A tour through Hilbert space

Exploring the Lieb-liniger model and the interaction quench

de Klerk, A.J.J.M.

A TOUR THROUGH HILBERT SPACE

EXPLORING THE LIEB-LINIGER MODEL AND THE INTERACTION QUENCH

BART DE KLERK

A Tour through Hilbert space: Exploring the Lieb-Liniger model and the interaction quench

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor

aan de Universiteit van Amsterdam

op gezag van de Rector Magnificus

prof. dr. P.P.C.C. Verbeek

ten overstaan van een door het College voor Promoties ingestelde commissie,

in het openbaar te verdedigen in de Aula

op vrijdag 14 juli 2023, te 14:00 uur

door

Albertus Johannes Jacobus Maria de Klerk

geboren te Noordwijkerhout

Promotiecommissie

Promotor:	prof. dr. J.-S. Caux	Universiteit van Amsterdam
Copromotor:	dr. N.J. Robinson	EPSRC
Overige leden:	prof. dr. C.J.M. Schoutens	Universiteit van Amsterdam
	dr. J. van Wezel	Universiteit van Amsterdam
	dr. P. R. Corboz	Universiteit van Amsterdam
	dr. V. Gritsev	Universiteit van Amsterdam
	prof. dr. B. Doyon	King's College London
	prof. dr. T. Giamarchi	Université de Genève
	prof. dr. J.M. Maillet	ENS Lyon

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

LIST OF PUBLICATIONS

This thesis is based on the following publications:

[1] *Improved Hilbert space exploration algorithms for finite temperature calculations*
A.J.J.M. de Klerk and J.-S. Caux
SciPost Phys. Core 6(2):039 (2023)

[2] *On computing non-equilibrium dynamics following a quench*
N.J. Robinson, A.J.J.M. de Klerk and J.-S. Caux
SciPost Phys. 11(6):104 (2021)

[3] F.H.L. Essler, A.J.J.M. de Klerk, N.J. Robinson, and J.-S. Caux
To be published.

Other publications by the author of this thesis:

[4] *Knotted optical vortices in exact solutions to Maxwell's equations.*
A.J.J.M. de Klerk, R.I. van der Veen, J.W. Dalhuisen and D. Bouwmeester
Phys. Rev. A 95(5) 053820 (2017)

Contents

Chapter 1

Introduction

At the start of any PhD, the student is given the noble goal of solving problems that push the boundary of our knowledge if only ever so slightly. The trouble with this goal is that, although it is easy to find problems beyond our understanding, it is frustratingly difficult to actually make any progress on them. A common approach that makes the life of a theoretical physicist more pleasant is to ignore some of the infinite complexity that nature presents us with and instead study simplified models that capture as much of the phenomology we are interested in whilst remaining tractable. This thesis concerns one such model that has been the center of my work over the last four years.

Before delving into the specifics of this thesis and the model at its centre, let us take a step back and see how the study of this model and others with similar properties has come to be. This story begins with Hans Bethe's solution to the one-dimensional Heisenberg model in 1931 [5], which showed how one can deal exactly with the nearest neighbour interactions in the spin chain. For some time however, the result lay mostly dormant until it was used to solve another at first sight rather different model: whereas in the Heisenberg chain the particles are situated at fixed positions, it was next applied to a continuum model instead. This work was done by Elliot Lieb and Werner Liniger in 1963, who showed how Bethe's Ansatz could be used to solve the one-dimensional model of spinless bosons interacting via a delta function interaction potential [6, 7]. This put the Bethe Ansatz on the map for good and it has since been used to calculate the partition function for the six-vertex model by Lieb [8, 9], the eight-vertex model in the zero field case by Rodney Baxter [10], and the solution of the one-dimensional model of fermions interacting through a delta function interaction potential by Chen Ning Yang [11] to name some of the most famous examples. The latter two examples also spawned what we now call the Yang-Baxter equation which plays a central role in the Algebraic Bethe Ansatz and spawned the mathematical research area of quantum groups.

Models that can be solved using the Bethe Ansatz have the special property that a single interaction between multiple particles can be described as mul-

tiple interactions between only two particles. This property can be taken to be the defining feature of quantum integrable systems, although some discussion remains as to how to best characterize integrability [12]. Beyond being a mathematical property thas has allowed for a wealth of exact results to be discovered, advances in the field of ultra-cold atoms have made some of these integrable models an experimental reality [13, 14]. One experiment in particular, describing the lack of equilibration of a one-dimensional Bose gas prepared in an out-of-equilbrium initial state, showcased the interesting behaviour integrable quantum systems can exhibit even in the presence of the integrability-breaking perturbations present in experiments [15]. Since then the study of integrable models has become an important pillar for the study of out-of-equilibrium physics and helped improve our understanding of strongly correlated physics in one dimension [16–25].

This thesis is centred around the Lieb-Liniger model, describing a one-dimensional gas of Bosons with ultra-local interactions which we introduce in Chap. 2. We show how to solve this model using the coordinate Bethe Ansatz and discuss its spectrum. Although the Lieb-Liniger model is a continuum model, the structure of its Bethe Ansatz solutions is the simplest among the integrable models. As a result, many of its equilibrium properties are well understood, including thermodynamic properties as well as correlation functions of local operators [26–29]. A crucial ingredient for the calculation of local correlation functions is produced by the Algebraic Bethe Ansatz, which allows for the efficient calculation of matrix elements of some physically important operators with respect to the eigenstates [30–34]. We give an overview of the Algebraic Bethe Ansatz and discuss how it can be leveraged to obtain efficient expressions for the matrix elements of physically important operators with respect to eigenstates of the model. Finally, we discuss how the thermodynamic limit of the model can be taken leading to the Thermodynamic Bethe Ansatz [35].

Having described the Lieb-Liniger model and its spectrum, we turn to the calculation of its correlation functions. We discuss how the Lehmann representation gives rise to a summation over the matrix elements of eigenstates in Hilbert space. Although there exist some partial analytical results, in general one has to resort to explicit numerical summations over intermediate states in order to evaluate such summations. In Chap. 3 we consider algorithms aimed at efficiently evaluating such summations numerically with as few eigenstates as possible. In order to find the right rare eigenstates important to the calculation at hand we develop Hilbert space exploration algorithms which leverage the properties of the integrable model. We start by discussing the general principles of the type of Hilbert space exploration algorithms we developed and then start building our algorithm step by step. Ultimately we arrive at an efficient algorithm matching the performance of existing algorithms at zero

temperature and providing additional advantages at finite temperature [1].

These Hilbert space exploration algorithms again play an important role when we consider a global quench in the interaction strength of the Lieb-Liniger model in Chap. 4. In this chapter we start by discussing how the time evolution of the ground state following such a quench is equivalent to diagonalizing a perturbed Hamiltonian. The resulting object is infinite-dimensional, so in order to make progress we introduce a truncation of the Hilbert space such that we can diagonalize the corresponding truncated matrix numerically. Furthermore, we show how the procedure can be recast into an iterative numerical renormalization group scheme allowing us to consider large truncated bases. What follows next is a discussion of how to cleverly choose the truncated basis in order to obtain optimal convergence. This approach is built upon Hilbert space exploration algorithms like the ones discussed in Chap. 3 and allows us to deal with far larger system sizes than one can consider using the coordinate Bethe Ansatz [2].

After this chapter, we switch gears and consider the statistical properties of the matrix elements of the Lieb-Liniger model in Chap. 5. We start by investigating the scaling of the matrix element distributions of the single and two particle-hole sectors with system size and its implications. Then we present a surprising result regarding the off-diagonal matrix elements between two different macrostates followed by a discussion on sampling microstates corresponding to a given thermal state. The sampling algorithm that follows from this is then used to study the statistical properties of typical states. One of the key insights of this chapter is the surprisingly slow convergence we observe with system size, often requiring hundreds of particles, as well as the subtleties that arise when sampling typical states [3]. We finish the thesis with a short summary.

Due to the numerical nature of much of the material in this thesis, a lot of work went into the development of the code used to generate the data. In order to allow the interested reader to study the implementations of the algorithms presented here, we have made the code available at

https://github.com/AJJMdeKlerk/MERG.

CHAPTER 2

THE LIEB-LINIGER MODEL

Having motivated the general themes of this thesis, in this chapter we review the properties of the Lieb-Liniger model relevant for what is to come. We start by introducing its Hamiltonian in Sec. 2.1, after which we solve it using the coordinate Bethe Ansatz in Sec. 2.2. In Sec. 2.3 we discuss properties of its spectrum and in Sec. 2.4 we summarize how Algebraic Bethe Ansatz can be used to obtain explicit expressions for the matrix elements of physically important operators. We give an overview of the thermodynamic limit of the Lieb-Liniger model in Sec. 2.5.

2.1 THE HAMILTONIAN

The Lieb-Liniger model [6, 7] describes one-dimensional bosons with a delta function interaction potential, so its Hamiltonian can, in first quantised form, be written as

$$H = \sum_{i=1}^{N} \left[-\frac{\hbar^2}{2m} \frac{\partial^2}{\partial x_i^2} + 2c \sum_{i<j} \delta(x_i - x_j) \right]. \tag{2.1}$$

where we have N particles and x_i is the position of the i^{th} particle. Furthermore, c represents the interaction strength and m the mass. We take $\hbar = 1 = 2m$ to define our units. In the following we will often consider the second quantised form of this Hamiltonian, which with $\hbar$ and $2m$ set to one reads

$$H = \int dx \left[\partial_x \Psi^\dagger(x) \partial_x \Psi(x) + c \Psi^\dagger(x) \Psi^\dagger(x) \Psi(x) \Psi(x) \right]. \tag{2.2}$$

Here $\Psi^\dagger$ is the bosonic creation operator which satisfies the canonical commutation relations, i.e.

$$\left[\Psi(x), \Psi^\dagger(y) \right] = \delta(x - y). \tag{2.3}$$

In the limit where the interaction strengh vanishes, i.e. where $c = 0$, we are dealing with free bosons so we can readily diagonalise the Hamiltonian using

5

a Fourier transform. In the opposite limit where the bosons are infinitely repulsive, i.e. for $c \to \infty$, we can map the problem to that of non-interacting spinless fermions in one dimension [36]. How to solve the model for the intermediate regime is the topic of the next section.

2.2 The eigenstates

One of the reasons for studying the Lieb-Liniger model is that we can solve its Schrödinger equation to obtain the complete spectrum. The spectrum was first solved using something we now call the coordinate Bethe Ansatz, inspired by Bethe's solution to the one-dimensional Heisenberg spin chain [5–7]. It is within the context of this original approach that we solve the model here.

2.2.1 Coordinate Bethe Ansatz

The Bethe ansatz approach to solving the Lieb-Liniger model begins by considering the fundamental domain D_N, which for N particles is defined as

$$D_N = \{x \in \mathbb{R}^N \mid x_1 < x_2 < \cdots < x_N\}. \tag{2.4}$$

Note that the restriction of the problem to this domain does not constitute a loss of generality as we can extend this solution to $\mathbb{R}^N$ by invoking the symmetry requirements of the wavefunction when exchanging particles.

The Bethe ansatz is then that the wavefunction is a superposition of all permutations of plane waves with quasi-momenta λ_j (often called rapidities) and amplitudes A_σ that are at this stage undetermined giving

$$\Psi_N(x) = \sum_{\sigma \in S_N} A_\sigma e^{\sum_{j=1}^{N} i\lambda_{\sigma(j)} x_j}. \tag{2.5}$$

where S_N is the group of permutations of the numbers $\{1, \ldots, N\}$. The coefficients A_σ can then be determined by the boundary conditions for the fundamental domain arising from the delta function interaction potential for two-particle collisions. This gives

$$A_\sigma = (-1)^{[\sigma]} \prod_{1 \leq l < j \leq N} (\lambda_{\sigma(j)} - \lambda_{\sigma(l)} + ic), \tag{2.6}$$

where $(-1)^{[\sigma]}$ is the sign of the permutation. Imposing periodic boundary conditions on the Bethe wave function gives rise to what are called the Bethe equations, which determine the values of the rapidities λ_j, which read

$$e^{i\lambda_j L} = \prod_{l \neq j} \frac{\lambda_j - \lambda_l + ic}{\lambda_j - \lambda_l - ic}. \tag{2.7}$$

Any set of rapidities $\{\lambda_i\}_{1 \leq i \leq N}$ satisfying Eq. (2.7) thus gives rise to an eigenstate of the Lieb-Liniger model.

2.2.2 THE QUANTUM NUMBER REPRESENTATION

Instead of considering the Bethe equations directly, it is useful to take the logarithm of Eq. (2.7) resulting in the logarithmic Bethe equations

$$\lambda_j + \frac{2}{L}\sum_{k=1}^{N}\operatorname{atan}\left(\frac{\lambda_j - \lambda_k}{c}\right) - \frac{2\pi I_j}{L} = 0. \tag{2.8}$$

where we introduced the quantum numbers $\{I_j\}_{1\leq j\leq N}$ which are integers when N is odd and half-odd integers when N is even. The quantum numbers are more than a mathematical necessity introduced by the logarithm, they turn out to be a convenient way of uniquely labelling the eigenstates. There is a one-to-one correspondence between the quantum numbers and the rapidities which respects the ordering, meaning that if $I_j > I_k$ then $\lambda_j > \lambda_k$ due to the monotonic nature of the second term on the left hand side of Eq. (2.8). Consequently, no two quantum numbers can be equal since then the wavefunction formally vanishes. How to solve the logarithmic coordinate Bethe Ansatz Equations is discussed in Appendix A.

The solvability of the model is a result of the fact that all interactions can be reduced to two-body interactions, a hallmark of integrability. Another special property of the Lieb-Liniger model, also sometimes used to characterize integrability, is that it has infinitely many non-trivial commuting conserved charges whose eigenvalues are given by

$$Q_n = \sum_{j=1}^{N}\lambda_j^n \tag{2.9}$$

for any $n \in \mathbb{N}$. The first conserved charge, which represents the momentum, can also be expressed in terms of the quantum numbers via

$$P = \sum_j \lambda_j = \frac{2\pi}{L}\sum_j I_j. \tag{2.10}$$

The second conserved charge Q_2 also has an important physical interpretation as it represents the energy.

2.3 THE SPECTRUM

The usefulness of the representation of an eigenstate by its quantum numbers goes beyond it being a tool to solve the Bethe equations or as a visualisation tool. It also allows us to obtain a better understanding of the excitation spectrum in terms of two types of elementary excitations. We turn to this
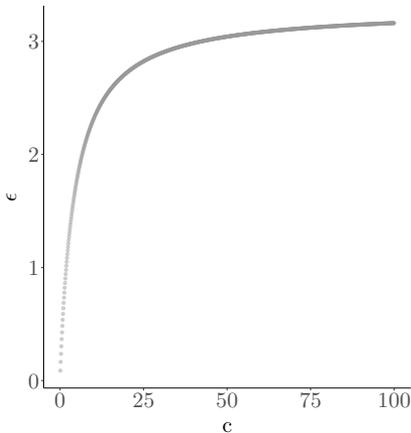
(a) The ground state energy density ϵ as a function of the interaction strength for $N = 128 = L$ for c from 0 to 100 in increments of 0.1.

(b) The kinetic and interaction energy density as a function of the interaction strength for $N = 128 = L$ for c from 0 to 100 in increments of 0.1.

Figure 2.1

characterization of the spectrum after discussing the structure of the ground state.

The ground state of the Lieb-Liniger model is the eigenstate for which the quantum numbers are as close to zero as possible. Since the quantum numbers are not allowed to coincide and they can only take on integer or half odd integer values, this results in a configuration like a Fermi sea. To be precise, the quantum numbers of the ground state are given by

$$\left\{ -\frac{N+1}{2}, -\frac{N+1}{2} + 1, \ldots, \frac{N-1}{2} \right\}. \tag{2.11}$$

We can therefore define a Fermi momentum in analogy with the Fermi sea, given by $k_F = \frac{\pi}{L}(N - \frac{1}{2})$ so that k_F is between the last occupied and first unoccupied mode. Note that although the quantum number configuration corresponding to the ground state does not depend on the interaction strength, the corresponding rapidities, which can be obtained by solving the logarithmic Bethe Ansatz equations, do. The dependence of the ground state energy as a function of the interaction strength is shown in Fig. 2.1a.

