value of the desired Hamiltonian can then be evaluated as the sum over the expectation values obtained using K different sets of vertical MPOs and gMPOs (for the K different values of λ). Since the only requirement of this technique is the representation of 1D Gaussian functions, this basis can be encoded directly within the MPO and gMPO tensors, which completely avoids the conceptual and computational complexity of introducing fictitious superlattices, as in Refs. [70, 117].

Unfortunately, there is no known exact, compact representation of a 1D MPO with pairwise Gaussian interactions. However, it can be generated in a nearly numerically exact manner using the method outlined in Section 4.3 to create \hat{W}_{gen} . Fig. 4.7(a) shows the required bond dimension for the Gaussian MPO for different values of λ . The result that $D_{op} = 14$ in the worst case for an accuracy of $\sim 10^{-10}$ is a modest bond dimension for an MPO, which is what makes the current approach of using an exact Gaussian basis amenable to the gMPO algorithm. Although this same scheme could, in principle, be implemented using PEPOs on the same lattice, it would require the use of PEPOs with $D_{op} = 28$ in the worst case. In practice, the

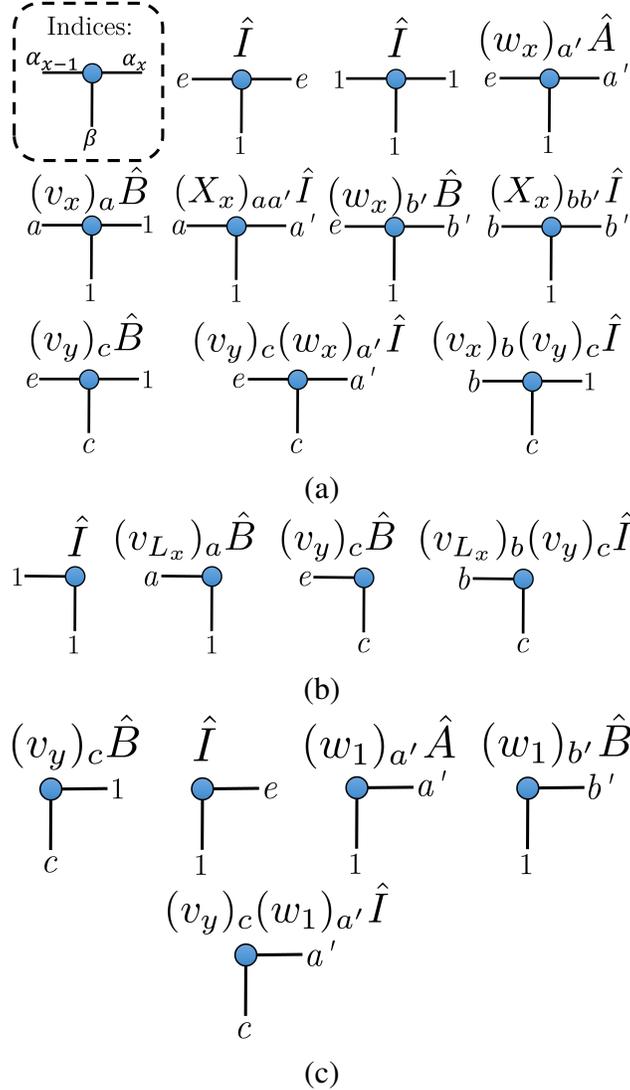


Figure 4.6: The gMPO tensors $\hat{M}[x, y > 1]$ for long-range Gaussian interactions (Eq. (4.28)). (a) Is for $1 < x < L_x$, (b) is for $x = L_x$, and (c) is for $x = 1$. Since the dimension of each bond α_{x-1} , α_x , and β can vary depending on x , y and the value of the exponential coefficient λ in the Gaussian interaction, we will use symbols to label specific values of the indices. For a horizontal (α) index of bond dimension $2g + 2$ that takes index values $\{1, 2, \dots, 2g + 2\}$, we label the first value by “1”, the next g values by a (if the bond points left) or a' (if the bond points right), the next g values by b (if the bond points left) or b' (if the bond points right), and the final value by e . This is the convention that is explained in Eq. (4.12) and is also used in Eq. (4.28). For a vertical (β) index of bond dimension $g + 1$ that takes index values $\{1, 2, \dots, g + 1\}$, we label the first element by “1” and the remaining g elements by c . This corresponds directly with Eqns. (4.27) & (4.28)

factor of D_{op}^7 in the computational cost of PEPO-based contractions makes a PEPO with a bond dimension of this size unusable. However, since the use of gMPOs reduces the dependence of the cost on the operator bond dimension to at most D_{op}^3 (in step 4), and D_{op}^1 in the most time intensive step (compression in step 5), using this bond dimension for the vertical MPOs and gMPOs is entirely feasible.

The explicit forms of the tensors in this case can be viewed as a direct generalization of the tensors from the previous Section (4.5), Eqs. (4.23) & (4.24). This follows from the discussion in Section 4.3 regarding \hat{W}_{gen} and $\hat{W}_{gen-sym}$ as direct generalizations of $\hat{W}_{uniform}$ and $\hat{W}_{uniform-sym}$, respectively. Since the tensors in Section 4.5 are derived from $\hat{W}_{uniform(-sym)}$ and in the current case we want to use tensors based on the $\hat{W}_{gen(-sym)}$ representation of a Gaussian MPO, the tensors in (4.23) and (4.24) generalize to the current case in an analogous way to the $\hat{W}_{uniform(-sym)} \rightarrow \hat{W}_{gen(-sym)}$ generalization of Section 4.3.

Specifically, the MPO for step 2 is the \hat{W}_{gen} representation of Gaussian interactions with exponential coefficient λ_k , which is determined from the algorithm in Ref. [97]. From this MPO, the data for each \vec{v}_i , \vec{w}_i , and X_i can be extracted (according to Eq. (4.11)). These can then be used to construct the other tensors for pairwise interactions mediated by a 2D Gaussian potential. The vertical MPO tensors are given by,

$$\begin{aligned} \hat{O}_{\beta_1}[x, 1] &= \begin{pmatrix} \hat{I} & (w_1)_c \hat{A} \end{pmatrix}, \\ \hat{O}_{\beta_{y-1}\beta_y}[x, L_y > y > 1] &= \begin{pmatrix} \hat{I} & (w_y)_c \hat{A} \\ \hat{0} & (X_y)_{c'c} \hat{I} \end{pmatrix}, \end{aligned} \quad (4.27)$$

where c and c' index through the vector \vec{w}_y and matrix X_y , like in Eq. (4.11). The

gMPO tensors are,

$$\begin{aligned}
\hat{M}_1[1, y > 1] &= \left(\hat{C}, (w_1)_{a'}\hat{A}, (w_1)_{b'}\hat{B}, \hat{I} \right), \\
\hat{M}_1[L_x, y > 1] &= \left(\hat{I}, (v_{L_x})_a\hat{B}, \hat{0}, \hat{C} \right)^T, \\
\hat{M}_1[L_x > x > 1, y > 1] &= \begin{pmatrix} \hat{I} & \hat{0} & \hat{0} & \hat{0} \\ (v_x)_a\hat{B} & (X_x)_{aa'}\hat{I} & \hat{0} & \hat{0} \\ \hat{0} & \hat{0} & (X_x)_{bb'}\hat{I} & \hat{0} \\ \hat{C} & (w_x)_{a'}\hat{A} & (w_x)_{b'}\hat{B} & \hat{I} \end{pmatrix}, \\
\hat{M}_c[1, y > 1] &= \left((v_y)_c\hat{B}, (v_y)_c \cdot (w_1)_{a'}\hat{I}, \hat{0}, \hat{0} \right), \\
\hat{M}_c[L_x, y > 1] &= \left(\hat{0}, \hat{0}, (v_y)_c \cdot (v_{L_x})_b\hat{I}, (v_y)_c\hat{B} \right)^T, \\
\hat{M}_c[L_x > x > 1, y > 1] &= \\
&\begin{pmatrix} \hat{0} & \hat{0} & \hat{0} & \hat{0} \\ \hat{0} & \hat{0} & \hat{0} & \hat{0} \\ (v_y)_c \cdot (v_x)_b\hat{I} & \hat{0} & \hat{0} & \hat{0} \\ (v_y)_c\hat{B} & (v_y)_c \cdot (w_x)_{a'}\hat{I} & \hat{0} & \hat{0} \end{pmatrix}, \\
&c \in \{2, 3, \dots, \text{len}(\vec{w}_{y-1}) + 1\}. \tag{4.28}
\end{aligned}$$

Here c is used consistently between Eqns. (4.27) & (4.28) to index the vertical MPO bond β . In an identical manner to Eq. (4.12), a, a', b, b' are used to index the coefficient vectors \vec{w}_x, \vec{v}_x and the coefficient matrix X_x . Note that within a given \hat{M}_c matrix, the value of c is fixed while the values of a, a', b, b' range appropriately over the dimensions of the matrix. This means that an expression such as $(v_y)_c \cdot (v_x)_b$ in Eq. (4.28) is a scalar multiplying a vector.

For these expressions to always make sense, we require $L_x \geq L_y$ so that the bottom MPO is long enough to extract all the necessary coefficient vectors and matrices for the vertical direction. The crucial component of this representation is how \vec{w} and \vec{v} appear in \hat{M}_c . For the \hat{I} operator in the bottom row of the matrix, which couples the action of \hat{A} (from below) to the action of \hat{B} (to the right) in the gMPO row, the ‘‘completion’’ interaction coefficients \vec{v}_y are encoded along the β index while the ‘‘beginning’’ interaction coefficients \vec{w}_x are encoded along the α_x index. Similarly for the \hat{I} operator in the first column of the matrix, which couples the action of \hat{A} (from below) to the action of \hat{B} (to the left) in the gMPO row, the ‘‘completion’’ interaction coefficients \vec{v}_y are encoded along the β index while the ‘‘completion’’ interaction coefficients \vec{w}_x are encoded along the α_{x-1} index. This formulation allows for

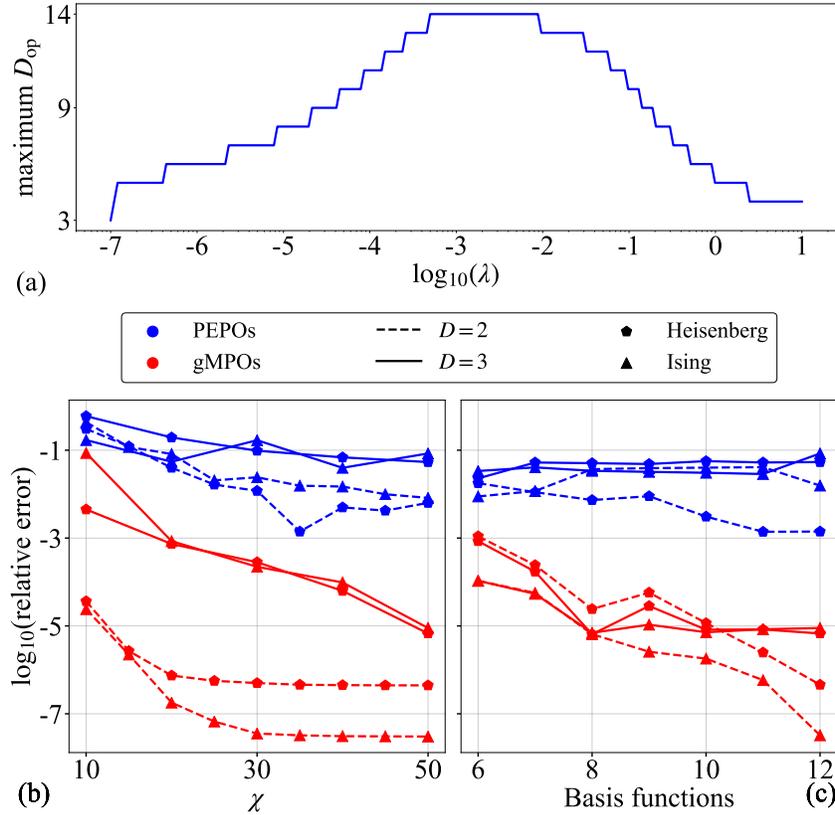


Figure 4.7: (a) The maximum bond dimension of a numerically exact $L = 250$ MPO representation of long-range pairwise Gaussian interactions for various values of the exponential coefficient λ . The algorithm from [97] was used with a singular value threshold of 10^{-10} . (b)-(c) The relative error in the computed expectation value of $\hat{H} = \sum_{i < j} \sigma_i^z \sigma_j^z / |\mathbf{r}_i - \mathbf{r}_j|$ for 8×8 , $D = 2$ (dashed) and 3 (solid) ground states of the AFM Heisenberg model (pentagons) and FM transverse field Ising model (triangles). We compare the Gaussian gMPO technique (red) with the CF-PEPO technique (blue) from Ref. [70]. In (b) we use 12 basis functions to fit the Coulomb potential and vary the boundary dimension χ of the contraction algorithm [35]. In (c) we fix $\chi = 35$ for $D = 2$, $\chi = 50$ for $D = 3$ and vary the number of basis functions K . The convergence of the gMPOs is rapid and strictly governed by K and χ (for a given trial state, either the curve in (b) or in (c) is always decreasing), while the PEPOs converge slowly and become saturated by other sources of numerical error.

vertical interactions of the form $\sum_{i=1}^y e^{-\lambda(y-i)^2} \hat{A}_i \hat{I}_y$, to be “completed” and thus scalar multiplied by “completed” horizontal interactions of the form $\sum_{i=x+1}^L e^{-\lambda(i-x)^2} \hat{I}_x \hat{B}_i$ and $\sum_{i=1}^x e^{-\lambda(x-i)^2} \hat{B}_i \hat{I}_x$, yielding the desired 2D Gaussian potential. The other entries of the tensors can be understood by their analogous form to the previous section and their direct correspondence with $\hat{W}_{gen-sym}$ (Eq. (4.12)).

Due to the inherent challenge of explicitly writing and interpreting the algebraic expressions for \hat{M} when the dimension of the β index is greater than 2, it can be more intuitive to understand the form of these gMPO tensors from a graphical presentation, which is given in Fig. 4.6. Additionally, a straightforward example implementation of these tensors can be found online [122].

The performance of this scheme relative to the PEPO-based scheme from Ref. [70] for evaluating the expectation value of $\hat{H} = \sum_{i<j}(\sigma_i^z \sigma_j^z)/|\mathbf{r}_i - \mathbf{r}_j|$ is shown in Fig. 4.7 and Table 4.1. One notable difference between this case and the previous sections is that there is no longer a generic, guaranteed speedup of the gMPOs over PEPOs because the two methods work differently. The PEPOs encode long-range coefficients by introducing a large auxiliary lattice, while the gMPOs do so by using an increased bond dimension. Since these things affect the computational scaling in different ways and their precise costs depend on specific numerical thresholds, one method is not strictly faster than the other.

However, in practice we observe that the gMPOs are many orders of magnitude more computationally efficient than the PEPOs. The simplest way to see this is to first note that for given values of K , χ , and $D > 2$, the CF-PEPO and gMPO schemes require similar levels of computational effort (Table 4.1). Yet with these same parameters, the gMPOs are approximately 4 orders of magnitude more accurate than the CF-PEPOs (Fig. 4.7). This can be extended to recognize that in order to obtain a given level of accuracy, the gMPOs will be many orders of magnitude faster than the PEPOs, or more generally that the gMPOs can obtain a more accurate answer than the PEPOs in less time.

Additionally, the convergence towards high accuracy is faster and more straightforward when using gMPOs than when using PEPOs. In the case of the gMPOs, the accuracy is systematically governed by χ and K (see Fig. 4.7(b)-(c)). This becomes clear by observing that, for a given trial state, its curve in either Fig. 4.7(b) or (c) is always decreasing. On the other hand, the convergence of the PEPO curves stall. The medium- and high-accuracy regimes are not bounded by errors due to the basis size or χ , but instead by larger numerical errors stemming from additional complicated parameters involved with making the basis radially symmetric [70, 117]. In fact, this is the inherent reason for the major accuracy difference. The gMPO Gaussian basis is radially symmetric up to $\sim 10^{-10}$ (the singular value threshold used in the approximation algorithm), whereas the PEPO bases are only radially symmetric up to significant numerical errors [70, 117].

As a final point, we note that a slightly faster implementation of this long-range gMPO scheme is possible. Since the bond dimensions D_{op} reported in Fig. 4.7(a) are only for \hat{W}_{gen} , the horizontal bond dimension of the gMPO tensors in Eq. (4.28) is almost twice as large. A factor of ~ 2 speedup can be gained in step 4 of the boundary gMPO algorithm if non-symmetric gMPO tensors of horizontal dimension D_{op} are used instead, so that the interactions $\hat{A}_{x,y}\hat{B}_{x,y+b} + \hat{A}_{x,y}\hat{B}_{x+a,y} + \hat{A}_{x,y}\hat{B}_{x+a,y+b}$ are encoded in one gMPO and the interactions $\hat{A}_{x,y}\hat{B}_{x-a,y+b}$ in another. The cost of this bond dimension reduction is an increase in the number of gMPOs that need to be independently evaluated from K to $2K$, but this still leaves a factor of 2 for the speedup because the cost of step 4 depends quadratically on the horizontal bond dimension of the gMPOs.

4.6 Summary

In this work we have presented an algorithm which can evaluate the expectation value of general 2D Hamiltonians without using a PEPO. To accomplish this, we introduced the formalism of a gMPO and showed how it can be used in combination with MPOs to efficiently compute the energy of a PEPO on-the-fly. In addition to the conceptual simplification of rewriting PEPOs in terms of the more familiar MPOs, we also showed that computing the energy using this strategy is 1-2 orders of magnitude faster while being equally as accurate as explicitly using a PEPO. The structure of the algorithm also allows for a new technique to be used for constructing and evaluating 2D Hamiltonians with physical long-range interaction potentials, which we demonstrated to be multiple orders of magnitude more accurate and efficient than existing strategies. We expect that this work will lower the computational and conceptual barriers to using tensor network operators in future PEPS calculations. We hope that this opens the door to the study of new, more complicated Hamiltonians in the tensor network community.

Finally, although this work focused on the specific case of finite systems, the fundamental requirement for the formulation of the algorithm to apply is that the contraction method starts from the boundary. Since much is known about infinite MPOs [28, 123, 124] and many prominent contraction methods for infinite PEPS [34] also begin from the boundary [39, 40, 87, 125], we expect that the concepts presented in this work can be generalized to the infinite case.

4.7 Short Appendix

In Section 4.5, we reported the exact construction of the vertical MPO matrices and the gMPO tensors for a Hamiltonian that had non-symmetric “linear” interactions up to distance $R = 1$, and non-symmetric diagonal interactions up to distance $R = \sqrt{2}$. Following the concepts in that example, and the general ideas behind MPO construction, in this Appendix we will give the exact construction for a Hamiltonian with non-symmetric linear interactions up to a general distance R , and non-symmetric diagonal interactions up to $\sqrt{2}R$. The interactions coefficients will be denoted $j_{x,y}$, where x is the horizontal distance between the local operators and y is the vertical distance.

The vertical MPO matrices are size $(R + 2) \times (R + 2)$ and they are given by,

$$\hat{O}_{\beta_1}[x, 1] = \left(\hat{I} \quad \hat{A} \quad \left[\hat{\hat{O}} \right]_{R-1} \right),$$

$$\hat{O}_{\beta_{y-1}\beta_y}[x, y > 1] = \left(\begin{array}{c|c|c} \hat{I} & \hat{A} & \hat{\hat{O}} \\ \hline \hat{\hat{O}} & \hat{\hat{O}} & \mathbb{1}_{(R-1) \times (R-1)} \hat{I} \\ \hline \hat{\hat{O}} & \hat{\hat{O}} & \hat{\hat{O}} \end{array} \right). \quad (4.29)$$

The gMPO tensors are size $(2R + 2) \times (2R + 2) \times (R + 1)$, where the third dimension

is the size of the β index. They are given by,

$$\begin{aligned}
& \hat{M}_1[1, y > 1] = \\
& \left(\hat{C}, j_{1,0}\hat{A}, \dots, j_{R,0}\hat{A}, \hat{B}, \left[\hat{\theta} \right]_{R-1} \hat{I} \right), \\
& \hat{M}_1[1, y > 1] = \left(\hat{I}, \hat{B}, \left[\hat{\theta} \right]_{2R-1} \hat{C} \right)^T, \\
& \hat{M}_1[L_x > x > 1, y > 1] = \\
& \left(\begin{array}{c|ccc|c|c} \hat{I} & \hat{\theta} & \hat{\theta} & \hat{\theta} & \hat{\theta} & \hat{\theta} \\ \hline \hat{B} & \hat{\theta} & \hat{\theta} & \hat{\theta} & \hat{\theta} & \hat{\theta} \\ \hline \hat{\theta} & \mathbb{1}_{(R-1) \times (R-1)} \hat{I} & \hat{\theta} & \hat{\theta} & \hat{\theta} & \hat{\theta} \\ \hline \hat{\theta} & \hat{\theta} & \hat{\theta} & \hat{\theta} & \mathbb{1}_{(R-1) \times (R-1)} \hat{I} & \hat{\theta} \\ \hline \hat{\theta} & \hat{\theta} & \hat{\theta} & \hat{\theta} & \hat{\theta} & \hat{\theta} \\ \hline \hat{C} & j_{1,0}\hat{A}, \dots, j_{R-1,0}\hat{A} & j_{R,0}\hat{A} & \hat{B} & \hat{\theta} & \hat{I} \end{array} \right), \\
& \hat{M}_k[1, y > 1] = \\
& \left(j_{0,k-1}\hat{B}, j_{1,k-1}\hat{I}, \dots, j_{R,k-1}\hat{I}, \left[\hat{\theta} \right]_{R+1} \right), \\
& \hat{M}_k[L_x, y > 1] = \\
& \left(\left[\hat{\theta} \right]_{R+1} j_{1,k-1}\hat{I}, \dots, j_{R,k-1}\hat{I}, j_{0,k-1}\hat{B} \right)^T, \\
& \hat{M}_k[L_x > x > 1, y > 1] = \\
& \left(\begin{array}{c|ccc|c} \left[\hat{\theta} \right]_{R+1} & \hat{\theta} & & \hat{\theta} \\ \hline j_{1,k-1}\hat{I} & & & \hat{\theta} \\ j_{2,k-1}\hat{I} & & \hat{\theta} & \hat{\theta} \\ \vdots & & & \\ j_{R,k-1}\hat{I} & & & \hat{\theta} \\ \hline j_{0,k-1}\hat{B} & j_{1,k-1}\hat{I}, j_{2,k-1}\hat{I}, \dots, j_{R,k-1}\hat{I} & & \left[\hat{\theta} \right]_{R+1} \end{array} \right), \\
& k \in \{2, 3, \dots, R+1\}. \tag{4.30}
\end{aligned}$$

In these expressions, $\mathbb{1}_{N \times N}$ is an $N \times N$ identity matrix. Additionally, when something is enclosed in square brackets and labelled with a subscript n , it means “repeat this n times”. Based on the dimensions of the other blocks, it should be clear which axis it should be expanded along. This is only used in places where it is not otherwise obvious to expand the blocks to match the dimensions of adjacent blocks.

ENTANGLEMENT IN THE QUANTUM PHASES OF AN UNFRUSTRATED RYDBERG ATOM ARRAY

This chapter is based on the following publication:

1. O'Rourke, M. J. & Chan, G. K. Entanglement in the quantum phases of an unfrustrated Rydberg atom array. *arXiv preprint arXiv:2201.03189* (2022).

5.1 Abstract

We report on the ground state phase diagram of interacting Rydberg atoms in the geometrically unfrustrated square lattice array. Using new tensor network algorithms, we scale to large systems in two dimensions while including all long-range interactions, revealing the phases in the bulk and their analogs in accessible finite arrays. We find a greatly altered phase diagram from earlier numerical and experimental studies, and in particular, we uncover an emergent entangled quantum nematic phase that appears in the absence of frustration. Broadly our results yield a conceptual guide for future experiments, while our techniques provide a blueprint for converging numerical studies in other lattices.

5.2 Introduction

Rydberg atom arrays, where cold atoms are trapped in an optical lattice and interact via excitation into Rydberg states [126, 127], have generated interest for quantum information processing and to realize exotic many-body states [128–142]. A recent experiment [143], backed by numerical studies [144, 145], has suggested a richness in 2D Rydberg atom array ground states on a square lattice. However, although the observed, non-disordered, phases are not all classical crystals, they contain little entanglement [144]. Thus it remains unclear whether such arrays realize non-trivial entangled quantum ground-states on simple lattices. Here, we describe new numerical techniques that greatly reduce finite size effects, allowing us to confidently converge the bulk phase diagram. We also showcase techniques that address large finite two-dimensional lattices realized in experiments, while incorporating all long-range interactions. Unexpectedly, we derive quite different physics from our simulations compared to both previous theoretical and experimental analyses—including the emergence of a non-trivial, entangled nematic phase, even on the

geometrically unfrustrated square lattice array.

The Rydberg atom array Hamiltonian is

$$\hat{H} = \sum_{i=1}^N \left[\frac{\Omega}{2} \hat{\sigma}_i^x - \delta \hat{n}_i \right] + \frac{1}{2} \sum_{i \neq j} \frac{V}{(|\vec{r}_i - \vec{r}_j|/a)^6} \hat{n}_i \hat{n}_j. \quad (5.1)$$

Here $\hat{\sigma}_i^x = |0_i\rangle\langle 1_i| + |1_i\rangle\langle 0_i|$ and $\hat{n}_i = |1_i\rangle\langle 1_i|$ ($\{|0_i\rangle, |1_i\rangle\}$ denote ground and Rydberg states of atom i). a is lattice spacing, Ω labels Rabi frequency, and δ describes laser detuning. V parameterizes the interaction strength between excitations. This can be re-expressed in terms of the Rydberg blockade radius R_b , with $V/(R_b/a)^6 \equiv \Omega$. We study the square lattice in units $a = \Omega = 1$ [144], yielding two free parameters δ and R_b .

The ground states of this Hamiltonian are simply understood in two limits. For $\delta/\Omega \gg 1$, $R_b \neq 0$, the system is classical and one obtains classical crystals of Rydberg excitations [146–149] whose spatial density is set by the competition between δ and R_b . For $\delta/\Omega \ll 1$, $R_b \neq 0$, Rydberg excitations are disfavored and the solutions are dominated by Rabi oscillations, leading to a trivial “disordered” phase [144, 150, 151]. In between these limits, it is known in 1D that no other density-ordered ground states exist besides the classical-looking crystals, with a Luttinger liquid appearing on the boundary between ordered and disordered phases [151].

In 2D, however, the picture is quite different. An initial study [144] using the density matrix renormalization group (DMRG) [22, 25, 26, 30] found additional quantum crystalline (or “density-ordered”) phases, where the local excitation density is not close to 0 or 1. A recent experiment on a 256 programmable atom array has realized such phases [143]. However, as also discussed there, the density-ordered phases are unentangled quantum mean-field phases, and thus not very interesting. In addition, more recent numerical results [134] highlight the sensitivity of the physics to the tails of the Rydberg interaction and finite size effects. Thus, whether Rydberg atom arrays on a simple unfrustrated lattice—such as the square lattice—support interesting quantum ground-states, remains an open question.

Here, we resolve these questions through high-fidelity numerical simulations. To do so, we employ variational tensor network methods. Tensor networks have led to breakthroughs in the understanding of 2D quantum many-body problems [41], and we rely on two techniques that address specific complexities of simulating interactions in Rydberg atom arrays. The first we term Γ -point DMRG, which captures interactions out to *infinite* range, while employing a traditional finite system

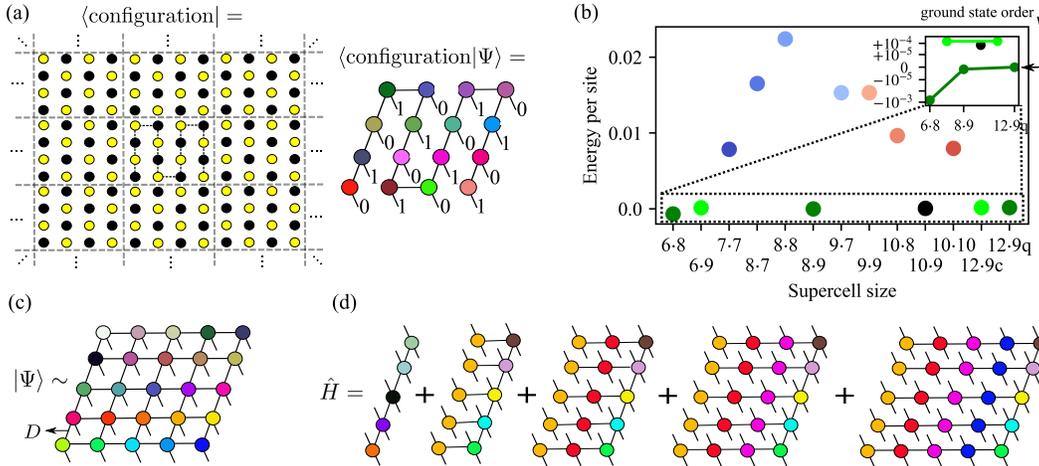


Figure 5.1: Numerical methods and strategy. (a) A schematic representation of Γ -point DMRG. A single infinite bulk configuration is given by periodic images of the central supercell configuration. The wavefunction coefficient for this infinite configuration is given by the contraction of a snake MPS, which is defined only within a single supercell. (b) By widely varying the size of the supercell, Γ -point DMRG obtains many different ground states. Identifying all accessible supercells which give the same ground state order (shown with identically colored points), we can ensure that all competing low-energy states are well converged w.r.t. finite size effects, and thus properly identify the true ground state (inset shows ground-state order (dark green) converged w.r.t. supercell size, separated from other low-energy orders by 10^{-4} energy units). (c) A PEPS wavefunction ansatz with bond dimension D for a finite system. Each tensor is a different color because they can all be unique. (d) A simplified diagrammatic representation of the long-range Hamiltonian construction for PEPS in Ref. [152]. All terms in the Hamiltonian are accounted for by a sum of L_x comb tensor network operators. Tensors of the same color are identical.

two-dimensional DMRG methodology [30], removing interaction truncations and boundary effects present in earlier studies [133, 134, 144, 145]. This allows us to controllably converge the bulk phase diagram. The second is a representation of long-range interactions [152] compatible with projected entangled pair states (PEPS) [23, 63–65]. With this, we use PEPS to find the ground states of a Hamiltonian with long-range interactions for the first time, and specifically here, model the states of finite arrays of large widths as used in experiment. Both techniques can be used for more faithful simulations of Rydberg atoms in other settings. We first describe the new numerical methods, before turning to the bulk and finite-size phase behavior of square lattice Rydberg arrays and the question of entangled quantum phases.