The ground state energy can be decomposed into the kinetic and an interaction energy part, whose energy densities are given by

$$\epsilon_{\text{pot}} = c\langle 0|\Psi^\dagger(0)\Psi^\dagger(0)\Psi(0)\Psi(0)|0\rangle \tag{2.12}$$

$$\epsilon_{\text{kin}} = \tfrac{1}{L}\langle 0|H_{LL}|0\rangle - \epsilon_{\text{pot}} \tag{2.13}$$

8

where $|0\rangle$ represents the ground state. The relative importance of either contribution to the ground state energy is shown in Fig. 2.1b. We see that as the interaction strength increases, the relative importance of the potential energy decreases. The value of the interaction strength where both are equal lies around 4.35. Although the exact value of the interaction strength at which the kinetic and potential energy contributions are equal is state and thus temperature dependent, we take $c = 4$ as the threshold separating the weakly and strongly interacting regimes.

Having identified the ground state, let us turn to the excited states. Any excited state can of course be characterized by its quantum numbers, but we can also introduce a language of excitations that organizes the excited states into families of states. This classification will prove useful for the discussions in Chapter 3 in particular. Before discussing this classification, let us discuss a way of visualizing the quantum numbers of an eigenstate as well as some common nomenclature.

In order to introduce a way of visualising the quantum numbers, consider a set of quantum numbers $\{I_i\}_{i \leq 5}$, which we assume to be ordered, an assumption we retain for the remainder of this thesis. Since no two quantum numbers can be the same, we can think of them as five particles on a one-dimensional lattice, which can be visualized as follows:

$$\cdots \; \circ \; \bullet \; \circ \; \bullet \; \bullet \; \bullet \; \circ \; \bullet \; \circ \; \cdots$$
$$\text{-4 -3 -2 -1 \; 0 \; 1 \; 2 \; 3 \; 4}$$

The numbers below the circles represent the quantum numbers corresponding to the position on the line, so this state represents $\{-3, -1, 0, 1, 3\}$. In this thesis we visualize the quantum numbers in order to illustrate the principles of our algorithms. For this purpose we are not interested in the absolute values of the quantum numbers, but rather only the differences between the quantum numbers allowing us to drop the numbers below the circles going forward.

For such eigenstates it is often convenient to talk about particles and holes. In this context we define them by fixing an eigenstate, the reference state, defined by some quantum numbers $\{I_i^{RS}\}_{i \leq N}$. Then for any other N particle state $\{I_i\}_{i \leq N}$, the quantum numbers I_l which do not occur in $\{I_i^{RS}\}_{i \leq N}$ are called particles whereas the lattice positions which are now empty as a result are called holes. As an example, let us take the five particle ground state which can be visualized as

$$\cdots \; \circ \; \circ \; \bullet \; \bullet \; \bullet \; \bullet \; \bullet \; \circ \; \circ \; \cdots$$

as the reference state. Then we can label the particles and holes for the
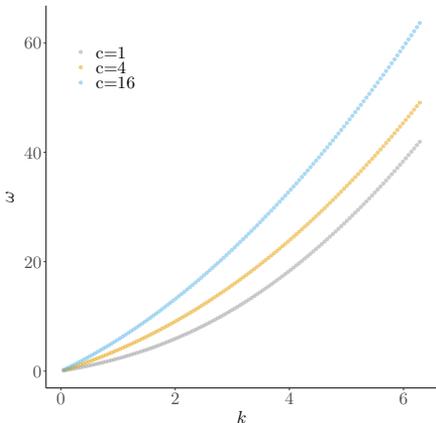
(a) The momentum k of the type I excited states versus the energy difference with the ground state $\omega = E - E_{gs}$ for $N = 128 = L$ and $c = 1, 4, 16$.
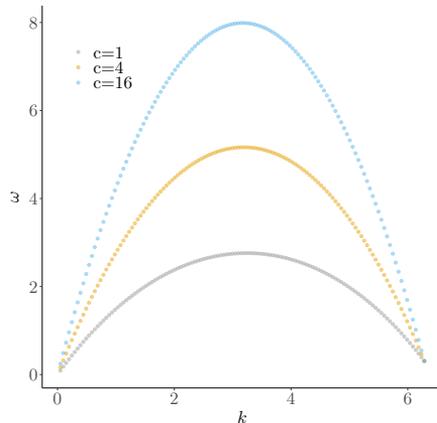
(b) The momentum k of the type II excited states versus the energy difference with the ground state $\omega = E - E_{gs}$ for $N = 128 = L$ and $c = 1, 4, 16$.

Figure 2.2

following three states with a p and h respectively.
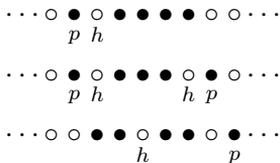
Note that we always have an equal number of particles and holes enabling us to talk about particle-hole pairs.

Now suppose we again take the ground state as the reference state and from it we wish to create some excited state. This can of course be done in many ways, but the simplest way is by hopping one of the outermost quantum numbers to the nearest vacant site. This creates a particle-hole pair at the edge of the Fermi sea which looks like

Now there are two possibilities to increase the momentum and energy further without creating additional particle-hole pairs. Either we move the particle further away from the Fermi sea, or we move the hole further into the bulk of the Fermi sea of quantum numbers.

Moving the particle further away from the edge of the Fermi sea gives rise to what are called the type I excitations, the first three of which can be visualized

as

$$\circ \; \circ \; \circ \; \circ \; \bullet \; \bullet \; \bullet \; \bullet \; \bullet \; \underset{h}{\circ} \; \underset{p}{\bullet} \; \bullet \; \circ \; \circ \; \circ$$

$$\circ \; \circ \; \circ \; \circ \; \bullet \; \bullet \; \bullet \; \bullet \; \bullet \; \underset{h}{\circ} \; \circ \; \underset{p}{\bullet} \; \circ \; \circ$$

$$\circ \; \circ \; \circ \; \circ \; \bullet \; \bullet \; \bullet \; \bullet \; \bullet \; \underset{h}{\circ} \; \circ \; \circ \; \underset{p}{\bullet} \; \circ$$
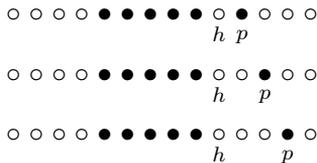
Under the restrictions that we can create only one particle-hole pair each of these represent the highest energy state we can create for a given momentum change. After all, the energy of an eigenstate is equal to the sum of the squares of its rapidities, so we maximize the energy for a given momentum transfer in the one-particle hole sector by creating the highest possible quantum number. Furthermore, the energy is strictly increasing as we increase the momentum transfer as shown in Fig. 2.2a.

Moving the hole further into the Fermi sea gives rise to what are called the type II excitations, of which the first three can visualized as

$$\circ \; \circ \; \circ \; \circ \; \bullet \; \bullet \; \bullet \; \bullet \; \bullet \; \underset{h}{\circ} \; \underset{p}{\bullet} \; \bullet \; \circ \; \circ \; \circ$$

$$\circ \; \circ \; \circ \; \circ \; \bullet \; \bullet \; \bullet \; \bullet \; \underset{h}{\circ} \; \bullet \; \underset{p}{\bullet} \; \bullet \; \circ \; \circ \; \circ$$

$$\circ \; \circ \; \circ \; \circ \; \bullet \; \bullet \; \bullet \; \underset{h}{\circ} \; \bullet \; \bullet \; \underset{p}{\bullet} \; \bullet \; \circ \; \circ \; \circ$$
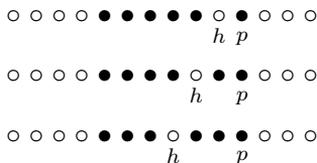
The type II states are the one particle-hole states with the lowest energy for a given momentum since for a given jump in momentum they keep the distribution as tightly packed around zero as possible. As the hole moves further inward, it initially increases the energy further and further until we hit the middle of the Fermi sea. After this point, we are taking negative quantum numbers and putting them on the right edge of the Fermi sea which increases the energy again as shown in Fig. 2.2b. The energies of the other eigenstates with a single particle-hole pair lie between the curves of the type I and type II excitations. States with more particle-hole pairs are of course not restricted to this region in (k, ω) space.

2.4 Matrix elements

Although the coordinate Bethe ansatz provides a convenient framework for finding the eigenstates of the Lieb-Liniger model, it is not suitable for determining the norms of eigenstates or more general overlaps between eigenstates in a computationally tractable manner. To this end, a set of tools has been developed that we refer to as the Algebraic Bethe Ansatz. The results obtained through the Algebraic Bethe Ansatz play an important role throughout the rest of this thesis, but since we do not use the formalism directly only a cursory outline is given here, for more information we refer the reader to [26].

Rather than starting from a Hamiltonian and finding the corresponding eigenstates, the Algebraic Bethe Ansatz starts from the more general problem of considering a Hilbert space $\mathcal{H}$ and asking if we can construct a complete set of local operators $\{\hat{Q}_n\}$ on this space which are in involution, i.e. that satsify $[\hat{Q}_n, \hat{Q}_m] = 0$ for all $m, n \in \mathbb{N}$. The problem of finding a complete set of local operators in involution can be rephrased by introducing an operator-valued function $\mathcal{T}(\lambda)$ of the form

$$\mathcal{T}(\lambda) = \sum_{n=0}^{\infty} \hat{Q}_n \lambda^n \tag{2.14}$$

called the transfer matrix. The idea is that if we can find such a transfer matrix satisfying

$$[\mathcal{T}(\lambda), \mathcal{T}(\mu)] = 0 \tag{2.15}$$

for all λ and μ, we have discovered a model where infinitely many charges in involution. Furthermore, by taking any of the charges as the Hamiltonian we ensure that the operators are constants of motion.

The task of finding an integrable model is now reduced to finding a transfer matrix satisfying Eq. (2.15). In order to simplify the solution to this problem, one introduces an auxilliary vector space $\mathcal{A}$ along with a new operator-valued function $T(\lambda)$ acting on $\mathcal{A} \otimes \mathcal{H}$, called the monodromy matrix, satisfying

$$\mathrm{Tr}_{\mathcal{A}} T(\lambda) = \mathcal{T}(\lambda) \tag{2.16}$$

which also satisfies the equivalent of Eq. (2.15), i.e.

$$[\mathrm{Tr}_{\mathcal{A}_1} T(\lambda), \mathrm{Tr}_{\mathcal{A}_2} T(\mu)] = 0. \tag{2.17}$$

The advantage of introducing the auxilliary space and the monodromy matrix is that now we can rewrite Eq. (2.17) to give

$$0 = [\mathrm{Tr}_{\mathcal{A}_1} T_1(\lambda), \mathrm{Tr}_{\mathcal{A}_2} T_2(\mu)] \tag{2.18}$$

$$= \mathrm{Tr}_{\mathcal{A}_1 \otimes \mathcal{A}_2}[T_1(\lambda) \otimes T_2(\mu) - T_2(\mu) \otimes T_1(\lambda)] \tag{2.19}$$

where the indices indicate on which auxilliary space the operators act. This implies that rather than having to find a transfer matrix which commutes with itself, we have to find an intertwining operator for the monodromy matrix, i.e. an operator R acting on the tensor product of auxilliary spaces satisfying

$$R_{12}(\lambda, \mu) T_1(\lambda) \otimes T_2(\mu) R_{12}^{-1}(\lambda, \mu) = T_2(\mu) \otimes T_1(\lambda). \tag{2.20}$$

Considering the intertwining relations for three monodromy matrices reveals a consistency condition that the intertwining operator has to satisfy called the Yang-Baxter equation given by

$$R_{12}(\lambda, \mu)R_{13}(\lambda, \nu)R_{23}(\mu, \nu) = R_{23}(\mu, \nu)R_{13}(\lambda, \nu)R_{12}(\lambda, \mu). \qquad (2.21)$$

Thus the problem of finding an integrable model reduces to defining a monodromy matrix and finding the corresponding intertwining operator satisfying the Yang-Baxter equation.

Given a monodromy matrix for $\mathcal{A} = \mathbb{C}^2$, as is the case for the Lieb-Liniger model, we can consider its matrix elements

$$T(\lambda) = \begin{pmatrix} A(\lambda) & B(\lambda) \\ C(\lambda) & D(\lambda) \end{pmatrix} \qquad (2.22)$$

whose commutation relations are fixed by the corresponding R matrix. These matrix elements, which themselves are operators acting on the Hilbert space, can in turn be used to construct the states in Hilbert space by assuming a psuedovacuum state $|0\rangle$ and acting on it with $B(\lambda)$, which acts as the creation operator. In order to ensure that the pseudovacuum is also an eigenstate of the conserved charges it is sufficient to require that it is an eigenstate of $A(\lambda)$ and $D(\lambda)$, with eigenvalues $a(\lambda)$ and $d(\lambda)$ respectively, while it is annihilated by $C(\lambda)$, which acts as the annihilation operator. Repeatedly applying $B(\lambda)$ to this pseudovaccum gives rise to new states given by

$$|\lambda\rangle = \prod_{j=1}^{M} B(\lambda_j)|0\rangle. \qquad (2.23)$$

which are eigenstates of the transfer matrix provided that the $\{\lambda_j\}_{j=1}^{M}$ satisfy the Bethe equations. This algebraic description of eigenstates forms the basis for the determination of the overlaps between two states of the form described in Eq. (2.23) of which at least one satisfies the Bethe Ansatz equations as detailed by Slavnov's theorem [30]. This theorem underpins all the matrix element expressions we present in the following section.

For completeness, we note that the Lieb-Liniger model can be constructed from the Algebraic Bethe Ansatz for $\mathcal{A} = \mathbb{C}^2$ as mentioned previously and by taking

$$R(\lambda, \mu) = \begin{pmatrix} f(\lambda, \mu) & 0 & 0 & 0 \\ 0 & g(\lambda, \mu) & 1 & 0 \\ 0 & 1 & g(\lambda, \mu) & 0 \\ 0 & 0 & 0 & f(\lambda, \mu) \end{pmatrix} \qquad (2.24)$$

where the functions f and g are given by

$$f(\lambda, \mu) = \frac{\mu - \lambda + ic}{\mu - \lambda} \quad \text{and} \quad g(\lambda, \mu) = \frac{ic}{\mu - \lambda}. \qquad (2.25)$$

13

The eigenvalues of A and B that ensue are, in turn, given by

$$a(\lambda) = e^{-i\lambda L/2} \quad \text{and} \quad d(\lambda) = e^{i\lambda L/2}. \tag{2.26}$$

The norm of eigenstates, which can also be obtained from taking the limit of Slavnov's theorem, confirming an earlier conjecture by Gaudin [37], can be expressed in terms of the determinant of a $N \times N$ matrix $\mathcal{N}$ giving

$$\langle \lambda | \lambda \rangle = c^N \prod_{j<l} \frac{(\lambda_j - \lambda_l)^2 + c^2}{(\lambda_j - \lambda_l)^2} \det \mathcal{N} \tag{2.27}$$

where

$$\mathcal{N}_{jl} = \delta_{jl} \left(L + \sum_{k=1}^{N} K(\lambda_j, \lambda_k) \right) - K(\lambda_j, \lambda_l) \tag{2.28}$$

with

$$K(\lambda, \mu) = \frac{2c}{c^2 + (\lambda - \mu)^2}. \tag{2.29}$$

2.4.2 Physically important operators

Although it is already rare that we can determine the full spectrum of an interacting integrable model, knowledge of the spectrum alone is not sufficient to do the calculations we present in this thesis. For the chapters that follow we need an additional ingredient, which are the matrix elements of physically important operators with respect to the eigenstates. Throughout this thesis we consider the density operator

$$\rho(x) = \Psi^\dagger(x)\Psi(x) \tag{2.30}$$

as well as the g_2 operator

$$g_2(x) = \Psi^\dagger(x)\Psi^\dagger(x)\Psi(x)\Psi(x). \tag{2.31}$$

For both these operators, efficient expressions for their matrix element with respect to the eigenstates of the Lieb-Liniger model have been determined using the Algebraic Bethe Ansatz [33, 34, 31].

The off-diagonal matrix elements of the density operator, up to normalization by the norm as given in Eq. (2.27), for states $|\lambda\rangle$ and $|\mu\rangle$ none of whose rapidities are equal are given by

$$\langle \mu | \rho(0) | \lambda \rangle = \prod_{a=1}^{N} \left(V_a^+ - V_a^- \right) \prod_{b,d=1}^{N} \left(\frac{\lambda_b - \lambda_d + ic}{\mu_b - \lambda_d} \right)$$

$$\times \frac{i(P_\lambda - P_\mu)}{V_p^+ - V_p^-} \det \left[\mathbb{1} + U(\lambda_p) \right] \tag{2.32}$$

14

where

$$P_\lambda = \sum_{i=1}^{N} \lambda_i \tag{2.33}$$

$$V_j^{\pm} = \prod_{m=1}^{N} \frac{\mu_m - \lambda_j \pm ic}{\lambda_m - \lambda_j \pm ic} \tag{2.34}$$

$$U_{jk}(\lambda_p) = \frac{i}{V_j^+ - V_j^-} \frac{\prod_{m=1}^{N}(\mu_m - \lambda_j)}{\prod_{m=1, m\neq j}^{N}(\lambda_m - \lambda_j)} [K(\lambda_j, \lambda_k) - K(\lambda_p, \lambda_k)] \tag{2.35}$$

$$K(a, b) = \frac{2c}{c^2 + (a - b)^2} \tag{2.36}$$

and $V_p^{\pm}$ is $V_j^{\pm}$ as defined above with λ_j replaced by an arbitrary complex number λ_p. So, like the equation for the norm of an eigenstate, it requires the computation of a matrix of size $N \times N$. The normalized matrix elements for the diagonal density operator on the other hand follows directly from translational symmetry and is given by the density of particles for the state under consideration.

The off-diagonal matrix elements of $g_2(0)$ with respect to two Bethe states $|\lambda\rangle$ and $|\mu\rangle$ which have no rapidities in common can again be expressed in terms of a single determinant of size $N \times N$ similar to the expression for the off-diagonal density operator. The equations read

$$\langle\mu|g_2(0)|\lambda\rangle = \prod_{a=1}^{N} \left(V_a^+ - V_a^-\right) \prod_{b,d=1}^{N} \left(\frac{\lambda_b - \lambda_d + ic}{\mu_b - \lambda_d}\right)$$

$$\times \frac{(-1)^N J}{6c(V_p^+ - V_p^-)(V_s^+ - V_s^-)} \det\left[\mathbb{1} + V(\lambda_p)\right] \tag{2.37}$$

where

$$E_\lambda = \sum_{i=1}^{N} \lambda_i^2 \tag{2.38}$$

$$J = (P_\lambda - P_\mu)^4 + 3(E_\lambda - E_\mu)^2 - 4(P_\lambda - P_\mu)(Q_{3,\lambda} - Q_{3,\mu}) \tag{2.39}$$

$$V_{jk}(\lambda_p) = \frac{i}{V_j^+ - V_j^-} \frac{\prod_{m=1}^{N}(\mu_m - \lambda_j)}{\prod_{m=1, m\neq j}^{N}(\lambda_m - \lambda_j)}$$

$$\times [K(\lambda_j, \lambda_k) - K(\lambda_p, \lambda_k)K(\lambda_s, \lambda_j)] \tag{2.40}$$

and $V_j^{\pm}$ is the same as defined in Eq. (2.34) and we have introduced another arbitrary complex number λ_s.

Expressions for the diagonal matrix elements of the g_2 operator have also been derived in [33], but rather than having to compute the determinant of a

single matrix of size N^2, it involves a sum over $\frac{1}{2}N(N-1)$ such matrices. As such, the computation of this matrix element is computationally costly, but it turns out that we can leverage the Feynman-Hellman theorem to compute these matrix elements in a more efficent manner. The Feynman-Hellman theorem states that for any normalised eigenstate $|\phi\rangle$ with energy E_ϕ and any parameter α the following identity holds

$$\partial_\alpha E_\Psi = \partial_\alpha \langle\phi|H|\phi\rangle = \langle\phi|\frac{\partial H}{\partial\alpha}|\phi\rangle \tag{2.41}$$

This result follows directly from expanding the middle expression and using

$$\partial_\alpha\langle\phi|\phi\rangle = 0. \tag{2.42}$$

In the case of the Lieb-Liniger model we can consider α to be the interaction strength, so that for any normalized Bethe state $|\lambda\rangle\rangle$ we get

$$\partial_c E_\lambda = \langle\lambda| \int_0^L dx \Psi^\dagger(x)\Psi^\dagger(x)\Psi(x)\Psi(x)|\lambda\rangle \tag{2.43}$$

$$= L\langle\lambda|\Psi^\dagger(0)\Psi^\dagger(0)\Psi(0)\Psi(0)|\lambda\rangle. \tag{2.44}$$

This means that the Feynman-Hellman theorem can be used to commpute the diagonal elements of g_2 provided that we can find a way to obtain $\partial_c E_\lambda$.

In order to compute $\partial_c E_\lambda$ we need to compute the derivatives of the rapidities with respect to the interaction strength. These can be computed by solving the equations obtained by taking the derivative of the Bethe equations with respect to the interaction strength. This gives

$$\partial_c\lambda_j + \frac{1}{L}\sum_{k=1}^N \frac{2c(\partial_c\lambda_j - \partial_c\lambda_k)}{c^2 + (\lambda_j - \lambda_k)^2} - \frac{1}{L}\sum_{k=1}^N \frac{2(\lambda_j - \lambda_k)}{c^2 + (\lambda_j - \lambda_k)^2} = 0. \tag{2.45}$$

The Jacobian of these equations is, like for the regular Bethe Equations, equal to the Gaudin matrix. Since the Gaudin matrix is positive definite a solution exists, it is unique, and can be efficiently determined using the Newton Rhapson-method described in Appendix A.

A property of the matrix elements crucial to the algorithms we discuss in Chapter 3 and 4, is that on average, the off-diagonal matrix elements are largest when the bra and ket states share the most quantum numbers. [1] To

---

[1] Here it is good to note that the value of the quantum numbers that they do not share is of secondary importance. For example, even when one of the quantum numbers I_j and the corresponding rapidity λ_j go off to $\pm\infty$, the term in the logarithmic Bethe equations, corresponding to the rapidity at infinity evaluates to $\pm\pi/L$ which effectively shifts I_k by $\mp 1/2$. Thus even a massive change in one quantum number shifts the other quantum numbers only a little.

see why, note that Slavnov's formula for the overlap $\langle \mu | \lambda \rangle$ between a Bethe state $|\lambda\rangle$ and an arbitrary set of rapidities $|\mu\rangle$ has poles for coinciding rapidities [26]. As such, overlaps between Bethe states are maximal when the number of (close to) coinciding rapidities is maximal. The matrix elements of the density and g_2 operators are derived from this formula for the overlaps by determining the action of the operator under consideration on the bra or ket and using the overlap formula. For example, acting with $\Psi(0)$ on a Bethe state results in a superposition of states where one of the rapidities is removed and the others are shifted due to the interactions. Therefore, the off-diagonal matrix elements of the density operator are maximal when the bra and ket differ by one quantum number. As the number of differing quantum numbers increases, the matrix element becomes smaller due to the smaller number of (close to) coinciding rapidities. Similarly, the off-diagonal matrix elements of the g_2 operator are maximal when the bra and ket differ by one or two quantum numbers. The importance of the number of differences between the quantum numbers of the bra and ket is also influenced by the interaction strength, with its importance diminishing as the interaction strength decreases.

2.5   The thermodynamic limit

Thus far we have considered the Lieb-Liniger model for some finite number of particles at some length L with periodic boundary conditions. In this section we discuss how to take the thermodynamic limit, i.e. the limit where $N, L \rightarrow \infty$ at a fixed finite density N/L, leading to what is called the Thermodynamic Bethe Ansatz [35]. Although the later chapters do not use the thermodynamic description of states directly, it is necessary to understand the language of thermal states we will employ.

Before we can take the thermodynamic limit of the logarithmic Bethe equations, we have to define a continuum version of the sets of quantum numbers we have been considering. In order to define this continuum limit we have to rescale the quantum numbers since the quantum numbers themselves diverge in the limit where $N, L \rightarrow \infty$ due to the fact that they are all distinct. We can define a density distribution for a set of quantum numbers $\{I_j\}_{j=1}^{N}$ as

$$\rho(x) = \frac{1}{L} \sum_{j=1}^{N} \delta\left(x - \frac{I_j}{L}\right). \tag{2.46}$$

This allows us to rewrite the logarithmic Bethe Ansatz equations from Eq. (2.8) as

$$\lambda(x) + \int_{-\infty}^{\infty} dy \, 2 \, \text{atan}\left(\frac{\lambda(x) - \lambda(y)}{c}\right) \rho(y) = 2\pi x \tag{2.47}$$

where $\lambda(x)$ is a real-valued function representing the continuum limit of the rapidities, which is defined for all $x \in \mathbb{R}$ by this relation. The complement of the density of the occupied quantum numbers is called the hole density, defined by

$$\rho_h(x) = \frac{1}{L} \sum_{\{n\}}^{N} \delta \left(x - \frac{n}{L} \right).$$ (2.48)

where the $\{n\}$ are the integers or half-odd integers unequal to any quantum number of the state under consideration. In the thermodynamic limit these densities can be described by smooth functions in x satisfying

$$\rho(x) + \rho_h(x) = \int_{-\infty}^{\infty} dy \delta(x - y) = 1$$ (2.49)

Note that in contrast to our conventions at finite size, any quantum number is called a particle and any vacant site is called a hole.

The densities of the particles and the holes can also be expressed in rapidity space using the relation between λ and x described by Eq. (2.47) by

$$\rho(\lambda) = \rho(x(\lambda)) \frac{dx(\lambda)}{d\lambda} \quad \text{and} \quad \rho_h(\lambda) = \rho_h(x(\lambda)) \frac{dx(\lambda)}{d\lambda}.$$ (2.50)

Here $\rho(\lambda)$ is called the root density. These definitions allow us to rewrite the Bethe Ansatz equations as

$$\lambda(x) + \int_{-\infty}^{\infty} d\mu \, 2 \operatorname{atan} \left(\frac{\lambda - \mu}{c} \right) \rho(\mu) = 2\pi x$$ (2.51)

Differentiating this equation with respect to λ gives rise to what is called the Lieb equation, reading

$$2\pi + 2\pi \int_{-\infty}^{\infty} d\mu \frac{2c}{(\lambda - \mu)^2 + c^2} \rho(\mu) = \rho(\lambda) + \rho_h(\lambda).$$ (2.52)

This relation determines $\rho(\lambda)$ given $\rho_h(\lambda)$ and vice versa. Analogously to the finite size case, the conserved quantities are given by the moments in terms of the rapidities, only now we need to take in to account the root density giving for example

$$p = \frac{P}{L} = \int_{-\infty}^{\infty} d\lambda \lambda \rho(\lambda)$$ (2.53)

$$e = \frac{E}{L} = \int_{-\infty}^{\infty} d\lambda \lambda^2 \rho(\lambda)$$ (2.54)

whereas the density is given by $n = \int_{-\infty}^{\infty} d\lambda \rho(\lambda)$.

At finite system size there are many distinct eigenstates whose distribution of rapidities approximate a given root distribution. To understand this, let us consider partitioning the real line into intervals of size $d\lambda$ called boxes. Then each box has a certain number of particles and holes fixed by the root distribution, which gives

$$L\rho(\lambda)d\lambda : \text{ number of particles in } [\lambda, \lambda + d\lambda] \qquad (2.55)$$

$$L\rho_h(\lambda)d\lambda : \text{ number of holes in } [\lambda, \lambda + d\lambda]. \qquad (2.56)$$

An eigenstate satisfying these constraints is said to have the right box occupation numbers. The number of such eigenstates for the given partition into boxes can be computed to be

$$\frac{[L(\rho(\lambda) + \rho_t(\lambda)])!}{[L\rho(\lambda)d\lambda]![L\rho_h(\lambda)d\lambda]!} = \exp(dS). \qquad (2.57)$$

Assuming that $d\lambda \gg L^{-1}$ such that each box can fit many rapidities at the system size L we are interested in, we can approximate this formula using Stirling's formula which gives

$$dS = Ld\lambda[(\rho(\lambda) + \rho_h(\lambda)) \ln(\rho_h(\lambda) + \rho_h(\lambda))$$
$$- \rho_h(\lambda) \ln \rho_h(\lambda) - \rho(\lambda) \ln \rho(\lambda)] \qquad (2.58)$$

where $S = \int dS$ is the entropy corresponding to the root distribution.

It should be pointed out that the correspondence between a root distribution, the macrostate, and the corresponding eigenstates at finite size, the microstates, with the definitions above depends on the partitioning of the real line we choose. In Chapter 5 we take a closer look at the implications this has. For now, we note that we can find the best approximation of the root distribution at a given finite size N by considering

$$z(\lambda) = L \int_{-\infty}^{\lambda} \mathrm{d}\mu\rho(\mu) \qquad (2.59)$$

to obtain a set of rapidities $\{\lambda_i\}$ corresponding to all the rapidities for which $z(\lambda)$ crosses an integer value. These rapidities can then through the Bethe Ansatz equations be turned into a set of integers after rounding. The resulting state is what we call the representative state corresponding to the macrostate, i.e. the root distribution.

In order to find the root distribution at a given temperature T one considers the saddle point equations for the partition function in the grand canonical ensemble which gives rise to the Yang-Yang equation which reads

$$\epsilon(\lambda) = \lambda^2 - \mu - \frac{T}{2\pi} \int_{-\infty}^{\infty} \mathrm{d}\mu K(\lambda, \mu) \ln(1 + e^{-\epsilon(\mu)/T}). \qquad (2.60)$$

Here μ represents the chemical potential and $\epsilon(\lambda)$ is related to the particle and hole densities through

$$\frac{\rho_h(\lambda)}{\rho(\lambda)} = e^{\epsilon(\lambda)/T}.$$ (2.61)

Simultaneously solving the Yang-Yang equation together with Eq. (2.61) and Eq. (2.52) while ensuring $\rho(\lambda) \geq 0$ and $\rho_h(\lambda) \geq 0$ as well as

$$\int_{-\infty}^{\infty} \rho(\lambda)d\lambda = \frac{N}{L}$$ (2.62)

gives rise to a thermal root distribution.

HILBERT SPACE EXPLORATION ALGORITHMS

To see how the need for Hilbert space scanning algorithms arises, consider the computation of the dynamical structure factor in the ground state of the Lieb-Liniger model, given by

$$S(k,\omega) = \int_0^L dx \int dt e^{-ikx+i\omega t} \langle \rho(x,t)\rho(0,0) \rangle \tag{3.1}$$

$$= \frac{2\pi}{L} \sum_\alpha |\langle 0|\rho_k|\alpha \rangle|^2 \delta(\omega - E_\alpha - E_0), \tag{3.2}$$

where ρ_k is the Fourier transform of the density operator. In order to numerically approximate Eq. (3.2), we need a set of eigenstates $|\alpha\rangle$ for which we compute the energies E_α, the matrix elements $\langle 0|\rho_k|\alpha \rangle$, and perform the summation. As the Lieb-Liniger model possesses an infinite number of eigenstates, an evaluation of the sum in Eq. (15) neccessitates a truncation and the accuracy of the calculation depends on the number of eigenstates and their matrix elements $\langle 0|\rho_k|\alpha \rangle$. The convergence is quantified by the f sum rule [38], which states that

$$\int \frac{d\omega}{2\pi} \omega S(k,\omega) = \frac{Nk^2}{L}. \tag{3.3}$$

Given Eq. (3.3), we can convert the contributions to the summation in Eq. (3.2) into a weighing function for a given eigenstate $|\alpha\rangle$ given by

$$w_f(\alpha) = \frac{L}{Nk^2}(E_\alpha - E_0) |\langle 0|\rho_k|\alpha \rangle|^2 \tag{3.4}$$

for $k \neq 0$ such that the summation over the weights of all eigenstates $|\alpha\rangle$ at a fixed momentum value gives 1. When considering the summation over the weights of Eq. (3.4), there are two contributing factors that determine the importance of an eigenstate, its energy and the matrix element. For the Lieb-Liniger model it turns out that most matrix elements except for a tiny portion of states are vanishingly small, dominating the effect the energy has.

Therefore our focus lies primarily on finding those states for which the matrix elements are large in order to get a good saturation of the sumrule. Hilbert space scanning algorithms are designed to preferentially generate eigenstates for which for example $w(|\alpha\rangle) = |\langle 0|\rho_k|\alpha\rangle|^2$ or $w(|\alpha\rangle) = (E_\alpha - E_0)|\langle 0|\rho_k|\alpha\rangle|^2$ are maximal [39, 40, 29].