5.3 Numerical strategy and techniques

Bulk simulations and Γ -point DMRG

A challenge in simulating Rydberg atom arrays is the long-range tails of the interaction. Because itinerancy only arises indirectly as an effective energy scale [150], the main finite size effects arise from interactions. Many previous studies have employed a cylindrical DMRG geometry common in 2D DMRG studies [30]. However, there the interaction is necessarily truncated to the cylinder half-width, while along the open direction, edge atoms experience different interactions than in the bulk; both choices produce strong finite size effects.

To avoid these problems, we perform 2D DMRG calculations in a site Bloch basis. Given the site basis $|n_{x,y}\rangle$, $n \in \{0, 1\}$, the Bloch basis states are periodic combinations, $|\tilde{n}_{x,y}\rangle = \sum_l |n_{(x,y)+R_l}\rangle$ where $R_l = (n \cdot L_x, m \cdot L_y)$, $n, m \in \mathbb{Z}$, are lattice vectors, and L_x, L_y are the supercell side lengths. The above correspond to Bloch states at the Γ -point of the supercell Brillouin zone. The corresponding matrix product state (MPS) is expressed as $|\Psi\rangle = \sum_{\{e\}} \prod_{x,y} A_{\{e_{x,y}\}}^{\tilde{n}_{x,y}} |\tilde{n}_{x,y}\rangle$ where $A^{\tilde{n}_{x,y}}$ is the MPS tensor associated with Bloch function $\tilde{n}_{x,y}$, $e_{x,y}$ denote its bonds, and a 2D snake ordering has been chosen through the lattice. Such a Γ -point 2D DMRG formally models an infinite lattice (Fig. 5.1a) with a wavefunction constrained by the supercell shape. This differs from using a periodic MPS as periodicity is enforced by the Bloch basis rather than the MPS. It is also different from a cylindrical/toroidal geometry, which are both finite; here the lattice remains infinite, and there is entanglement between supercells induced by the Bloch basis. The interactions are then expressed as a lattice sum,

$$\hat{H} = \sum_i \left[\frac{1}{2} \hat{\sigma}_i^x - \delta \hat{n}_i \right] + \frac{1}{2} \sum_{i \neq j+R_l, R_l} \frac{R_b^6}{|\vec{r}_i - \vec{r}_{j+R_l}|^6} \hat{n}_i \hat{n}_j, \quad (5.2)$$

where the operators now act in the Bloch basis. Further details of this approach are discussed in the supplementary information (SI) in Appendix A.

The only finite size parameter is the supercell size $L_x \times L_y$. We thus perform exhaustive scans over L_x, L_y to identify competing orders. We systematically converge the energy per site of low-energy orders by increasing the commensurate supercell sizes to contain many copies of the order (up to 108 sites). The finite size effects converge rapidly because no interactions are truncated and there are no edge effects even in the smallest cells, allowing us to converge the energy per site to better than 10^{-5} , compared to the smallest energy density difference we observe between competing phases of $\sim 10^{-4}$ (see Fig. 5.1b and Appendix A).

Finite simulations and PEPS with long-range interactions

To simulate ground-states of finite arrays, we consider finite systems (with open boundaries) of sizes 9×9 up to 16×16 atoms. This resembles capabilities of near-term experiments [141, 143]. The width of the largest arrays challenges what can be confidently described with MPS and DMRG for more entangled states. Consequently, we employ PEPS wavefunctions which capture area law entanglement in 2D, and can thus be scaled to very wide arrays (Fig. 5.1c). Together with DMRG calculations on moderate width finite lattices, the two methods provide complementary approaches to competing phases and consistency between the two provides strong confirmation. However, PEPS are usually combined with short-range Hamiltonians. We now discuss a way to combine long-range Hamiltonians efficiently with PEPS without truncations.

For this, we rely on the representation we introduced in Ref. [152]. This encodes the long-range Hamiltonian as a sum of “comb” tensor network operators (Fig. 5.1d). As discussed in Ref. [152], arbitrary isotropic interactions can be efficiently represented in this form, which mimics the desired potential via a sum of Gaussians, i.e. $\frac{1}{r^6} = \sum_{k=1}^{k_{\max}} c_k e^{-b_k r^2}$ (where $k_{\max} \sim 7$ for the desired accuracy in this work). The combs can be efficiently contracted much more cheaply than using a general tensor network operator.

While Ref. [152] described the Hamiltonian encoding, here we must also find the ground-state. We variationally minimize $\langle \Psi | \hat{H} | \Psi \rangle$ using automatic differentiation [119]. Combined with the comb-based energy evaluation, this allows for both the PEPS energy and gradient to be evaluated with a cost linear in lattice size. (Stably converging the PEPS optimization involves some challenges. Further details can be found in Appendix A).

5.4 Bulk phases

Summary of the phase diagram. Fig. 5.2a shows the bulk phase diagram from Γ -point DMRG with infinite-range interactions. We first discuss the orders identified by their density profiles (orders of some phase transitions are briefly discussed in Appendix A). Where we observe the same phases as in earlier work [144] we use the same names, although there are very substantial differences with earlier phase diagrams.

$R_b < 1.8$. With weaker interactions, the ground states progress through densely-packed, density-ordered phases starting from checkerboard (pink) ($R_b \sim 1.2$), to

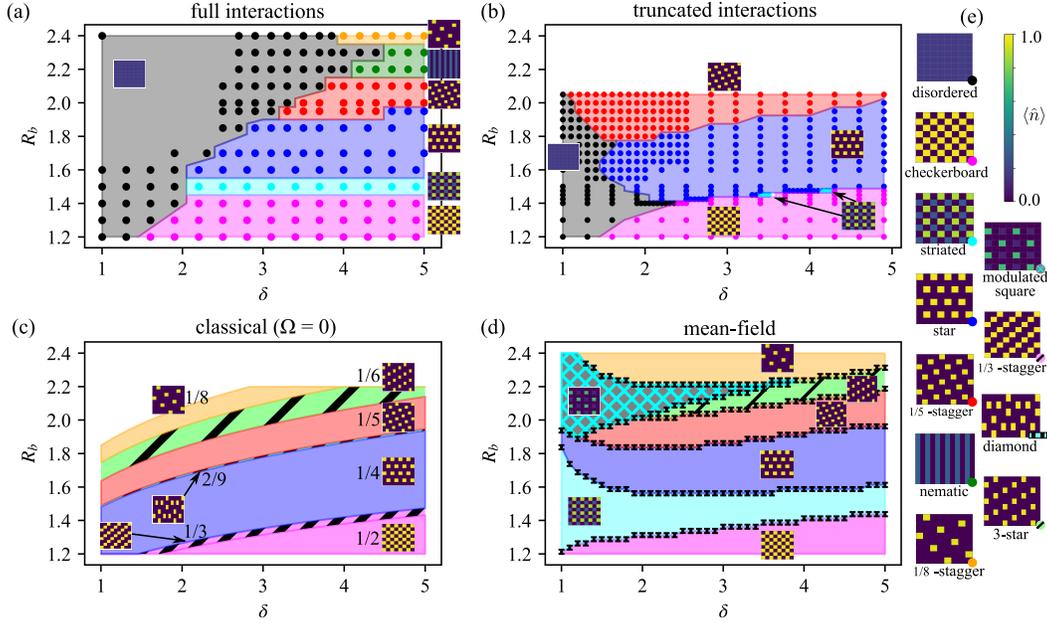


Figure 5.2: Phase diagrams of the bulk system under various assumptions. The color of a dot/region identifies the ground state order. The density profiles for each color are given in (e) and shown near each phase domain. (a) The phase diagram given by Γ -point DMRG including all long-range interactions. (b) The phase diagram from Γ -point DMRG when interactions are truncated to 0 beyond a distance of $|\vec{r}_i - \vec{r}_j| = 2$. (c) The classical phase diagram (when all sites are either fully occupied or empty) including all long-range interactions. (d) The mean-field phase diagram, including all long-range interactions. Error bars display the uncertainty of the computed phase boundaries. (e) Representative density profiles for all phases in (a)-(d), identified by the colored dot in each lower right corner. All profiles have Γ -point boundary conditions on all edges. In (a)-(b) dots denote computed data, while shading is a guide for the eye. (c),(d) are computed with very fine resolution/analytically, and thus no dots are shown.

striated (cyan) ($R_b \sim 1.5$), to star (blue) ($R_b \sim 1.6$). While the checkerboard and star phases are classical-like crystals, the striated state is a density-ordered quantum phase, seen previously [144].

$R_b > 1.8$. Here, the phases look very different from earlier work, which truncated the interactions [144]. Ordered ground states start with the $\frac{1}{5}$ -“staggered” phase (red) ($R_b \sim 1.95$), then progress to a “nematic” phase (dark green) ($R_b \sim 2.2$) and the $\frac{1}{8}$ -“staggered” phase (gold) ($R_b \sim 2.4$). There is also a small region at larger δ (not shown) where the nematic phase and a “3-star” classical-like crystal appear to be essentially degenerate, with an energy difference per site of $\Delta e < 3 \cdot 10^{-5}$ (see Appendix A).

Effects of interactions. In Fig. 5.2b we show the phase diagram computed using Γ -point DMRG with interactions truncated to distance 2. This approximation resembles earlier numerical studies [144], but here bulk boundary conditions are enforced by the Bloch basis, rather than cylindrical DMRG. Comparing Figs. 5.2a,b we see the disordered and striated phases are greatly stabilized using the full interaction, and new longer-range orders are stabilized at larger R_b . Comparing Fig. 5.2b and Ref. [144], we see that having all atoms interact on an equal footing (via the Bloch basis) destroys some quantum ordered phases seen in [144] at larger R_b .

Classical, mean-field, and entangled phases. Without the Rabi term Ω , one would obtain classical Rydberg crystals without a disordered phase. Fig. 5.2c shows the classical phase diagram. For the δ values here, the 1D classical phase diagram has sizable regions of stability for all accessible unit fraction densities [147, 151]. However, the connectivity of the square lattice in 2D changes this. For example, only a tiny part of the phase diagram supports a $\frac{1}{3}$ -density crystal, and we do not find a stable $\frac{1}{7}$ -density crystal within unit cell sizes of up to 10×10 . All ordered quantum phases in Fig. 5.2a appear as classical phases except for the striated and nematic phases, while there are small regions of classical phases at densities $\frac{1}{3}$ and $\frac{2}{9}$ with no quantum counterpart. The striated and nematic phases emerge near the $\frac{1}{3}$ and $\frac{1}{7}$ density gaps respectively, however the nematic phase also supersedes the large region of the $\frac{1}{6}$ density “3-star” crystal.

Ref. [143] suggested that quantum density-ordered phases are qualitatively mean-field states of the form $\prod_i \alpha_i |0_i\rangle + \sqrt{1 - |\alpha_i|^2} |1_i\rangle$. Fig. 5.2d shows the mean-field phase diagram. The disordered phase does not appear, as it emerges from defect hopping and cannot be described without some entanglement for $\delta \geq 1$ [150]. The mean-field phase diagram contains features of both the classical and quantum phase diagrams. The striated quantum phase indeed appears as a mean-field state, confirmed by the match between the mean-field and exact correlation functions (Fig. 5.3a). However, the nematic phase does *not* appear, and in its place is the same $\frac{1}{6}$ -density crystal stabilized in the classical phase diagram. This shows that a treatment of entanglement is required to describe the nematic phase.

Nature of the entangled nematic phase. Fig. 5.3b shows the density correlation function of the nematic phase, which does not display mean-field character. To reveal the phase structure, Fig. 5.3c shows the lowest energy classical states in the same region of the phase diagram. Due to the Rydberg blockade radius ($R_b = 2.3$), excitations are spaced by 3 units within a column, giving 3 column configurations

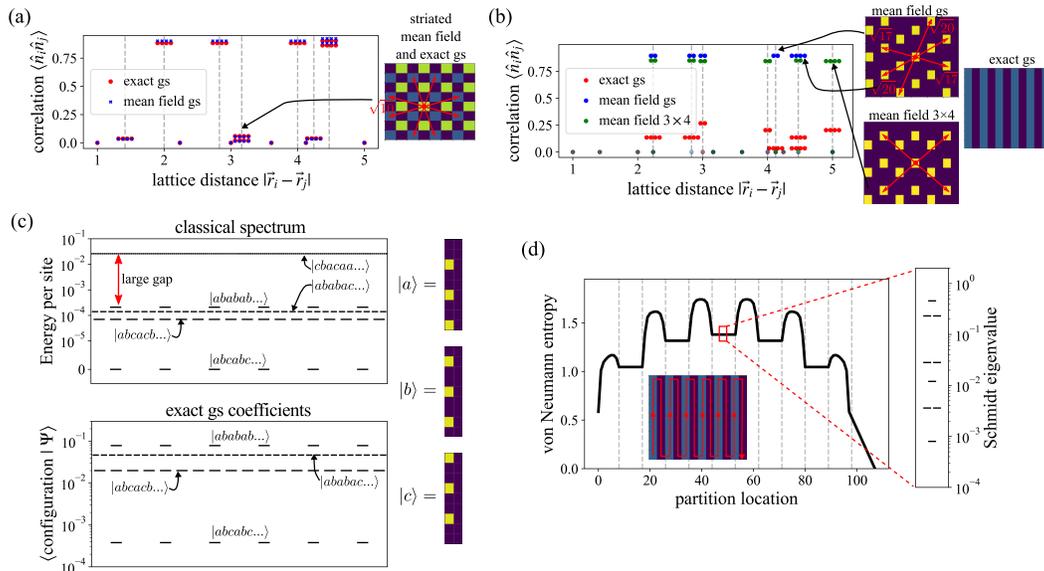


Figure 5.3: Mean-field striated versus entangled nematic phase. (a) Density-density correlation functions of the mean-field and exact striated ground state, both at $(\delta, R_b) = (3.1, 1.5)$; these agree, confirming the mean-field nature of the striated phase. (b) Density-density correlation functions for the entangled nematic phase ground state and two different mean-field ground states (from a 6×3 unit cell and a 3×4 unit cell) at $(\delta, R_b) = (5.0, 2.3)$. In (a)-(b), 2-fold/4-fold degeneracy of a peak is indicated by 2/4 horizontal dots distributed around the proper distance coordinate. 8-fold degeneracy in (a) is shown as two rows of 4 dots. The non-mean-field (entangled) character of the nematic phase is evident. (c) Structure of the nematic state in terms of classical configurations constructed via compositions of 3 individual column states $|a\rangle$, $|b\rangle$, $|c\rangle$. In the classical limit, there are distinct sets of low-energy configurations, all characterized by the absence of adjacent columns in the same state (e.g. $|aa\dots\rangle$) and large degeneracies due to permutational symmetry between $|a\rangle$, $|b\rangle$, and $|c\rangle$. The lowest in energy is 6-fold degenerate, corresponding to the 3-star state. However, in the quantum nematic state the configurations that are slightly higher in energy have much larger wavefunction coefficients. The most relevant classical states in the wavefunction are those with the greatest amount of single-excitation itinerancy without introducing unfavorable states like $|aa\dots\rangle$. (d) Bipartite entanglement entropy for each possible bipartition of the 12×9 supercell nematic ground state. One inset shows the “path” that the “partition location” axis follows through the supercell MPS, while the other shows the entanglement spectrum at a central cut.

$|a\rangle, |b\rangle, |c\rangle$. Column-column interactions, however, prevent adjacent columns from being in the same configuration (with excitations separated by 2 units); thus states such as $|abcb\dots\rangle$ are “allowed”, but $|accb\dots\rangle$ are not. Without long-range interactions, these column constraints give rise to an exponential classical degeneracy. Long-range interactions partially lift the classical degeneracy, yielding the $|abc\dots\rangle$ crystal (3-star phase) and its 6-fold degenerate permutations. However, after including quantum fluctuations and entanglement in each of the distinct crystals through a 4th order perturbative treatment of σ_x (giving rise to defect itinerancy), $|abab\dots\rangle$ and related configuration energies are lowered below those of the $|abc\dots\rangle$ configurations; entanglement stabilizes a different order (see Appendix A).

Fig. 5.3c gives the weights of the configurations in the quantum ground-state, which are distributed across the exponentially numerous $|abab\dots\rangle, |abcbab\rangle$ etc., configurations, with the classical crystal $|abc\dots\rangle$ configurations strongly disfavored. The bi-partite entanglement entropy and entanglement spectrum are further shown in Fig. 5.3d. The entanglement spectrum carries 3 large Schmidt values across every cut along the DMRG snake MPS, showing the state is macroscopically entangled in a non-trivial way (see Appendix A for discussion), and well approximated by an MPS of bond dimension 3. The size of our current calculations cannot conclusively resolve the fate of this macroscopically entangled state in the true thermodynamic limit, due to the small energy scales involved. However, the stability of the simpler locally entangled quantum crystal $|abab\dots\rangle$ is certain, and it is generally clear that quantum fluctuations are much stronger in this phase than in any of the surrounding ordered phases (further discussion in Appendix A).

5.5 Finite phase diagram

Current experiments are limited to lattices with open boundary conditions consisting of a few hundred atoms [141, 143]. To investigate how this modifies the bulk behavior, we computed the phase diagram of selected finite lattices from size 9×9 to 16×16 , using DMRG for the smaller sizes and our PEPS methodology for the larger ones.

We first focus (in Fig. 5.4a) on understanding the fate of the ordered phases on the 15×15 lattice along three slices: $\delta = 2.7, 4.0$, and 5.0 (16×16 lattice phases, as well as other lattice sizes, are discussed in Appendix A). Here, many finite lattice ground state orders resemble those in the bulk. However, their regions of stability are substantially reduced and their patterns are broken by frustration. Out

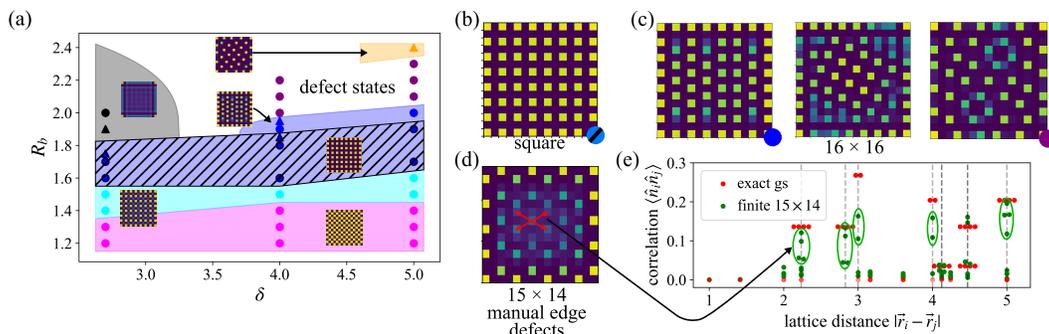


Figure 5.4: Phase diagram of the 15×15 finite system and finite lattice orders. (a) The phase diagram, where colors correspond to the same phase classifications as Fig. 5.2. Triangles represent tentative classification of points showing inconsistent PEPS convergence, see Appendix A. A new “square” order is specified in (b) and various examples of boundary-bulk frustrated ground states in (c). (d) The density profile for a nematic-like ground state that can be stabilized on a 15×14 lattice at $(\delta, R_b) = (3.4, 2.1)$ with manually tailored edge excitations (see text). (e) Comparing the correlations of the finite nematic phase to the “exact” bulk phase. The degeneracy of the peaks is split by the boundary excitations, but the number of peaks is generally conserved between the two (green ovals), which provides a clear distinction from mean-field states (see Fig. 5.3b).

of the density-ordered quantum phases, the striated mean-field phase remains due to its commensurate boundary-bulk configurations, while in the region of strongest interactions the nematic phase is destabilized. A new region of classical order, called here the square phase (Fig. 5.4b), emerges across much of the $R_b = 1.5 - 1.8$ region where the star phase was stable in the bulk [145]. We distinguish the square order from the striated order in the sense that the former has negligible quantum fluctuations on the $(1, 1)$ -sublattice, although it is unclear if the square and striated orders constitute truly distinct phases (in the bulk phases the square order is not stable, only the striated order appears).

In Fig. 5.5, we directly compare the experimental results on the 13×13 lattice to our calculations on the same lattice. The analysis of the experiments in Ref. [143] was based on simulations on the 9×9 lattice using truncated interactions. This assigned only part of the experimental non-zero order parameter space to a square/striated phase (see Fig. 5.5a left panel, note, the order parameter does not distinguish between square/striated orders). However, our simulations (Fig. 5.5b left panel) in fact reproduce the full region of the non-zero order parameter, and thus the whole region seen experimentally should be assigned to a square/striated phase, with the square order appearing in the upper part of the region. Similarly, the

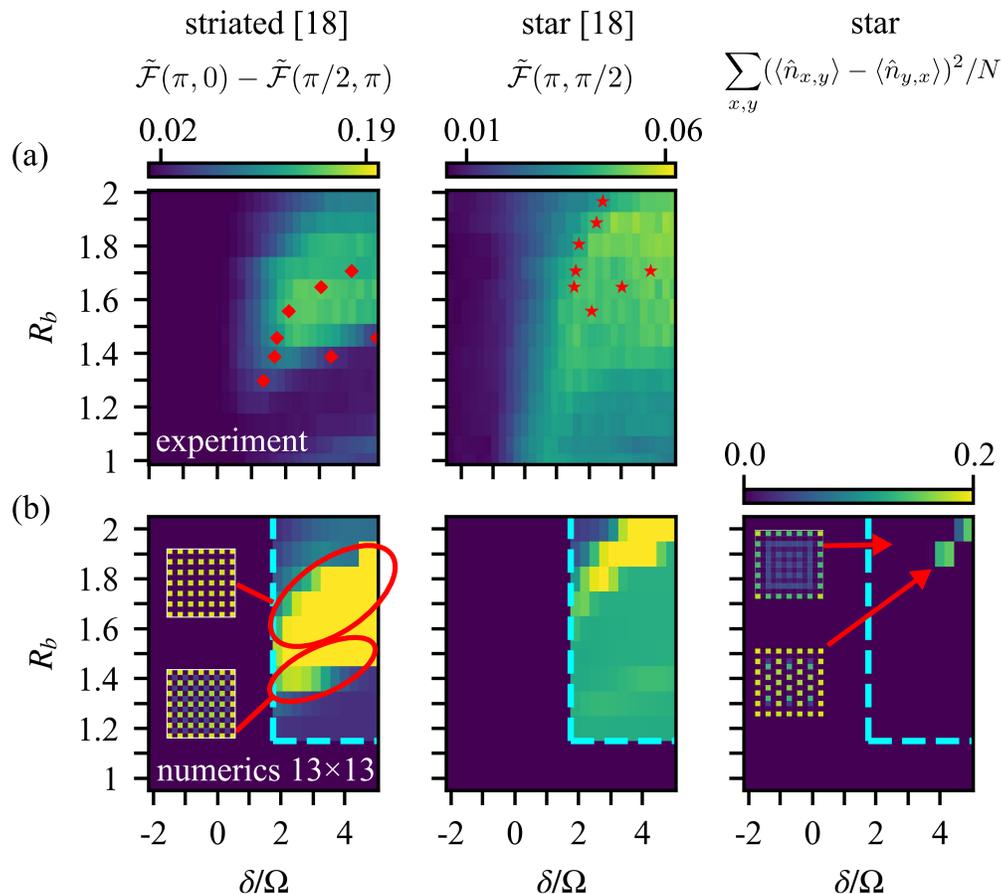


Figure 5.5: Comparison to experiment. The (a) row directly reproduces experimental phase diagram data on the 13×13 lattice (data extracted from Ref. [143] Fig. 4), while the (b) row is 13×13 numerical data computed in this work. The first two columns show the order parameters used in [143] to identify the square/striated and star phases, while the third column shows a new, more sensitive order parameter for the star phase. Red dots in (a) denote the phase boundaries assigned in [143], while the cyan dotted lines in (b) indicate the subset of parameter space that was computed. Our calculations support a re-interpretation of the experimental data with a significantly larger square/striated region and much smaller star phase.

experimental analysis identified a large region of star order (Fig. 5.5a middle panel). This assignment is complicated by the edge ordering before the bulk [153], which prevents the order parameter used from cleanly distinguishing the star phase from other phases. However, our simulations suggest that the region of the star phase should be considered to be much smaller, located at the very top of the non-zero order region, and this is confirmed using a different, more sensitive order parameter (Fig. 5.5b, right panel). Overall, the measured data corresponds more closely to our numerics than earlier simulations, giving confidence in our more precise interpretation (more discussion in Appendix A).

Stabilizing entangled ground-state order. Generally, additional impacts of boundary physics can be understood in terms of frustration of the bulk order by the boundary order, where excitations concentrate more densely due to the lower energetic penalty from fewer long-range interactions on the edge. Examples of the effects of this frustration, ranging from modified bulk orders, to defect dominated states, are shown in Fig. 5.4b-c (see also Appendix A).

We searched for conditions to stabilize the entangled nematic ground-state on a finite lattice by manipulating boundary effects. We scanned various rectangular sizes and explicitly “removed” patterns of atoms from the edges to induce different bulk orders. We found the best conditions to stabilize the nematic phase occur near $(\delta, R_b) = (3.4, 2.1)$, on a 15×14 lattice, while removing edge atoms to create a spacing of 4 on two edges and 3 on the other two edges (Fig. 5.4d) ¹. Although there are strong finite size effects, the density profile and correlation functions (Fig. 5.4d-e) reveal qualitative similarities to the bulk nematic phase, in particular, the presence of 4-fold correlation peaks at distance $\sqrt{5}$ and $\sqrt{8}$, which are also a feature of the bulk entangled phase (Fig. 5.3b).

5.6 Summary

Using new tensor network simulation methods, we have obtained a converged understanding of the phase diagram of Rydberg atom arrays in both bulk and finite simple square lattices. Surprisingly, our bulk phase diagram is quite different from that predicted in earlier numerical studies, while on finite lattices, our results support a reinterpretation of previous experimental analysis. Theoretically, this is due to the subtle effects of the long-range interactions that are addressed by our techniques, while experimentally, it brings into focus the challenge of more accurate theoretic-

¹Note that the location of this state in phase space cannot be directly compared to the locations of states in Fig 5.4a due to the significant difference in the treatment of the boundary.

cal models to interpret increasing experimental capabilities in quantum many-body physics [153–156]. Perhaps most intriguingly, we find strong evidence that the geometrically unfrustrated square lattice supports an entangled quantum nematic phase, brought about by the competition between emergent itinerancy and the constraints of the Rydberg interaction.

A primary focus of Rydberg atom array experiments has been to realize well-studied short-range Hamiltonians, for example, on frustrated lattices. However, we find that lattice frustration is not necessary to produce interesting entanglement effects in Rydberg systems. In fact our work highlights the richness and complexity intrinsic to Rydberg atom arrays, due to the non-trivial effects of their native interactions.

ULTRACOLD HYPERMETALLIC POLAR MOLECULES FOR PRECISION MEASUREMENTS

This chapter is based on the following publication:

1. O'Rourke, M. J. & Hutzler, N. R. Hypermetallic polar molecules for precision measurements. *Phys. Rev. A* **100**, 022502. doi:10.1103/PhysRevA.100.022502 (2 Aug. 2019).

6.1 Abstract

Laser cooling is a powerful method to control molecules for applications in precision measurement, as well as quantum information, many-body physics, and fundamental chemistry. However, many optically-active metal centers in valence states which are promising for these applications, especially precision measurement, are difficult to laser cool. In order to extend the control afforded by laser cooling to a wider array of promising atoms, we consider the use of small, hypermetallic molecules that contain multiple metal centers. We provide a detailed analysis of YbCCCa and YbCCAl as prototypical examples with different spin multiplicities, and consider their feasibility for precision measurements making use of the heavy Yb atom. We find that these molecules are linear and feature metal-centered valence electrons, and study the complex hybridization and spin structures that are relevant to photon cycling and laser cooling. Our findings suggest that this hypermetallic approach may be a versatile tool for experimental control of metal species that do not otherwise efficiently cycle photons, and could present a new polyatomic platform for state-of-the-art precision measurements.

6.2 Introduction

Cold molecules have applications in diverse areas, from quantum information and many-body physics to searches for new physics beyond the Standard Model [157–160]. All of these applications benefit from (or rely on) the ability to efficiently scatter a large number of photons from a molecule, which allows for laser-cooling and trapping as well as effective quantum state preparation and readout. However, due to their inherently complex internal structures, finding molecules that can efficiently cycle photons is often difficult. A primary concern is that since there are no

selection rules for vibration during an electronic decay, the population of excitations can rapidly become diluted over a large number of internal states.