Another problem where the need for generating appropriate eigenstates arises, is when choosing a basis for truncated spectrum methods [41–43]. Consider a quench of the Lieb-Liniger model where we change the interaction strength at $t = 0$ from c_i to c_f. Truncated spectrum methods can be used to compute the time evolution of the initial state $|\Psi_0\rangle$ in terms of a set of eigenstates of the Lieb-Liniger model at interaction strength c_f, i.e.

$$|\Psi_0(t)\rangle = \sum_\alpha b_\alpha e^{-iE_\alpha t}|\alpha\rangle \qquad (3.5)$$

where $b_\alpha = \langle\alpha|\Psi_0\rangle$ are the coefficients being approximated by the truncated spectrum methods. However, the accuracy of the expansion in Eq. (3.5) depends on the choice of basis states $|\alpha\rangle$. By choosing a weighing function that approximates $\langle\alpha|\Psi_0\rangle$ we can leverage our scanning algorithms to generate a close to ideal basis for this quench [2].

In this chapter, we develop novel Hilbert space exploration algorithms, inspired by those implemented in ABACUS, which we compare to ABACUS by considering the dynamical structure factor. We show that at zero temperature the performance of our algorithms is virtually identical to those in ABACUS, whereas we note that our algorithm performs more optimally for the finite temperature calculation. We also compare ABACUS to our algorithms for generating an optimal basis for computing local observables following a quench in the interaction strength [2]. In this case we notice a more dramatic increase in performance at finite temperature, highlighting the advantages of our approach for problems where states with multiple particle-hole excitations relative to the thermal state play an important role.

We start by reviewing the general principles of Hilbert space scanning algorithms in Sec. 3.1. We then continue by writing down the simplest algorithm that satisfies the general rules we require of all scanning algorithms in Sec. 3.2. This initial algorithm does not target a specific momentum sector, but with a simple addition to the rules of the algorithm it can be made to do so, as we show in 3.3. In Sec. 3.4 we show how this simple algorithm can be improved upon by adding additional rules relating to the preservation of the number of particle-hole pairs. In Sec. 3.5 we show how the algorithm can be improved even more not by adding additional rules, but by cleverly prioritizing certain parts of the calculation. The resulting algorithm is then compared to the existing state of the art in Sec. 3.6, where we show the advantages of our algorithm for finite temperature computations. We conclude and discuss

how our insights are relevant to scanning algorithms for the spin chain in Sec. 3.7. The work in this chapter was done in collaboration with J.-S. Caux [1].

3.1 Scanning Algorithms

The algorithms for exploring Hilbert space that we present in this chapter can all be generally described as generating a single-rooted tree where each node represents an eigenstate. The algorithms differ in the rules that determine which new eigenstates are generated from a given node, resulting in trees with different topologies even when they are generated from the same initial eigenstate, which we will herein call the seed state. In some cases, the seed state for an algorithm will be directly related to the observable we are trying to compute (for example, for the calculation of the dynamical structure factor, Eq. (3.2), the seed state will be the ground state) while for other problems it may not be (this is the case for the quench problem, which we discuss later). This approach using tree-building algorithms, was shown to be very successful for the computation of correlation functions in integrable models [40]. The purpose of this chapter is to introduce new algorithms for scanning and comparing their properties for different problems.

There are two properties we require of all of our Hilbert space exploration algorithms:

1. *Uniqueness:* No eigenstate should come up more than once when generating a tree of eigenstates.

2. *Completeness*: All eigenstates in a pre-defined sector of Hilbert space must occur in the tree if we give the algorithm infinite computation time.

Property 1 ensures that we can use the eigenstates generated in a tree to perform a summation such as the ones in Eq. (3.2) and Eq. (3.5) without having to separately keep track of which eigenstates we already included. Property 2 comes from the fact that we want to be able to approximate summations such as the one in Eq. (3.2) arbitrarily well given infinite computational resources. If some states would not be generated by the algorithm, this would not be possible.

In the absence of a UV cut-off, any momentum sector of the Lieb-Liniger model is infinite dimensional even in a finite volume system. Therefore we can never truly generate the full corresponding tree with finite computational resources. As such, it is not only relevant what the final tree would look like, but also how it is built, i.e. what it looks like after some finite time. In order to ensure that we spend our time wisely, we can pause the generation of new eigenstates from certain branches of the tree. The goal here is to pause the

generation of eigenstates in the algorithm until they are the most important ungenerated eigenstates that are remaining. The result is that after some finite time we end up with a truncated tree where most of the nodes on the outside of the tree could be used to generate additional new eigenstates.

The way in which we determine which branches to pause at a given time and the interplay between this and the rules for growing the tree as determined by our algorithm can have a strong influence on the way the tree is grown and therefore the quality of our algorithm. To see this, consider an algorithm that satisfies the completeness and overcounting criteria. Now suppose we run this algorithm for some time which gives rise to a growing tree as depicted in Fig. 3.1. The dot-filled nodes represent eigenstates whose weight is below some threshold value whereas the weight of the stripe-filled nodes is bigger, allowing us to visually distinguish between high-weight and low-weight nodes. The algorithm showcased here is not ideal since the high-weight nodes 7 and 8 are generated after the low-weight nodes 5 and 6. When scaled up this means we can, at finite runtime, miss important states because they are effectively locked behind low-weight states. In an ideal algorithm the descendents of a node would therefore always have a lower weight than the parent node.

The extent to which the descendents of nodes have a weight lower than their parents depends on an interplay between the rules of the algorithm and the weighing function considered. Therefore there may be some algorithms which are more compatible with certain weighing functions and others that are more compatible with others. For the Lieb-Liniger model, which we consider in this thesis, we can identify a criterion that is common to the weighing functions we want to consider. This criterion is the number of particle-hole excitations as it is closely related to the matrix element value of the operators we want to consider as explained in Chapter. 2. This allows us to create close to optimal algorithms for this system.

3.2   A basic scanning algorithm

One of the simplest ways of generating a new set of quantum numbers from a given set is by changing one of the quantum numbers by ± 1, the minimal possible value provided that this does not render two quantum numbers equal. Such a change can be identified with a particle hopping to the left or right on the lattice, where a particle hopping to the right looks like

$$\cdots \circ \bullet \circ \bullet \bullet \bullet \circ \bullet \circ \cdots$$
$$\downarrow$$
$$\cdots \circ \circ \bullet \bullet \bullet \bullet \circ \bullet \circ \cdots$$

24

(a) Zeroth step

(b) First step

(c) Second step

(d) Third step

(e) Fourth step

(f) Fifth step
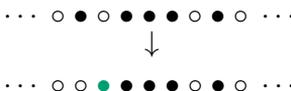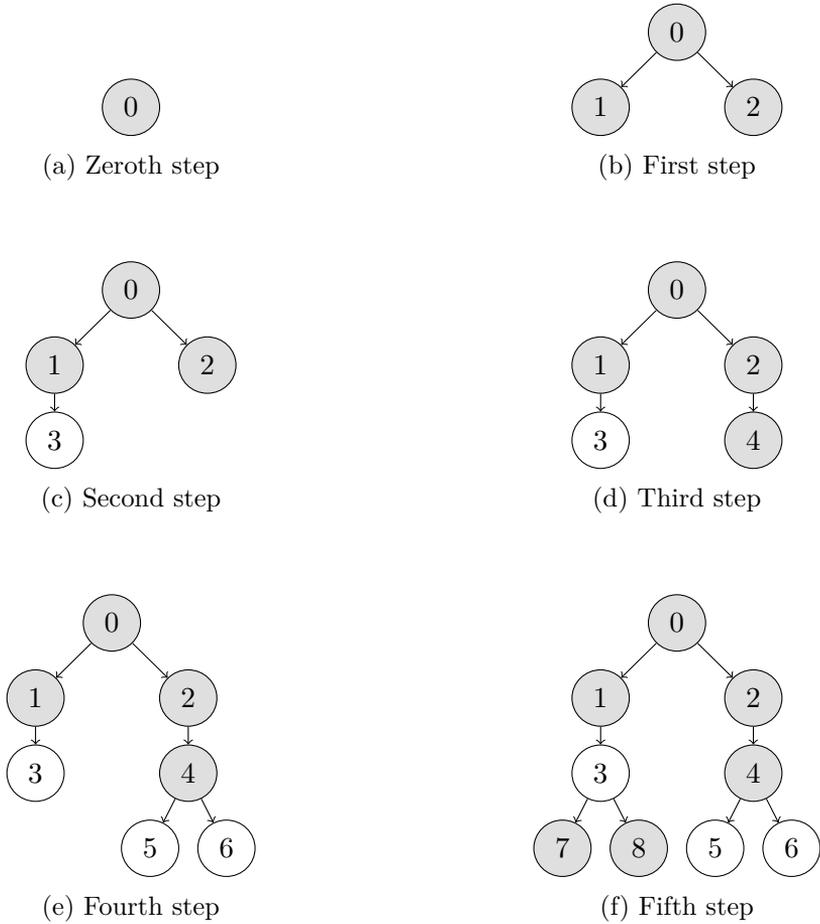
Figure 3.1: Illustration of the state of the tree after the first six steps of a non-optimal algorithm. The dot-filled circles represent eigenstates whose weight is smaller than some threshold value, whereas the stripe-filled circles represent eigenstates whose weight is larger than the same threshhold. An optimal algorithm would therefore not generate stripe-filled circles after dot-filled circles, which we see occuring in the fifth step of the algorithm.

and a particle hopping to the left looks like

$$\cdots \circ \bullet \circ \bullet \bullet \bullet \circ \bullet \circ \cdots$$
$$\downarrow$$
$$\cdots \bullet \circ \circ \circ \bullet \bullet \bullet \circ \bullet \circ \cdots$$
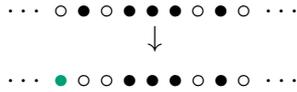
In both cases the first line represents the quantum numbers of the initial state and the line following that represents the quantum numbers of its descendent where one coloured particle has hopped to a neighbouring lattice site.

Such particle hops can be used to formulate rules for generating descendents and building a tree, but we need to impose rules to avoid overcounting. We start with an algorithm where the descendents of a node are those where a single particle has hopped one position to the right or left provided that this does not result in a collision of particles. In the following we identify the issues with this algorithm and propose solutions, the result of which will be our first real scanning algorithm which does not overcount states.

The first issue we consider is that when a particle first moves to the right and then back it will result in the same state we started with. Without any restrictions we can generate subtrees that look like

$$\cdots \circ \bullet \circ \bullet \bullet \bullet \circ \bullet \circ \cdots$$
$$\downarrow$$
$$\cdots \circ \bullet \circ \bullet \bullet \circ \bullet \bullet \circ \cdots$$
$$\downarrow$$
$$\cdots \circ \bullet \circ \bullet \bullet \bullet \circ \bullet \circ \cdots$$
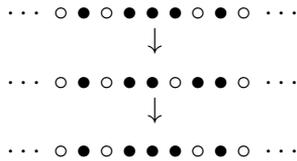
The first rule of our scanning algorithms is therefore that once a particle has moved to either the left or the right, it can only continue moving in that direction. We refer to particles that have moved to the right as rightmovers whereas we refer to particles that have moved to the left as leftmovers. In our visualizations we colour the leftmovers blue and the rightmovers orange.

The second issue is due to there being no preferred ordering of moving particles to the left or right, as with the current rules the following subtree could be generated
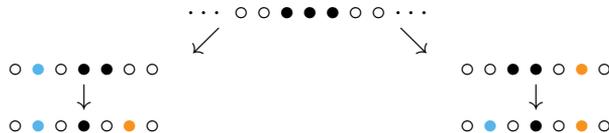
Here we see that the bottom two states are the same even though they are not the same node in the tree and no particle has changed direction. In order to avoid overcounting we impose the rule that if a state has a rightmover then its descendents can not have additional leftmovers. In the subtree we just

considered this means that the state on the bottom right would not have been generated as there is already a rightmover (the right most particle in the above figure).

The final issue that arises is similar to the previous one, only now it involves only moves to the left or right. To understand the issue, note that with our current rules the following subtree could be generated
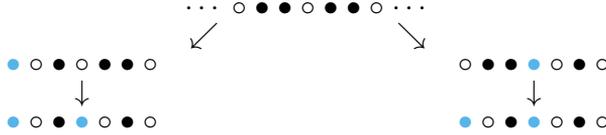
Again, we see that the two lowest configurations of integers are identical, despite being on different branches of the tree. This issue can be avoided by only allowing particles to hop to the left if they are the rightmost leftmover or to its right. This would rule out the configuration on the bottom right, as it is generated by creating a leftmover to the left of the existing leftmover. We can run into the same problem when considering rightmovers giving rise to

which can be avoided by only allowing particles to hop to the right if they are the leftmost rightmover or to its left.

The rules we have imposed thus far constitute the first scanning algorithm that is complete and does not count the same state twice as we will show shortly. For now, let us summarise the rules of the algorithm.

Let $\{I_i\}_{i \leq N}$ be the set of quantum numbers at a node. We denote the leftmoving and rightmoving quantum numbers by $\{I_j^L\}_{j \leq N_L}$ and $\{I_j^R\}_{j \leq N_R}$ respectively.

- **Move quantum numbers to the right:** Generate a descendent for every $I_l \in \{I_i\}_{i \leq N}$ for which $I_l \notin \{I_j^L\}_{j \leq N}$, $I_l \leq I_0^R$, and $I_l + 1 \neq I_{l+1}$ if $l \neq N$ with quantum numbers given by $\{I_i + \delta_{l,i}\}_{i \leq N}$.

- **Move quantum numbers to the left:** If $\{I_k^R\} = \emptyset$ generate a descendent for every $I_l \in \{I_i\}_{i \leq N}$ such that $I_l \geq I_{N_L}^L$, and $I_{l-1} \neq I_l - 1$ if $l \neq 0$ with quantum numbers given by $\{I_i - \delta_{l,i}\}_{i \leq N}$.

The maximal number of descendents is therefore $2N$. What remains is the proof that this algorithm generates a tree containing every set of allowed quantum numbers exactly once regardless of the seed state chosen.

Consider an arbitrary seed state $\{I_i^{SS}\}_{i \leq N}$ and an arbitrary target state $\{\tilde{I}_i\}_{i \leq N}$. In order to show completeness and the absence of overcounting it is sufficient to show that the target state occurs in the generated tree precisely once. This is equivalent to there being one way of applying the rules of the algorithm to get from the seed state to the target state. The quantum numbers of the node we consider at an intermediate step of the algorithm will be denoted by $\{I_i\}_{i \leq N}$.

Since our algorithm does not allow particles to hop to the left in the presence of rightmovers, we first have to move all particles for which $\tilde{I}_l < I_l^{SS}$. We first consider the leftmost particle of this type with index l and claim that we can move it all the way to its target position without collisions. To show this, assume the contrary, which implies that $l \neq 1$ and $I_{l-1}^{SS} \geq \tilde{I}_l$. However, since l was the index of the leftmost particle that had to move to the left, we know that $\tilde{I}_{l-1} \geq I_{l-1}^{SS}$. Combining these statements gives $\tilde{I}_{l-1} \geq I_{l-1}^{SS} \geq \tilde{I}_l$ which is a contradiction. Repeating this argument for all quantum numbers of the seed state that have to be decreased to reach the target state starting with the smallest one, we can move all such quantum numbers to the right position in a unique way.

Having fixed the quantum numbers that have to be decreased, we are left with quantum numbers for which $\tilde{I}_l \geq I_l^{SS}$. This time we claim that we can start from the rightmost particle for which this inequality holds and the rules of our algorithm allow it to be put in its place without collisions. To prove this, we again assume the contrary which implies that $l \neq N$ and that $I_{l+1} \leq \tilde{I}_l$. However, since I_l was the rightmost quantum number that had to be increased, and since we already fixed the leftmovers we have that $I_{l+1} = \tilde{I}_{l+1}$. Together this gives $\tilde{I}_{l+1} = I_{l+1} \leq \tilde{I}_l$ giving the contradiction we require. This finishes the proof of completeness as well as showing that there is no overcounting.

Throughout the remainder of this chapter, we refer to the scanning algorithm developed in this section as stepwise scanning (SWS).

3.3 IMPOSING MOMENTUM CONSERVATION

In the previous section we introduced a scanning algorithm that would generate every state in Hilbert space exactly once, given infinite computational resources. However, often we are interested in a particular momentum sector, so for such problems it is not particularly well-suited. After all, it would mean that we are only interested in a tiny subset of the states that we generate. In this section we introduce a variant of the previous algorithm which restricts itself to a given momentum sector.

To turn stepwise scanning into an algorithm that generates descendents whose momentum is equal to that of their parent, we combine the rules we have for moving particles to the left and right. The momentum of a state

is proportional to the sum of the quantum numbers, Eq. (2.10), so moving one particle one step to the right and another one step to the left ensures that momentum remains preserved. Furthermore, this approach eliminates the second problem we encountered in the previous section due to which we imposed the rule that no particle can hop left in the presence of rightmovers. The resulting algorithm goes as follows.