Molecules with particular electronic structure and bonding properties avoid this problem [161–163], which has made laser cooling and trapping of molecules a reality in the last few years [164–169]. Suitable molecules typically have a very simple structure of non-bonding *s* electrons localized on a metal center [162, 163, 170], such as SrF [167], SrOH [171], CaF [168, 169], YO [166, 172], TlF [173], YbF [174], BaF [175, 176], and isoelectronic analogues. This decouples the electronic and vibrational motion, resulting in highly diagonal Franck-Condon (FC) matrices [177, 178]. This property allows these molecules to be laser cooled with a reasonable number of “repump” lasers to return excited vibrational states back to the photon cycling process. Many difficulties remain, including avoiding rotational branching, “dark states” that don’t scatter photons [179, 180], and Renner-Teller effects, but a diagonal FC matrix is a necessary condition for laser cooling methods that rely on spontaneous decay.

Despite the success of this scheme in recent years, molecules useful for precision measurements of fundamental symmetry violations to search for physics beyond the Standard Model pose a number of additional challenges [159, 181]. First, the molecule must feature a heavy atom that has core-penetrating valence electrons, since sensitivity to CP-violating physics relies on relativistic motion of electrons near the nucleus. The rapid scaling of this feature with proton number, typically Z^{2-3} , effectively restricts the choice of atom to those on the bottom two rows of the periodic table and having valence *s* or *p* electrons. The requirement of efficient photon cycling restricts even further, and generally requires that the atoms have *only* valence *s* and *p* electrons. Thus diatomic species like BaF [176], RaF [182], TlF [173], and YbF [174] are promising candidates for laser cooling and are sensitive to new physics beyond the Standard Model.

However, the requirement of simple electronic structure effectively precludes the advantageous Ω -doublets that arise from electronic orbital angular momentum [183–185]. These nearly-degenerate states of opposite parity can be fully polarized in the lab, leading to “internal co-magnetometer” states that are important for rejection of systematic effects. Species such as ThO [183] and HfF⁺ [184], which are used in the most sensitive experiments to search for the electron EDM, and other species with experimentally useful Ω -doubled states such as TaN [186] or WC [187], would be extremely challenging to laser cool using current techniques.

However, polyatomic molecules can offer both photon cycling and fully-polarizable states through their unique mechanical modes [185]. The electronic structure that enables certain diatomic molecules to cycle photons is largely independent of the bonding partner, provided that it has similar valence and ionic nature [162, 163, 170, 188]. For example, SrOH has similar properties to isoelectronic SrF and was recently laser cooled [171]. Since sensitivity to CP-violating physics arises from electronic structure at the metal center, molecules like YbOH [185, 189–191] and RaOH [188] have sensitivity comparable to their fluoride analogues, but with significant experimental advantages. Molecules with at least three atoms have nearly-degenerate mechanical modes of opposite parity, such as linear bending modes or symmetric top rotations about the symmetry axis [185]. Molecules of the type MOH, MCCH, MOCH₃, and others, where M is a suitable metal such as Yb or Ra, are therefore promising candidates for photon cycling *with* a robust mechanism for rejection of systematic effects. This could enable efficient state preparation/readout along with the possibility of laser cooling and trapping to achieve long coherence times and perform extremely sensitive searches for CP-violation. For these types of searches, this ability to simultaneously have laser cooling and robust systematic error rejection through parity doublets is unique to polyatomic molecules.

In some sense, the metal atom in these molecules is providing the photon cycling functionality as well as sensitivity to new physics, while the bonding partner is providing the polarization. A question then arises: can we attach *multiple* metals with interesting properties to a molecule to realize them simultaneously [185]? For example, a molecule like YbCCCa could provide enhanced scattering rates and advanced co-magnetometry, or a molecule like TaCOCa could be used to laser cool and trap a Ta-containing molecule via photon cycling on the Ca center, thereby enabling access to the advantages of the deformed Ta nucleus for precision measurements [186].

In the limit where the two metal centers are infinitely far apart, they will truly be independent and their unique properties can be accessed individually. However, smaller molecules are more advantageous for practical applications. In this work, we therefore consider the molecules YbCCCa [185] and YbCCAl to explore whether the “smallest possible” molecules where the metals do not bond to the same atom can be thought of as two more-or-less independent centers.

Hypermetallic oxides of the form MOM have been studied both experimentally and theoretically [192–194], and recently the mixed hypermetallic BaOCa⁺ was created

and studied in an ion trap [194]. We consider molecules with additional separation between the metals to provide more flexibility in choosing metals and bonding partners. Additionally, the added distance between the centers should reduce their couplings to each other.

YbCCH [195], CaCCH [196], and AlCCH [197] have all been studied spectroscopically, and are linear with low-lying electronic excitations centered on the metal. The species were created via gas-phase chemical reactions of the ablated metal with a reactive gas such as acetylene (HCCH), which suggests a production mechanism for the molecules discussed here. Yb ($Z = 70$) is sensitive to a range of leptonic and hadronic CP-violating physics while Ca ($Z = 20$) and Al ($Z = 13$) are not, yet they tend to create bonds with higher Frank-Condon factors (FCF). For example, the 0-0 FCF for the $A \rightsquigarrow X$ transition is $\approx 99\%$ [198] in CaF and $> 99.9\%$ in AlF [161, 199], compared to $\sim 93\%$ for YbF [200]. Thus, the hope for these molecules is that the Ca and Al centers will provide better laser cooling than YbOH or YbCCH, but still with similar mass and while maintaining nearly-degenerate states of opposite parity. As we shall discuss, the Ca and Al metal centers are distinct due to the different sets of possible spin configurations that they permit in the molecules that we consider.

The primary goal of this work is to study the validity of the simple expectation of multiple, quasi-independent cycling centers on these small molecules. We find that it is indeed the case that the two metal centers can be considered as reasonably independent, and can cycle photons. We also find that the hybridization and spin structure of these exotic molecules plays a critical role in their utility for laser cooling. Our work highlights the potential utility of this hypermetallic approach, and illuminates possibilities for future theoretical and experimental investigations with different metals and molecular bonds which could explore molecules with heavy metal centers that cannot cycle photons.

6.3 Electronic Structure

A clear trend can be seen in previous works [170, 201] on molecules such as CaOH [162], CaNC [162], SrOH [171], and YbOH [185] that metals with alkaline earth-like valence electronic configurations tend to form bonds and hybridized orbitals which are beneficial for creating highly diagonal FC matrices. Thus, a molecule such as YbCCCa, in which both metal centers have an alkaline earth-like valence structure, is a natural place to begin the investigation of small molecules

with more than one optically active metal.

YbCCCa

Using various methodologies from computational quantum chemistry (described in detail in Section 6.5), we find that the linear geometry of YbCCCa is lower in energy than various bent and trigonal configurations in both the ground and excited states of interest, which is supported by spectroscopy on similar molecules [195, 196]. Additionally, all the excited states of interest lie below the ionization energy. The molecule is open-shell and has the desired bonding pattern, which causes the ground state to have an unpaired $4s\sigma$ electron on the Ca and an unpaired $6s\sigma$ electron on the Yb as the highest occupied molecular orbital (HOMO) and HOMO-1, respectively (see Fig. 6.1a-b). The ground state spin structure of these two electrons is characterized by close competition between singlet $X^1\Sigma$ and triplet $X^3\Sigma$ states. This appears to be a significant piece of evidence that the cycling electrons on the two metal centers are highly independent, because in the limit where they are truly independent we expect the singlet and triplet states to be exactly degenerate. The computed splitting between these two states in YbCCCa is approximately 10^{-3} eV, with the singlet lying lower than the triplet. This value is considered quite small to resolve with high certainty using standard quantum chemistry methods, but even if one assumes a large error of $\pm 50\%$ it is still easily resolved experimentally as it corresponds to a frequency on the order of ~ 100 GHz. Fortunately, neither the precise size of this splitting nor the ordering of the states are critical to our conclusions, provided that the gap between the states is larger than typical radiative widths of ~ 10 MHz.

The structure of the lowest lying excited states primarily consists of $4s\sigma \rightarrow 4p\pi$ transitions on the Ca atom (which we will call the Ca *A* state), $6s\sigma \rightarrow 6p\pi$ on the Yb atom, (which we will call the Yb *A* state), and $4s\sigma \rightarrow 3d\sigma$ on the Ca atom (which we will call the Ca *B* state), which are the potential laser cooling transitions. We use this nomenclature because these transitions are analogous to the $X^2\Sigma \rightarrow A^2\Pi$ and $X^2\Sigma \rightarrow B^2\Sigma$ transitions in the single-center molecules [170]. In both the $X\Sigma$ ground state and $A\Pi$ excited states, we observe similar molecular orbital hybridization to the molecules studied in Ref. [162], for which those authors gave a detailed discussion about the likely advantages of this structure for diagonal FC matrices. The molecular orbitals of the promoted electrons and their corresponding holes are shown in Figure 6.1 for all of these excited states. These orbitals are constructed as the eigenvectors of the difference of the density matrices for the

State	Yb-C		C≡C	C-Ca		$ \mu $ (Debye)
	L_0 (Å)	E_0 (eV)	L_0 (Å)	L_0 (Å)	E_0 (eV)	
$X^3\Sigma$	2.351	6.5209	1.243	2.290	4.5179	0.9715
Ca $A^1\Pi$	2.352	4.6355	1.243	2.282	2.6336	–
Yb $A^1\Pi$	2.306	4.1175	1.241	2.298	2.1145	–
Ca $B^3\Sigma$	2.347	4.1918	1.242	2.295	2.1888	–

Table 6.1: YbCCCa bond lengths L_0 , bond energies E_0 , and molecular frame permanent dipole moment (μ) for the ground state, along with bond lengths in the excited states of interest.

excited state and the ground state, $\rho_{ES} - \rho_{GS}$, which is dominated by two nonzero eigenvalues ($\approx \{+1, -1\}$) corresponding to the electron and hole, respectively. The equilibrium bond lengths for the ground and excited states of interest, as well as the Yb-C and C-Ca bond energies and the permanent dipole moment of the molecule are given in Table 6.1.

Similar to the X ground state, close competition between singlet and triplet spin configurations is also observed in the A and B states. A full treatment of these excited states including spin-orbit coupling effects reveals that the small energy gap between $A^1\Pi$ and $A^3\Pi$ states induces strong intersystem crossing, causing some of the true $A\Pi$ spin-orbit sub-levels to be linear combinations of both $^1\Pi$ and $^3\Pi$ configurations¹. Such mixing does not occur between the B state sub-levels due to their Σ symmetry. Table 6.2 gives the strength of the mixing for the A states, along with the excitation energies and transition dipole moments (μ_{tr}) for all the states of interest. Although the mixing reported in the Table is derived from the very small computed value of the energy gap between $^1\Pi$ and $^3\Pi$ configurations, our general conclusion of “strong spin-orbit mixing” is robust to large relative errors in the exact value of the computed gap. Further discussion of the spin-orbit splitting and mixing, including this important point, can be found in Appendix 6.8.

This strong mixing between different spin configurations is a complicating factor for any laser cooling scheme that would use $\text{Ca } X \rightarrow A$ as the main transition. Electrons which are pumped into the mixed A sub-levels have a roughly equal chance of decaying back into the singlet or triplet manifold of the ground state.

¹The $m_s=0$ sectors of the singlet and triplet manifolds combine to form highly mixed spin-orbit sub-levels of the $A^1\Pi$ state. The other sub-levels are linear combinations of states in the $m_s = \pm 1$ sectors, meaning they are entirely within the triplet manifold. See Appendix 6.8 for a more detailed discussion.

Since these two manifolds are not exactly degenerate, they each have their own set of vibrational modes which are essentially identical to each other, but are split by ~ 100 GHz. This doubles the number of accessible vibrational states during an $A\Pi \rightsquigarrow X$ decay. To avoid doubling the number of repump lasers, one would need to selectively pump the $X^3\Sigma$ ground state and use the exclusively triplet $X^3\Sigma \rightarrow A^3\Pi$ transition for laser cooling. The number of additional repump lasers needed for this selective pumping of $X^3\Sigma$, on top of those necessary for the $X^3\Sigma$ vibrational states, depends heavily on the precise values of the main FCFs for the $A\Pi \rightsquigarrow X^1\Sigma$ decay. Note, however, that these strong couplings in the A states could be useful for engineering couplings between the centers.

The Ca $X \rightarrow B$ transition provides a simpler laser cooling scheme, as the Σ symmetry of both X and B prevents their spin-orbit sub-levels from mixing. One could therefore drive the $X^1\Sigma \rightarrow B^1\Sigma$ or $X^3\Sigma \rightarrow B^3\Sigma$ transition with reduced risk of leaking into the other spin manifold. Additionally, since the energy gap between the Ca B and A states is relatively small compared to the $B - X$ gap, the radiative decay rate for $B \rightsquigarrow A$ is suppressed by a factor of ~ 2000 compared the decay rate for $B \rightsquigarrow X$.

Note that for both molecules examined in this work, we consider only single excitations. Excitation of one of the metal centers will shift the transitions, likely on the order of the change in spin-orbit splitting between ground and excited states. Given that this shift is considerably larger than the radiative width ², single excitations therefore likely “blockade” a second excitation, making simultaneous excitations difficult to achieve in the laboratory without the addition of even more lasers. On the other hand, these effects likely have interesting applications on their own.

YbCCAI

Given the desirable electronic transition structure and orbital hybridization of YbCCa, but the complicated spin structure, the molecule YbCCAI has also been investigated. Since this molecule only has a doublet spin configuration in both the ground and excited states, the intersystem crossing issues in YbCCa are avoided in YbCCAI. Similar to YbCCa, we find that the energy of the linear geometry of YbCCAI is lower than bent and trigonal structures for both the $X^2\Sigma$ ground and $A^2\Pi$ excited states, which is supported by spectroscopy on similar molecules [195, 197]. Also, the excited states of interest again lie below the ionization energy. As

²Even the energy of interaction between the spin and orbital magnetic moments of the different metal-centered electrons is likely larger than the radiative width.

Transition			Energy	$ \mu_{tr} $	${}^1\Pi/{}^3\Pi$	
			(eV)	(nm)	Admix (%)	
$X {}^1\Sigma$	\rightarrow	Ca $A \Pi$	1.8843	658	4.020	52%/48%
Ca $A \Pi$	\rightsquigarrow	$X {}^3\Sigma$	1.8834	658	4.208	52%/48%
$X {}^3\Sigma$	\rightarrow	Ca $A {}^3\Pi$	1.8824	657	4.031	0%/100%
$X {}^3\Sigma$	\rightarrow	$B {}^3\Sigma$	2.3314	532	4.747	–
$X {}^1\Sigma$	\rightarrow	$B {}^1\Sigma$	2.3291	532	4.718	–
$X {}^1\Sigma$	\rightarrow	Yb $A \Pi$	2.4034	516	4.560	50%/50%
Yb $A \Pi$	\rightsquigarrow	$X {}^3\Sigma$	2.4025	516	4.576	50%/50%
$X {}^3\Sigma$	\rightarrow	Yb $A {}^3\Pi$	2.4021	515	4.558	0%/100%
$X {}^3\Sigma$	–	$X {}^1\Sigma$	0.0009	–	–	–
$B {}^3\Sigma$	–	$B {}^1\Sigma$	0.0032	–	–	–

Table 6.2: Computed excitation energies and transition dipole moments for the lowest-lying excitations of YbCCCa. Here X denotes the ground state and A , B are the excited states. Since the spin-orbit effects are strong in this molecule, some of the $A \Pi$ spin-orbit sub-levels cannot be identified as purely singlet or triplet (they are left without a spin label). The other $A \Pi$ sub-levels are purely triplet states, so they are identified with the proper label. Column 4 details the magnitude of the singlet and triplet components of the mixed $A \Pi$ states. Detailed discussion of spin-orbit mixing is provided in Appendix 6.8.

State	Yb–C		C \equiv C	C–Al		$ \mu $ (Debye)
	L_0 (Å)	E_0 (eV)	L_0 (Å)	L_0 (Å)	E_0 (eV)	
$X {}^2\Sigma$	2.448	4.4757	1.228	1.935	5.8613	4.3541
Al $A {}^2\Pi$	2.455	2.0163	1.218	1.895	3.4019	–
Yb $A {}^2\Pi$	2.401	2.2828	1.221	1.952	3.6684	–

Table 6.3: YbCCAl bond lengths L_0 , bond energies E_0 , and molecular frame permanent dipole moment (μ) for the ground state, along with bond lengths in the excited states of interest.

before, the bond lengths, bond energies and permanent dipole moment are given in Table 6.3, the excited state energies and transition dipole moments are given in Table 6.4, and the electron/hole orbitals for the ground state and low lying excited states are given in Figure 6.2.

We find that YbCCAl is also open-shell and has the desired bonding pattern, which causes the ground state to have a single, unpaired, non-bonding $6s\sigma$ electron on the Yb as the HOMO and two non-bonding $3s\sigma$ electrons on the Al as the dominant feature of the HOMO-1 (see Fig. 6.2a-b). The $X \rightarrow A$ excited state on the Yb atom

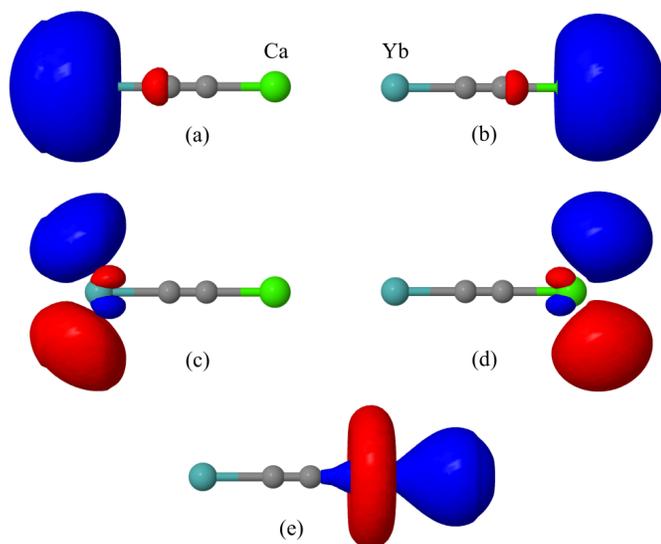


Figure 6.1: Electron and hole orbitals for the ground state and low lying excited states of YbCCCa, computed as the eigenvectors of the difference of density matrices $\rho_{ES} - \rho_{GS}$. *Top:* The hole orbitals, each of which is occupied in the ground state and unoccupied in the Π excited state that is depicted directly below it. (a) shows the unpaired $6s\sigma$ HOMO-1 on the Yb atom, while (b) shows the unpaired $4s\sigma$ HOMO on the Ca atom. *Middle:* The lowest-lying excited state ($A\Pi$) for the Yb Σ electron (c), and the Ca Σ electron (d). *Bottom:* The Ca $4s\sigma \rightarrow 3d\sigma$ B excited state (e).

Transition	Energy		$ \mu_{tr} $ (Debye)
	(eV)	(nm)	
$X \ ^2\Sigma \rightarrow \text{Al } A \ ^2\Pi$	2.4594	504	0.0735
$X \ ^2\Sigma \rightarrow \text{Yb } A \ ^2\Pi$	2.1929	565	6.345

Table 6.4: Computed excitation energies and transition dipole moments for the lowest-lying excitations of YbCCAl. Here X denotes the ground state and A denotes an excited state.

(Fig. 6.2c) is highly similar to the $X \rightarrow A$ excited state in YbCCCa (Fig. 6.1c), as expected.

However, the $X \rightarrow A$ excitation on the Al atom (Fig. 6.2d) has noticeably worse features than the corresponding Ca excitation in YbCCCa (Fig. 6.1d). The lack of $s-p$ hybridization of the Al orbitals leaves a significant portion of the electron density in the Al A state within the C-Al bonding region. This leads to a relatively large

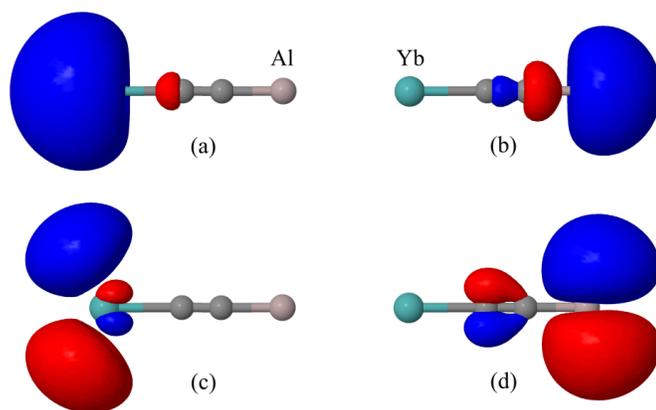


Figure 6.2: Electron and hole orbitals for the ground state and low lying excited states of YbCCAl, computed as the eigenvectors of the difference of density matrices $\rho_{ES} - \rho_{GS}$. *Top:* The hole orbitals, each of which is occupied in the ground state and unoccupied in the Π excited state that is depicted directly below it. (a) shows the unpaired $6s\sigma$ HOMO-1 on the Yb atom, while (b) shows the $3s\sigma$ HOMO on the Al atom. *Bottom:* The lowest-lying excited state ($A\Pi$) for the Yb Σ electron (c), and the Al Σ electron (d).

change in the C-Al bond length (Table 6.3) when compared to the corresponding C-Ca bond length in the Ca A state of YbCCCa (Table 6.1). Further, the Al A state is not completely localized on the Al atom, as there is some additional density in the π -bonding region of the $C\equiv C$ bond. This induces a shortening of the $C\equiv C$ bond, similar to what was observed experimentally in AlCCH [197], which we do not observe in the excited states of YbCCCa. These features of the Al A state are not optimal for photon cycling, as geometry changes tend to reduce the diagonality of the FC matrix.

As a final note, we point out that for both YbCCCa and YbCCAl our calculations reveal that there are no Δ excited states on the metal atoms which are intermediate in energy between the Σ and Π states we have reported. This result is expected for the Al atom, but the level structures of atomic Yb^+ and Ca^+ would suggest intermediate Δ states. This “reordering” of the Π and Δ levels has been understood for alkaline earth monohalides using a ligand field model [202, 203], as well as for CaOH using electronic structure techniques [204]. It has also been observed experimentally [205, 206]. We expect that the qualitative nature of these results

hold for our slightly more complicated YbCCM species. However, there may still be perturbative correlated effects between the low-lying Π states and the displaced Δ states which have physical implications that are not captured in this study [207].

6.4 Vibrational Structure

FC matrix elements have been computed for both YbCCCa and YbCCAl, using methods described in Section 6.5. There we also discuss that the typical level of error in these calculations is no less than a few percent, and thus unequivocal assessment of the cycling properties of these molecules always requires experimental measurements. Nonetheless, the numbers presented in this section serve as a useful guide for future experimental and theoretical investigation of these exotic hypermetallic small molecules.

Due to their linear geometry, YbCCCa and YbCCAl both have five vibrational modes, two of which are doubly-degenerate, for a total of 7 modes. The assignments and energies of these modes are given for the triplet ground states of both molecules in Table 6.5. ν_1 , ν_2 , and ν_3 label the number of vibrational quanta populating the C–M, Yb–C, and C \equiv C stretching modes of the molecule YbCCM, respectively (where M is either Ca or Al). These modes are non-degenerate and have the same σ symmetry as the vibronic ground state. In reality the physical modes are superpositions of these possible modes, though these simple descriptions are reasonably accurate due to the mass differences between the constituent atoms [185, 208]. ν_4 and ν_5 label the population of the doubly-degenerate “anti-symmetric” and “symmetric” bending modes of the C atoms about the symmetry axis, respectively³. Linear combinations of these π -symmetric degenerate pairs can be formed which correspond to states with definite angular momentum $\ell_k = \nu_k, \nu_k - 2, \nu_k - 4, \dots, 1/0$ about the symmetry axis [185, 209]. This quantum number ℓ characterizes the different sub-levels which can occur for the degenerate modes.

We will denote the vibrational wavefunction of the ground electronic state $\chi_0(S)$, where $S = \{\nu_1\nu_2\nu_3\nu_4\nu_5\}$ is some vibrational state. Similarly, $\chi_1(0)$ will denote the excited electronic state. We only consider decays from the ground vibrational state in the excited electronic state since molecules will be excited to this state selectively during the laser cooling process, though other excitations are relevant for

³ Since the YbCCM molecules have no inversion symmetry, the anti-symmetric and symmetric nomenclature to distinguish the bending modes technically does not make sense. However, these names are well-defined for 4 atom molecules with inversion symmetry, such as HCCH, and there is a very clear correspondence between the normal modes of YbCCM and those of HCCH, so we use this nomenclature anyways for the sake of clarity.

Mode	Assignment	Frequency (cm ⁻¹)	
		YbCCCa	YbCCAl
C–M Stretch	$\nu_1(\sigma)$	484.92	612.78
Yb–C Stretch	$\nu_2(\sigma)$	185.57	216.77
C≡C Stretch	$\nu_3(\sigma)$	2006.15	2047.00
Asymm. Bend	$\nu_4(\pi)$	43.15	58.34
Symm. Bend	$\nu_5(\pi)$	140.89	197.60

Table 6.5: Vibration frequencies for each of the vibrational modes of the triplet ground state of YbCCM, where M is either Ca or Al. σ and π represent the symmetry of the vibration, and whether it is non-degenerate or doubly-degenerate, respectively.

laser cooling schemes since repumping is invariably required. We seek to compute the FCFs $|\langle \chi_0(S) | \chi_1(0) \rangle|^2$ for the relevant $|\chi_1\rangle$ states discussed earlier in order to understand if photon cycling on the two metal atoms is possible with only a small number of repumping lasers ⁴.

In general, the restrictions on S are simply that $|\chi_0(S)\rangle$ must have a total symmetry of σ , because the $|\chi_1(0)\rangle$ state is σ -symmetric and therefore any other spontaneous decays $|\chi_1(0)\rangle \rightsquigarrow |\chi_0(S)\rangle$ are forbidden. More specifically, this means that the stretching modes ν_{1-3} can be arbitrarily populated, but the bending modes ν_{4-5} are subject to selection rules during radiative decay of $|\chi_1(0)\rangle$ [209],

$$\begin{aligned} \Delta \ell_k &= 0, \quad k \in \{4, 5\} \\ \sum_k \Delta \nu_k &= 0, \pm 2, \pm 4, \pm 6, \dots \end{aligned} \quad (6.1)$$

Finally, for any single excited state $|\chi_1(0)\rangle$ under consideration in this work we have the useful property,

$$\sum_S |\langle \chi_0(S) | \chi_1(0) \rangle|^2 = 1, \quad (6.2)$$

which provides a normalized scale with which to assess the branching ratios.

YbCCCa

The computed FCFs for the $X^3\Sigma \rightarrow \text{Ca } A^3\Pi$, $X^3\Sigma \rightarrow \text{Yb } A^3\Pi$, and $X^3\Sigma \rightarrow \text{Ca } B^3\Sigma$ transitions are shown in Table 6.6. The FC matrix for the Ca $X \rightarrow A$ transition was calculated to be highly diagonal, with a 0-0 FCF of 0.99 and only 3 other transitions with an FCF greater than $4 \cdot 10^{-4}$. However, quantitative estimates of the systematic

⁴Note that the original literature calls $\langle \chi_0(S) | \chi_1(0) \rangle$ the Franck-Condon Factor, but that definition is less useful and increasingly uncommon.

$X^3\Sigma \rightarrow \text{Ca } A^3\Pi$	FCF	Sum
$\chi_0(\{00000\}) \rightarrow \chi_1(0)$	0.99	0.99
$\chi_0(\{01000\}) \rightarrow \chi_1(0)$	0.005	0.994
$\chi_0(\{10000\}) \rightarrow \chi_1(0)$	0.003	0.997
$\chi_0(\{0002^0_0\}) \rightarrow \chi_1(0)$	0.002	0.999
$\chi_0(\{00011\}) \rightarrow \chi_1(0)$	0.0004	0.9994
$\chi_0(\{20000\}) \rightarrow \chi_1(0)$	0.0003	0.9997
$\chi_0(\{11000\}) \rightarrow \chi_1(0)$	0.0002	0.9999
Sum:	0.9999	
$X^3\Sigma \rightarrow \text{Yb } A^3\Pi$	FCF	Sum
$\chi_0(\{00000\}) \rightarrow \chi_1(0)$	0.75	0.75
$\chi_0(\{01000\}) \rightarrow \chi_1(0)$	0.16	0.91
$\chi_0(\{10000\}) \rightarrow \chi_1(0)$	0.06	0.97
$\chi_0(\{02000\}) \rightarrow \chi_1(0)$	0.01	0.98
$\chi_0(\{11000\}) \rightarrow \chi_1(0)$	0.01	0.99
$\chi_0(\{20000\}) \rightarrow \chi_1(0)$	0.004	0.995
$\chi_0(\{03000\}) \rightarrow \chi_1(0)$	0.0008	0.996
$\chi_0(\{00011\}) \rightarrow \chi_1(0)$	0.0003	0.996
Sum:	0.996	
$X^3\Sigma \rightarrow \text{Ca } B^3\Sigma$	FCF	Sum
$\chi_0(\{00000\}) \rightarrow \chi_1(0)$	0.995	0.995
$\chi_0(\{10000\}) \rightarrow \chi_1(0)$	0.003	0.998
$\chi_0(\{00002^0_0\}) \rightarrow \chi_1(0)$	0.001	0.999
$\chi_0(\{0002^0_0\}) \rightarrow \chi_1(0)$	0.0004	0.9994
$\chi_0(\{20000\}) \rightarrow \chi_1(0)$	0.0003	0.9997
$\chi_0(\{11000\}) \rightarrow \chi_1(0)$	0.0002	0.9999
Sum:	0.9999	

Table 6.6: Franck-Condon factors for the metal-centered electronic transitions of YbCCCa. Multiply populated vibrational modes ν which are degenerate bending modes are additionally labeled by their symmetry-projected angular momentum quantum number ν^ℓ . Effects of systematic errors in the calculations are not included in these numbers (see Sections 6.4 & 6.5 for estimates of these effects).

error in these calculations suggest a reduced level of diagonality. Accounting for our “worst case” error estimates for this transition, the 0-0 FCF is reduced to 0.9, but only 3 additional states are required to reach a total efficiency of 0.997, which is comparable with the “diagonal” results in Table 6.6. However, 8 total states (1 main transition + 7 repump states) are required for an efficiency of 0.999. which is significantly worse than the results in Table 6.6. The details of these error estimates

are discussed in Section 6.5.

Without considering the estimates of systematic error or the complications that arise due to intersystem crossing, the results in Table 6.6 suggest that the Ca atom can scatter thousands of photons with only 3-4 lasers (1 main transition and 2-3 repumps) and tens of thousands of photons with 7 lasers. When considering the worst case systematic error, these numbers are increased to 4-8 lasers just for the ability to scatter thousands of photons. Further, the consideration of intersystem crossing effects requires the addition of even more lasers in order to have a “closed” laser cooling cycle that accounts for the non-degenerate vibrational manifolds of the $X^1\Sigma$ and $X^3\Sigma$ states (as discussed in Section 6.3).

The Ca $X \rightarrow B$ displays significantly nicer properties. The FC matrix was computed to be even more diagonal than the $X \rightarrow A$ transition, with a 0-0 FCF of 0.995 and only 2 other transitions with an FCF greater than $4 \cdot 10^{-4}$. Additionally, the effects of the “worst case” systematic error estimates for this transition are smaller: the 0-0 FCF is reduced to 0.993, 4 total states are required for an efficiency of 0.9991, and 7 total states are required for an efficiency of 0.9999. Further, intersystem crossing effects are suppressed. This allows the Ca center to scatter thousands of photons with only 3-4 lasers, and tens of thousands of photons with 6-7 lasers.

The Yb-centered transition is less diagonal, showing more expansive branching than YbOH [185]. The main 0-0 FCF is only 0.75 and there are 5 other FCFs with values larger than 10^{-3} . This limits the scattering efficiency of the Yb atom to ~ 500 photons with a reasonable number of repump lasers, without even considering systematic errors or the additional lasers that are necessary due to intersystem crossing. This decreased efficiency of the Yb atom compared to YbOH is likely due to the fact that the Yb–C bond in YbCCa is significantly longer and “floppier” than the Yb–O bond in YbOH. This allows for a more significant off-diagonal vibrational decay channel through the Yb–C stretch, as we can see in Table 6.6. Note that for each metal center M, the two most dominant off-diagonal decays for the A states are the M–C stretch and the C–M’ stretch. The former is not surprising, but the latter may seem unusual since the metal centers are rather far apart. However, the descriptions of the mode assignments shown in Table 6.5 are only an approximation, and the true physical normal modes are admixtures of the idealized stretching modes described in the first column.

$X^2\Sigma \rightarrow \text{Al } A^2\Pi$	FCF	Sum
$\chi_0(\{00000\}) \rightarrow \chi_1(0)$	0.74	0.74
$\chi_0(\{10000\}) \rightarrow \chi_1(0)$	0.14	0.88
$\chi_0(\{01000\}) \rightarrow \chi_1(0)$	0.06	0.95
$\chi_0(\{11000\}) \rightarrow \chi_1(0)$	0.01	0.96
$\chi_0(\{00002^0\}) \rightarrow \chi_1(0)$	0.01	0.97
$\chi_0(\{20000\}) \rightarrow \chi_1(0)$	0.01	0.98
$\chi_0(\{00011\}) \rightarrow \chi_1(0)$	0.009	0.984
$\chi_0(\{00100\}) \rightarrow \chi_1(0)$	0.005	0.989
$\chi_0(\{02000\}) \rightarrow \chi_1(0)$	0.002	0.991
Sum:	0.991	
$X^2\Sigma \rightarrow \text{Yb } A^2\Pi$	FCF	Sum
$\chi_0(\{00000\}) \rightarrow \chi_1(0)$	0.74	0.74
$\chi_0(\{01000\}) \rightarrow \chi_1(0)$	0.18	0.91
$\chi_0(\{10000\}) \rightarrow \chi_1(0)$	0.05	0.96
$\chi_0(\{02000\}) \rightarrow \chi_1(0)$	0.02	0.98
$\chi_0(\{11000\}) \rightarrow \chi_1(0)$	0.009	0.988
$\chi_0(\{00100\}) \rightarrow \chi_1(0)$	0.005	0.993
$\chi_0(\{20000\}) \rightarrow \chi_1(0)$	0.002	0.995
$\chi_0(\{10100\}) \rightarrow \chi_1(0)$	0.001	0.996
Sum:	0.996	

Table 6.7: Franck-Condon factors for the metal-centered electronic transitions of YbCCAl. Multiply populated vibrational modes ν which are degenerate bending modes are additionally labeled by their symmetry-projected angular momentum quantum number ν^ℓ .

YbCCAl

The FCFs for the $X^2\Sigma \rightarrow \text{Al } A^2\Pi$ and $X^2\Sigma \rightarrow \text{Yb } A^2\Pi$ transitions are shown in Table 6.7. The Yb $X \rightarrow A$ excitation shows similar branching ratios to the Yb-centered excitation in YbCCCa, although the populated modes differ slightly. Despite the same branching ratios, the YbCCAl excitation will have a higher optical efficiency in practice because it avoids the intersystem crossing of the YbCCCa A state. On the other hand, the Al-centered excitation is significantly less diagonal than the Ca $X \rightarrow A$ transition in YbCCCa. The large number of significant FCFs cause its optical efficiency to be too low for successful laser cooling. This result is expected based on the non-ideal electronic density in the Al $A^2\Pi$ excited state (Fig. 6.2d) and its relatively large impact on the geometry of the molecule, as discussed in Section 6.3.

6.5 Computational Details

The molecules YbCCCa and YbCCAl were predominantly investigated with the complete active space self-consistent field (CASSCF) [210, 211] and multireference configuration interaction (MRCI) [212, 213] methods from the MOLPRO quantum chemistry package [214]. Since we were interested in the nature of the low-lying excited states, all calculations involving excited states were performed using state averaging in CASSCF. These methods were chosen in order to address the electronic levels (including spin-orbit effects) as accurately as possible, despite the inherent difficulty of using them to compute more challenging quantities such as the FC matrix elements (due to very high computational cost).

Despite their typical levels of accuracy, using active space-based methods with heavy atoms such as Yb poses a challenge. As an example, consider the ideal chemical active space for YbCCCa: it should likely include all doubly occupied valence f orbitals on the Yb, 4 doubly occupied bonding orbitals, 2 singly occupied valence $s\sigma$ orbitals, 4 virtual $p\pi$ orbitals, and valence d orbitals on both the Yb and Ca. This active space (24 electrons, 28 orbitals) is far too large for MRCI, and even if the virtual d orbitals are removed, the (24e, 18o) reduced active space is still too large for MRCI.

To test whether this challenge would prohibit us from using the MRCI methodology (in favor of a more approximate, less computationally expensive method), we examined the importance of including the occupied f orbitals and virtual d orbitals in the active space by running MRCI on further reduced “test” active spaces that were constructed by exploiting orbital symmetries. Specifically, some essential components of the test active spaces always remained the same: the 4 bonding orbitals, 2 singly occupied valence $s\sigma$ orbitals (a-b in Figs. 6.1 & 6.2), the 4 virtual $p\pi$ orbitals (c-d in Figs. 6.1 & 6.2), and the Ca $3d\sigma$ orbital (Fig. 6.1e). However, added on top of these were permutations of the occupied f orbitals in 0, 1, or 2 symmetry sectors and the valence d orbitals in 0, 1, or 2 additional symmetry sectors (not necessarily the same as the f sectors). This allowed us to perform well-defined MRCI tests (i.e. the exact orbitals in the active space were unambiguously known) on active spaces of more reasonable sizes such as (14e, 14o) - (16e, 16o). For example, a prototypical (14e, 14o) test space included 2 occupied f orbitals and one virtual d orbital on each metal, along with the 4 bonding orbitals, 2 singly occupied valence $s\sigma$ orbitals, and the 4 virtual $p\pi$ orbitals.

Across many different permutations of such active spaces, the MRCI tests revealed

that configurations containing occupied valence d orbitals maximally contributed $\sim 0.02\%$ to the multi-determinant ground state and excited A states of interest. The B state in YbCCCa was completely dominated by the occupation of the σ -symmetric Ca $3d$ orbital in Fig. 6.1e, and had equally small contributions from the $3d$ orbitals in other symmetry groups. This was likely due to the relatively large energy gap between the $3d\sigma$ orbital and the $3d\pi, \delta$ orbitals, arising from the significant $s-d$ and p_z-d hybridization seen in Fig. 6.1e (and the lack of any hybridization in the other d orbitals). A completely negligible contribution arose from all configurations in which an electron vacated the doubly occupied f or bonding orbitals. This suggests that, at least in this specific case, molecular Yb has significantly simpler electronic structure than atomic Yb⁺ [215, 216], as mentioned at the end of Section 6.3.

The excited states were examined in these tests by including 12 baseline states in the state averaging procedure: $2^1\Sigma$, $2^3\Sigma$, $4^1\Pi$, and $4^3\Pi$. Additional states were added to this average based on the symmetry of the f and d orbitals included in a given active space permutation. These tests allowed us to search for all the low-lying excited states and revealed no additional allowed states in the energy range of the A and B states of interest. Most importantly, the results allowed us to define much more tractable (yet still physically realistic) active spaces for our MRCI studies of YbCCCa and YbCCAl.

YbCCCa

All electronic orbitals, density matrices, transition dipole moments, excitation energies, and bond energies were obtained using MRCI. The full ANO-RCC basis was used for Ca and C [217, 218] while a contracted ANO basis [219, 220] was used in combination with a 28 electron small core pseudopotential [221] for the treatment of scalar relativistic effects in Yb. An active space of (2e, 7o) was used, including the two singly occupied valence $s\sigma$ electrons (Fig. 6.1a-b), 4 virtual $p\pi$ orbitals (Fig. 6.1c-d), and the one relevant Ca $3d\sigma$ virtual orbital shown in Fig. 6.1e. Given the active space, we used 12 states in the state averaging procedure: $2^1\Sigma$, $2^3\Sigma$, $4^1\Pi$, and $4^3\Pi$. The spin-orbit coupling analysis was done in MOLPRO via the state interaction formalism [222–225] with all 12 of the MRCI wavefunctions from the state averaging.

Optimized geometries, normal modes, and vibrational frequencies for the ground and excited states were obtained at the level of unrestricted Kohn-Sham (UKS) DFT using the B3LYP functional in the Q-Chem quantum chemistry package [226].

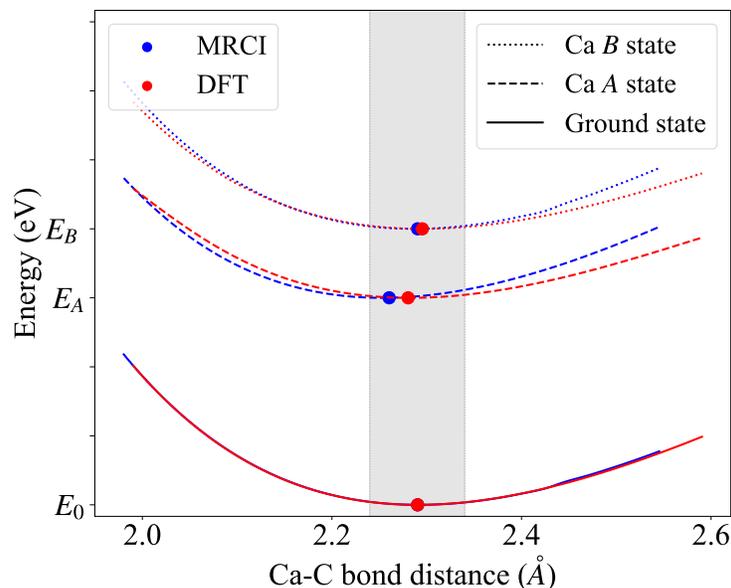


Figure 6.3: Potential energy surfaces obtained with DFT and MRCI along the Ca-C bond coordinate in YbCCCa. The solid lines are the ground state, dashed lines are the Ca A state, and dotted lines are the Ca B state. Red curves are from DFT and blue curves are from MRCI. The dots show the minima of each curve. The curves have been shifted by scalar values in order to make them easier to compare; hence the energy axis does not have numerical tick labels because the gaps between the curves are not to scale. However, for a sense of scale of each curve individually, the tick marks are placed at intervals of 0.3 eV.

Excited states were obtained using TDDFT. The $X^3\Sigma$ ground state was easier to reliably isolate than the $X^1\Sigma$ ground state in the UKS procedure, so we report FC matrix elements for the triplet manifold. Given that the splitting between the singlet and triplet manifolds is on the order of ~ 100 GHz for the X , A , and B states, we do not expect a meaningful difference in the FCFs between the singlet and triplet sub-levels of these states, up to the level of accuracy that can be expected from these calculations. This approach was chosen instead of continuing to use MRCI due to the numerical difficulty and computational cost associated with computing vibrational frequencies with CASSCF+MRCI.

We justify the validity of this use of DFT in Figure 6.3 by examining the DFT and MRCI energy landscapes along the relevant C-Ca bond coordinate. We use the MRCI curves to estimate the level of systematic error in the DFT frequencies and geometry changes for each state, and then assess the effects of these errors on the final FCF values. The curves in Figure 6.3 have been shifted by scalars on both the

x- and y-axes in order to make the comparison of the curves easier. Shifting in the y coordinate is necessary because DFT and MRCI do not predict the same excitation energies. Shifting in the x coordinate is required because DFT and MRCI do not predict the exact same Ca-C bond length in the ground state; however this alone has no effect on the values of the FCFs because they depend on *changes* in geometry and vibrational frequencies between ground and excited states (and these quantities are preserved by scalar shifts).

The ground state curves are virtually identical in the window of importance (highlighted in gray). Their second derivatives (which are directly related to vibrational frequencies) at the equilibrium geometry differ by less than 0.5% between MRCI and DFT. In the *A* state, the $X \rightarrow A$ bond length change from the MRCI curve is 0.02 Å greater than is predicted by DFT and the second derivative from MRCI is $\sim 15\%$ larger than DFT. The effects of “worst case” systematic errors, in which we assume all elements of the DFT Hessian are underestimated by $\sim 15\%$ (along with the 0.02 Å Ca-C geometry change error), on the final FCF values for the Ca $X \rightarrow A$ transition are discussed in Section 6.4. In the *B* state, the $X \rightarrow B$ bond length change from the MRCI curve is only ~ 0.005 Å greater than is predicted by DFT, and the second derivative from MRCI is $\sim 8\%$ larger than DFT. Similarly, the effects of the “worst case” of these errors on the FCFs are discussed in Section 6.4.

YbCCAl

All electronic orbitals, density matrices, transition dipole moments, excitation energies, and bond energies were obtained with MRCI, while the optimized geometries, normal modes, and their frequencies were obtained with CASSCF (all in MOLPRO). The def2-TZVPP basis was used for all atoms [227–229], which included the same 28 electron small core pseudopotential as above for scalar relativistic effects. Given that the Al-centered excitation contains some $C\equiv C$ π -bonding density (see Fig. 6.2d), the $C\equiv C$ π bonding orbitals were included in the active space for these calculations, along with the valence $s\sigma$ orbitals (Fig. 6.2a-b) and the 4 virtual $p\pi$ orbitals (Fig. 6.2c-d) for a (7e,8o) active space.

For this molecule, we included 5 states in the state averaging procedure: 1 $^2\Sigma$ and 4 $^2\Pi$. For the normal mode frequency calculations, we noticed less agreement between DFT and MRCI than we saw in Fig. 6.3 for YbCCCa. This may have been due to the larger magnitude of geometry changes in the YbCCAl excited states (Section 6.3). Thus, we instead chose to use CASSCF to compute the normal modes

and their frequencies.

For both molecules, the FCFs were computed from the ab initio molecular data using the ezSpectrum software [230]. Multiple approximations enter into these calculations. Firstly, all the ab initio molecular data was obtained within the Born-Oppenheimer approximation, so certain effects such as Renner-Teller are not accounted for. Additionally, the FCFs are computed by assuming that the potential energy surfaces in the immediate vicinity of the equilibrium geometries for both the ground and excited state can be approximated by a harmonic potential. Based on Fig. 6.3, this does not appear to be a strong assumption in our case. Finally, the FCFs were computed analytically, including hot bands and Duchinsky rotations [231–233], so no assumptions were made about the normal modes of the ground and excited state being parallel.

As a concluding remark to this section, it is worth mentioning explicitly that no quantum chemical calculations for molecules this heavy have sufficient accuracy to serve as a replacement for spectroscopic measurements. The first experimental step toward using such molecules is measurement of the energy levels and branching ratios via broadband optical spectroscopy.

6.6 Discussion

These results suggest that both YbCCCa and YbCCAl do indeed show many of the desired properties for polar molecules with multiple cycling centers. Both the ground and excited states are bound and linear, the valence s electrons of interest for laser cooling are localized on their respective metal centers, and the Franck-Condon factors are reasonably diagonal (especially for the Ca center). However, YbCCAl reveals that metals without an alkaline earth-like valence structure may not have strong enough $s - p$ hybridization of the excited state orbitals to remove the electron density from the bonding region, though this is worth investigating further with other species, particularly other group IIIA elements. It seems likely that this generic feature will often cause significant geometrical changes in the excited state of polyatomic molecules, which reduces the diagonal nature of the FC matrix and allows branching into a significant number of vibrational modes. These effects have also been discussed elsewhere in the context of molecules with a single metal center [170].

In YbCCCa, both metal atoms have alkaline earth-like valence structures and the $s - p$ hybridization of the excited state orbitals significantly improves FCFs for the

Ca center over the Al center. The two singly-bonded alkaline earth-like atoms create a diradical with two singly occupied s orbitals, one on each metal. Since we seek a molecule which has two metal centers that are as independent as possible, this electronic structure creates a feature of critical importance: the more independent the two open-shell electrons are, the smaller the energy gap between singlet and triplet configurations of the molecule in both ground and excited states. This quasi-degeneracy of singlet and triplet states gives rise to complexities in cycling with the $X \rightarrow A$ transition. The singlet-triplet gap will likely never be small enough that it is not resolved by a laser ⁵, but also not large enough that strong spin-orbit mixing of the $A^1\Pi$ and $A^3\Pi$ excited states can be avoided (even for light molecules, as is seen for Ca in Table 6.2). This could effectively double the number of vibronic states with significant FCFs, though there may be routes to avoid leakage between the singlet and triplet manifolds as discussed in Section 6.3.

Regardless, the Ca B state avoids the challenges caused by spin-orbit coupling because both the X and B states are Σ -symmetric. This allows the highly diagonal, spin-pure $X^1\Sigma \rightarrow B^1\Sigma$ or $X^3\Sigma \rightarrow B^3\Sigma$ transitions on the Ca center to be considered as potentially feasible laser cooling transitions. These transitions may only require $\sim 4 - 5$ repump lasers in order to cycle tens of thousands of photons, and the Ca $B \rightsquigarrow A$ decay has a radiative rate which is suppressed by a factor of ~ 2000 compared to the desired Ca $B \rightsquigarrow X$ decay, due to the difference in their respective energies. No such B state was investigated in detail on the Yb center due to the lack of existing experimental data on B states in Yb-containing molecules. Additionally, configurations with holes in Yb f orbitals lie energetically below any potential Yb B state, which makes an accurate study of its properties significantly more challenging.

Our results also suggest a compelling alternative approach for precision measurements utilizing hypermetallic MCCM' molecules. For metals that do not make cycling centers in molecules, such as Th, Ta, U, Pa, *etc.*, such a molecular scheme should make it possible to cycle photons, apply optical forces, and potentially implement laser cooling to perform precision measurement on these species, while maintaining the ability to realize full polarization and internal co-magnetometer states for robust systematic error rejection [185]. Given that most atoms do not make laser-coolable molecules, yet may contain interesting properties such as highly deformed nuclei with extreme sensitivity to nuclear CP-violation [234], our approach could make available all of the tools of ultracold molecular precision measurement

⁵Due to the relative magnitude of various Breit-Pauli terms such as the spin-spin interaction compared to the spectral width of modern lasers.

to these exotic species. Utilizing the diagonal transitions of a Ca (or analogous) metal center in a molecule containing Th, Ta, Pa, etc. may offer *extremely favorable* coherence times compared to molecules such as ThO [183, 235, 236] or TaN [237], or polyatomic analogues such as ThOH⁺ [238] or TaCH [239].

The additional benefits of this molecular design for precision measurements may be numerous, especially for CP violation searches. First, the cycling center could be used for enhanced state detection to yield improved statistics. By using state-dependent optical pumping (or coherent transfer) of spin states to internal states, for example excited vibrational states, the cycling center can be used for efficient detection of the initial spin state. Second, the cycling center offers an additional co-magnetometer that can be used to diagnose stray fields and other systematic effects. Third, the requirements are more relaxed for a molecule in which the non-cycling center is the focus of the precision measurement, compared to the dual-cycling molecules examined in this work. A slight perturbation to the FCFs of an optical cycling precision measurement atom (such as Yb) can destroy experimental efficiency through loss of photon cycling, but a slight perturbation to a CP-violation sensitivity parameter of a non-cycling measurement atom (such as Th) will still result in a promising molecule. Such a molecule would also avoid the “excitation blockade” discussed earlier since simultaneous excitation would be undesirable in the first place. Lastly, the polyatomic structure allows us to use diamagnetic species with sensitivity to nuclear CP violation and good robustness against magnetic effects, such as the ¹Σ states of divalent Th or Ra [185, 238] or monovalent Tl [185, 240] while still maintaining strong systematic error rejection and providing optical readout schemes. All of these areas are worth considering in further theoretical studies.

6.7 Summary

In summary, we have explored the vibronic structure of two prototypical hypermetallic small molecules for precision measurement experiments, YbCCCa and YbCCAl. Despite the small size of the molecules, the electronic properties of each of the metal centers remain quite independent. This allows for photon cycling that is localized on each metal, exploiting their different advantages. Although these two molecules do not possess all of the desired properties for precision measurement experiments that specifically utilize a Yb atom, they suggest a more general class of promising molecules which contain Ca and a heavy metal that does not make use of photon cycling for precision measurement. This general recipe for hypermetallic small

molecules likely allows for the laser cooling of a wide variety of heavy metal atoms via a Ca center, which is one potential path towards ultra precise next generation experiments.

6.8 Short Appendix

In this section we will discuss in detail the magnitude of spin-orbit coupling effects in YbCCCa. The results in Table 6.2 show some spin-orbit sub-levels of the A states which are heavily mixed between $^3\Pi$ and $^1\Pi$ configurations, along with some sub-levels which are purely triplets. In total, there are 8 spin-orbit eigenstates on each metal atom (2 degenerate π orbitals \times (1 singlet state + 3 triplet states)) which group into pairs that are exactly degenerate, leaving 4 distinct sub-levels of each A state. The mixed sub-levels, of which there are two degenerate pairs per metal atom, arise from linear combinations of the $m_s = 0$ sectors in the singlet and triplet manifolds. No such linear combinations between singlet and triplet manifolds can be made for the $m_s = \pm 1$ triplet sectors (due to symmetry), so there are two additional degenerate pairs which are purely triplet sub-levels of the A state. Since the energy gap between the mixed sub-levels and the pure sub-levels is only ~ 100 MHz, the radiative decay lifetime between the pure sub-levels and the mixed ones is essentially infinite relative to experimental timescales. For this reason, only the lower energy mixed and triplet sub-levels are reported in Table 6.2. We do not include the Ω quantum number in the term symbols since we have two distinct metal-centered electrons that couple to the internuclear axis independently, as opposed to coupling as a single total spin as in a diatomic molecule. This means that Ω is not necessarily meaningful, especially for larger molecules where the spins are coupled even more weakly.

For a detailed understanding of why the $m_s = 0$ sectors mix so strongly, we can perform an explicit analysis on a small part of the spin-orbit matrix $H_{\text{SOC}} = \sum_i \alpha_i \vec{L}_i \cdot \vec{S}_i$, where \vec{L}_i is the orbital angular momentum, \vec{S}_i is the electron spin, and α_i is a spin-orbit constant. We will examine here only the Ca-centered A state, although identical reasoning extends to the mixing of the Yb-centered sub-levels as well. In the spin-pure basis $\{A^1\Pi_{\text{Ca}}, A^3\Pi_{\text{Ca}}\}$ (and in units of cm^{-1}), we computed,

$$H_0 = \begin{bmatrix} -418,852,888 & 0 \\ 0 & -418,852,890 \end{bmatrix} \quad (6.3)$$

$$H_{\text{SOC}} = \begin{bmatrix} 15198.43 & 20.74i \\ -20.74i & 15196.52 \end{bmatrix}. \quad (6.4)$$

Note that the off-diagonal terms are of the same order as the spin-orbit constant in CaOH [241, 242]. Thus the “full” Hamiltonian is given by $H_{\text{tot}} = H_0 + H_{\text{SOC}}$,

$$H_{\text{tot}} = \begin{bmatrix} -418837689.57 & 20.74i \\ -20.74i & -418837693.48 \end{bmatrix}. \quad (6.5)$$

This matrix can be analyzed by a simplified matrix of the form,

$$\tilde{H}_{\text{tot}} = \begin{bmatrix} N + \epsilon & \chi \\ -\chi & N - \epsilon \end{bmatrix}, \quad (6.6)$$

where 2ϵ is the splitting between between $A^1\Pi$ and $A^3\Pi$ due to standard correlation effects, while the mixing χ is due to spin-orbit effects. Note that the very large energies on the diagonals of H_0 and H_{SOC} mostly come from the quantum chemical “background” (i.e. the core electrons), which is computed to very high accuracy. Typically only small components of the valence energy are subject to significant possible errors, so error in the splitting ϵ should be considered as a percentage of ϵ , not a percentage of N .

In the limit $\epsilon \rightarrow 0, |\chi| > 0$, the eigenvectors of \tilde{H}_{tot} approach the fully mixed $[i/\sqrt{2}, \pm 1/\sqrt{2}]^T$, while in the limit $\chi \rightarrow 0, |\epsilon| > 0$, the eigenvectors approach the completely unmixed $[1, 0]^T, [0, 1]^T$. In the case of H_{tot} , we have $\epsilon \sim 2 \text{ cm}^{-1}$ and $\chi \sim 20i \text{ cm}^{-1}$, which gives highly mixed eigenvectors: $[0.673i, -0.740]^T, [0.740i, 0.673]^T$.

Thus, the strong mixing effect emerges because the singlet-triplet energy gap is small compared to the magnitude of the spin-orbit coupling term. However, as we just discussed, ϵ and χ are the parts of the computation which are highly sensitive to the electronic structure and are subject to possible errors based on the computational methodology. Nonetheless, we expect our conclusion of “strong spin-orbit mixing” to be valid because even if we assume our calculated values of ϵ and χ both have very large errors of $\sim 90\%$ each, we still have $|\epsilon/\chi| \approx 1$. Using this ratio to compute the eigenvectors of \tilde{H}_{tot} , we see that they are still quite mixed: $[0.924i, 0.383]^T, [-0.383i, 0.924]^T$.

Furthermore, in the context of potential experiments the spin-orbit mixing is only negligible in the limit when $|\chi/\epsilon| \approx 10^{-4}$. Assuming χ is approximately correct due to its similarity with the CaOH results, this limit could only be reached if the error in ϵ is $10^5 - 10^6\%$, which we deem unlikely based on our computational methodology (discussed in Section 6.5).

BIBLIOGRAPHY

1. Klitzing, K. v., Dorda, G. & Pepper, M. New method for high-accuracy determination of the fine-structure constant based on quantized Hall resistance. *Physical review letters* **45**, 494 (1980).
2. Laughlin, R. B. Anomalous quantum Hall effect: an incompressible quantum fluid with fractionally charged excitations. *Physical Review Letters* **50**, 1395 (1983).
3. Bednorz, J. G. & Müller, K. A. Possible highT_c superconductivity in the Ba-La-Cu-O system. *Zeitschrift für Physik B Condensed Matter* **64**, 189–193 (1986).
4. Wu, M.-K. *et al.* Superconductivity at 93 K in a new mixed-phase Y-Ba-Cu-O compound system at ambient pressure. *Physical review letters* **58**, 908 (1987).
5. Maeda, H., Tanaka, Y., Fukutomi, M. & Asano, T. A new high-T_c oxide superconductor without a rare earth element. *Japanese Journal of Applied Physics* **27**, L209 (1988).
6. Born, M. & Oppenheimer, R. Zur Quantentheorie der Molekeln. *Annalen der Physik* **389**, 457–484 (1927).
7. Hohenberg, P. & Kohn, W. Inhomogeneous electron gas. *Physical review* **136**, B864 (1964).
8. Kohn, W. & Sham, L. J. Self-consistent equations including exchange and correlation effects. *Physical review* **140**, A1133 (1965).
9. Mardirossian, N. & Head-Gordon, M. Thirty years of density functional theory in computational chemistry: an overview and extensive assessment of 200 density functionals. *Molecular Physics* **115**, 2315–2372. doi:10.1080/00268976.2017.1333644 (2017).
10. Slater, J. C. A simplification of the Hartree-Fock method. *Physical review* **81**, 385 (1951).
11. Čížek, J. On the correlation problem in atomic and molecular systems. Calculation of wavefunction components in Ursell-type expansion using quantum-field theoretical methods. *The Journal of Chemical Physics* **45**, 4256–4266 (1966).
12. McClain, J., Sun, Q., Chan, G. K.-L. & Berkelbach, T. C. Gaussian-based coupled-cluster theory for the ground-state and band structure of solids. *Journal of chemical theory and computation* **13**, 1209–1218. doi:10.1021/acs.jctc.7b00049 (2017).