Let $\{I_i\}_{i \leq N}$ be the set of quantum numbers at a node. We denote the leftmoving and rightmoving quantum numbers by $\{I_j^L\}_{j \leq N_L}$ and $\{I_j^R\}_{j \leq N_R}$ respectively.

- **Generate rightmovers:** Generate an intermediate descendent C_r for every $I_l \in \{I_i\}_{i \leq N}$ such that $I_l \notin \{I_j^L\}_{j \leq N_L}$, $I_l \leq I_0^R$, and $I_l + 1 \neq I_{l+1}$ if $l \neq N$.

- **Generate leftmovers:** For every intermediate descendent C_r generate a descendent for every $I_l \in \{I_i^{C_r}\}_{i \leq N}$ such that $I_l \notin \{I_j^{C_r,R}\}_{j \leq N_{C_r,R}}$, $I_l \geq I_{N_{C_r,L}}^L$, and $I_l - 1 \neq I_{l-1}$ if $l \neq 0$.

The scanning routine described here, which we call momentum preserving stepwise scanning (SWS-MP), generates at most N^2 descendents in contrast to the at most $2N$ descendents in regular stepwise scanning. The difference arises due to the fact that we imposed momentum conservation, which led us to essentially apply first the first step of regular stepwise scanning, generating at most N intermediate descendents, and then applying the second step of the stepwise scanning algorithm to these intermediate descendents.

To show that this algorithm is also complete in the sense that it can generate any state whose momentum is equal to that of the seed state, consider a random seed state $\{I_i^{SS}\}_{i \leq N}$ and a random target state $\{\tilde{I}_{i \leq N}\}$ with the same momentum. There are indices $k \in \{k_1, \ldots, k_{N_L}\}$ such that $\tilde{I}_k \leq I_k^{SS}$ as well as indices $l \in \{l_1, \ldots, l_{N_R}\}$ such that $\tilde{I}_l \geq I_l^{SS}$ which represent the indices of what will become the leftmovers and rightmovers respectively. Note that the size of these sets of indices can be different, as momentum preservation only requires the number of hops to the right and left to be preserved. In order to reach the target state from the seed state we again have to start by moving the leftmost leftmover, i.e. I_{k_1}, and the rightmost rightmover, i.e. I_{N_R}. The proof that these particles can hop to their target positions without collisions is exactly the same as the proof of completeness for regular stepwise scanning, so for that we refer the reader to the previous section. The same argument holds for the rightmovers being able to move to their target positions. Since there is still a unique order of moves by which we can reach the target state given the seed state, we also have no overcounting.

To understand the problems that arise when we try to apply this algorithm to situations where the momentum of the states we want to scan for is not

equal to the natural candidate for the seed state, consider the example of the dynamical structure factor, see Eq. (3.2). Here the reference state we are interested in is the ground state $|0\rangle$ and we want to use a scanning algorithm to find the states $|\alpha\rangle$ at some fixed value of momentum most important to the summation. However, since the states in the intermediate summation do not belong to the same momentum sector as the ground state we cannot use a momentum preserving stepwise scanning algorithm in order to find them. In principle we can choose some other state from the targeted momentum sector and use it as a seed state to the algorithm, but this raises the question of which state to choose. Even though in the limit where we have infinite computational resources this does not matter, it does affect what the tree looks like after a finite amount of time. After all, a bad choice can lead to the situation where important contributions are only generated after a long time because they are far down in the tree or are "hidden" as descendents of unimportant states. In fact, it is unclear if there even exists a seed state that would not lead to a Bethe tree with undesirable properties in this case.

In order to avoid having to choose a seed state for the target momentum sector, we choose to tweak the rules of the scanning algorithm such that we can use a given reference state (in the case of the dynamical structure factor, the ground state) despite its momentum not being equal to that of the momentum sector we are interested in. After all, we want to generate the states with few particle-hole pairs first as these are the states we generally expect to have the largest matrix elements for local operators. Taking the reference state to be the seed state ensures that these few particle-hole states are generated because the rules of our algorithm generate descendents with at most two additional particle-hole pairs. In order to generate states in the targeted momentum sector from a given reference state we add the rule that if the momentum of a node is smaller than the target momentum, its descendents are generated by following the first step of the algorithm, whereas if it is larger the rules of the second step are used. Note that applying only one of the two steps of momentum preserving stepwise scanning generates descendents whose momentum is changed by the minimal amount compared to their parent state. As such, there can be no overshooting of the target momentum sector. Furthermore, every branch created from the seed state reaches the target momentum sector in $\frac{k_{target}L}{2\pi}$ steps or dies off before then.

In order to assess the quality of this algorithm, we consider again the dynamical structure factor introduced in Eq. (3.2). Starting from the ground state, we can generate a tree of eigenstates where we use the weighing function that selects for states which contribute most strongly for the saturation of the f sum rule as defined in Eq. (3.3). For now, we generate descendents node by node starting each time with one of the highest weight nodes as we elaborate on more in section 3.5. The results for this calculation for three different

values of the interaction strength are shown in Fig. 3.2a, b, c We see that in all three cases we reach close to optimal convergence with very few states meaning the most important states are generated first.

At finite temperature, we can do the same calculation provided that we replace the ground state with the representative state of the thermal state we wish to consider [26]. The calculation of the dynamical structure factor at finite temperature is inherently more difficult, but it also turns out that the simple scanning algorithm we have developed is not optimal for the finite temperature case. To understand why, consider Fig. 3.2d, e, f where we again consider the the f sumrule convergence. Besides the significantly poorer rate of convergence, we observe clear plateaus where convergence stagnates. This indicates that the algorithm is not successfully generating the states that contribute to the finite-temperature correlation function most strongly first. In the next section we propose changes to our algorithm that lead to better convergence.

3.4 Imposing additional constraints

Thus far we have proposed algorithms where we generate descendents by moving quantum numbers by the minimal amount at each step of the algorithm, assuming this minimal change would result in the most important states first being generated first. However, it turns out that there is a property of the quantum numbers defining an eigenstate which is an important indicator of its importance for most of the weighing functions we are interested in that we have ignored thus far. This property is the number of particle-hole pairs of a state with respect to a given state of importance to the calculation at hand, herein the reference state. In this section we present algorithms that use knowledge of this property to produce more efficient algorithms.

What makes the number of particle-hole pairs, introduced in Sec. 2.3, an important property to consider is that generally the fewer particle-hole pairs a state has, the bigger the off-diagonal matrix element between it and the reference state for local operators. This means that the states with few particle-hole pairs are generally those with the largest weights, meaning we should generate them first. For the regular and momentum preserving stepwise scanning algorithms, however, a descendent can have fewer particle-hole pairs than its parent. For example, consider a seed state given by

○ ● ○ ● ○ ● ○ ○ ○

then it generates the following subtree where a particle and hole annihilate

(a) $c = 1$, $T = 0$ (b) $c = 4$, $T = 0$ (c) $c = 16$, $T = 0$

(d) $c = 1$, $T = 1$ (e) $c = 4$, $T = 1$ (f) $c = 16$, $T = 1$
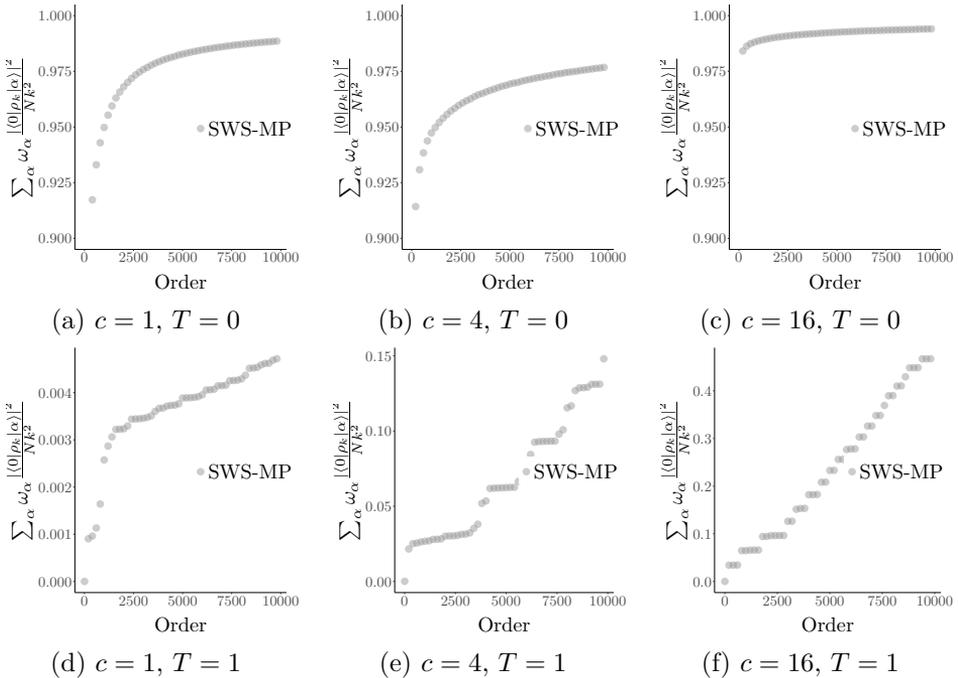
Figure 3.2: Saturation of the f sumrule with the number of states included in the summation. Starting from the ground state for (a)-(c) and the representative thermal state at $T = 1$ for (d)-(f), we generated 10,000 states using momentum preserving stepwise scanning for a target momentum of $k = \pi$, and $N = 128 = L$. We plot the sum rule saturation after every 200 states for $c = 1$ in (a) and (d), for $c = 4$ in (b) and (e), and for $c = 16$ in (c) and (f). Convergence is near perfect after very few states in the zero temperature case, whereas in the finite temperature case convergence is poor for the number of states considered. Furthermore, the interaction strength is seen to be an important variable in the finite temperature case, with smaller interaction strengths corresponding to poorer convergence.

one another
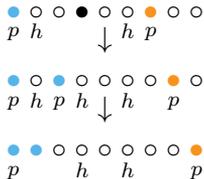
Taking on board that matrix elements of local operators depend upon the number of particle-hole excitations in a structured way, we can modify the previously proposed algorithms to more efficiently explore the Hilbert space.

One way to proceed is to simply forbid the annihilation of particle-hole pairs within a modified algorithm. However, this approach comes at a cost since adding this restriction to the rules of the regular and momentum preserving stepwise scanning algorithms breaks completeness. In order to regain this necessary feature, we allow quantum numbers that have not moved yet to hop into the position of a hole if it lies between it and one of the neighbouring quantum numbers. For example, this allows the following subtree to be generated:
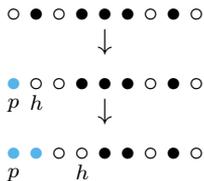
Note that this means that this allows for quantum number jumps of more than one position. Applying these changes to stepwise scanning scanning leads to the following rules, which constitute the leapwise scanning algorithm (LWS).

Let $\{I_i\}_{i \leq N}$ be the set of quantum numbers at a node. We again denote the leftmoving and rightmoving quantum numbers by $\{I_j^L\}_{j \leq N_L}$ and $\{I_j^R\}_{j \leq N_R}$ respectively. Furthermore, we label the positions of the holes as $\{I_j^h\}_{j \leq N_h}$ and the particles by $\{I_j^p\}_{j \leq N_p}$.

- **Generate higher momentum descendents:** Generate a descendent for every $I_l \in \{I_i\}_{i \leq N}$ for which $I_l \notin \{I_j^L\}_{j \leq N_L}$, $I_l \leq I_0^R$, and either

  - $I_l + 1 \neq I_{l+1}$ if $l \neq N$ with quantum numbers given by $\{I_i + \delta_{l,i}\}_{i \leq N}$, or

  - there exists a k such that $I_l < I_k^h < I_{l+1}$ with quantum numbers given by $\{I_i + \delta_{i,l}(I_k^h - I_i)\}_{i \leq N}$

- **Generate lower momentum descendents:** If $\{I_k^R\} = \emptyset$ generate a descendent for every $I_l \in \{I_i\}_{i \leq N}$ such that

33

- ○ $I_l \geq I_{N_L}^L$, $(I_l - 1) \notin \{I_j^h\}_{j \leq N_h}$, and $I_{l-1} \neq I_l - 1$ if $l \neq 0$ with quantum numbers given by $\{I_i - \delta_{l,i}\}_{i \leq N}$

- ○ there exists a k such that $I_{l-1} < I_k^h < I_l$ with quantum numbers given by $\{I_i + \delta_{i,l}(I_k^h - I_i)\}$

The proof that this algorithm is complete and does not overcount carries over directly from the proof for the stepwise scanning algorithm.

Like stepwise scanning, leapwise scanning is not suitable for targeting a fixed momentum sector. One step towards a solution of this problem is to, like before, combine the first and second step of the leapwise scanning algorithm. However, in the current case the resulting algorithm does not preserve momentum since the momentum increasing move from the first step and the momentum decreasing momentum from the second step may not cancel. The solution to this issue is the same as the solution to the problem we had with momentum preserving leapwise scanning algorithm when we wanted to consider a target momentum sector whose momentum was different from that of the reference state. We generate the momentum increasing descendents for states whose momentum is smaller than the target momentum and momentum decreasing descendents for states whose momentum is larger. For states at the right momentum we combine the steps as we do in momentum preserving stepwise scanning. This algorithm, which we refer to as momentum preserving leapwise scanning (LWS-MP), can be summarised as follows.

Let $\{I_i\}_{i \leq N}$ be the set of quantum numbers at a node. We denote the leftmoving and rightmoving quantum numbers by $\{I_j^L\}_{j \leq N_L}$ and $\{I_j^R\}_{j \leq N_R}$ respectively. Furthermore, we label the positions of the holes as $\{I_j^h\}_{j \leq N_h}$ and the particles by $\{I_j^p\}_{j \leq N_p}$.

- **Generate rightmovers:** Generate an intermediate descendent C_r for every $I_l \in \{I_i\}_{i \leq N}$ such that $I_l \notin \{I_j^L\}_{j \leq N}$, $I_l \leq I_0^R$ and either

  - ○ $I_l + 1 \notin \{I_j^h\}_{j \leq N_h}$, and $I_l + 1 \neq I_{l+1}$ if $l \neq N$ with quantum numbers given by $\{I_i + \delta_{l,i}\}_{i \leq N}$, or

  - ○ there exists a k such that $I_l < I_k^h < I_{l+1}$ with quantum numbers given by $\{I_i + \delta_{i,l}(I_k^h - I_i)\}_{i \leq N}$

- **Generate leftmovers:** For every intermediate descendent C_r generate a descendent for every $I_l \in \{I_i^{C_r}\}_{i \leq N}$ such that $I_l \notin \{I_j^R\}_{j \leq N_R}$, $I_l \geq I_{N_L}^L$, and either

  - ○ $(I_l - 1) \notin \{I_j^h\}_{j \leq N_h}$, and $I_{l-1} \neq I_l - 1$ if $l \neq 0$ with quantum numbers given by $\{I_i - \delta_{l,i}\}_{i \leq N}$, or

  - ○ there exists a k such that $I_{l-1} < I_k^h < I_l$ with quantum numbers given by $\{I_i + \delta_{i,l}(I_k^h - I_i)\}$
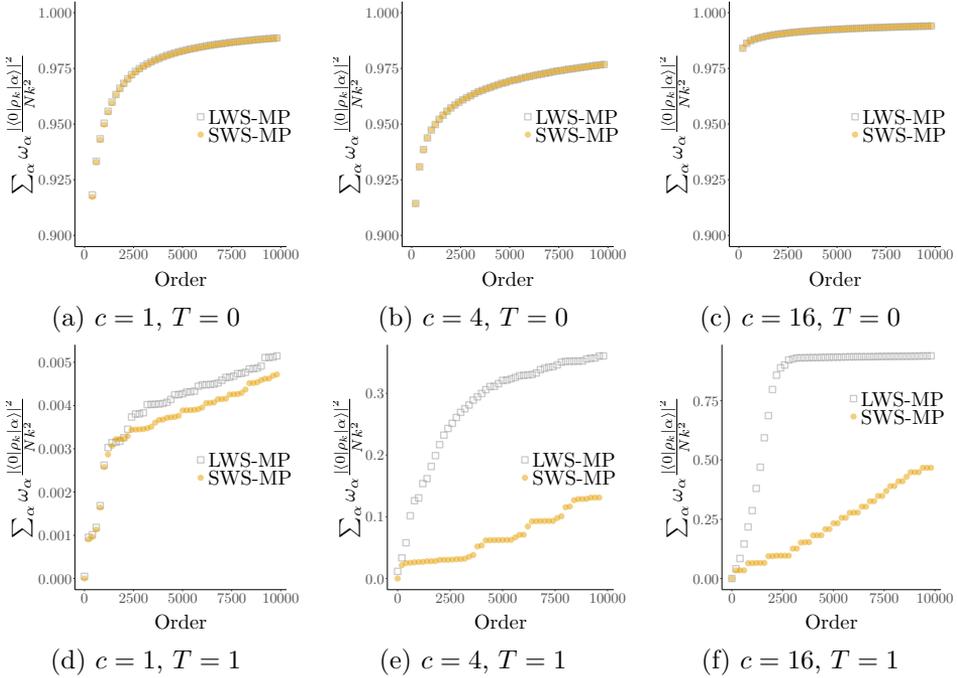
Figure 3.3: Comparison of the saturation of the f sumrule with the number of states included in the summation between momentum preserving stepwise (SWS-MP) and momentum preserving leapwise (LWS-MP) scanning. Starting from the ground state for (a)-(c) and the representative thermal state at $T = 1$ for (d)-(f), we generate 10,000 states for a target momentum of $k = \pi$, and $N = 128 = L$. We plot the sum rule saturation after every 200 states for $c = 1$ in (a) and (d), for $c = 4$ in (b) and (e), and for $c = 16$ in (c) and (f). Convergence is identical for both momentum preserving stepwise and leapwise scanning zero temperature, but not for the finite temperature case where momentum preserving leapwise scanning performs better. The plateaus at large and intermediate interaction strengths are barely noticeable for momentum preserving leapwise scanning whereas they remain pronounced in the case where $c = 1$.

If the momentum of the node is smaller than the target momentum, we only do the first step and the intermediate descendents become the descendents. On the other hand, if the momentum of the node is smaller than the target momentum we apply the second step to the node. The proof that the resulting algorithm is complete and does not overcount is the same as the proof for the momentum preserving stepwise scanning algorithm.

At zero temperature momentum preserving leapwise scanning gives the practically the same results as momentum preserving stepwise scanning as can be seen in Fig. 3.3a, b, c. To understand this, note that at $T = 0$ we start from the ground state which consists of a single block of neighbouring quantum numbers without holes. The only initial excitations that are then allowed by our rules are ones that move the outermost quantum numbers out leaving behind holes. This in turn gives space for more particle-hole pairs to be created, but no jumps over vacant positions that are not holes will occur during momentum preserving leapwise scanning. As a result, the descendents generated by momentum preserving stepwise and leapwise scanning as well as the resulting trees are virtually the same.

At finite temperature momentum preserving leapwise scanning outperforms momentum preserving stepwise scanning as can be seen in part Fig. 3.3d, e, f. In this case, we start from a seed state where there is no longer a zero temperature Fermi sea, but instead there are vacant positions between the quantum numbers, i.e. the quantum numbers of the seed state are no longer of the form $\{a, a+1, a+2, \dots\}$ but rather something like $\{a, a+3, a+4, a+7, \dots\}$ for example. Therefore, when descendents are generated where quantum numbers neighbouring vacant positions are moved, they leave behind holes surrounded by vacant positions. For momentum preserving leapwise scanning, states with the same number of particle-hole excitations can be created in such a scenario by letting another quantum number jump over the vacant positions to this newly created hole position, whereas in the case of momentum preserving stepwise scanning a bunch of intermediate states would have to be created. Not having to generate these intermediate states is what leads to the increase in efficiency.

3.5 Beyond the topology of the tree

The algorithms we have introduced thus far determine the topology of the tree corresponding to a given seed state and target momentum sector, but they do not determine the order in which the nodes of the tree are generated. For the numerics displayed in the previous sections we adopted the rule that, by default, we pause all branches and after each step of generating new descendents we find the highest weight state all of whose descendents we then generate. However, in Fig. 3.3d, e, f, we saw that this approach produces

plateaus which indicate that less important states are generated before their more important counterparts. The main reason for this is that at a given node there are descendents with different levels of importance (as, for example, they have different numbers of particle-hole excitations). By generating all descendents of a node at the same time, states with more particle-hole pairs can be generated before states with fewer particle-hole pairs further down the tree. In this section we show how splitting the descendents of the momentum preserving leapwise scanning algorithm into three groups and treating them separately results in a more efficient algorithm.

The division of the descendents of a given node into three groups is done based on the number of additional particle-hole pairs they have compared to their parent, which is either zero, one, or two. Given a node of the tree, we expect the states in the group of descendents with the same number of particle-hole pairs as the parent node to be of similar importance as the parent. Since we always consider the paused node with the highest weight, its descendents with the same number of particle-hole pairs are the ones we expect to be the most important unexplored eigenstates. Therefore we always start by generating these descendents if they have not yet been generated. In contrast to regular momentum preserving leapwise scanning, we do not also generate the descendents with additional particle-hole pairs at the same time.

Once we have generated the descendents with the same number of particle-hole pairs for a given node, we can generally not yet forget about this node as we could previously, since we still have to generate some of its descendents with additional particle-hole pairs. Therefore, we store it in a different list which tracks the states for which the descendents with no additional particle-hole pairs have been generated. Furthermore, we give it a weight by choosing a random descendent with an additional particle-hole pair and computing its weight, which is then used to weigh the parent node in this secondary list. Having populated both the initial list of paused nodes and this secondary list, we can choose to generate new descendents by either considering a state from the original list and generating more states with the same number of particle-hole pairs as their parent, or by considering a state from the secondary list, and generating the descendents of the node with one additional particle-hole pair. Which choice we make depends on where the state with the highest weight is, as it could be in the first list where its weight is that of the eigenstate, whereas in the second list it would be the weight of a randomly chosen descendent with an additional particle-hole pair.

In some cases, the full list of descendents would also have included states with two additional particle-hole pairs with respect to their parent node. In this case we follow a similar procedure and move the node from the secondary list to a third where we keep nodes whose descendents with zero and one additional particle-hole pair have been generated and associate to it another
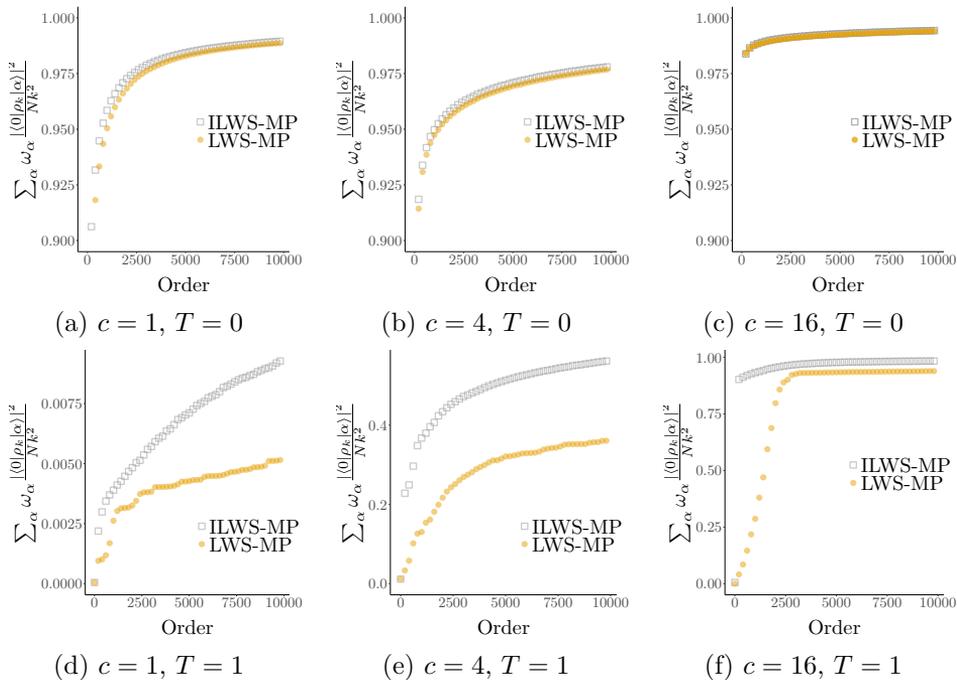
Figure 3.4: Comparison of the saturation of the f sumrule with the number of states included in the summation between momentum preserving leapwise scanning (LWS-MP) and improved momentum preserving leapwise scanning (ILWS-MP). Starting from the ground state for (a)-(c) and the representative thermal state at $T = 1$ for (d)-(f), we generate 10,000 states for a target momentum of $k = \pi$, and $N = 128 = L$. We plot the sum rule saturation after every 200 states for $c = 1$ in (a) and (d), for $c = 4$ in (b) and (e), and for $c = 16$ in (c) and (f). Convergence at zero temperature sees a tiny improvement for improved momentum preserving leapwise scanning to regular momentum preserving leapwise scanning and remains near perfect for very few states. At finite temperature, we see a dramatic increase in performance in both overall convergence as well as the number of states required to achieve this convergence for improved momentum preserving leapwise scanning, eliminating all plateaus. Still, absolute convergence remains a challenge at lower values of the interaction strength for the number of states considered regardless of the algorithm used.

representative weight obtained by computing the weight of one of the states with two additional particle-hole pairs. We then repeat the procedure outlined above with three lists. The resulting procedure is what we refer to as improved momentum preserving leapwise scanning (ILWS-MP).

The advantage of this approach is that, at the cost of computing at most two additional matrix elements per node, we can ensure that we are generating the states most important to the calculation under consideration. Despite this representing an additional computational effort, it can still represent a net gain as it can allow us to generate far fewer states for a given accuracy of the calculation under consideration. Whether the additional cost of computing these matrix elements is worth it can depend, for example, on the seed state under consideration. To illustrate this, consider the zero temperature dynamical structure factor calculation whose results are illustrated in Fig. 3.4a, b, c. Here we see that the performance the of regular and improved momentum preserving leapwise scanning is virtually identical. On the other hand, when considering the finite temperature equivalent as illustrated in Fig. 3.4d, e, f, we see a big difference in performance. The algorithm where we allow some descendents to be generated at a later point in time outperforms the algorithm where all descendents are generated. Not only does it generate the same states quicker, but it also appears to achieve higher sum rule saturations in part because the algorithm where all descendents are generated gets stuck generating unimportant states.

3.6 Comparison to the state of the art

In this section, we compare the improved momentum preserving leapwise scanning algorithm to the most recent version of the state of the art software for the computation of correlation functions in the Lieb-Liniger model called ABACUS [40]. In Sec. 3.6.1 we consider the dynamical structure factor at zero and finite temperature. Since ABACUS was developed in part to compute the dynamical structure factor, it makes it the ideal candidate for a fair comparison. In Sec. 3.6.2 we consider the generation of an optimal basis for the interaction quench. In both cases, we use commit 08C85CF590 of ABACUS as available at https://jscaux.org/git/jscaux/ABACUS. This version of ABACUS generates a tree topologically equivalent to the one generated by stepwise scanning, generating all its descendents by moving quantum numbers by at most one position. The difference between ABACUS and our momentum preserving stepwise scanning is the way the tree is built. Most crucially, ABACUS deals with the need to generate states with few particle-hole pairs by forcing the generation of branches of the tree where a particle and hole recombine. Therefore we can view it as stepwise scanning with forced recombinations, which we refer to with its abbreviation (SWS-FR) throughout this section. As we shall see, this

is not the optimal way to deal with finite temperature states.

3.6.1 THE DYNAMICAL STRUCTURE FACTOR

For the ground state dynamical structure factor calculation, illustrated in Fig. 3.5a, b, c, we see that the performance of SWS-FR and our improved momentum preserving leapwise scanning routine are virtually identical. Similarly to the discussion comparing momentum preserving stepwise and leapwise scanning, there is not a lot of freedom on how to generate the states when starting from the ground state. The differences which do exist are caused by states being generated in a different order due to a different way of building the same tree.

At finite temperature, we observe an increase in performance for the improved momentum preserving leapwise scanning compared to SWS-FR for all values of the interaction strength as illustrated in Fig. 3.5. At finite temperature, the seed states we consider no longer consist of a contiguous block of quantum numbers with empty positions between neighbouring quantum numbers being introduced. In this case, creating a new particle-hole pair leaves a hole that may not neighbour another quantum number that has not moved yet. For the improved momentum preserving leapwise scanning algorithm, states with the same number of particle-hole pairs can be direct descendents of this state by allowing the quantum number to move by more than one position. For example, we saw that we can have subtrees like

$$\circ \bullet \circ \bullet \bullet \bullet \circ \bullet \circ$$
$$\downarrow$$
$$\bullet \circ \circ \bullet \bullet \bullet \circ \bullet \circ$$
$$p \; h \qquad \downarrow$$
$$\bullet \bullet \circ \circ \bullet \bullet \circ \bullet \circ$$
$$p \qquad h$$
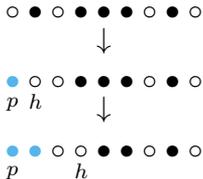
For the stewpwise scanning algorithm with forced recombinations such jumps are not allowed, forcing it to first create an additional particle-hole pair and move it to annihilate the initial hole leading to the following subtree:

$$\circ \bullet \circ \bullet \bullet \bullet \circ \bullet \circ$$
$$\downarrow$$
$$\bullet \circ \circ \bullet \bullet \bullet \circ \bullet \circ$$
$$p \; h \qquad \downarrow$$
$$\bullet \circ \bullet \circ \bullet \bullet \circ \bullet \circ$$
$$p \; h \; p \; h \quad \downarrow$$
$$\bullet \bullet \circ \circ \bullet \bullet \circ \bullet \circ$$
$$p \qquad h$$
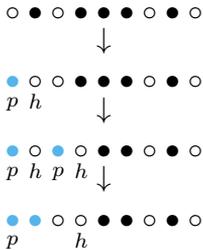
The difference may look innocent in this simple example, but as the distances between quantum numbers in the initial state grow and the number of such
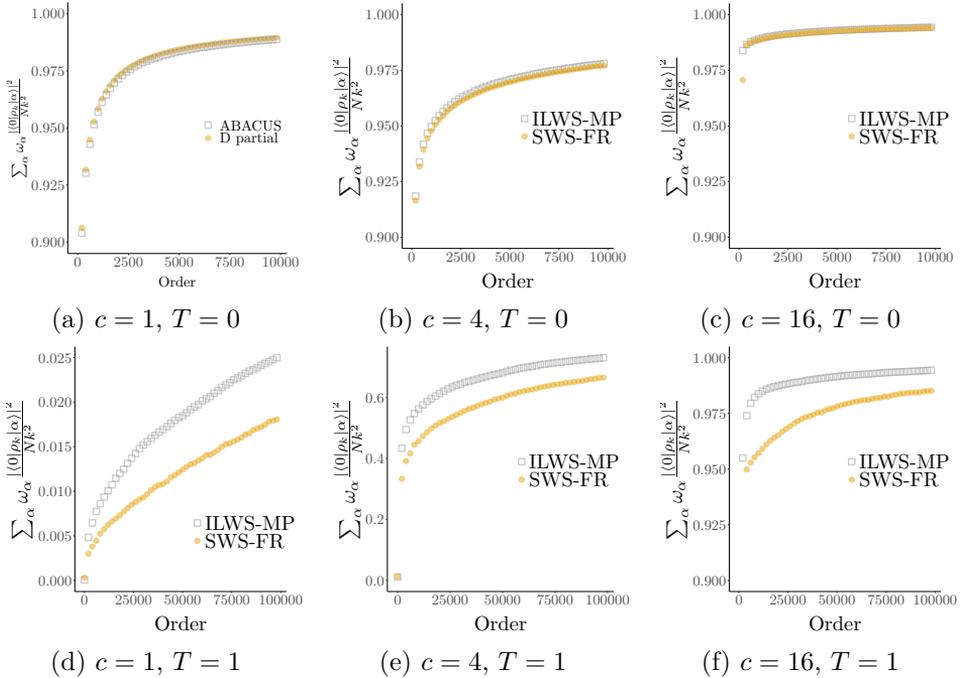
Figure 3.5: Comparison of the saturation of the f sumrule with the number of states included in the summation between the improved momentum preserving leapwise scanning (ILWS-MP) and momentum preserving stepwise scanning with forced recombinations. Starting from the ground state for (a)-(c) and the representative thermal state at $T = 1$ for (d)-(f), we generate 10,000 states for (a)-(c) and 100,000 states for (d)-(f) for a target momentum of $k = \pi$, and $N = 128 = L$. We plot the sum rule saturation after every 200 states for $c = 1$ in (a) and (d), for $c = 4$ in (b) and (e), and for $c = 16$ in (c) and (f). Convergence at zero temperature is virtually identical, whereas the improved momentum preserving scanning outperforms SWS-FR at finite temperature.