13. Gao, Y. *et al.* Electronic structure of bulk manganese oxide and nickel oxide from coupled cluster theory. *Physical Review B* **101**, 165138. doi:10.1103/PhysRevB.101.165138 (2020).
14. Zhu, T., Cui, Z.-H. & Chan, G. K.-L. Efficient formulation of ab initio quantum embedding in periodic systems: Dynamical mean-field theory. *Journal of Chemical Theory and Computation* **16**, 141–153. doi:10.1021/acs.jctc.9b00934 (2019).
15. Cui, Z.-H., Zhu, T. & Chan, G. K.-L. Efficient implementation of ab initio quantum embedding in periodic systems: Density matrix embedding theory. *Journal of Chemical Theory and Computation* **16**, 119–129. doi:10.1021/acs.jctc.9b00933 (2019).
16. De Boer, J. H. & Verwey, E. J. Semi-conductors with partially and with completely filled 3d-lattice bands. *Proceedings of the Physical Society (1926-1948)* **49**, 59 (1937).
17. Mott, N. & Peierls, R. Discussion of the paper by de Boer and Verwey. *Proceedings of the Physical Society (1926-1948)* **49**, 72 (1937).
18. Hubbard, J. Electron correlations in narrow energy bands. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences* **276**, 238–257 (1963).
19. Yang, J. *et al.* Ab initio determination of the crystalline benzene lattice energy to sub-kilojoule/mole accuracy. *Science* **345**, 640–643. doi:10.1126/science.1254419 (2014).
20. Li, Z., Guo, S., Sun, Q. & Chan, G. K. Electronic landscape of the P-cluster of nitrogenase as revealed through many-electron quantum wavefunction simulations. *Nature chemistry* **11**, 1026–1033. doi:10.1038/s41557-019-0337-3 (2019).
21. Larsson, H. R., Zhai, H., Umrigar, C. J. & Chan, G. K.-L. The chromium dimer: closing a chapter of quantum chemistry. *Journal of the American Chemical Society* **144**, 15932–15937. doi:https://doi.org/10.1021/jacs.2c06357 (2022).
22. Schollwöck, U. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics* **326**, 96–192. doi:10.1016/j.aop.2010.09.012 (2011).
23. Orús, R. A practical introduction to tensor networks: Matrix product states and projected entangled pair states. *Annals of Physics* **349**, 117–158. doi:10.1016/j.aop.2014.06.013 (2014).
24. Hastings, M. B. An area law for one-dimensional quantum systems. *Journal of statistical mechanics: theory and experiment* **2007**, P08024. doi:10.1088/1742-5468/2007/08/P08024 (2007).

25. White, S. R. Density matrix formulation for quantum renormalization groups. *Physical review letters* **69**, 2863. doi:10.1103/PhysRevLett.69.2863 (1992).
26. White, S. R. Density-matrix algorithms for quantum renormalization groups. *Physical Review B* **48**, 10345. doi:10.1103/PhysRevB.48.10345 (1993).
27. *Papers Using the ITensor Library* <http://itensor.org/docs.cgi?page=papers>.
28. McCulloch, I. P. Infinite size density matrix renormalization group, revisited. *arXiv preprint arXiv:0804.2509* (2008).
29. Sharma, S., Sivalingam, K., Neese, F. & Chan, G. K. Low-energy spectrum of iron–sulfur clusters directly from many-particle quantum mechanics. *Nature chemistry* **6**, 927–933. doi:10.1038/nchem.2041 (2014).
30. Stoudenmire, E. & White, S. R. Studying two-dimensional systems with the density matrix renormalization group. *Annu. Rev. Condens. Matter Phys.* **3**, 111–128. doi:10.1146/annurev-conmatphys-020911-125018 (2012).
31. Barthel, T., Chung, M.-C. & Schollwoeck, U. Entanglement scaling in critical two-dimensional fermionic and bosonic systems. *Physical Review A* **74**, 022329. doi:10.1103/PhysRevA.74.022329 (2006).
32. Li, W., Ding, L., Yu, R., Roscilde, T. & Haas, S. Scaling behavior of entanglement in two-and three-dimensional free-fermion systems. *Physical Review B* **74**, 073103. doi:10.1103/PhysRevB.74.073103 (2006).
33. Schuch, N., Wolf, M. M., Verstraete, F. & Cirac, J. I. Computational complexity of projected entangled pair states. *Physical review letters* **98**, 140506. doi:10.1103/PhysRevLett.98.140506 (2007).
34. Jordan, J., Orús, R., Vidal, G., Verstraete, F. & Cirac, J. I. Classical simulation of infinite-size quantum lattice systems in two spatial dimensions. *Physical review letters* **101**, 250602. doi:10.1103/PhysRevLett.101.250602 (2008).
35. Lubasch, M., Cirac, J. I. & Banuls, M.-C. Unifying projected entangled pair state contractions. *New Journal of Physics* **16**, 033014. doi:10.1088/1367-2630/16/3/033014 (2014).
36. Lubasch, M., Cirac, J. I. & Banuls, M.-C. Algorithms for finite projected entangled pair states. *Physical Review B* **90**, 064425. doi:10.1103/PhysRevB.90.064425 (2014).
37. Corboz, P., Rice, T. M. & Troyer, M. Competing States in the t - J Model: Uniform d -Wave State versus Stripe State. *Phys. Rev. Lett.* **113**, 046402. doi:10.1103/PhysRevLett.113.046402 (4 July 2014).

38. Phien, H. N., Bengua, J. A., Tuan, H. D., Corboz, P. & Orús, R. Infinite projected entangled pair states algorithm improved: Fast full update and gauge fixing. *Phys. Rev. B* **92**, 035142. doi:10.1103/PhysRevB.92.035142 (3 July 2015).
39. Corboz, P. Variational optimization with infinite projected entangled-pair states. *Physical Review B* **94**, 035133. doi:10.1103/PhysRevB.94.035133 (2016).
40. Vanderstraeten, L., Haegeman, J., Corboz, P. & Verstraete, F. Gradient methods for variational optimization of projected entangled-pair states. *Physical Review B* **94**, 155123. doi:10.1103/PhysRevB.94.155123 (2016).
41. Zheng, B.-X. *et al.* Stripe order in the underdoped region of the two-dimensional Hubbard model. *Science* **358**, 1155–1160. doi:10.1126/science.aam7127 (2017).
42. Liu, W.-Y. *et al.* Gapless spin liquid ground state of the spin- $\frac{1}{2}$ $J_1 - J_2$ Heisenberg model on square lattices. *Phys. Rev. B* **98**, 241109. doi:10.1103/PhysRevB.98.241109 (24 Dec. 2018).
43. He, L. *et al.* Peps++: towards extreme-scale simulations of strongly correlated quantum many-particle models on sunway taihulight. *IEEE Transactions on Parallel and Distributed Systems* **29**, 2838–2848. doi:10.1109/TPDS.2018.2848618 (2018).
44. Liu, W.-Y., Huang, Y.-Z., Gong, S.-S. & Gu, Z.-C. Accurate simulation for finite projected entangled pair states in two dimensions. *Physical Review B* **103**, 235155. doi:10.1103/PhysRevB.103.235155 (2021).
45. Fröwis, F., Nebendahl, V. & Dür, W. Tensor operators: Constructions and applications for long-range interaction systems. *Physical Review A* **81**, 062337. doi:10.1103/PhysRevA.81.062337 (2010).
46. Crosswhite, G. M. & Bacon, D. Finite automata for caching in matrix product algorithms. *Physical Review A* **78**, 012356. doi:10.1103/PhysRevA.78.012356 (2008).
47. Vidal, G. Entanglement renormalization. *Physical review letters* **99**, 220405. doi:10.1103/PhysRevLett.99.220405 (2007).
48. Changlani, H. J., Kinder, J. M., Umrigar, C. J. & Chan, G. K.-L. Approximating strongly correlated wave functions with correlator product states. *Physical Review B* **80**, 245116. doi:10.1103/PhysRevB.80.245116 (2009).
49. Mezzacapo, F., Schuch, N., Boninsegni, M. & Cirac, J. I. Ground-state properties of quantum many-body systems: entangled-plaquette states and variational Monte Carlo. *New Journal of Physics* **11**, 083026. doi:10.1088/1367-2630/11/8/083026 (2009).

50. Verstraete, F., Murg, V. & Cirac, J. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Advances in Physics* **57**, 143–224. doi:10.1080/14789940801912366 (2008).
51. Vidal, G., Latorre, J. I., Rico, E. & Kitaev, A. Entanglement in Quantum Critical Phenomena. *Phys. Rev. Lett.* **90**, 227902. doi:10.1103/PhysRevLett.90.227902 (22 June 2003).
52. Calabrese, P. & Cardy, J. Entanglement entropy and quantum field theory. *Journal of Statistical Mechanics: Theory and Experiment* **2004**, P06002. doi:10.1088/1742-5468/2004/06/P06002 (2004).
53. Fannes, M., Nachtergaele, B. & Werner, R. F. Finitely correlated states on quantum spin chains. *Communications in mathematical physics* **144**, 443–490. doi:10.1007/BF02099178 (1992).
54. Östlund, S. & Rommer, S. Thermodynamic limit of density matrix renormalization. *Physical review letters* **75**, 3537. doi:10.1103/PhysRevLett.75.3537 (1995).
55. Fannes, M., Nachtergaele, B. & Werner, R. Finitely correlated pure states. *Journal of functional analysis* **120**, 511–534. doi:10.1006/jfan.1994.1041 (1994).
56. Schollwöck, U. The density-matrix renormalization group. *Rev. Mod. Phys.* **77**, 259–315. doi:10.1103/RevModPhys.77.259 (1 Apr. 2005).
57. Dorando, J. J., Hachmann, J. & Chan, G. K.-L. Analytic response theory for the density matrix renormalization group. *The Journal of chemical physics* **130**, 184111. doi:10.1063/1.3121422 (2009).
58. Haegeman, J. *et al.* Time-dependent variational principle for quantum lattices. *Physical review letters* **107**, 070601. doi:10.1103/PhysRevLett.107.070601 (2011).
59. Nakatani, N., Wouters, S., Van Neck, D. & Chan, G. K.-L. Linear response theory for the density matrix renormalization group: Efficient algorithms for strongly correlated excited states. *The Journal of chemical physics* **140**, 024108. doi:10.1063/1.4860375 (2014).
60. Vanderstraeten, L., Mariën, M., Verstraete, F. & Haegeman, J. Excitations and the tangent space of projected entangled-pair states. *Phys. Rev. B* **92**, 201111. doi:10.1103/PhysRevB.92.201111 (20 Nov. 2015).
61. Pollmann, F. & Turner, A. M. Detection of symmetry-protected topological phases in one dimension. *Phys. Rev. B* **86**, 125441. doi:10.1103/PhysRevB.86.125441 (12 Sept. 2012).
62. Huang, C.-Y., Chen, X. & Pollmann, F. Detection of symmetry-enriched topological phases. *Phys. Rev. B* **90**, 045142. doi:10.1103/PhysRevB.90.045142 (4 July 2014).

63. Nishino, T. & Okunishi, K. Corner transfer matrix renormalization group method. *Journal of the Physical Society of Japan* **65**, 891–894. doi:10.1143/JPSJ.65.891 (1996).
64. Verstraete, F. & Cirac, J. I. Renormalization algorithms for quantum-many body systems in two and higher dimensions. *arXiv preprint cond-mat/0407066* (2004).
65. Verstraete, F., Wolf, M. M., Perez-Garcia, D. & Cirac, J. I. Criticality, the area law, and the computational power of projected entangled pair states. *Physical review letters* **96**, 220601. doi:10.1103/PhysRevLett.96.220601 (2006).
66. Murg, V., Verstraete, F. & Cirac, J. I. Exploring frustrated spin systems using projected entangled pair states. *Phys. Rev. B* **79**, 195119. doi:10.1103/PhysRevB.79.195119 (19 May 2009).
67. Wang, L., Gu, Z.-C., Verstraete, F. & Wen, X.-G. Tensor-product state approach to spin- $\frac{1}{2}$ square J_1 - J_2 antiferromagnetic Heisenberg model: Evidence for deconfined quantum criticality. *Phys. Rev. B* **94**, 075143. doi:10.1103/PhysRevB.94.075143 (7 Aug. 2016).
68. Haghshenas, R., Lan, W.-W., Gong, S.-S. & Sheng, D. N. Quantum phase diagram of spin-1 J_1 - J_2 Heisenberg model on the square lattice: An infinite projected entangled-pair state and density matrix renormalization group study. *Phys. Rev. B* **97**, 184436. doi:10.1103/PhysRevB.97.184436 (18 May 2018).
69. Haghshenas, R. & Sheng, D. N. $U(1)$ -symmetric infinite projected entangled-pair states study of the spin-1/2 square J_1 - J_2 Heisenberg model. *Phys. Rev. B* **97**, 174408. doi:10.1103/PhysRevB.97.174408 (17 May 2018).
70. O'Rourke, M. J., Li, Z. & Chan, G. K.-L. Efficient representation of long-range interactions in tensor network algorithms. *Phys. Rev. B* **98**, 205127. doi:10.1103/PhysRevB.98.205127 (20 Nov. 2018).
71. Haghshenas, R., Gong, S.-S. & Sheng, D. N. Single-layer tensor network study of the Heisenberg model with chiral interactions on a kagome lattice. *Phys. Rev. B* **99**, 174423. doi:10.1103/PhysRevB.99.174423 (17 May 2019).
72. Phien, H. N., McCulloch, I. P. & Vidal, G. Fast convergence of imaginary time evolution tensor network algorithms by recycling the environment. *Phys. Rev. B* **91**, 115137. doi:10.1103/PhysRevB.91.115137 (11 Mar. 2015).
73. Evenbly, G. Gauge fixing, canonical forms, and optimal truncations in tensor networks with closed loops. *Phys. Rev. B* **98**, 085155. doi:10.1103/PhysRevB.98.085155 (8 Aug. 2018).
74. Pollmann, F., Khemani, V., Cirac, J. I. & Sondhi, S. L. Efficient variational diagonalization of fully many-body localized Hamiltonians. *Phys. Rev. B* **94**, 041116. doi:10.1103/PhysRevB.94.041116 (4 July 2016).

75. Wahl, T. B., Pal, A. & Simon, S. H. Efficient Representation of Fully Many-Body Localized Systems Using Tensor Networks. *Phys. Rev. X* **7**, 021018. doi:10.1103/PhysRevX.7.021018 (2 May 2017).
76. Evenbly, G. & Vidal, G. Algorithms for entanglement renormalization. *Phys. Rev. B* **79**, 144108. doi:10.1103/PhysRevB.79.144108 (14 Apr. 2009).
77. Kao, Y.-J., Hsieh, Y.-D. & Chen, P. Uni10: an open-source library for tensor network algorithms. *Journal of Physics: Conference Series* **640**, 012040. doi:10.1088/1742-6596/640/1/012040 (2015).
78. Evenbly, G. Algorithms for tensor network renormalization. *Phys. Rev. B* **95**, 045117. doi:10.1103/PhysRevB.95.045117 (4 Jan. 2017).
79. Higham, N. J. Stable iterations for the matrix square root. *Numerical Algorithms* **15**, 227–242. doi:10.1023/A:1019150005407 (Sept. 1997).
80. Corboz, P., Orús, R., Bauer, B. & Vidal, G. Simulation of strongly correlated fermions in two spatial dimensions with fermionic projected entangled-pair states. *Physical Review B* **81**, 165104. doi:10.1103/PhysRevB.81.165104 (2010).
81. Pi žorn, I., Wang, L. & Verstraete, F. Time evolution of projected entangled pair states in the single-layer picture. *Phys. Rev. A* **83**, 052321. doi:10.1103/PhysRevA.83.052321 (5 May 2011).
82. Xie, Z. *et al.* Optimized contraction scheme for tensor-network states. *Physical Review B* **96**, 045128. doi:10.1103/PhysRevB.96.045128 (2017).
83. Zaletel, M. P. & Pollmann, F. Isometric tensor network states in two dimensions. *Physical review letters* **124**, 037201. doi:10.1103/PhysRevLett.124.037201 (2020).
84. Carleo, G. & Troyer, M. Solving the quantum many-body problem with artificial neural networks. *Science* **355**, 602–606. doi:10.1126/science.aag2302 (2017).
85. Corboz, P. Improved energy extrapolation with infinite projected entangled-pair states applied to the two-dimensional Hubbard model. *Physical Review B* **93**, 045116. doi:10.1103/PhysRevB.93.045116 (2016).
86. Murg, V., Verstraete, F. & Cirac, J. I. Variational study of hard-core bosons in a two-dimensional optical lattice using projected entangled pair states. *Physical Review A* **75**, 033605. doi:10.1103/PhysRevA.75.033605 (2007).
87. Orús, R. & Vidal, G. Simulation of two-dimensional quantum systems on an infinite lattice revisited: Corner transfer matrix for tensor contraction. *Physical Review B* **80**, 094403. doi:10.1103/PhysRevB.80.094403 (2009).
88. Corboz, P., Rice, T. M. & Troyer, M. Competing States in the t - J Model: Uniform d -Wave State versus Stripe State. *Phys. Rev. Lett.* **113**, 046402. doi:10.1103/PhysRevLett.113.046402 (4 July 2014).

89. Levin, M. & Nave, C. P. Tensor renormalization group approach to two-dimensional classical lattice models. *Physical review letters* **99**, 120601. doi:10.1103/PhysRevLett.99.120601 (2007).
90. Evenbly, G. & Vidal, G. Tensor network renormalization. *Physical review letters* **115**, 180405. doi:10.1103/PhysRevLett.115.180405 (2015).
91. Martin, R. M. *Electronic structure: basic theory and practical methods* (Cambridge university press, 2004).
92. Szabo, A. & Ostlund, N. S. *Modern Quantum Chemistry: Intro to Advanced Electronic Structure Theory* (Dover publications, 1996).
93. White, S. R. & Martin, R. L. Ab initio quantum chemistry using the density matrix renormalization group. *The Journal of chemical physics* **110**, 4127–4130. doi:10.1063/1.478295 (1999).
94. Chan, G. K.-L., Keselman, A., Nakatani, N., Li, Z. & White, S. R. Matrix product operators, matrix product states, and ab initio density matrix renormalization group algorithms. *The Journal of chemical physics* **145**, 014102. doi:10.1063/1.4955108 (2016).
95. Stoudenmire, E., Wagner, L. O., White, S. R. & Burke, K. One-dimensional continuum electronic structure with the density-matrix renormalization group and its implications for density-functional theory. *Physical review letters* **109**, 056402. doi:10.1103/PhysRevLett.109.056402 (2012).
96. Wagner, L. O., Stoudenmire, E., Burke, K. & White, S. R. Reference electronic structure calculations in one dimension. *Physical Chemistry Chemical Physics* **14**, 8581–8590. doi:10.1039/C2CP24118H (2012).
97. Stoudenmire, E. M. & White, S. R. Sliced basis density matrix renormalization group for electronic structure. *Physical review letters* **119**, 046401. doi:10.1103/PhysRevLett.119.046401 (2017).
98. Dolfi, M., Bauer, B., Troyer, M. & Ristivojevic, Z. Multigrid algorithms for tensor network states. *Physical review letters* **109**, 020604. doi:10.1103/PhysRevLett.109.020604 (2012).
99. Mardirossian, N., McClain, J. D. & Chan, G. K.-L. Lowering of the complexity of quantum chemistry methods by choice of representation. *The Journal of chemical physics* **148**, 044106. doi:10.1063/1.5007779 (2018).
100. Xie, Z. Y. *et al.* Tensor Renormalization of Quantum Many-Body Systems Using Projected Entangled Simplex States. *Phys. Rev. X* **4**, 011025. doi:10.1103/PhysRevX.4.011025 (1 Feb. 2014).
101. Corboz, P. & Mila, F. Crystals of Bound States in the Magnetization Plateaus of the Shastry-Sutherland Model. *Phys. Rev. Lett.* **112**, 147203. doi:10.1103/PhysRevLett.112.147203 (14 Apr. 2014).

102. Picot, T. & Poilblanc, D. Nematic and supernematic phases in kagome quantum antiferromagnets under the influence of a magnetic field. *Phys. Rev. B* **91**, 064415. doi:10.1103/PhysRevB.91.064415 (6 Feb. 2015).
103. Picot, T., Ziegler, M., Orús, R. & Poilblanc, D. Spin- S kagome quantum antiferromagnets in a field with tensor networks. *Phys. Rev. B* **93**, 060407. doi:10.1103/PhysRevB.93.060407 (6 Feb. 2016).
104. Fano, G., Ortolani, F. & Ziosi, L. The density matrix renormalization group method: Application to the PPP model of a cyclic polyene chain. *The Journal of chemical physics* **108**, 9246–9252. doi:10.1063/1.476379 (1998).
105. Chan, G. K.-L. & Head-Gordon, M. Highly correlated calculations with a polynomial cost algorithm: A study of the density matrix renormalization group. *The Journal of chemical physics* **116**, 4462–4476. doi:10.1063/1.1449459 (2002).
106. Chan, G. K.-L. & Sharma, S. The density matrix renormalization group in quantum chemistry. *Annual review of physical chemistry* **62**, 465–481. doi:10.1146/annurev-physchem-032210-103338 (2011).
107. Pirvu, B., Murg, V., Cirac, J. I. & Verstraete, F. Matrix product operator representations. *New Journal of Physics* **12**, 025012. doi:10.1088/1367-2630/12/2/025012 (2010).
108. Crosswhite, G. M., Doherty, A. C. & Vidal, G. Applying matrix product operators to model systems with long-range interactions. *Physical Review B* **78**, 035116. doi:10.1103/PhysRevB.78.035116 (2008).
109. Zhao, H. *et al.* Renormalization of tensor-network states. *Physical Review B* **81**, 174411. doi:10.1103/PhysRevB.81.174411 (2010).
110. Baxter, R. J. *Exactly solved models in statistical mechanics* chap. 2,7 (Academic Press, 1982).
111. Braess, D. & Hackbusch, W. Approximation of $1/x$ by exponential sums in $[1, \infty]$. *IMA journal of numerical analysis* **25**, 685–697. doi:10.1093/imanum/dri015 (2005).
112. Verstraete, F., Garcia-Ripoll, J. J. & Cirac, J. I. Matrix product density operators: Simulation of finite-temperature and dissipative systems. *Physical review letters* **93**, 207204. doi:10.1103/PhysRevLett.93.207204 (2004).
113. McCulloch, I. P. From density-matrix renormalization group to matrix product states. *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P10014. doi:10.1088/1742-5468/2007/10/P10014 (2007).
114. Verstraete, F., Murg, V. & Cirac, J. Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems. *Adv. Phys.* **57**, 143–224. doi:10.1080/14789940801912366 (2008).

115. Fishman, M., White, S. R. & Stoudenmire, E. M. The ITensor Software Library for Tensor Network Calculations. *SciPost Phys. Codebases*, 4. doi:10.21468/SciPostPhysCodeb.4 (2022).
116. Haghshenas, R. & Sheng, D. U (1)-symmetric infinite projected entangled-pair states study of the spin-1/2 square J_1 - J_2 Heisenberg model. *Physical Review B* **97**, 174408. doi:10.1103/PhysRevB.97.174408 (2018).
117. Li, Z., O'Rourke, M. J. & Chan, G. K.-L. Generalization of the exponential basis for tensor network representations of long-range interactions in two and three dimensions. *Phys. Rev. B* **100**, 155121. doi:10.1103/PhysRevB.100.155121 (15 Oct. 2019).
118. Lin, L. & Tong, Y. Low-rank representation of tensor network operators with long-range pairwise interactions. *SIAM Journal on Scientific Computing* **43**, A164–A192. doi:10.1137/19M1287067 (2021).
119. Liao, H.-J., Liu, J.-G., Wang, L. & Xiang, T. Differentiable Programming Tensor Networks. *Phys. Rev. X* **9**, 031041. doi:10.1103/PhysRevX.9.031041 (3 Sept. 2019).
120. Beylkin, G. & Monzón, L. On approximation of functions by exponential sums. *Applied and Computational Harmonic Analysis* **19**, 17–48. doi:10.1016/j.acha.2005.01.003 (2005).
121. Beylkin, G. & Monzón, L. Approximation by exponential sums revisited. *Applied and Computational Harmonic Analysis* **28**, 131–149. doi:10.1016/j.acha.2009.08.011 (2010).
122. https://gitlab.com/mattorourke41/lr_gmpo_public.
123. Zauner-Stauber, V., Vanderstraeten, L., Fishman, M. T., Verstraete, F. & Haegeman, J. Variational optimization algorithms for uniform matrix product states. *Phys. Rev. B* **97**, 045145. doi:10.1103/PhysRevB.97.045145 (4 Jan. 2018).
124. Parker, D. E., Cao, X. & Zaletel, M. P. Local matrix product operators: Canonical form, compression, and control theory. *Physical Review B* **102**, 035147. doi:10.1103/PhysRevB.102.035147 (2020).
125. Fishman, M. T., Vanderstraeten, L., Zauner-Stauber, V., Haegeman, J. & Verstraete, F. Faster methods for contracting infinite two-dimensional tensor networks. *Phys. Rev. B* **98**, 235148. doi:10.1103/PhysRevB.98.235148 (23 Dec. 2018).
126. Bloch, I., Dalibard, J. & Zwerger, W. Many-body physics with ultracold gases. *Rev. Mod. Phys.* **80**, 885–964. doi:10.1103/RevModPhys.80.885 (3 July 2008).

127. Saffman, M. Quantum computing with atomic qubits and Rydberg interactions: progress and challenges. *Journal of Physics B: Atomic, Molecular and Optical Physics* **49**, 202001. doi:10.1088/0953-4075/49/20/202001 (Oct. 2016).
128. Endres, M. *et al.* Atom-by-atom assembly of defect-free one-dimensional cold atom arrays. *Science* **354**, 1024–1027. doi:10.1126/science.aah3752 (2016).
129. Pichler, H., Wang, S.-T., Zhou, L., Choi, S. & Lukin, M. D. Computational complexity of the Rydberg blockade in two dimensions. *arXiv preprint arXiv:1809.04954* (2018).
130. Bernien, H. *et al.* Probing many-body dynamics on a 51-atom quantum simulator. *Nature* **551**, 579–584. doi:10.1038/nature24622 (Nov. 2017).
131. Keesling, A. *et al.* Quantum Kibble–Zurek mechanism and critical dynamics on a programmable Rydberg simulator. *Nature* **568**, 207–211. doi:10.1038/s41586-019-1070-1 (2019).
132. De Léséleuc, S. *et al.* Observation of a symmetry-protected topological phase of interacting bosons with Rydberg atoms. *Science* **365**, 775–780. doi:10.1126/science.aav9105 (2019).
133. Samajdar, R., Ho, W. W., Pichler, H., Lukin, M. D. & Sachdev, S. Quantum phases of Rydberg atoms on a kagome lattice. *Proceedings of the National Academy of Sciences* **118**. doi:10.1073/pnas.2015785118 (2021).
134. Verresen, R., Lukin, M. D. & Vishwanath, A. Prediction of Toric Code Topological Order from Rydberg Blockade. *Physical Review X* **11**, 031005. doi:10.1103/PhysRevX.11.031005 (3 July 2021).
135. Barredo, D., De Léséleuc, S., Lienhard, V., Lahaye, T. & Browaeys, A. An atom-by-atom assembler of defect-free arbitrary two-dimensional atomic arrays. *Science* **354**, 1021–1023. doi:10.1126/science.aah3778 (2016).
136. Lee, W., Kim, H. & Ahn, J. Defect-free atomic array formation using the Hungarian matching algorithm. *Physical Review A* **95**, 053424. doi:10.1103/PhysRevA.95.053424 (5 May 2017).
137. Brown, M. O., Thiele, T., Kiehl, C., Hsu, T.-W. & Regal, C. A. Gray-Molasses Optical-Tweezer Loading: Controlling Collisions for Scaling Atom-Array Assembly. *Physical Review X* **9**, 011057. doi:10.1103/PhysRevX.9.011057 (1 Mar. 2019).
138. Ohl de Mello, D. *et al.* Defect-Free Assembly of 2D Clusters of More Than 100 Single-Atom Quantum Systems. *Physical Review Letters* **122**, 203601. doi:10.1103/PhysRevLett.122.203601 (20 May 2019).
139. Barredo, D., Lienhard, V., De Léséleuc, S., Lahaye, T. & Browaeys, A. Synthetic three-dimensional atomic structures assembled atom by atom. *Nature* **561**, 79–82. doi:10.1038/s41586-018-0450-2 (2018).

140. Kumar, A., Wu, T.-Y., Giraldo, F. & Weiss, D. S. Sorting ultracold atoms in a three-dimensional optical lattice in a realization of Maxwell's demon. *Nature* **561**, 83–87. doi:10.1038/s41586-018-0458-7 (2018).
141. Scholl, P. *et al.* Quantum simulation of 2D antiferromagnets with hundreds of Rydberg atoms. *Nature* **595**, 233–238. doi:10.1038/s41586-021-03585-1 (2021).
142. Pichler, H., Wang, S.-T., Zhou, L., Choi, S. & Lukin, M. D. Quantum optimization for maximum independent set using Rydberg atom arrays. *arXiv preprint arXiv:1808.10816* (2018).
143. Ebadi, S. *et al.* Quantum phases of matter on a 256-atom programmable quantum simulator. *Nature* **595**, 227–232. doi:10.1038/s41586-021-03582-4 (2021).
144. Samajdar, R., Ho, W. W., Pichler, H., Lukin, M. D. & Sachdev, S. Complex Density Wave Orders and Quantum Phase Transitions in a Model of Square-Lattice Rydberg Atom Arrays. *Physical Review Letters* **124**. doi:10.1103/PhysRevLett.124.103601 (Mar. 2020).
145. Felser, T., Notarnicola, S. & Montangero, S. Efficient Tensor Network Ansatz for High-Dimensional Quantum Many-Body Problems. *Physical Review Letters* **126**, 170603. doi:10.1103/PhysRevLett.126.170603 (17 Apr. 2021).
146. Bak, P. Commensurate phases, incommensurate phases and the devil's staircase. *Reports on Progress in Physics* **45**, 587 (1982).
147. Bak, P. & Bruinsma, R. One-dimensional Ising model and the complete devil's staircase. *Physical Review Letters* **49**, 249 (1982).
148. Fendley, P., Sengupta, K. & Sachdev, S. Competing density-wave orders in a one-dimensional hard-boson model. *Physical Review B* **69**, 075106. doi:10.1103/PhysRevB.69.075106 (7 Feb. 2004).
149. Schauß, P. *et al.* Crystallization in Ising quantum magnets. *Science* **347**, 1455–1458. doi:10.1126/science.1258351 (2015).
150. Weimer, H. & Büchler, H. P. Two-Stage Melting in Systems of Strongly Interacting Rydberg Atoms. *Physical Review Letters* **105**, 230403. doi:10.1103/PhysRevLett.105.230403 (23 Nov. 2010).
151. Rader, M. & Läuchli, A. M. Floating Phases in One-Dimensional Rydberg Ising Chains. *arXiv preprint arXiv:1908.02068* (2019).
152. O'Rourke, M. J. & Chan, G. K.-L. Simplified and improved approach to tensor network operators in two dimensions. *Phys. Rev. B* **101**, 205142. doi:10.1103/PhysRevB.101.205142 (20 May 2020).

153. Kalinowski, M. *et al.* Bulk and boundary quantum phase transitions in a square Rydberg atom array. *Physical Review B* **105**, 174417. doi:10.1103/PhysRevB.105.174417 (2022).
154. Merali, E., De Vlugt, I. J. & Melko, R. G. Stochastic series expansion quantum monte carlo for rydberg arrays. *arXiv preprint arXiv:2107.00766* (2021).
155. Miles, C. *et al.* Machine learning discovery of new phases in programmable quantum simulator snapshots. *arXiv preprint arXiv:2112.10789* (2021).
156. Czischek, S., Moss, M. S., Radzihovsky, M., Merali, E. & Melko, R. G. Data-enhanced variational Monte Carlo simulations for Rydberg atom arrays. *Physical Review B* **105**, 205108. doi:10.1103/PhysRevB.105.205108 (2022).
157. Carr, L. D., DeMille, D., Krems, R. V. & Ye, J. Cold and ultracold molecules: science, technology and applications. *New J. Phys.* **11**, 55049. doi:10.1088/1367-2630/11/5/055049 (May 2009).
158. Quémener, G. & Julienne, P. S. Ultracold molecules under control! *Chem. Rev.* **112**, 4949–5011. doi:10.1021/cr300092g (Sept. 2012).
159. Safronova, M. S. *et al.* Search for new physics with atoms and molecules. *Reviews of Modern Physics* **90**, 025008. doi:10.1103/RevModPhys.90.025008 (June 2018).
160. Bohn, J. L., Rey, A. M. & Ye, J. Cold molecules: Progress in quantum engineering of chemistry and quantum matter. *Science* **357**, 1002–1010. doi:10.1126/science.aam6299 (Sept. 2017).
161. Di Rosa, M. D. Laser-cooling molecules. *Eur. Phys. J. D* **31**, 395–402. doi:10.1140/epjd/e2004-00167-2 (Nov. 2004).
162. Isaev, T. A. & Berger, R. Polyatomic Candidates for Cooling of Molecules with Lasers from Simple Theoretical Concepts. *Phys. Rev. Lett.* **116**, 063006. doi:10.1103/PhysRevLett.116.063006 (Feb. 2016).
163. Kozyryev, I., Baum, L., Matsuda, K. & Doyle, J. M. Proposal for Laser Cooling of Complex Polyatomic Molecules. *ChemPhysChem* **17**, 3641–3648. doi:10.1002/cphc.201601051 (Nov. 2016).
164. Tarbutt, M. R. Laser cooling of molecules. *Contemp. Phys.*, 1–21. doi:10.1080/00107514.2018.1576338 (Feb. 2019).
165. Shuman, E. S., Barry, J. F. & Demille, D. Laser cooling of a diatomic molecule. *Nature* **467**, 820–823. doi:10.1038/nature09443 (Sept. 2010).
166. Hummon, M. T. *et al.* 2D Magneto-Optical Trapping of Diatomic Molecules. *Phys. Rev. Lett.* **110**, 143001. doi:10.1103/PhysRevLett.110.143001 (Apr. 2013).

167. Barry, J. F., McCarron, D. J., Norrgard, E. B., Steinecker, M. H. & DeMille, D. Magneto-optical trapping of a diatomic molecule. *Nature* **512**, 286–289. doi:10.1038/nature13634 (Aug. 2014).
168. Truppe, S. *et al.* Molecules cooled below the Doppler limit. *Nat. Phys.* **13**, 1173–1176. doi:10.1038/nphys4241 (Aug. 2017).
169. Anderegg, L. *et al.* Radio Frequency Magneto-Optical Trapping of CaF with High Density. *Phys. Rev. Lett.* **119**, 103201. doi:10.1103/PhysRevLett.119.103201 (Sept. 2017).
170. Ellis, A. M. Main group metal-ligand interactions in small molecules: New insights from laser spectroscopy. *Int. Rev. Phys. Chem.* **20**, 551–590. doi:10.1080/01442350110062542 (Oct. 2001).
171. Kozyryev, I. *et al.* Sisyphus Laser Cooling of a Polyatomic Molecule. *Phys. Rev. Lett.* **118**, 173201. doi:10.1103/PhysRevLett.118.173201 (Apr. 2017).
172. Collopy, A. L. *et al.* 3D Magneto-Optical Trap of Yttrium Monoxide. *Physical Review Letters* **121**, 213201. doi:10.1103/PhysRevLett.121.213201 (Nov. 2018).
173. Norrgard, E. B. *et al.* Hyperfine Structure of the B3Pi1 State and Predictions of Optical Cycling Behavior in the X-B transition of TlF. *Phys. Rev. A* **95**, 062506. doi:10.1103/PhysRevA.95.062506 (2017).
174. Smallman, I., Wang, F., Steimle, T., Tarbutt, M. & Hinds, E. Radiative branching ratios for excited states of ¹⁷⁴YbF: Application to laser cooling. *J. Mol. Spectrosc.* **300**, 3–6. doi:10.1016/j.jms.2014.02.006 (June 2014).
175. Aggarwal, P. *et al.* Measuring the electric dipole moment of the electron in BaF. *Eur. Phys. J. D* **72**, 197. doi:10.1140/epjd/e2018-90192-9 (Nov. 2018).
176. Chen, T., Bu, W. & Yan, B. Radiative deflection of a BaF molecular beam via optical cycling. *Phys. Rev. A* **96**, 053401. doi:10.1103/PhysRevA.96.053401 (Nov. 2017).
177. Franck, J. & Dymond, E. Elementary processes of photochemical reactions. *Transactions of the Faraday Society* **21**, 536–542 (1926).
178. Condon, E. A theory of intensity distribution in band systems. *Physical Review* **28**, 1182 (1926).
179. Stuhl, B. K., Sawyer, B. C., Wang, D. & Ye, J. Magneto-optical Trap for Polar Molecules. *Phys. Rev. Lett.* **101**, 243002. doi:10.1103/PhysRevLett.101.243002 (Dec. 2008).
180. Tarbutt, M. R. Magneto-optical trapping forces for atoms and molecules with complex level structures. *New J. Phys.* **17**, 15007. doi:10.1088/1367-2630/17/1/015007 (Jan. 2015).

181. Khriplovich, I. B. & Lamoreaux, S. K. *CP Violation Without Strangeness* doi:10.1007/978-3-642-60838-4 (Springer-Verlag Berlin Heidelberg, 1997).
182. Isaev, T., Hoekstra, S. & Berger, R. Laser-cooled RaF as a promising candidate to measure molecular parity violation. *Phys. Rev. A* **82**, 052521. doi:10.1103/PhysRevA.82.052521 (Nov. 2010).
183. V. Andreev *et al.* Improved limit on the electric dipole moment of the electron. *Nature* **562**, 355–360. doi:10.1038/s41586-018-0599-8 (Oct. 2018).
184. Cairncross, W. B. *et al.* Precision Measurement of the Electron's Electric Dipole Moment Using Trapped Molecular Ions. *Phys. Rev. Lett.* **119**, 153001. doi:10.1103/PhysRevLett.119.153001 (Oct. 2017).
185. Kozyryev, I. & Hutzler, N. R. Precision Measurement of Time-Reversal Symmetry Violation with Laser-Cooled Polyatomic Molecules. *Phys. Rev. Lett.* **119**, 133002. doi:10.1103/PhysRevLett.119.133002 (Sept. 2017).
186. Flambaum, V. V., DeMille, D. & Kozlov, M. G. Time-Reversal Symmetry Violation in Molecules Induced by Nuclear Magnetic Quadrupole Moments. *Phys. Rev. Lett.* **113**, 103003. doi:10.1103/PhysRevLett.113.103003 (Sept. 2014).
187. Lee, J. *et al.* Optical spectroscopy of tungsten carbide for uncertainty analysis in electron electric-dipole-moment search. *Phys. Rev. A* **87**, 022516. doi:10.1103/PhysRevA.87.022516 (Feb. 2013).
188. Isaev, T. A., Zaitsevskii, A. V. & Eliav, E. Laser-coolable polyatomic molecules with heavy nuclei. *Journal of Physics B* **50**, 225101. doi:10.1088/1361-6455/aa8f34 (Nov. 2017).
189. Gaul, K. & Berger, R. Ab initio study of parity and time-reversal violation in laser-coolable triatomic molecules. *Phys. Rev. A* **101**, 012508. doi:10.1103/PhysRevA.101.012508 (1 Jan. 2020).
190. Denis, M. *et al.* Enhancement factor for the electric dipole moment of the electron in the BaOH and YbOH molecules. *Phys. Rev. A* **99**, 042512. doi:10.1103/PhysRevA.99.042512 (4 Apr. 2019).
191. Prasanna, V. S., Shitara, N., Sakurai, A., Abe, M. & Das, B. P. Enhanced sensitivity of the electron electric dipole moment from YbOH: The role of theory. *Phys. Rev. A* **99**, 062502. doi:10.1103/PhysRevA.99.062502 (June 2019).
192. Antonov, I. O., Barker, B. J. & Heaven, M. C. Pulsed-field ionization zero electron kinetic energy spectrum of the ground electronic state of BeOBe⁺. *J. Chem. Phys.* **134**, 044306. doi:10.1063/1.3541255 (Jan. 2011).
193. Ostojić, B., Jensen, P., Schwerdtfeger, P. & Bunker, P. R. Singlet-triplet interaction in Group 2 M₂O hypermetallic oxides. *J. Mol. Spectrosc.* **301**, 20–24. doi:10.1016/j.jms.2014.05.003 (2014).

194. Puri, P. *et al.* Synthesis of mixed hypermetallic oxide BaOCa+from laser-cooled reagents in an atom-ion hybrid trap. *Science* **357**, 1370–1375. doi:10.1126/science.aan4701 (2017).
195. Looock, H.-P., Bérces, A., Simard, B. & Linton, C. Laser spectroscopy of the $\tilde{A}^2\Pi \leftarrow \tilde{X}^2\Sigma^+$ transition of ytterbium monoacetylide. *J. Chem. Phys.* **107**, 2720–2727. doi:10.1063/1.474629 (Aug. 1997).
196. Marr, A. J., Perry, J. & Steimle, T. C. A molecular beam optical/Stark study of calcium monoacetylide. *J. Chem. Phys.* **103**, 3861–3863. doi:10.1063/1.470038 (Sept. 1995).
197. Apetrei, C., Ding, H. & Maier, J. P. Gas phase electronic spectrum of linear AlCCH. *Physical Chemistry Chemical Physics* **9**, 3897–3901. doi:10.1039/B706384A (2007).
198. Wall, T. *et al.* Lifetime of the A(v'=0) state and Franck-Condon factor of the A-X(0-0) transition of CaF measured by the saturation of laser-induced fluorescence. *Phys. Rev. A* **78**, 062509. doi:10.1103/PhysRevA.78.062509 (Dec. 2008).
199. Wells, N. & Lane, I. C. Electronic states and spin-forbidden cooling transitions of AlH and AlF. *Phys. Chem. Chem. Phys.* **13**, 19018. doi:10.1039/c1cp21313j (2011).
200. Zhuang, X. *et al.* Franck–Condon factors and radiative lifetime of the $A^2\Pi_{1/2} - X^2\Sigma^+$ transition of ytterbium monofluoride, YbF. *Phys. Chem. Chem. Phys.* **13**, 19013. doi:10.1039/c1cp21585j (Aug. 2011).
201. Brazier, C. R., Ellingboe, L. C. & Bernath, P. F. Laser Spectroscopy of Alkaline Earth Monoalkoxide Free Radicals. *J. Am. Chem. Soc.* **108**, 2126–2132. doi:10.1021/ja00269a002 (1986).
202. Rice, S. F., Martin, H. & Field, R. W. The electronic structure of the calcium monohalides. A ligand field approach. *The Journal of Chemical Physics* **82**, 5023–5034. doi:10.1063/1.448676 (1985).
203. Allouche, A., Wannous, G. & Aubert-Frécon, M. A ligand-field approach for the low-lying states of Ca, Sr and Ba monohalides. *Chemical physics* **170**, 11–22. doi:10.1016/0301-0104(93)80087-P (1993).
204. Taylor, C. M., Chaudhuri, R. K. & Freed, K. F. Electronic structure of the calcium monohydroxide radical. *The Journal of Chemical Physics* **122**, 044317. doi:10.1063/1.1834511 (2005).
205. Verges, J. *et al.* The B' 2 Delta state of CaF. *J. Phys. B At. Mol. Opt. Phys.* **26**, 279–284. doi:10.1088/0953-4075/26/2/011 (Jan. 1993).
206. Jarman, C. N. & Bernath, P. F. High resolution laser spectroscopy of the $\tilde{C}^2\Delta - \tilde{X}^2\Sigma^+$ transition of CaOH and CaOD: Vibronic coupling and the Renner–Teller effect. *J. Chem. Phys.* **97**, 1711–1718. doi:10.1063/1.463158 (Aug. 1992).

207. Melville, T. C. & Coxon, J. A. The visible laser excitation spectrum of YbOH: The $A^2\Pi - X^2\Sigma^+$ transition. *J. Chem. Phys.* **115**, 6974–6978. doi:10.1063/1.1404145 (2001).
208. Oberlander, M. D. *Laser excited fluorescence studies of reactions of group 2 metals with oxygen containing molecules and of heavy group 15 clusters with fluorine* PhD thesis (The Ohio State University, 1995).
209. Herzberg, G. *Molecular Spectra and Molecular Structure, Volume 3: Electronic Spectra and Electronic Structure of Polyatomic Molecules* (D. Van Nostrand, 1967).
210. Werner, H.-J. & Knowles, P. J. A second order multiconfiguration SCF procedure with optimum convergence. *The Journal of Chemical Physics* **82**, 5053–5063. doi:10.1063/1.448627 (1985).
211. Knowles, P. J. & Werner, H.-J. An efficient second-order MC SCF method for long configuration expansions. *Chemical Physics Letters* **115**, 259–267. doi:10.1016/0009-2614(85)80025-7 (1985).
212. Werner, H.-J. & Knowles, P. J. An efficient internally contracted multiconfiguration-reference configuration interaction method. *The Journal of Chemical Physics* **89**, 5803–5814. doi:10.1063/1.455556 (1988).
213. Knowles, P. J. & Werner, H.-J. An efficient method for the evaluation of coupling coefficients in configuration interaction calculations. *Chemical Physics Letters* **145**, 514–522. doi:10.1016/0009-2614(88)87412-8 (1988).
214. Werner, H.-J., Knowles, P. J., Knizia, G., Manby, F. R., Schütz, M., *et al.* *MOLPRO, version 2015.1, a package of ab initio programs* see <https://www.molpro.net/>. 2015.
215. Porsev, S. G., Safronova, M. S. & Kozlov, M. G. Correlation effects in Yb⁺ and implications for parity violation. *Phys. Rev. A* **86**, 022504. doi:10.1103/PhysRevA.86.022504 (2 Aug. 2012).
216. Feldker, T., Fürst, H., Ewald, N. V., Joger, J. & Gerritsma, R. Spectroscopy of the $^2S_{1/2} \rightarrow ^2P_{3/2}$ transition in Yb ii: Isotope shifts, hyperfine splitting, and branching ratios. *Phys. Rev. A* **97**, 032511. doi:10.1103/PhysRevA.97.032511 (3 Mar. 2018).
217. Roos, B. O., Veryazov, V. & Widmark, P.-O. Relativistic atomic natural orbital type basis sets for the alkaline and alkaline-earth atoms applied to the ground-state potentials for the corresponding dimers. *Theoretical Chemistry Accounts* **111**, 345–351. doi:10.1007/s00214-003-0537-0 (2004).
218. Roos, B. O., Lindh, R., Malmqvist, P.-Å., Veryazov, V. & Widmark, P.-O. Main group atoms and dimers studied with a new relativistic ANO basis set. *The Journal of Physical Chemistry A* **108**, 2851–2858. doi:10.1021/jp031064+ (2004).

219. Cao, X. & Dolg, M. Valence basis sets for relativistic energy-consistent small-core lanthanide pseudopotentials. *The Journal of Chemical Physics* **115**, 7348–7355. doi:10.1063/1.1406535 (2001).
220. Cao, X. & Dolg, M. Segmented contraction scheme for small-core lanthanide pseudopotential basis sets. *Journal of Molecular Structure: THEOCHEM* **581**, 139–147. doi:10.1016/S0166-1280(01)00751-5 (2002).
221. Dolg, M., Stoll, H. & Preuss, H. Energy-adjusted abinitio pseudopotentials for the rare earth elements. *The Journal of Chemical Physics* **90**, 1730–1734. doi:10.1063/1.456066 (1989).
222. Malmqvist, P. Å., Roos, B. O. & Schimmelpfennig, B. The restricted active space (RAS) state interaction approach with spin-orbit coupling. *Chemical physics letters* **357**, 230–240. doi:10.1016/S0009-2614(02)00498-0 (2002).
223. Sjøvoll, M., Gropen, O. & Olsen, J. A determinantal approach to spin-orbit configuration interaction. *Theoretical Chemistry Accounts* **97**, 301–312. doi:10.1007/s002140050265 (1997).
224. Berning, A., Schweizer, M., Werner, H.-J., Knowles, P. J. & Palmieri, P. Spin-orbit matrix elements for internally contracted multireference configuration interaction wavefunctions. *Molecular Physics* **98**, 1823–1833. doi:10.1080/00268970009483386 (2000).
225. Yabushita, S., Zhang, Z. & Pitzer, R. M. Spin-Orbit Configuration Interaction Using the Graphical Unitary Group Approach and Relativistic Core Potential and Spin-Orbit Operators. *The Journal of Physical Chemistry A* **103**, 5791–5800. doi:10.1021/jp9901242 (1999).
226. Shao, Y. *et al.* Advances in molecular quantum chemistry contained in the Q-Chem 4 program package. *Molecular Physics* **113**, 184–215. doi:10.1080/00268976.2014.952696 (Sept. 2014).
227. Weigend, F., Häser, M., Patzelt, H. & Ahlrichs, R. RI-MP2: optimized auxiliary basis sets and demonstration of efficiency. *Chemical physics letters* **294**, 143–152. doi:10.1016/S0009-2614(98)00862-8 (1998).
228. Weigend, F. & Ahlrichs, R. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for H to Rn: Design and assessment of accuracy. *Physical Chemistry Chemical Physics* **7**, 3297–3305. doi:10.1039/B508541A (2005).
229. Gulde, R., Pollak, P. & Weigend, F. Error-Balanced Segmented Contracted Basis Sets of Double- ζ to Quadruple- ζ Valence Quality for the Lanthanides. *Journal of chemical theory and computation* **8**, 4062–4068. doi:10.1021/ct300302u (2012).
230. Mozhayskiy, V. & Krylov, A. *ezSpectrum* <http://iopshell.usc.edu/downloads>.

231. Kupka, H. & Cribb, P. Multidimensional Franck–Condon integrals and Duschinsky mixing effects. *The Journal of chemical physics* **85**, 1303–1315. doi:10.1063/1.451216 (1986).
232. Berger, R., Fischer, C. & Klessinger, M. Calculation of the vibronic fine structure in electronic spectra at higher temperatures. 1. Benzene and pyrazine. *The Journal of Physical Chemistry A* **102**, 7157–7167. doi:10.1021/jp981597w (1998).
233. Duschinsky, F. The importance of the electron spectrum in multi atomic molecules. Concerning the Franck-Condon principle. *Acta Physicochim. USSR* **7**, 551–566 (1937).
234. Auerbach, N., Flambaum, V. V. & Spevak, V. Collective T - and P -Odd Electromagnetic Moments in Nuclei with Octupole Deformations. *Physical Review Letters* **76**, 4316–4319. doi:10.1103/PhysRevLett.76.4316 (June 1996).
235. Wentink, T. & Spindler, R. J. The isoelectronic series ScF through ThO, I notes on the band spectra of TiO, HfO and ThO. *J. Quant. Spectrosc. Radiat. Transf.* **12**, 1569–1590. doi:10.1016/0022-4073(72)90131-8 (Nov. 1972).
236. Kokkin, D. L., Steimle, T. C. & DeMille, D. Branching ratios and radiative lifetimes of the U, L, and I states of thorium oxide. *Phys. Rev. A* **90**, 062503. doi:10.1103/PhysRevA.90.062503 (Dec. 2014).
237. Bouchard, J. L., Steimle, T., Kokkin, D. L., Sharfi, D. J. & Mawhorter, R. J. Branching ratios, radiative lifetimes, and transition dipole moments for tantalum nitride, TaN. *J. Mol. Spectrosc.* **325**, 1–6. doi:10.1016/j.jms.2016.05.006 (2016).
238. Flambaum, V. V. Enhanced nuclear Schiff moment and time-reversal violation in 229Th-containing molecules. *Phys. Rev. C* **99**, 035501. doi:10.1103/PhysRevC.99.035501 (Mar. 2019).
239. Kingston, C. T. *Laser spectroscopy of some transition metal-containing free radicals* PhD thesis (University of British Columbia, 2001). doi:10.14288/1.0061350.
240. Hunter, L. R., Peck, S. K., Greenspon, A. S., Alam, S. S. & DeMille, D. Prospects for laser cooling TlF. *Phys. Rev. A* **85**, 012511. doi:10.1103/PhysRevA.85.012511 (Jan. 2012).
241. Kovács, I. On the Spin-Orbit Interaction in Diatomic Molecules. I. *Can. J. Phys.* **36**, 309–328. doi:10.1139/p58-034 (Mar. 1958).
242. Bernath, P. F. & Brazier, C. R. Spectroscopy of CaOH. *Astrophys. J.* **288**, 373. doi:10.1086/162800 (Jan. 1985).

243. Ronca, E., Li, Z., Jimenez-Hoyos, C. A. & Chan, G. K.-L. Time-step targeting time-dependent and dynamical density matrix renormalization group algorithms with ab initio Hamiltonians. *Journal of chemical theory and computation* **13**, 5560–5571. doi:10.1021/acs.jctc.7b00682 (2017).
244. Xing, X., Li, X. & Lin, L. Unified analysis of finite-size error for periodic Hartree-Fock and second order Moller-Plesset perturbation theory. *arXiv preprint arXiv:2108.00206* (2021).
245. McLean, W. in *Contemporary Computational Mathematics - A Celebration of the 80th Birthday of Ian Sloan* (eds Dick, J., Kuo, F. Y. & Woźniakowski, H.) 911–930 (Springer International Publishing, Cham, 2018). doi:10.1007/978-3-319-72456-0_40.
246. Nocedal, J. & Wright, S. *Numerical optimization* (Springer Science & Business Media, 2006).
247. Vidal, G. Efficient classical simulation of slightly entangled quantum computations. *Physical Review Letters* **91**, 147902. doi:10.1103/PhysRevLett.91.147902 (2003).
248. Jiang, H.-C., Weng, Z.-Y. & Xiang, T. Accurate determination of tensor network state of quantum lattice models in two dimensions. *Physical Review Letters* **101**, 090603. doi:10.1103/PhysRevLett.101.090603 (2008).

Appendix A

SUPPLEMENTARY INFORMATION: ENTANGLEMENT IN THE QUANTUM PHASES OF AN UNFRUSTRATED RYDBERG ATOM ARRAY (CHAPTER 5)

This appendix is based on the following publication:

1. O'Rourke, M. J. & Chan, G. K. Entanglement in the quantum phases of an unfrustrated Rydberg atom array. *arXiv preprint arXiv:2201.03189* (2022).

A.1 Numerical methods

This section gives details for the numerical simulations in this work. Principally, it will focus on algorithmic subtleties and sources of error, as well as the strategies employed to resolve the physics of the Rydberg atom system.

Γ -point DMRG

Theoretical details and relation to other methods

In this work we chose to perform 2D DMRG in a site Bloch basis at the Γ -point in the Brillouin zone. Let us define the computational supercell of the DMRG calculation to be of dimension $L_x \times L_y$ sites. Then, such a Γ -point site Bloch basis state $|\tilde{n}_{x,y}\rangle$ is related to the normal site basis state $|n_{x,y}\rangle$ at site $r_{x,y}$ by

$$|\tilde{n}_{x,y}\rangle = \sum_l |n_{(x,y)+R_l}\rangle, \quad (\text{A.1})$$

where $R_l = (n \cdot L_x, m \cdot L_y)$; $n, m \in \mathbb{Z}$. In other words, each single particle basis state is a superposition of the original site basis states separated by lattice vectors of the supercell. The Bloch function has unit norm per supercell.

The many-particle 2D DMRG wavefunction is then

$$|\Psi\rangle = \sum_{\{e\}} \prod_{x,y} A_{\{e_{x,y}\}}^{\tilde{n}_{x,y}} |\tilde{n}_{x,y}\rangle, \quad (\text{A.2})$$

where $\mathbf{A}^{\tilde{n}_{x,y}}$ is the MPS tensor associated with Bloch function $\tilde{n}_{x,y}$, $e_{x,y}$ denote its bonds, and an ordering has been chosen through the lattice.

As explained in Chapter 5, this representation is different from the cylindrical boundary condition MPS employed in previous studies [133, 134, 144]. The primary

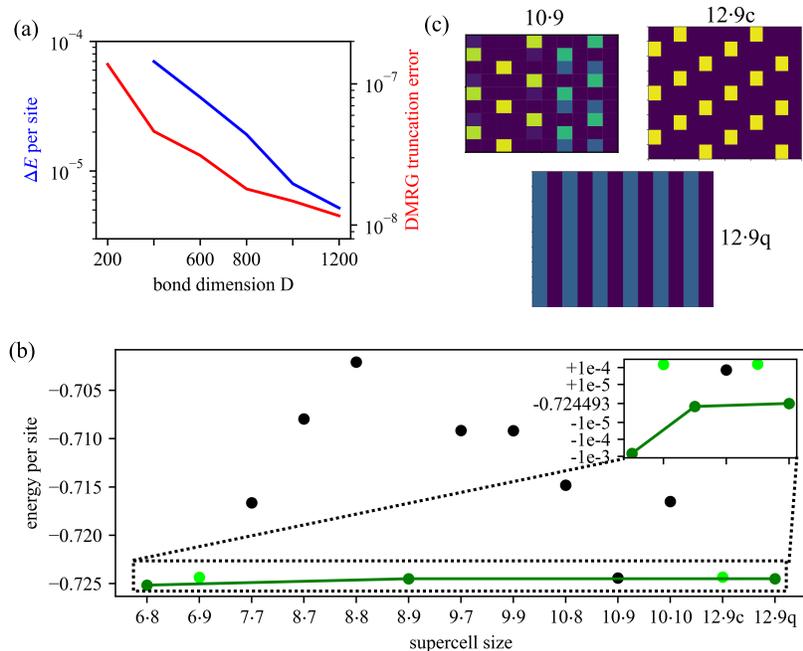


Figure A.1: Convergence of Γ -point DMRG in the most difficult region of the phase diagram $(\delta, R_b) = (5.0-6.0, 2.3)$. (a) Shows the convergence w.r.t. bond dimension of the largest truncated DMRG singular value (red) and the change in energy per site relative to the energy obtained with bond dimension $D = 200$ (blue). (b) The energies per site of a large variety of supercell sizes. This is adapted from Fig. 5.1 to highlight the relevant points. The connected dark green points are the nematic phase, and lime green points are the low energy 3-star $\frac{1}{6}$ -density crystalline phase. The inset shows the convergence of the nematic phase energy w.r.t. supercell size and gaps to the other low energy solutions, whose density profiles are shown in (c). Note that, between (a) and (b), the nematic phase is converged to below 10^{-5} accuracy while the competing states differ in energy by at least 10^{-4} .

advantage of the current approach is that regardless of supercell size, the 2D DMRG state in the Bloch basis represents an *infinite* system in 2D (rather than a finite system in at least one direction in prior cylindrical studies) simply because the underlying single-particle basis is a discrete periodic function on the infinite 2D square lattice. Thus there is no need to truncate the Rydberg interactions unlike in cylinder studies. We note that this type of Bloch basis is widely used in electronic structure theory partly for similar reasons, namely it allows one to treat the infinite range Coulomb interaction. For an example of a DMRG calculation of an infinite system using such Bloch bases (known as crystalline atomic orbitals) in electronic structure, see e.g. Ref. [243].