(a) $c = 1$, $T = 1$ (b) $c = 4$, $T = 1$ (c) $c = 16$, $T = 1$
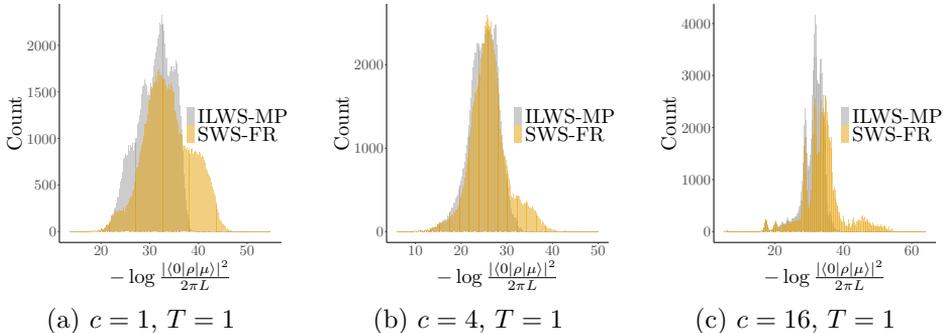
Figure 3.6: Comparison of the histograms of the f sumrule weights between the improved momentum preserving leapwise scanning (ILWS-MP) and SWS-FR. Starting from the representative thermal state at $T = 1$, we generate 10,000 states for a target momentum of $k = \pi$, and $N = 128 = L$. We plot the resulting histogram for $c = 1$ in (a), for $c = 4$ in (b), and for $c = 16$ in (c). We see that the improved momentum preserving leapwise scanning and SWS-FR find the same states with very large weights (those on the left), but SWS-FR generates more less important states (the states on the right).

isolated quantum numbers increases, the number of additional states that have to be generated in order to reach all states with the same number of particle-hole pairs grows rapidly.

The fact that SWS-FR generates such less important intermediate states can be seen from Fig. 3.6. where we consider the histograms of the first 100,000 states generated by each algorithm. In these histograms, the x axis measures the importance of an eigenstate for the saturation of the f-sum rule, where the importance decreases as x increases. The rightmost gray part of the histograms represent the intermediate states that are generated too soon due to the sub-optimal topology of the tree that SWS-FR is building. Since this problem is related to the topology of the tree, it occurs independently of the interaction strength considered.

3.6.2 Generating a basis for the interaction quench

Another way to benchmark our algorithm is by considering the problem of a quench in the interaction strength from c_i to c_f [2]. In this case, we want to find an approximate expansion of some initial eigenstate $|\Psi_0\rangle$ of $H(c_i)$ in terms of eigenstates of $H(c_f)$. Truncated spectrum methods can be used to obtain the expansion coefficients b_α in

$$|\Psi_0(t)\rangle = \sum_\alpha b_\alpha e^{-iE_\alpha t}|\alpha\rangle. \tag{3.6}$$

(a) $c = 1$, $T = 0$ (b) $c = 4$, $T = 0$ (c) $c = 16$, $T = 0$

(d) $c = 1$, $T = 1$ (e) $c = 4$, $T = 1$ (f) $c = 16$, $T = 1$
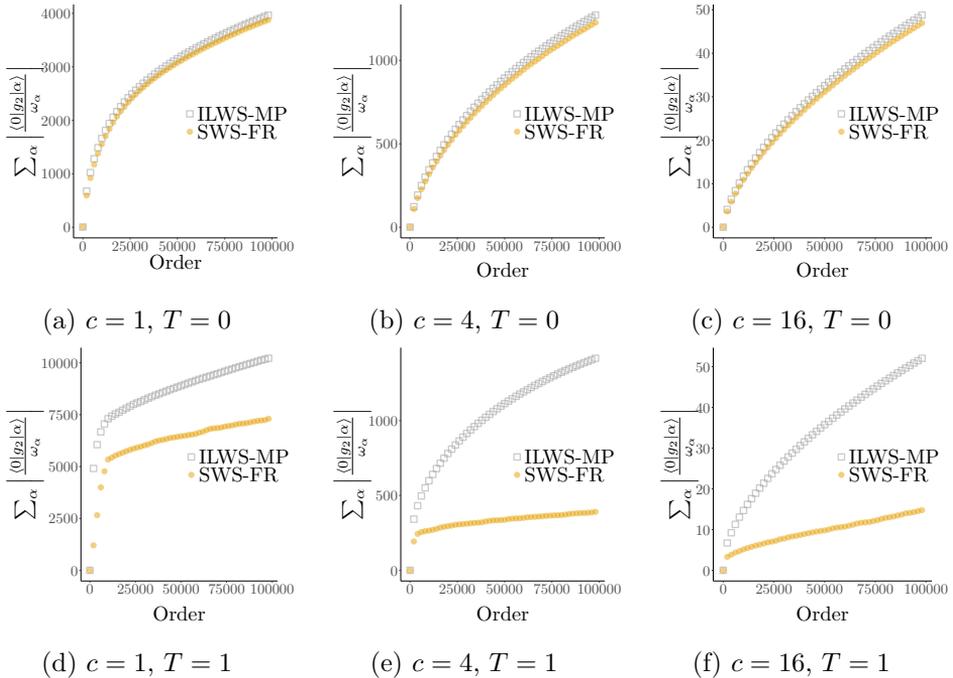
Figure 3.7: Comparison of the sum of eigenstate weights with the number of states included in the summation between the improved momentum preserving leapwise scanning (ILWS-MP) and momentum preserving stepwise scanning with forced recombinations (SWS-FR). Starting from the ground state for $(a)-(c)$ and the representative thermal state at $T = 1$ for $(d)-(f)$, we generate 10,000 states and 100,000 states respecitvely for a target momentum of $k = \pi$, and $N = 128 = L$. We plot the sum rule saturation after every 200 states for $c = 1$ in (a) and (d), for $c = 4$ in (b) and (e), and for $c = 16$ in (c) and (f). The differences between SWS-FR and the improved momentum preserving leapwise scanning for the sums at zero temperature are small whereas at finite temperature the latter outperforms SWS-FR by a large margin.

(a) $c = 1$, $T = 1$ (b) $c = 4$, $T = 1$ (c) $c = 16$, $T = 1$
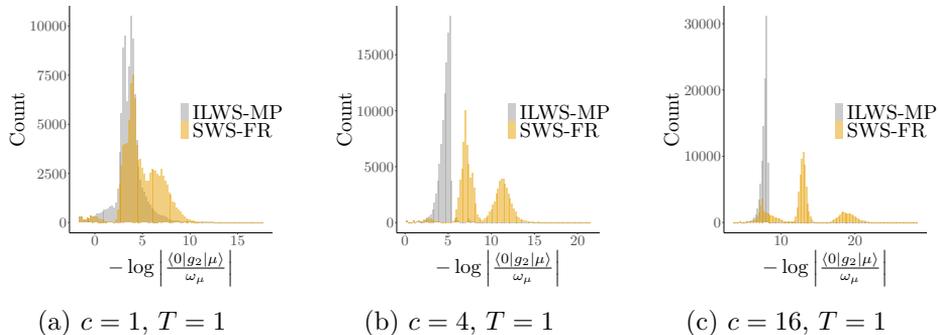
Figure 3.8: Comparison of the histograms of the first 10,000 states generated by SWS-FR and the improved momentum preserving leapwise scanning (ILWS-MP) for the basis generation problem at $k = 0$ and $N = 128 = L$ at $T = 1$. We see that there is some overlap in which states are generated by both algorithms for the large weight states (on the left) but we also see that SWS-FR generates far more low weight states. At larger interaction strengths, we see that these lower weight states generated by SWS-FR are grouped in two distinct bumps.

In order to obtain an expansion that captures the time evolution following the quench accurately, we need to choose our basis states $|\alpha\rangle$ wisely. In [2] we showed that a good estimate of the importance of an eigenstate $|\alpha\rangle$ is given by

$$w(|\alpha\rangle) = \left| \frac{\langle \Psi_0 | g_2 | \alpha \rangle}{E_{\Psi_0} - E_\alpha} \right|. \tag{3.7}$$

The task of our scanning algorithms is therefore to generate the eigenstates with the largest weights.

In order to compare our algorithm to the momentum preserving stepwise scanning algorithm with recombinations, we compare the sums of the weights generated having generated i eigenstates as shown in Fig. 3.7. We see that at zero temperature, the results from the improved momentum preserving leapwise scanning are only marginally better than those of SWS-FR, which can again be explained by the fact that at zero temperature the topology of trees generated by both algorithms is identical. At finite temperature however, we see a much more pronounced difference than we saw for the finite temperature calculation of the dynamical structure factor.

So why is the difference between SWS-FR and the improved momentum preserving leapwise scanning for this basis generation problem so much more pronounced compared to the finite temperature dynamical structure factor calculation? Note that the f sum rule for the dynamical structure factor is dominated by the one particle-hole sector, which both algorithms are perfectly

capable of generating in full, as it contains at most N states at fixed momentum. For the basis generation problem considered here however, the states contributing most strongly are those in the two particle-hole sector (states in the one particle-hole sector cannot have the same momentum as the seed state). The process of generating a given state in the two particle-hole sector is however very different between the improved momentum preserving leapwise scanning and SWS-FR. As we mentioned, SWS-FR is forced to generate intermediate states, of which there are increasingly many as temperature increases. The improved momentum preserving leapwise scanning on the other hand allows for jumps of quantum numbers by more than one position allowing it to generate these states without intermediate states resulting in a more efficient calculation. The choice of topology of the tree combined with the versatile way of building the tree is thus what allows the algorithm to focus on the contributions that are most important at a given point in the calculation, as illustrated by the histograms of contributions in Fig. 3.8

In order to substantiate our claim that our algorithm outperforms SWS-FR due to the different topology of the tree, consider Fig. 3.9. Here we see a breakdown of the histograms of weights generated by either the improved momentum preserving leapwise scanning (a) through (c) or SWS-FR (d) through (f) based on the number of particle-hole sectors the states whose weights are displayed are from previously shown in Fig. 3.8 based on the number of particle-hole pairs of a given contribution. We see that SWS-FR spends its time generating a lot of unimportant states from the three and four particle-hole sectors in order to generate the important contributions from the two particle-hole sector. Since for the improved momentum preserving leapwise scanning any state can be generated without intermediate states with more particle-hole pairs, it is able to avoid such problems. We therefore conclude that the improved momentum preserving leapwise scanning algorithm is better suited for dealing with problems at finite temperature than algorithms that are not able to strictly preserved the number of particle-hole pairs like stepwise scanning or SWS-FR. We have seen that this is especially true if the calculation under consideration is not one domtinated by the one particle-hole sector.

3.7 Conclusions

Despite the powerful analytical tools that have been developed for integrable systems, performing a numerical evaluation over eigenstates is still often an inevitable step required to compute correlation functions. It is therefore crucial that we have good Hilbert space exploration algorithms that allow us to accurately and efficiently approximate such summations. Here we have reviewed the basic principles that such an algorithm has to satisfy and developed a
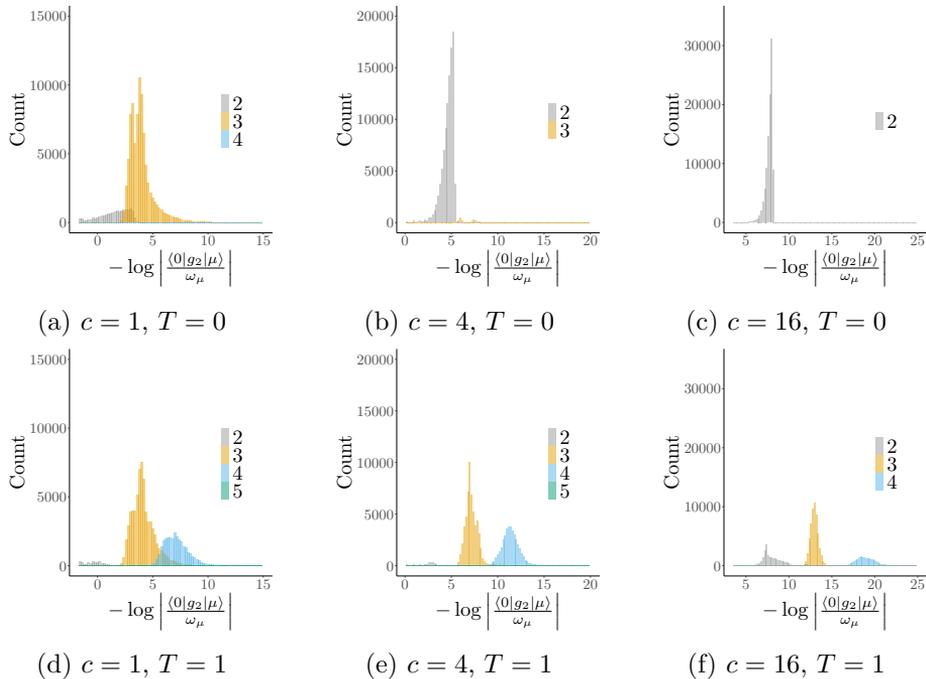
Figure 3.9: Breakdown of the histograms shown in Fig. 3.8 based on the number of particle-hole pairs for improved momentum preserving leapwise scanning (ILWS-MP) (a)-(c) and SWS-FR (d)-(f). We see that the bumps of unimportant states generated by SWS-FR are states from the three and four particle-hole sectors whereas improved momentum preserving leapwise scanning is capable of sticking to lower particle-hole sectors.

number of concrete examples that satisfy these criteria. Starting from the most basic algorithm satisfying these basic principles, we considered its shortcomings one by one and proposed solutions resulting in incremental changes that led us to the final algorithm. Finally, we compared this algorithm to the state of the art for the dynamical structure factor at zero and finite temperature as well as the problem of generating an efficient basis for a quench in the interaction strength.

Overall, our algorithms can be viewed as an algorithm for building a single-rooted tree where every node of the tree represents an eigenstate and the algorithm specifies the topology of the tree. Since the tree is infinite and we only have finite computational resources it is also important in what order the nodes of the tree are generated as this determines what the tree looks like after some finite time. A key characteristic for the importance of a state is the number of particle-hole pairs it has with respect to the seed state. In

order to avoid having to consider less relevant states with more particle-hole pairs we chose our final algorithm to have a topology where any node with n particle-hole pairs can be generated without having to generate nodes with more particle-hole pairs. This topology, in combination with a clever way of choosing the order in which to generate new states is what led to the most efficient algorithm.

Comparing our final algorithm to the current version of the state of the art library ABACUS showed that we outperform the latter when considering the dynamical structure factor at finite temperature as well as the problem of generating an efficient basis for a quench in the interaction strength at finite temperature. In the latter case the difference is particularly pronounced, emphasizing the importance of the choice of topology. After all, the reason that our algorithm outperforms ABACUS is primarily due to the fact that if ABACUS wants to generate certain states with n particle-hole pairs it has to go through states with more particle-hole pairs. The examples considered show that the improved momentum preserving leapwise scanning algorithm offers significant advantages for finite temperature calculations. One of the main advantages is due to the topology of the tree being generated, which allows any state from the n particle-hole sector to be generated without generating states from higher particle-hole sectors. This is especially important when considering calculations not dominated by the one particle-hole sector as the number of intermediate states required to construct a state from the n particle-hole sector grows dramatically with increasing n at finite temperature. Improved momentum preserving leapwise scanning is therefore more suited to dealing with problems involving for example the g_2 and g_3 operators, the latter being relevant to modelling three-body losses. Because of these advantages, future versions of ABACUS will incorporate this approach to be better equiped to deal with finite temperature calculations.

In this chapter we considered calculations where the number of particles of the states we explore is equal to the number of particles of the reference state. However, for some problems this is not the case (e.g. when calculating the Green's function). In these cases additional difficulties arise for finite temperature states since in this case the number of particle-hole pairs is no longer well-defined. Future research is required to come up with strategies to efficiently deal with these problems. Another interesting direction for future research direction would be the extension of the ideas in this paper to spin chains. Despite their differences, such as the existence of string solutions, the spin chain and the Bose gas also share key characteristics. For example, the quantum numbers of the spin chain can be viewed as a multi-level version of those of the Bose gas, each lattice representing the quantum numbers of a

given string sector which can be visualized as:

$$\vdots$$

2-strings: $\cdots \circ \circ \circ \circ \bullet \circ \bullet \circ \circ \cdots$

1-strings: $\cdots \circ \circ \bullet \bullet \bullet \circ \bullet \circ \cdots$

An algorithm like the one we developed can then be applied to the quantum numbers of every string sector. However, additional complications due to the differences are bound to arise. For example, in the spin chain more care has to be taken to avoid overcounting since a given eigenstate can be represented by multiple quantum number configurations. Furthermore, additional constraints such as the magnetization, as well as the periodicity of the momentum will require careful consideration.