Systematic convergence to the correct bulk behaviour in the Bloch representation

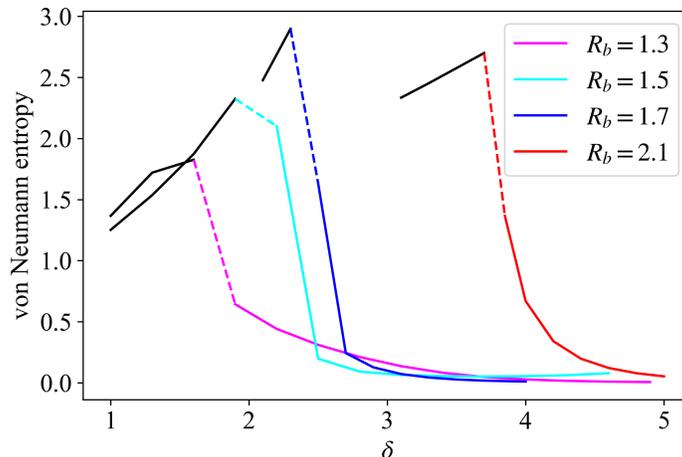


Figure A.2: Bipartite entanglement entropy of various crystalline phases as δ increases. Each line is a slice over δ values for a fixed R_b value. Black line segments denote when the ground state is in the disordered phase. Solid colored line segments denote when the ground state is an ordered crystalline phase (same color classifications as the phase diagram in Chapter 5). Dotted line segments denote the “transition zone” of a given line between the disordered phase and an ordered phase. These are simply a result of the finite resolution used to sample phase space in the phase diagrams.

is controlled by two parameters: the DMRG bond dimension and sampling of the Brillouin zone. To increase sampling of the Brillouin zone, we can introduce basis states away from the Γ -point, i.e. $|\tilde{n}_{x,y}^{\vec{k}}\rangle = \sum_l e^{-i\vec{k}\cdot R_l} |\tilde{n}_{(x,y)+R_l}\rangle$, and convergence with Brillouin zone sampling corresponds to choosing a sufficiently large numbers of vectors \vec{k} . Alternatively, we can simply increase the supercell size and work with the Γ -point functions of the increasing supercells. We choose the latter strategy in the current work. Thus, examining convergence with bond dimension and supercell size is fully sufficient to establish convergence to the thermodynamic limit.

The current 2D DMRG calculation in the Bloch representation can be formally related to calculations on a finite system as follows. Within the Bloch basis, the 1-particle components of the Rydberg Hamiltonian (Eq. 5.1) remain unaltered. The interactions are slightly transformed, becoming an infinite lattice sum over the real space lattice (Eq. 5.2). This Hamiltonian (Eq. 5.2) encodes the per supercell energy of the infinite bulk system. The infinite lattice sum can also be viewed as arising from interactions that loop around a torus infinitely many times, with the proper decaying form.

Another way to understand the 2D DMRG calculation in the Bloch basis is to examine the form of the correlation functions it predicts for an infinite system. Because the Bloch states at the Γ -point are periodic, all correlation functions are implicitly periodic across supercells. For example, transformed to the site basis, the density-density correlation function satisfies

$$\langle n_{x_1, y_1} n_{x_2, y_2} \rangle = \langle n_{x_1, y_1} n_{(x_2, y_2) + R_l} \rangle. \quad (\text{A.3})$$

Particles in adjacent supercells are thus entangled and correlated with each other, but in a highly constrained fashion. (This can be seen from the entanglement of a single particle state in the Bloch basis, which has the maximum entanglement entropy of $\log 2$ for a cut in the site basis). Note that a 2D infinite tensor network, such as an iPEPS, also introduces a constrained form of correlations between particles, but the constraint there is different and controlled solely by the bond dimension. In the 2D DMRG calculations in the Bloch basis, the full flexibility of long-range correlations is restored by increasing the Brillouin sampling, or equivalently, the supercell size.

Finite size errors

There are two types of finite size errors in the energy in the Γ -point formulation of the bulk Rydberg system. These can be associated with the Rydberg interaction energy and the (emergent) kinetic energy.

As discussed above, the use of the Γ -point basis induces a periodicity in the density-density correlation function, Eq. (A.3), and thus in the numerator of the Rydberg interaction term. This relation is exact for classical crystals and it is also exact for quantum states with such correlations (those that can be expressed exactly in the supercell Bloch basis, which obviously need not be classical crystals). However, one can imagine that such periodic correlations are inaccurate for certain quantum phases, such as the disordered phase.

As a metric for the energy error *per site* induced by the constrained form of the correlations, we compute the quantity

$$\Delta e = \frac{2 \cdot R_b^6}{\rho_{ex} \cdot \min(L_x, L_y)^6} (\langle \hat{n}_i \rangle - \langle \hat{n}_i \rangle^2). \quad (\text{A.4})$$

Here, $\langle \hat{n}_i \rangle$ is the expectation of the local Rydberg excitation for a single characteristic excited site, while ρ_{ex} is the density of sites which have the characteristic excitation of $\langle \hat{n}_i \rangle$. This is a measure of error for quantum crystals whose correlations do not

match those induced by the Bloch basis. Note that Δe is always positive, and it can be systematically reduced by increasing the supercell size.

The other source of systematic error comes from the effective itinerancy of the Rydberg atoms arising from the $\hat{\sigma}^x$ operator [150]. The error in the kinetic energy of fermionic systems when using a Bloch basis is well studied and understood to converge superalgebraically with Brillouin zone samples, in this case, the supercell size $L_x \times L_y$ (see e.g. Ref. [244] and references within). We expect a similarly rapid convergence here, although the precise quantitative effect can only be directly assessed through simulations. We have carried out such checks extensively to ensure convergence of our calculations, as discussed in Section A.1 and Fig. A.1.

Convergence and physical strategy

Despite the finite size effects discussed above, we find that we can converge our calculations to sufficiently high accuracy with reasonable bond dimensions and manageable supercell sizes. Even in the very complicated region of the phase diagram near $\delta = 5.0 - 6.0$, $R_b = 2.3$, we can distinguish the ground-state orders using a bond dimension of $D = 1200$, as shown in Fig. A.1. However, although this is enough to identify the ground state order, higher bond dimensions would be needed to capture the phase transitions with high precision; given the large region of phase space explored here, we leave such detailed calculations to future work.

The strategy used to generate the bulk phase diagram in Fig. 5.2a, as well as the truncated interaction phase diagram Fig. 5.2b, is as follows.

- For a given point in phase space (δ, R_b) , run a $D_{\max} = 1000$ simulation for all reasonable supercell sizes between 4×4 and 10×10 , as well as 12×9 .
- Identify all supercells for which the ground state has an energy per site within 10^{-2} of the lowest energy.
- If there are competing orders, ensure these solutions are all sufficiently converged by requiring (i) the largest singular value truncated during the final DMRG sweeps is less than 10^{-8} , and (ii) corrections to the energy when increasing supercell size (up to 12×9 maximally) are smaller than the energy gap between competing states (Fig. A.1).

- The ground state phase is then identified by evaluating simple density-based order parameters on the largest supercell size which hosts the ground state order.

The only time this convergence criteria is not satisfied is for disordered phase solutions near the order-disorder phase transition (largest truncated DMRG singular value is $\sim 10^{-6}$), for which all large supercells show a disordered solution. The classification of the phase in this region is supplemented by analyzing the ground state entanglement entropy, which shows a distinctive “drop” when the phase becomes ordered (see Fig. A.2).

Importantly, this strategy completely neglects possible orders with unit cells larger than 10×10 or 12×9 , as well as non-periodic solutions. Although orders with unit cells of this large size are not expected in the region of the phase diagram under investigation in this work due to the relatively high crystal densities (and thus close spatial packing) [144, 150, 151], our study cannot definitively rule out the stability of such solutions.

PEPS

The PEPS simulations in this work combine recent advances in optimizing PEPS wavefunctions using automatic differentiation [119] and 2D operator representations of long-range interactions [152]. This combination illuminated many new challenges for PEPS optimization with respect to complicated Hamiltonians. This section will detail the various challenges and the technical solutions used in this work. The instability of PEPS optimization remains an open problem and it is an area of future research to determine a PEPS optimization pipeline (using automatic differentiation) that is fully robust to problem instance. In this section, D will refer to the PEPS bond dimension and χ will refer to the maximum bond dimension allowed during contraction before approximations (via SVD) are performed.

Operator representation

The method proposed in Ref. [152] to represent Hamiltonians with long-range interactions writes the interaction potential as a sum of Gaussians,

$$\frac{1}{(\sqrt{x^2 + y^2})^6} \approx \sum_{k=1}^K c_k e^{-\lambda_k(x^2+y^2)} \equiv V_{\text{fit}}(\vec{r}). \quad (\text{A.5})$$

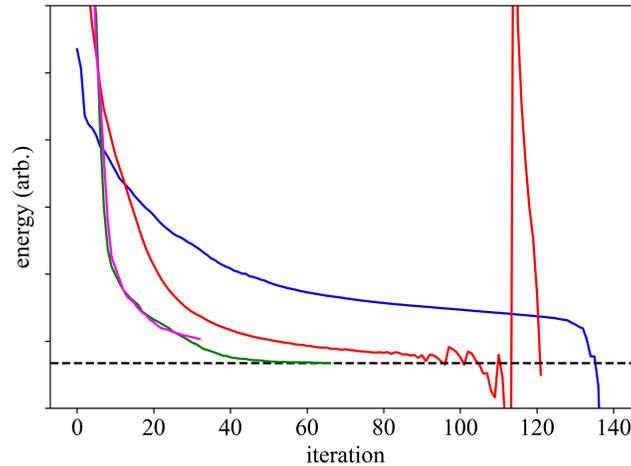


Figure A.3: Examples of typical optimization trajectories for long-range PEPS using different automatic differentiation schemes. The blue line often occurs with a naive implementation of the energy evaluation algorithms and use of a line search which does not minimize gradient norm. The red line can occur even when using a more sophisticated energy evaluation including local norms and/or a multi-evaluation cost function. The stable magenta and green lines result from combining the four techniques discussed in Sec. A.1. The difference between the magenta and green curves reflects the quality of the initial guess.

Using the methods in Ref. [245], we can obtain a $K = 7$ fit with error $\epsilon = \max_i |1/\vec{r}_i^6 - V_{\text{fit}}(\vec{r}_i)| = 10^{-5}$ on the domain $\vec{r} \in [1, 16\sqrt{2}]$, which is used throughout the work.

Essential computational techniques

As originally discussed in Ref. [119], when trying to use automatic differentiation to optimize a PEPS there are a few essential techniques that must be employed, which are not typically “default” in standard automatic differentiation libraries. They are “essential” in the sense that without them the computation of the energy expectation value and its derivative will typically not run to completion due to out-of-memory errors or numerical infinities. These techniques are:

- Numerical stabilization of the gradient of SVD, by adding Lorentzian broadening to the inverse singular values.
- Significant usage of “checkpointing” when evaluating the energy to reduce the memory load of computing gradients.

Both of these techniques are explained in significant detail in Ref. [119].

Stabilizing the optimization

A straightforward implementation of the energy expectation value as described in [152], with optimization via automatic differentiation including the above techniques, typically fails to find the ground state PEPS for the Rydberg Hamiltonian (see Fig. A.3). This failure can be generally attributed to the fact that in the quantity under optimization $E = \frac{\langle \psi | H | \psi \rangle}{\langle \psi | \psi \rangle}$, both the numerator and denominator are evaluated approximately and thus the computation is not strictly bound by the variational principle. Consequently, the optimization can find pathological regions of the PEPS parameter values which make the PEPS contractions inaccurate for the chosen χ , even when starting from an accurately contractible PEPS. Unfortunately, in this problem we find that simply raising the value of χ does not prevent this behavior until χ is impractically large.

In order to mitigate this problem we use the following four techniques in tandem:

- We employ line search methods that minimize the gradient norm as well as the energy. In this work, we use the BFGS algorithm [246] in conjunction with such a line search, as suggested in [119].
- We use the cost function $E_1/2 + E_2/2 + \lambda|E_2 - E_1|$ where E_1 and E_2 are the energies of PEPS on lattices rotated by 180 degrees and λ is a penalty factor. This strongly penalizes the optimization from entering parameter space with large contraction error (where E_1 and E_2 would be very different).
- During the first iterations of the gradient optimization we only update small patches of tensors at a time, which are chosen to break spatial symmetries that may be contained in the initial guess. After this has pushed the optimization towards the symmetries of the true ground state order, then all tensors can be updated at each optimization step.
- We evaluate the numerator and denominator of E in a consistent way by using “local normalization” during the computation of $\langle \psi | H | \psi \rangle$. This means that, writing H as a comb tensor sum $H = \sum_{i=1}^{L_x} h_i$, then for each comb tensor numerator $\langle \psi | h_i | \psi \rangle$, the associated denominator uses the identical contraction, but with h_i replaced by the identity (the environments are not recomputed).

Combining all four of these techniques removes the most egregious instabilities in the optimization trajectory (see Fig. A.3), at the cost of a slightly larger computational burden. However, as in more standard DMRG calculations with small bond

dimension, convergence to the correct ground-state (rather than a local minimum) still requires a reasonable initial guess.

Initial guess

Obtaining an accurate ground state PEPS typically relies on starting with an accurate initial guess. The predominant algorithms to generate such a guess for problems with a local Hamiltonian are simple update [80, 247, 248] or imaginary time projection of a converged small D solution to a larger D guess. However, in the presence of long-range interactions it becomes challenging to generalize either of these methods in an efficient and/or accurate way. We therefore used the following simple scheme to generate initial guesses in this work.

- Sum n manually constructed $D = 1$ PEPS to obtain an initial PEPS of bond dimension $D = n$. The configurations of these $D = 1$ PEPS were set to reproduce specific low energy Rydberg crystals and defects within them.
- For small R_b : truncate the long-range interactions in H to next-nearest, or next-next-nearest, neighbor interactions (distance of $\sqrt{2}$ or 2), and run conventional simple update starting from the above manually summed PEPS. This fails once the ground state excitations are spaced by more than 2.
- For large R_b : add positive random noise to the manually summed PEPS, and then run a highly approximate, first-order gradient optimization for ~ 25 iterations using a large step size when updating the parameters.

Convergence and physical strategy

Despite the simple procedure to generate initial guesses, we were usually able to systematically converge PEPS solutions according to the conventional protocol of increasing D and χ until the energies corresponding to multiple increasing (D, χ) pairs all vary by less than 0.01% relative to each other (e.g. see Fig. A.5). In this study, we used maximal values of $D = 5$, $\chi = 100$. However, for a small number of phase points (δ, R_b) we encountered inconsistent convergence of PEPS solutions (see Fig. A.4), where increasing D and χ did not systematically result in finding a PEPS with a lower energy, instead getting stuck in various local minima. We attribute this to the low quality of the initial guesses for larger D and R_b .

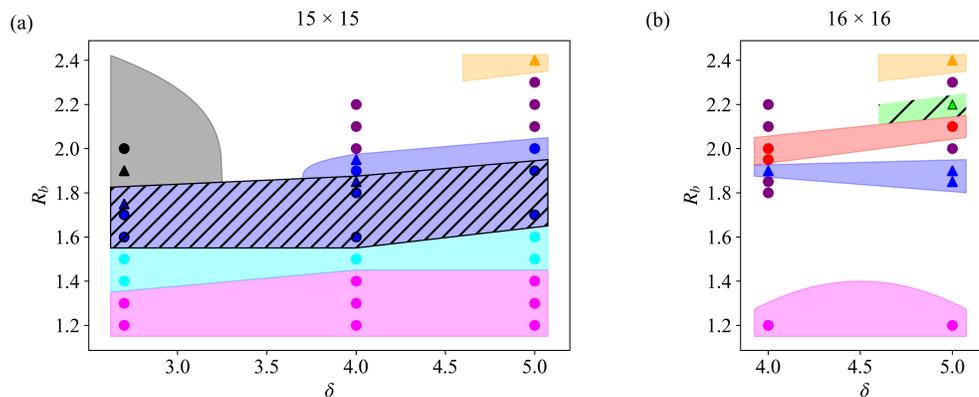


Figure A.4: Phase diagrams of 15×15 (a) and 16×16 (b) arrays, detailing convergence. Circular points indicate systematic convergence with PEPS up to bond dimension $D = 5$, while triangles indicate intermittent convergence with PEPS, requiring supplemental convergence checks using 2D DMRG. More details are available in Sections A.5 and A.1. Colors used in these plots correspond identically to the colors used in Chapter 5 to identify phases.

In these cases when PEPS *energies* could not be systematically converged to within 0.01%, the observed *order* of the various low-energy solutions was nonetheless the same. The differing energies arose due small quantitative differences such as single-site defects and variations in the local density $\langle \hat{n}_i \rangle$. To further increase certainty in the observed order, we also compared the PEPS solutions to the results of 2D DMRG on the same finite lattice, since the convergence properties of DMRG are much more well-understood. In all cases, the low-energy PEPS solutions had similar energies to the approximate DMRG (relative difference $< 1\%$), and they all showed the same generic low-energy ground state order. The energy gap between phases appeared to be sufficiently large to allow for a tentative classification of the order of this small number of phase points, even though the DMRG was not necessarily converged to high precision (due to the wide lattices) and the PEPS convergence could not be definitively confirmed. The uncertainty in convergence highlights remaining challenges in simulating complex large 2D interacting problems with competing phases using tensor network techniques. The relevant points in the finite lattice phase diagram are labelled by triangles in Fig. A.4 above, and in Fig. 5.4.

Finite 2D DMRG

Standard 2D DMRG calculations with open boundaries were used to study the 9×9 system, a low-entanglement region of the 13×13 system, and to supplement

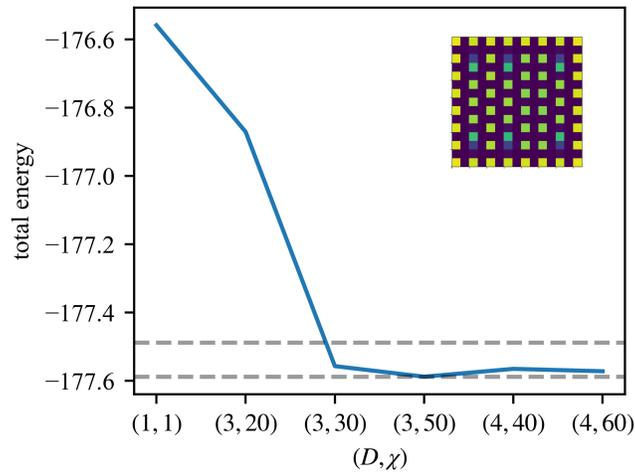


Figure A.5: An example of systematic convergence of PEPS on the 15×15 lattice for the frustrated star phase at $\delta = 4.0$, $R_b = 1.9$. The region between the horizontal lines indicates a change in energy of 0.01% relative to the lowest obtained value. The PEPS is deemed converged because many simulations with increasing (D, χ) return energies that fall within this region. Note that the star phase, like most ordered phases, is sufficiently converged by $D = 3$ due to the predominant mean-field character of most ordered phases (discussed in Chapter 5).

convergence of PEPS on the larger 15×14 , 15×15 , and 16×16 lattices. Like the PEPS calculations, these too included all long-range interactions (according to Eq. 5.1). The maximal bond dimension used for the 9×9 and 13×13 simulations was $D_{\max} = 1200$, which we found was more than enough to accurately study the regions of interest in Fig. 5.5 for these lattices (see Fig. A.6). For supplementing PEPS convergence on the larger lattices, we used $D_{\max} = 750$. Although this bond dimension is not large enough to capture the ground state energy or entanglement of such large systems with high precision, we found it sufficient to capture the first 3-4 digits of the ground state energy and to help with distinguishing between the different low-entanglement ordered phases present in the finite phase diagram, which have substantially larger gaps than the bulk system due to edge effects.

Mean field and classical

The mean field phase diagram for the bulk system (including all long-range interactions) in Fig. 5.2d was generated by the following procedure.

- Parameterize the single site wavefunction as $|\phi_i\rangle = \sin^2(\theta_i)|0\rangle + \cos^2(\theta_i)|1\rangle$, where $|0\rangle$ is the atomic ground state and $|1\rangle$ is the excited Rydberg state.

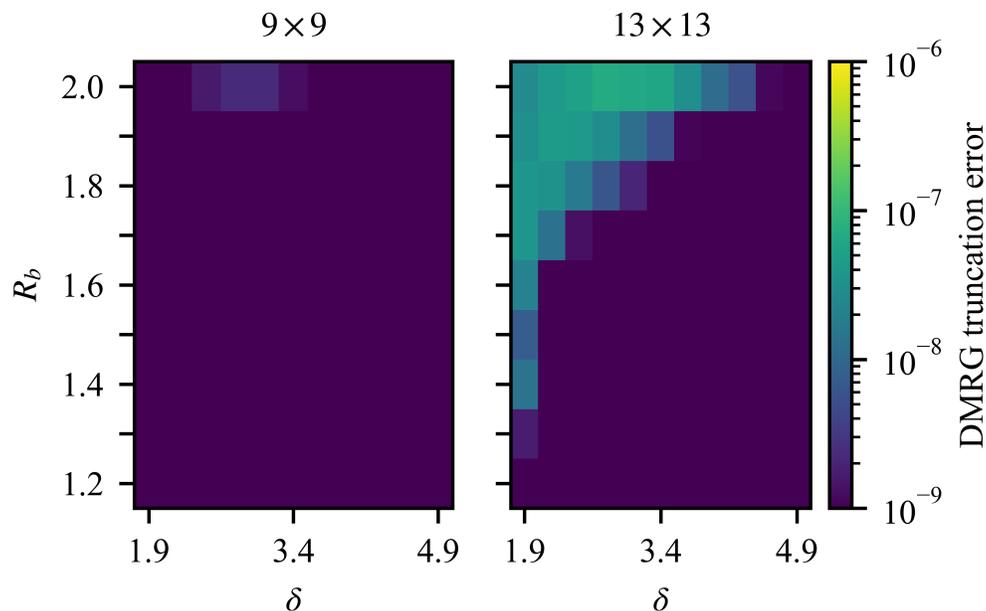


Figure A.6: Accuracy of 2D DMRG on the 9×9 and 13×13 finite lattices (open boundaries). The displayed regions of parameter space correspond exactly to the computed regions in Fig. 5.5 and Fig. A.12. The reported error is the largest truncated singular value during the final DMRG sweeps (i.e. once converged). Note that in the ordered regions the error is $\sim 10^{-9}$, and it grows to $\sim 10^{-7}$ as the ground state becomes disordered on the 13×13 lattice due to increasing entanglement.

- Construct a completely un-entangled many-body wavefunction as a typical product of these single-site states according to all reasonable unit cells between size 2×2 and 8×10 (supercells are not necessary for mean-field convergence).
- Initialize all possibly relevant configurations for each unit cell as initial guesses. These can be obtained from classical algebraic arguments or classical Monte Carlo.
- Minimize the Γ -point energy for all guesses with respect to the $\{\theta_i\}$ using gradient descent. Analytic gradients are easily derived, or automatic differentiation can be employed.
- Classify the phase of the lowest energy state using the same density-based order parameters as the Γ -point DMRG calculations.

The phase space was scanned with a δ -resolution of 0.1 and a R_b -resolution of 0.025. Importantly, these calculations are subject to the same limitation as the

Γ -point DMRG - they do not capture any possible low energy states with a unit cell larger than 8×10 . Although such states are not expected in the phase space under examination, this study cannot definitively rule them out.

The classical phase diagram for the bulk system (including all long-range interactions) in Fig. 5.2c was generated by the following procedure.

- Run classical Monte Carlo minimization of the Γ -point energy for every unit cell size between 2×2 and 10×10 at phase space points spaced by $\Delta\delta = 0.3$, $\Delta R_b = 0.1$.
- For all low energy configurations obtained at all phase points, derive their continuous functional form $E(\delta, R_b)$ by numerically integrating the interactions.
- Analytically solve for the intersection line between each adjacent pair of configurations in phase space that have minimal energy.

These calculations are also subject to the same limitation as above—any states with unit cells larger than 10×10 are not captured, and we cannot rule out their possible existence.

A.2 Entangled nematic phase

In this section we will focus on the nature of the nematic phase and its stability in the thermodynamic limit (TDL). The qualitatively mean-field nature of all other ordered phases permits a variety of straightforward arguments for their TDL stability.

The structure of the ground state nematic wavefunction in Fig. 5.3 suggests that itinerancy of Rydberg excitations plays an important role in stabilizing this state. Since itinerancy of excitations and defects emerges perturbatively [150], we will rewrite the 2D Hamiltonian like,

$$\begin{aligned} \hat{H} &= \hat{H}_D + \lambda \hat{H}_Q \\ &= \frac{1}{2} \sum_{i \neq j} \frac{R_b^6}{|\vec{r}_i - \vec{r}_j|^6} \hat{n}_i \hat{n}_j - \sum_i \delta \hat{n}_i \\ &\quad + \lambda \sum_i \hat{\sigma}_i^x. \end{aligned}$$

Here the eigenstates of \hat{H}_D are classical crystals, while $\lambda \equiv \frac{\Omega}{2} = \frac{1}{2}$. To investigate the energy scales involved in the itinerant processes, we can begin by performing non-

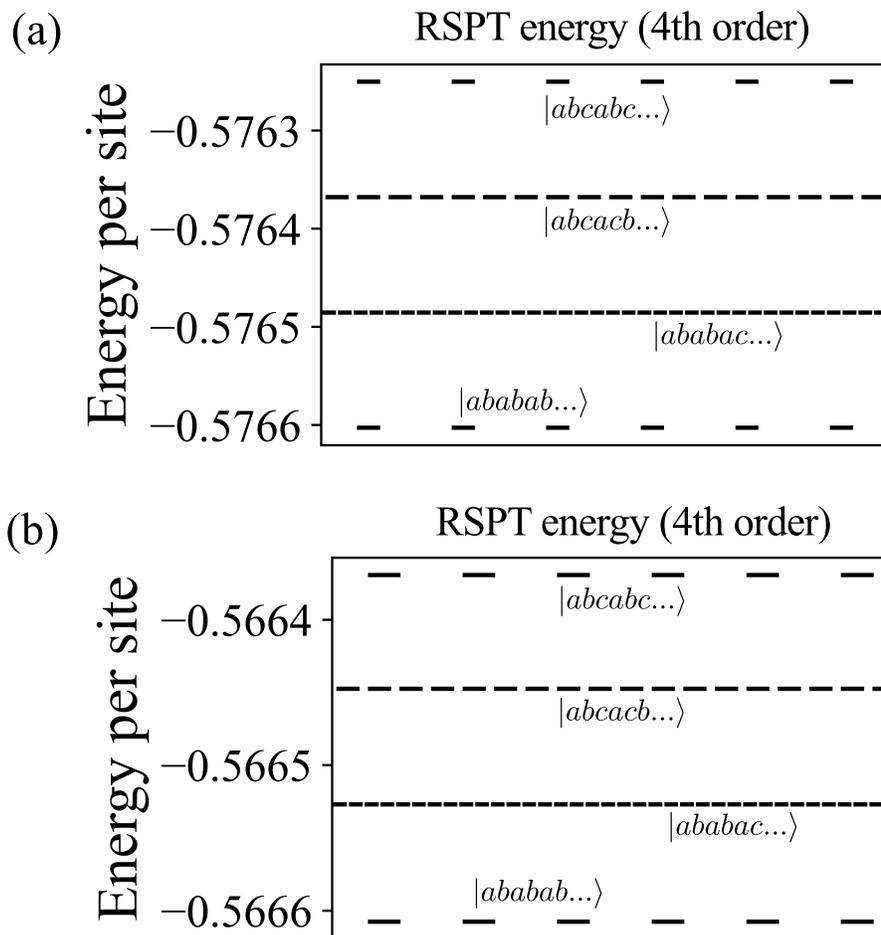


Figure A.7: Energies of various low-energy classical crystals that are corrected up to 4th-order Rayleigh-Schrodinger perturbation theory (RSPT), which includes effects of single excitation itinerancy. Results are shown for (a) interactions up to nearest neighbor columns ($|x_i - x_j| = 2$ in the 2D lattice), and (b) full long-range interactions. Comparing to the 0th-order classical energies with long-range interactions (Fig. 5.3c), we see that the classically unfavorable states (due to longer-range terms) are stabilized by the perturbations. The energy scale of single-excitation itinerancy is therefore larger than the long-range terms.

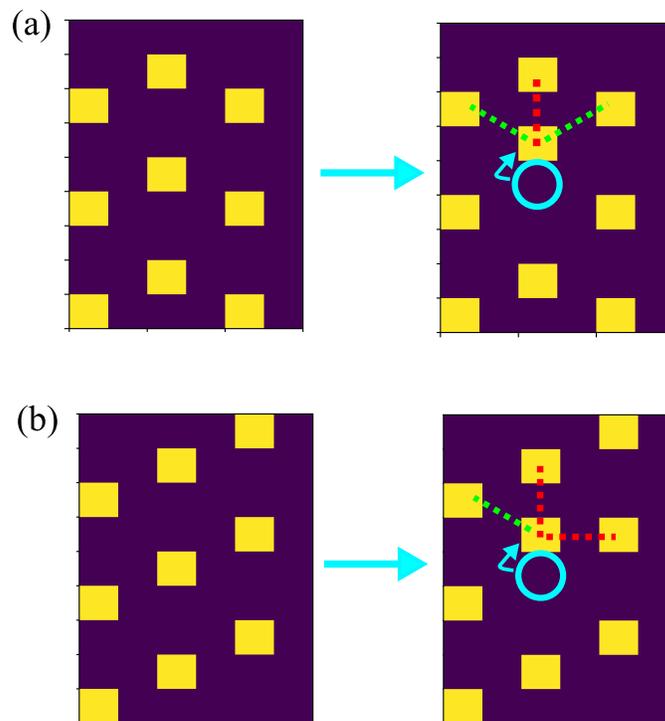


Figure A.8: Example of a perturbative hopping process for a single excitation that emerges in the second-order Rayleigh-Schrodinger correction to a classical initial wavefunction (fourth order energy). In (a) a single excitation in an $|ababab\rangle$ -type state can hop while only violating a single Rydberg blockade constraint (denoted by a red line). In (b), a single excitation in an $|abcabc\rangle$ -type state must violate two blockade constraints when hopping. These quantum fluctuations generate entanglement that preferentially stabilizes the $|ababab\rangle$ -type states, despite their higher classical and mean-field energy.

degenerate Rayleigh-Schrodinger perturbation theory (RSPT) starting from different initial column state crystals such as $|ababab\dots\rangle$, $|abcabc\dots\rangle$, etc (see Fig. 5.3).

Since the initial states that diagonalize \hat{H}_D are classical, with all sites having exactly $\langle \hat{n}_i \rangle = 0$ or 1, the first-order RSPT correction to the wavefunction (second-order energy) allows local superpositions of $|0_i\rangle$ and $|1_i\rangle$. This order of RSPT captures the energies of the mean-field ordered phases with high accuracy. The second-order RSPT correction to the wavefunction (4th-order energy) allows for effective “hopping” of a single excitation or defect from site i to another site j . The first- and third- order corrections to the energy are zero when starting from a single column state crystal. In Fig. A.7 we compute the energies up to the 4th order RSPT correction for various starting states, using all long-range interactions as well as

interactions truncated to only be between neighboring columns ($|x_i - x_j| \leq 2$ on the 2D lattice; in this limit the classical energies of all low-energy column states are exactly degenerate). The states of type $|ababab\dots\rangle$ have the lowest RSPT energies in both interaction schemes. At the level of single-particle hopping, this is because hops within a single column from excited site y to site $y \pm 1$ can be chosen to only violate 1 Rydberg blockade ($R_b \sim 2.3$) constraint (the new excitation is 2 sites away from a single excitation within its column, but still at least $\sqrt{5}$ away from all excitations in other columns). On the other hand, states with more $|abc\dots\rangle$ character have some single excitation hops which must violate two instances of the blockade (they can only hop to a position at $y \pm 1$ which is 2 sites from an excitation in its own column and 2 sites from an excitation in the adjacent column). See Fig. A.8.

By comparing the results (Fig. A.7) with truncated interactions and full interactions, we can see that the energetic contributions of the quantum fluctuation induced itinerancy are larger than those of the long-range interactions, since the interaction favors $|abc\dots\rangle$ states while itinerancy favors $|abab\dots\rangle$ states. These fluctuations create entanglement, because the in-column direction of low-energy hopping is highly dependent on the state of the adjacent columns. This suggests that the ground state should be populated by $|abab\dots\rangle$ states instead of the $|abcabc\dots\rangle$ mean-field ground state, which is what we observe in the nematic ground state obtained from DMRG (see Fig. 5.3). This analysis only requires a low-order of perturbation theory to stabilize the nematic quantum crystal $|ababab\dots\rangle$, regardless of the system size, which provides clear evidence for its stability in the TDL.

However, if we take this non-degenerate RSPT style of analysis for each separate column state crystal to its extreme, we will notice that it diverges at high orders. Specifically, when the perturbation series reaches an order comparable to the number of excitations in a single column, the itinerant perturbation term connects different column states such as $|ababab\dots\rangle$ and $|acabab\dots\rangle$. These classical configurations are quasi-degenerate (exactly degenerate with truncated interactions) and the energy differences in the denominators go to zero. Unlike a typical cat state, where this degeneracy is a small constant, the present case accesses a manifold of exactly- and quasi-degenerate states that grows exponentially in the number of columns (see next subsection). This prevents a straightforward application of degenerate perturbation theory and hints at the possibility of a more exotic, non-perturbative entangled order.

In the results reported for the 12×9 supercell in Fig. 5.3, we do in fact observe

non-trivial coupling between the exponentially large manifold of low-energy column states. It is clear the structure of the wavefunction (Fig. 3c) does not correspond to the single dressed $|ababab\dots\rangle$ state discussed earlier, since the wavefunction coefficients are distributed across the 6-fold permutations of the $|abab\dots\rangle$ configurations as well as the exponential manifold of other states around them. The structure of the entanglement spectrum between columns further supports the presence of a non-trivial (i.e. non-cat state), macroscopic entangled order (see next subsection) *for the solution on the 12×9 supercell.*

The stability of this order in the TDL is a much more complicated question than the single nematic quantum crystal discussed above. As we mentioned, a perturbative understanding will fail us. In a straightforward perturbative treatment (degenerate or non-degenerate) of the low-energy classical states, a size-extensive order of the series is needed to see coupling between the different quasi-degenerate states. For the finite supercells we are able to study, the energetic benefit of coupling the exponential classical manifold exceeds the finite cost to do so (due to finite effective column height). However, in the thermodynamic limit the coupling may go to zero due to the size-extensive order. The true TDL stability is not clear from this picture. Numerically, a detailed finite-size scaling analysis is needed to resolve the fate of this macroscopically entangled order, but this is beyond the scope of our current calculations, and perhaps beyond current methods in general due to the small energy scales involved. Analytically, we note that the itinerant-crystalline nature of this order, along with its non-perturbative behavior, is similar to the incommensurate floating solid phase that emerges near the order-disorder transition in 1D Rydberg chains [150, 151]. Future efforts may seek to understand the nematic state through that lens. Nevertheless, it is undoubtedly clear from our analysis that fluctuations and entanglement are much larger in this nematic phase than in the other ordered phases.

D=3 low-energy projectors and the entanglement spectrum

The character of the nematic phase has been extensively discussed in terms of the classical configurations that make up the quantum wavefunction. It has been pointed out that all the low-energy (and thus the most relevant) classical configurations can be described in a succinct notation like $|abcabc\dots\rangle$ in terms of compositions of 3 individual column states $|a\rangle$, $|b\rangle$, and $|c\rangle$ which are defined in Fig. 5.3. This notation is very suggestive of the idea that a **qualitative** model for the 2D state can be written as a 1D MPS with a local Hilbert space of dimension 3, spanning $|a\rangle$, $|b\rangle$,

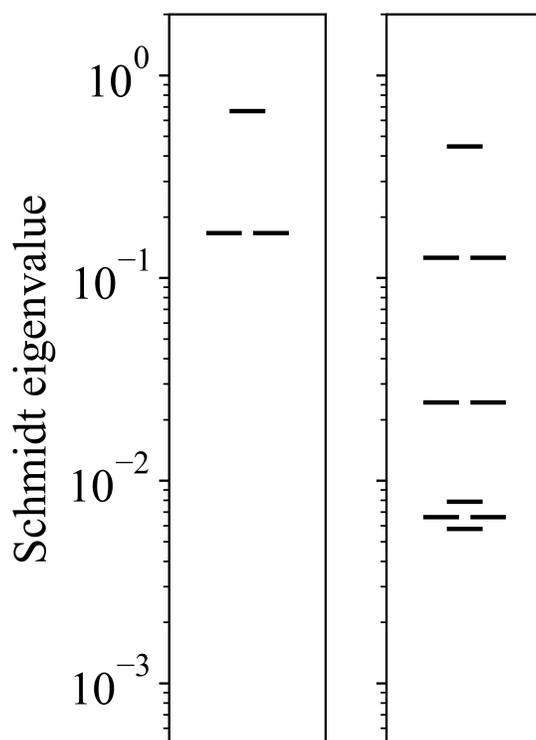


Figure A.9: Structure of the entanglement spectrum for the 1D low-energy projector model of the nematic order, as described in Section A.2. *Left*: The entanglement structure when projectors are applied on every bond of the MPS, in an open-boundary conditions style. *Right*: Entanglement structure when an additional periodic projector is applied between the first and last site in the MPS. This type of interaction is generally accounted for in the Γ -point basis calculations. Note the similarity of these Schmidt spectra to the nematic DMRG ground state shown in Fig. 5.3.

and $|c\rangle$.

A striking feature of the entanglement spectrum presented in Chapter 5 is the presence of 3 large eigenvalues, with a 1 : 2 degeneracy structure. To illustrate a possible origin of this pattern in the nematic phase, we create a simple model state with a similar entanglement spectrum.

The strength of the interactions where the nematic phase emerges ($R_b = 2.3$) is such that configurations with adjacent columns in the same state (e.g. $|abcabb\dots\rangle$) are much higher in energy than all configurations without any identical adjacent columns. The projector into the low energy subspace thus removes all configurations with adjacent columns in the same state. It can be written as a product of commuting

two-site operators,

$$\mathcal{P} = \prod_i \hat{P}_{i,i+1}, \quad (\text{A.6})$$

where $\hat{P}_{i,i+1} = \mathbb{1} - |aa\rangle\langle aa| - |bb\rangle\langle bb| - |cc\rangle\langle cc|$.

First, we will briefly comment on the eigenspectrum of \mathcal{P} . It's eigenvalues are all positive integers, with the smallest being 0. The number of 0 eigenvalues grows as $\sim 2^L$, where L is the length of the 1D chain. These eigenvalues correspond to all the possible arrangements of the individual column states that do not violate any constraints in Eq. (A.6). This reveals the origin of the exponential classical degeneracy that has been previously discussed in the 2D system.

Concerning the entanglement spectrum, we can apply the operator \mathcal{P} to a simple product state $|\psi_0\rangle$ (a $D = 1$ MPS) containing an equal mixture of all possible column states i.e. $|\psi_0\rangle = \prod_i (\lambda_a |a_i\rangle + \lambda_b |b_i\rangle + \lambda_c |c_i\rangle)$ where $|a_i\rangle, |b_i\rangle, |c_i\rangle$ represent states on column i . As long as $|\lambda_a| = |\lambda_b| = |\lambda_c|$, then $\mathcal{P}|\psi_0\rangle$ has the entanglement structure shown in Fig. A.9, which is very similar to that seen in the 2D nematic phase computed by DMRG.

A.3 Bulk phase diagram degenerate region

In Chapter 5 it was briefly mentioned that there is a small region of the bulk phase diagram where the nematic phase and 3-star phase become essentially degenerate. By this we mean that their gap becomes too small to resolve within the estimated finite size error in the Γ -point DMRG numerics. Using the Δe finite size error measure defined above, for $[R_b = 2.3, L_x = 12, L_y = 9]$, we have $\Delta e \approx 3 \cdot 10^{-5}$ in the nematic phase and $\Delta e \lesssim 8 \cdot 10^{-6}$ in the 3-star phase. An expanded view (in δ) of the upper part of the bulk phase diagram is shown in Fig. A.10. The degenerate region emerges between the nematic phase and the $\frac{1}{5}$ -staggered phase near $\delta = 7.0$, as indicated by the lime green color.

A.4 Bulk phase transitions

The disorder \rightarrow order phase transitions that occur throughout the bulk phase diagram have been characterized as continuous phase transitions in previous work [144]. Although full, precise characterization of all bulk phase transitions is beyond the scope of this work, we are able to estimate the order of some transitions using straightforward numerical differentiation of the energies. Figure A.11 shows the first and second derivatives of the energy as a function of δ , for various values of R_b , from simulations on an 8×8 supercell. The peaks in the second derivatives near

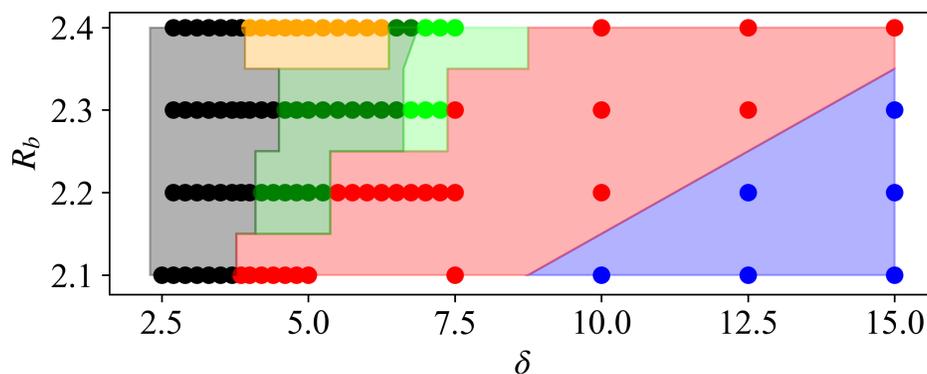


Figure A.10: Expanded view of the large- R_b part of the bulk phase diagram, computed with Γ -point DMRG. For $\delta \leq 5.0$, this data is identical to Fig. 5.2a. All colors correspond to the same phases as in the Fig. 5.2. The small lime green region indicates the degenerate zone where the gap between the 3-star and nematic phases becomes very small.

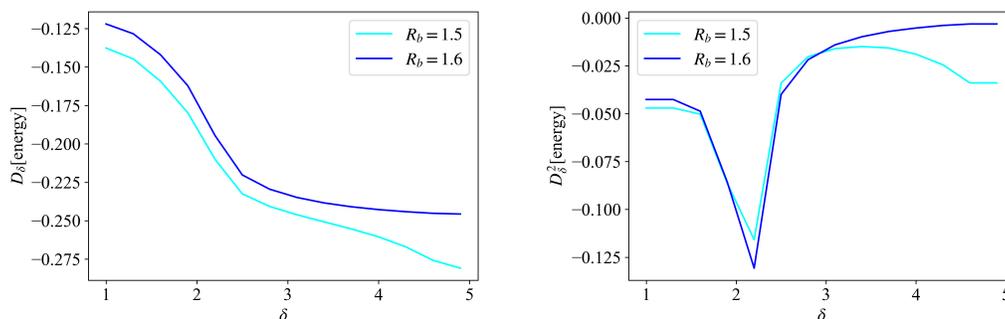


Figure A.11: Numerical evidence of second-order phase transitions between the disordered phase and the star (blue) and striated (cyan) phases, from simulations on an 8×8 supercell. *Top*: First derivative of the energy with respect to δ . *Bottom*: Second derivative of the energy with respect to δ . Both are estimated using standard finite difference formulas.

the critical values of δ tentatively support the conclusions that the disorder \rightarrow star and disorder \rightarrow striated phase transitions are indeed second-order. However, this data is also consistent with very recent work [153] which shows that although these two transitions appear continuous using an 8×8 simulation cell, they adopt strong characteristics of first-order transitions once the cell size is increased to 16×16 .

A.5 Finite phase diagram: 15×15 and 16×16

The phase diagram of the 15×15 lattice reported in Chapter 5 contained many of the ground state orders seen in the bulk phase diagram, but it also revealed the strong finite-size effects induced by the boundary. Due to the long-range van der Waals interactions, Rydberg excitations at the edge of the array incur roughly half of the energetic penalty that excitations in the interior do, but lower the energy by an equal amount (δ). Except at small values of R_b , this induces excitations along the edge of the array to be more densely packed than what would be expected from the bulk phase diagram at a given point (δ, R_b) . This generic effect causes frustration between the boundary and interior of the finite lattices, which gives rise to the square classical order and many defect-dominated states at large R_b , as discussed in Chapter 5. In these defect states, the optimal bulk density becomes so small relative to the optimal edge density that the ground states are permeated by edge-induced defects, leaving only small regions of any discernible order and making the precise configuration very sensitive to small changes in R_b and δ .

In addition to the 15×15 lattice, we also studied two slices ($\delta = 4.0, 5.0$) of the phase diagram of the 16×16 lattice to probe for bulk-like ordered phases where the 15×15 system is dominated by defects. Specifically, we focused on the $R_b > 1.8$ region, for which the results are shown in Fig. A.4b. We find a clear region of the stability for the boundary-bulk frustrated $\frac{1}{5}$ -stagger phase (red), for which the density profile is shown in Fig. 5.4c. Along with a small region of the 3-star phase (green and black), these regions are unique to the 16×16 lattice (i.e. they are not seen in 15×15). There are also some common features between the two array sizes, namely regions of the star and $\frac{1}{8}$ -stagger (gold) phase as well as many defect states. This suggests that the defect states are an intrinsic part of the physics of medium-sized arrays.

As reflected by the triangular markers in Fig. A.4b (which reflect inconsistent convergence) we found it more challenging than the 15×15 lattice to systematically converge the PEPS calculations with respect to (D, χ) , especially in the star phase (blue). In part, this was due to the boundary itself being frustrated; on an even-sided lattice it is not possible to place excitations in all corners and also along all edges spaced by a distance of 2. Because the corner excitations are strongly pinned due to their reduced interaction penalty, this causes the boundary to be frustrated and makes it more difficult to prepare a good initial guess with our rudimentary strategies.

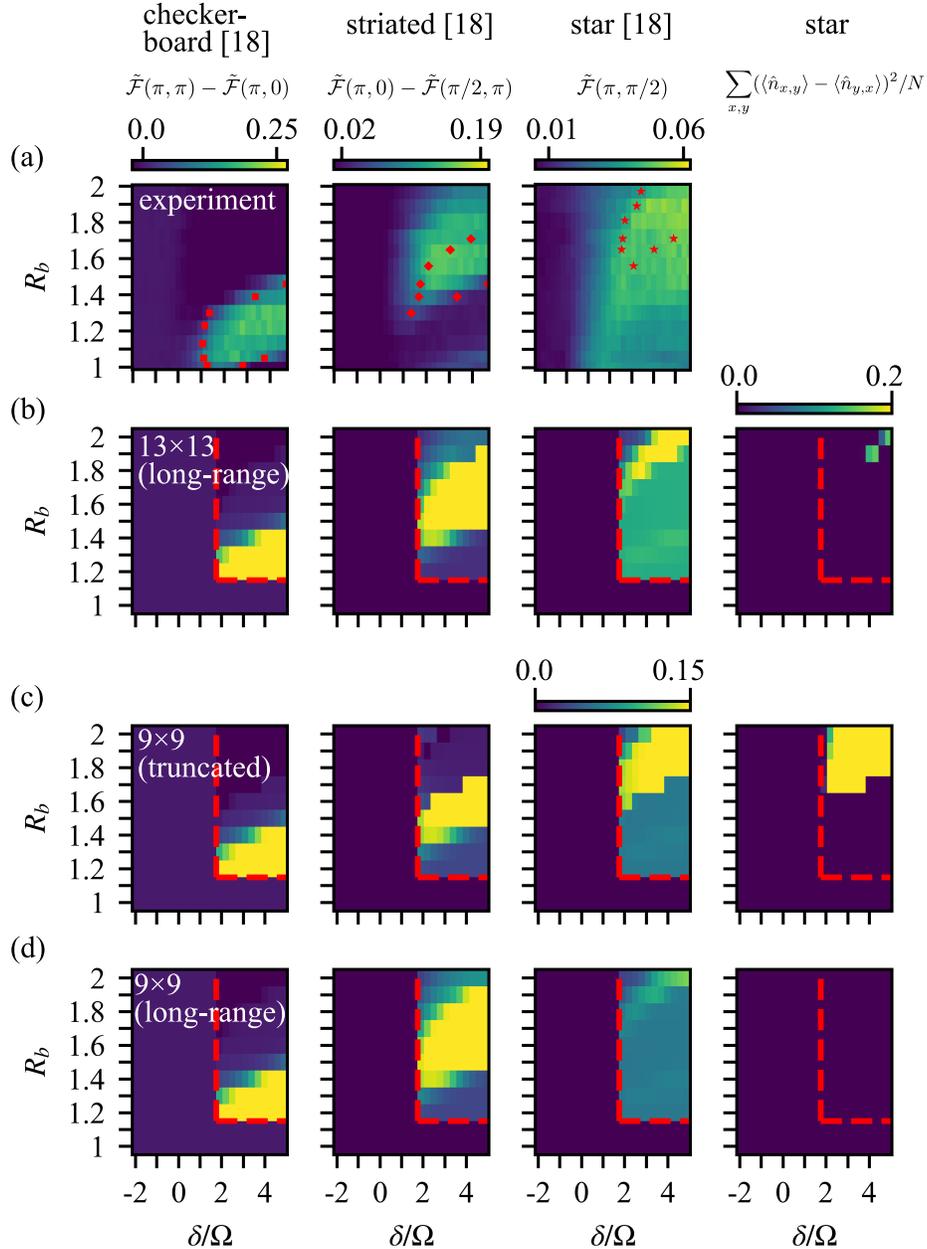


Figure A.12: Detailed comparison to experimental phase diagram. The (a) row directly reproduces the experimental phase diagram on the 13×13 lattice (data extracted from Ref. [143] Fig. 4). Rows (b)-(d) show analogous numerical data on 9×9 and 13×13 lattices, where (b) and (d) are results from simulations containing all long-range interactions and (c) shows results using interactions truncated to zero beyond distance 2. This is identical to the truncation scheme used in numerics in Ref. [143]. The first three columns show all three order parameters used in [143] to distinguish the phase diagram, while the fourth column shows a new, more precise order parameter for the star phase. Red dots in (a) denote the phase boundaries assigned in [143], while the cyan dotted lines in (b)-(d) indicate the subset of parameter space that was computed.

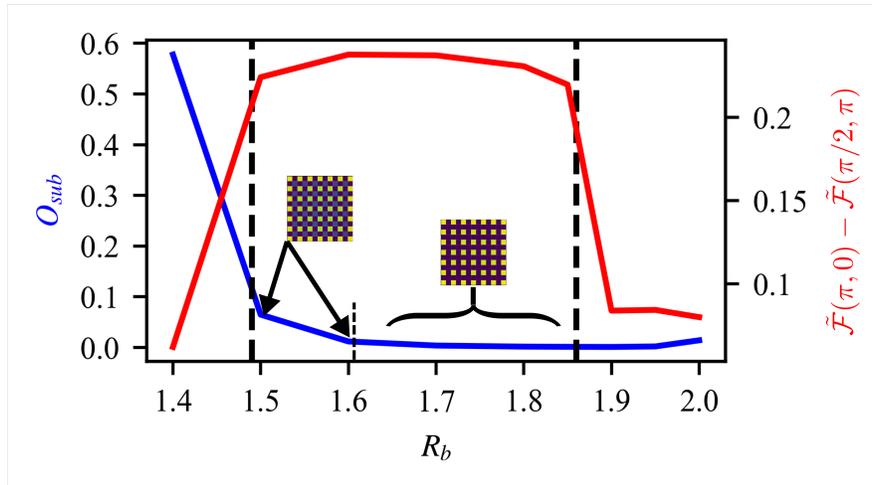


Figure A.13: Distinguishing the striated and square orders on the 13×13 lattice at the slice $\delta = 4.0$. The striated order parameter $\tilde{\mathcal{F}}(\pi, 0) - \tilde{\mathcal{F}}(\pi/2, \pi)$ (red) is large across the range $R_b = 1.5 - 1.85$, but the density of quantum fluctuations on the $(1, 1)$ -sublattice O_{sub} (blue) decays to ~ 0 by $R_b = 1.65$, revealing the square order.

A.6 Comparing to experiment: 9×9 and 13×13 lattices

Chapter 5 discussed discrepancies between our numerical results on the 13×13 lattice and analysis reported in a recent experiment [143], specifically concerning the striated, square and star phases. It was noted that the actual experimental data appears to agree with our numerics, but the interpretation of the data offered in Ref. [143] is inconsistent with ours. This section details the effect of the approximations made in the numerics of Ref. [143] on the interpretation of the data, and how relaxing those approximation leads to the interpretation described in our work (Chapter 5).

Context

In Ref. [143], the experimental data on the 13×13 square lattice was primarily understood with respect to DMRG calculations performed on the 9×9 lattice (all open boundaries), in which interactions were truncated to zero beyond a distance of 2. The experimental results of Ref. [143] are reproduced in Fig. A.12a, and they are compared to our numerical results on 9×9 and 13×13 lattices (Fig. A.12b-d). The region of the phase diagram that was studied included domains of stability for the disordered, checkerboard, striated, and star phases. The square “phase” was not separately reported, although it may be considered the classical limit of the striated phase.

We also introduce here a useful order parameter for detecting the star phase,

$$O_{star} = \sum_{x,y} (\langle \hat{n}_{x,y} \rangle - \langle \hat{n}_{y,x} \rangle)^2 / N, \quad (\text{A.7})$$

where $N = L_x \cdot L_y$. O_{star} detects a symmetry breaking that occurs in the star phase but not in the disordered, checkerboard, striated, or square phases. On a finite lattice, this provides a clean way to define the star phase separate from the other orders in this set. We also recapitulate the definition of the order parameters defined in [143] and used in Fig. A.12,

$$\tilde{\mathcal{F}}(k_1, k_2) = (\mathcal{F}(k_1, k_2) + \mathcal{F}(k_2, k_1)) / 2 \quad (\text{A.8})$$

$$\mathcal{F}(k_1, k_2) = \left| \sum_{x,y} \exp(i(k_1x + k_2y)) \langle \hat{n}_{x,y} \rangle \right| / N. \quad (\text{A.9})$$

Star phase stability

In Fig. A.12c, we recompute the main 9×9 phase diagram numerical results used in [143], which use truncated interactions. The bright region in $\tilde{\mathcal{F}}(\pi, \pi/2)$ predicts a large domain of stability for the star phase, which is corroborated by the value of O_{star} . This data was used in [143] to draw the expected phase boundary in the 13×13 experimental data seen in Fig A.12a. However, Fig. A.12d shows the analogous results on the 9×9 lattice when including all long-range interactions. Surprisingly, the star phase gets completely destabilized! This illustrates the hazard of interpreting the experimental data from smaller lattice simulations.

Unlike the 9×9 lattice, we observe that the 13×13 lattice phase diagram has a qualitative difference: it hosts a nonzero domain of star phase even when accounting for all long-range interactions. As pointed out in Chapter 5, $\tilde{\mathcal{F}}(\pi, \pi/2)$ is not a sensitive order parameter for the star phase as it appears on finite lattices, but O_{star} does reveal the tiny stable region of the star phase (see Fig. A.12b).

Square and striated phases

The overestimation of the extent of the star phase by using numerics from the 9×9 lattice with truncated interactions also results in an underestimation of the extent of the striated order parameter, $\tilde{\mathcal{F}}(\pi, 0) - \tilde{\mathcal{F}}(\pi/2, \pi)$, since $\tilde{\mathcal{F}}(\pi/2, \pi)$ is the star order parameter used in [143] (see Fig A.12c). These 9×9 results were used in [143] to interpret the striated phase domain in the experimental data, so the boundary drawn in Fig. A.12a is too small. In fact, the extent of the experimental data for $\tilde{\mathcal{F}}(\pi, 0) - \tilde{\mathcal{F}}(\pi/2, \pi)$ (Fig. A.12a) is significantly larger than the drawn boundary,

corresponding much more closely to the numerical data on the 13×13 including long-range interactions (Fig. A.12b), as mentioned in Chapter 5.

In this work, we distinguish a region of classical square order from the striated phase where the square order contains (almost) no quantum fluctuations on the $(1, 1)$ -sublattice, which are an essential feature of the striated phase in the bulk. $\tilde{\mathcal{F}}(\pi, 0) - \tilde{\mathcal{F}}(\pi/2, \pi)$ does not help distinguish between square and striated orders, and no classical square order was reported in Ref. [143]. In Fig. A.13 we show that a large part of the bright region in $\tilde{\mathcal{F}}(\pi, 0) - \tilde{\mathcal{F}}(\pi/2, \pi)$ on the 13×13 lattice should be interpreted as a classical square order by plotting,

$$O_{sub} = \begin{cases} \frac{4}{N} \sum_{x,y} \langle n_{x,y} \rangle & \text{if } x \bmod 2 = 1, y \bmod 2 = 1 \\ 0 & \text{else} \end{cases}$$

which detects the deformation of the density on the $(1, 1)$ -sublattice. This sublattice is defined in terms of a 2×2 cell, as in [143].

Numerical accuracy

All numerical results in Figs.A.12-A.13 were computed using DMRG. It was possible to study the 13×13 lattice using DMRG because we only investigated a low-entanglement region of the phase diagram. The level of accuracy for these calculations is shown in Fig. A.6 in terms of the largest truncated singular value during the DMRG sweep. In the ordered regions of the results, the largest truncated singular value is below 10^{-9} , which is generally considered accurate.