Chapter 4

Quenching the interaction strength

Having reviewed the solution of the Lieb-Liniger model and discussed the Hilbert space exploration algorithms required to compute correlation functions, we turn our attention to out-of-equilibrium dynamics in this chapter. In particular, we consider the dynamics following a quench in the interaction strength of the Lieb-Liniger model. The interest in the dynamics of isolated quantum systems is related to fundamental questions regarding the nature of equilibration in isolated quantum systems. The investigation of such questions however is made difficult by the absence of general methods by which we can compute the time evolution accurately up to long times. Even for the Lieb-Liniger model, where a lot of tools are available due to its integrability, analytical results are limited to situations where the model before the quench is at $c = 0$ or $c = \infty$, which are the effectively non-interacting limits [44, 45].

In this chapter we develop a new numerical approach for computing the overlaps based on the truncated spectrum approach [46]. Our approach focusses on finding the optimal set of computational basis states consisting of eigenstates of the final Hamiltonian to expand the initial state in. Computing the overlaps of the computational basis states numerically allows us to calculate the real-time dynamics and access the long time limit using the diagonal ensemble [47]. The explicit problem we consider is that of the interaction quench in the Lieb-Liniger model, for which we provide proof-of-principle calculations for small system sizes, albeit beyond what can be achieved with coordinate Bethe asnsatz calculations [48, 49].

Recently, the truncated spectrum approach and derived methods have been successfully used to compute overlaps and dynamics for the Ising field theory [50–52, 46, 53] and the sine-Gordon model [54–56, 46, 53]. The difference between the approach taken in these cases and the one we adopt here, is that we work explicitly with a computational basis of eigenstates of an interacting integrable system. As such, strong correlations are already built into the basis states, whereas the aforementioned results are obtained by using non-interacting bases of the Ising field theory and the sine-Gordon model. Even

more importantly, we show that the conventional energy-ordered truncated spectrum approach is inappropriate for the problem we consider and introduce alternative ordering metrics which reduces the computational cost by orders of magnitude.

We start by discussing how the problem of the interaction quench can be formulated in terms of a perturbed Hamiltonian in Sec. 4.1 and we discuss what an ideal truncation scheme would look like. Following this, we start by considering the conventional approach and show how we can overcome its limitations leading us to the high overlap states truncation scheme in Sec. 4.2. In the process, we discover that high energy states can be important showing that truncating the perturbed Hamiltonian by choosing an energy cutoff limits the accuracy of the results that can be obtained. Furthermore, we find a better metric for selecting and ordering the computational basis states based on matrix elements of the perturbing operator. These insights, combined in the high overlap states truncation scheme allows us to construct the initial state to better accuracy and at a lower computational cost. With this newly minted algorithm, we compute the real-time dynamics of correlation functions following the quench in Sec. 4.3 and compute the long time limit using the diagonal ensemble.

In Sec. 4.4 we study strongly non-perturbative quenches, where numerical renormalization group approaches within the high overlap truncation scheme need some modification. A modified algorithm, the matrix element renormalization group, is detailed in this section. We illustrate problems of the high overlap states truncation scheme and the success of the matrix element renormalization group in computing non-equilibrium dynamics following strongly non-perturbative quenches. Furthermore, we introduce a general version of the matrix element renormalization group algorithm able to deal with excited states as well as ground states. We conclude in Sec. 4.5, where we also suggest a number of future directions for studies. The work in this chapter was done in collaboration with J.-S. Caux and N.J. Robinson [2].

4.1 Formulating the problem

In this chapter we consider a quench of the repulsive Lieb-Liniger model in the interaction strength from some initial interaction strength $c_i > 0$ to another interaction strength $c_f > 0$. This means that we initialize the system in an eigenstate of the initial Hamiltonian $H(c_i)$, unless mentioned otherwise we choose the ground state, and time evolve it using the final Hamiltonian $H(c_f)$. The time evolution of the ground state of the initial Hamiltonian, which is not an eigenstate of the final Hamiltonian, is therefore non-trivial.

Analytical approaches, such as for example the quench action [57, 58], are centred around the overlaps between the initial state and the eigenstates of

the final Hamiltonian under which it evolves. After all, given these overlaps, we can expand the initial eigenstate as

$$|\Psi_0\rangle = \sum_{n=0}^{\infty} |\lambda\rangle \underbrace{\langle\lambda|\Psi_0\rangle}_{\text{the overlaps}} \tag{4.1}$$

in terms of the Bethe states $|\lambda\rangle$ of $H(c_f)$. This expansion in turn determines the time evolution of the initial state by adding the appropriate phase factors, which gives

$$|\Psi_0(t)\rangle \equiv e^{-iH_{LL}(c_f)t}|\Psi_0\rangle = \sum_{n=0}^{\infty} e^{-iE_\lambda t}|\lambda\rangle\langle\lambda|\Psi_0\rangle, \tag{4.2}$$

where E_λ is the energy of $|\lambda\rangle$ with respect to $H(c_f)$. Unfortunately, it is not generally known how to compute these overlaps even in cases where the powerful tools of the Algebraic Bethe Ansatz are applicable, despite progress in a few cases [44, 59–73]

Due to the limited success of analytical methods for computing the overlaps, it is a natural question to ask if numerical methods can be used instead. Since the expression for the time evolved state depends on the expansion of the initial state in terms of the eigenstates of the final Hamiltonian, we should aim to determine the corresponding overlaps numerically. In cases where the initial Hamiltonian can be constructed in terms of the eigenbasis of the final Hamiltonian, these overlaps can be determined by diagonalization of the Hamiltonian. However, for the interaction quench under consideration here, we are dealing with a continuum bosonic model, rendering the Hilbert spaces infinite dimensional. This necessitates the truncation of the Hilbert space of the final Hamiltonian in order to be able to represent the Hamiltonian as a finite matrix we can diagonalize. This process of truncating the Hilbert space is called the truncated spectrum approach.

It should be noted that our knowledge of the spectrum of the final Hamiltonian allows us to work directly with strongly correlated basis states. This is different to how the truncated spectrum appraoch has previously been used to tackle out-of-equilibrium dynamics [50, 51, 74, 75, 53, 76, 55, 56], as usually the computational basis used consists of non-interacting bosons or fermions. In such cases, both the initial state and eigenstates of the final Hamiltonian have to be expanded in the computational basis.

Knowledge of the spectrum of the final Hamiltonian in itself is however not enough to successfully use the truncated spectrum approach. We also have to be able to construct the initial Hamiltonian in terms of this set of basis states. For the case at hand here, this is achieved through knowledge of

matrix elements obtained using the Algebraic Bethe Ansatz [33, 34]. To see this, note that we can write relate the initial and final Hamiltonian by

$$H(c_i) = H(c_f) + (c_i - c_f) \int_0^L \mathrm{d}x \, \Psi^\dagger(x)\Psi^\dagger(x)\Psi(x)\Psi(x). \qquad (4.3)$$

From this formulation, we see that we can view the initial Hamiltonian as a perturbation of the final Hamiltonian. We emphasize however, that we put no restrictions on the difference $(c_i - c_f)$, so the perturbation need not be small in any sense. In the zero momentum sector, which is the relevant sector to the problem under because g_2 is a momentum-preserving operator and the initial state we consider has zero momentum, the matrix elements of the initial Hamiltonian are given by

$$\langle \mu | H(c_i) | \lambda \rangle = \delta_{\mu,\lambda} E_\lambda + (c_i - c_f) L \langle \mu | \left(\Psi^\dagger(0) \right)^2 \left(\Psi(0) \right)^2 | \lambda \rangle. \qquad (4.4)$$

Here the $|\mu\rangle$, and $|\lambda\rangle$ on the right hand side are again the eigenstates of $H(c_f)$ for some finite number of particles N. For such finite size systems, a determinant expression for the matrix element on the right hand side has been derived, as discussed in Sec. 2.4.2. The basis we choose for the expansion we also call the computational basis states.

4.2 Developing a High Overlap States Truncation Scheme

Having discussed how an expansion of the initial state in terms of the computational basis can be found by diagonalization of a truncated Hamiltonian, the question that presents itself is which states of the computational basis we should include in order to obtain optimal convergence of physical quantities. Suppose we can order the states of the computational basis according to the magnitude of their overlaps with the initial state, i.e. according to $|\langle \lambda | \Psi_0 \rangle|^2$. In this case, we can obtain an approximate expansion of the time evolved wave function by included the N_{tot} states with the highest weight, giving

$$|\Psi(t)\rangle_{\text{approx}} = \sum_{n=0}^{N_{\text{tot}}} e^{-iE_\lambda t} |\lambda\rangle \langle \lambda | \Psi_0 \rangle. \qquad (4.5)$$

In this expansion, the norm of the time evolved wave function is saturated up to

$$s(N_{tot}) = 1 - \sum_{n=0}^{N_{\text{tot}}} |\langle \lambda | \Psi_0 \rangle|^2. \qquad (4.6)$$

As a result, the errors of bounded operators $A(t)$ evaluated for the approximate time evolved state are themselves bounded by $\epsilon_{\max}$ given by

$$\epsilon_{\max}\left[A(t) - A(t)_{\text{approx}} \right] = s(N_{\text{tot}})\max_{m,n}(A_{m,n}). \qquad (4.7)$$

Here $A_{m,n} = \langle m|A|n \rangle$ are the matrix elements of A with respect to the m^{th} and n^{th} states of the computational basis. As a result, we can make the maximal error $\epsilon_{\max}$ arbitrarily small by making $s(N_{\text{tot}})$ arbitrarily small.

Unfortunately, as mentioned previously, we do not a priori know the overlaps. The remainder of this section is concerned with developing a method that comes as close as possible to this optimal truncated basis. The method we arrive at after several steps we call the high overlap state truncation scheme and draws inspiration from the conventional truncated spectrum aproach. Rather than present directly the final algorithm and show it works well, we show also the process that led us there since we deem it insightfull.

4.2.1 The truncated spectrum approach

Since the Hilbert space of the Lieb-Liniger model is infinite dimensional, its Hamiltonian is represented by an infinite dimensional matrix. Furthermore, the perturbed Hamiltonian we introduced in the formulation is dense in the computational basis of eigenstates of the final Hamiltonian. To come to an expansion of the initial state and in order to find the corresponding approximate numerical values of the overlaps $\langle \lambda|\Psi_0 \rangle$ we have to truncate the Hilbert space in some manner.

To start, let us consider the standard approach [41–43, 46] in the case where the perturbing operator is renormalization group relevant. In this case, the following holds:

1. The diagonalization of the Hamiltonian leads to a strong mixing of the low-energy states of the computational basis, i.e. the low-energy eigenstates are superpositions of many states.

2. The diagonalization of the Hamiltonian leads to a weak mixing of the high-energy states, i.e. the computational basis states are approximate eigenstates of the perturbed Hamiltonian.

3. A consequence of the previous points is that low-energy and high-energy states in the computational basis cannot strongly couple.

Given the presumed decoupling of the low and high energy states of the computational basis, the simplest truncation method would be to introduce an energy cutoff Λ. This is the approach taken by Yurov and Zamalodchikov in their original papers where they considered perturbed conformal field theories [41, 42]. Given such a cutoff, we can consider the convergence of quantities we are interested in, such as the energy of the approximate ground state, as a function of the cutoff energy. An example of this is shown in Fig. 4.1 where we consider the interaction quench for the ground state starting from from $c_i = 20$ to $c_f = 10$ for the case of ten particles at unit density. What allows us
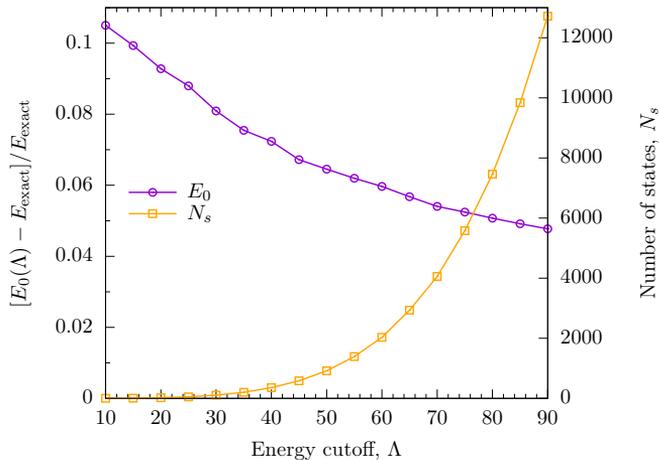
Figure 4.1: A combined plot of the number of states N_s below a given energy cutoff Λ and the relative error $(E(\Lambda) - E_{\text{exact}})/E_{\text{exact}}$ of the approximate ground state energy $E(\Lambda)$ obtained by including all N_s states below the cutoff. The system parameters considered for this plot are $c_i = 20$, $c_f = 10$, $N = 10 = L$.

to make this comparison is that we know the ground state energy of $H(c_i)$ so that we can measure convergence by considering $(E_0 - E_{\text{exact}})/E_{\text{exact}}$. For the data we were able to generate, the convergence of the energy of the ground state can be fitted with an exponential function.

Although a very natural approach to take, convergence with the cutoff energy Λ is slow for many models [46]. Furthermore, the number of states can grow very rapidly with the cutoff energy, which quickly renders us unable to diagonalize the truncated Hamiltonian. For the case at hand here, this can clearly be seen from Fig. 4.1, which we remind the reader concerns a calculation for a mere 10 particles. In order to deal with these limitations, various techniques have been developed [46]. In the next section, we show how the numerical renormalization group can be utilized to deal with the rapidly growing number of states as the cutoff increases.

4.2.2 The numerical renormalization group extension

To overcome the issue with the rapidly growing number of states as a function of the cutoff energy Λ, we adopt an iterative procedure proposed by Konik and Adamov [43]. This procedure enriches the ideas from Zamalodchikov's truncated spectrum approach [41, 42] with the ideas from Wilsons's numerical renormalization group [77]. The resulting procedure can be summarized as

follows:

1. Generate a computational basis $\{|\lambda_i\rangle\}_{1 \leq i \leq N_{\text{tot}}}$ and order the computational basis states in order of increasing energy as measured by the final Hamiltonian.

2. Construct the truncated perturbed Hamiltonian for the first $N_s + \Delta N_s$ states of the reordered computational basis and diagonalize it to obtain $N_s + \Delta N_s$ approximate[1] eigenstates $\{i\}_{1 \leq i \leq N_s + \Delta N_s}$ and corresponding approximate energies $\{\tilde{E}_1, \ldots, \tilde{E}_{N_s + \Delta N_s}\}$.

3. Discard the ΔN_s approximate eigenstates with the largest energies from the set of approximate eigenstates obtained by diagonalizing the truncated perturbed Hamiltonian.

4. Add the first ΔN_s states from the computational basis that have thus far not been part of the truncated basis and add them to the N_s approximate eigenstates kept from step three forming a new truncated basis.

5. Construct the truncated Hamiltonian in terms of this new truncated basis and diagonalize it to generate new approximate eigenstates and energies.

6. Return to step three.

Steps three through five are repeated, continually updating a set of N_s approximate eigenstates and energies until we exhaust the computational basis or until we are satisfied with the convergence reached. Note that for every round of the procedure we diagonalize a truncated Hamiltonian of size $(N_s + \Delta N_s)^2$. Despite having to repeat this procedure many times in order to exhaust the computational basis, this can be more efficient than having to diagonalize a single N_{tot}^2 matrix as would be required for the standard truncated spectrum approach. In fact, using the renormalization group extension allows computational bases of hundreds of thousands or even millions of states to be exhausted [78, 79], for which it would be impossible to diagonalize the full truncated perturbed Hamiltonian.

To illustrate the performance of the numerical renormalization group extension, see Fig. 4.2 where we consider the quench from $c_i = 20$ to $c_f = 10$ for $N = 10 = L$ where the parameters for the numerical renormalization group are $N_s = 600$ and $\Delta N_s = 200$. Here we consider the convergence of the ground state energy E_0 as a function of the number of computational basis states.

---

[1] They are of course the numerically exact eigenstates of the truncated pertrubed Hamiltonian we consider here, but this makes them only approximate eigenstates of the full perturbed Hamiltonian.
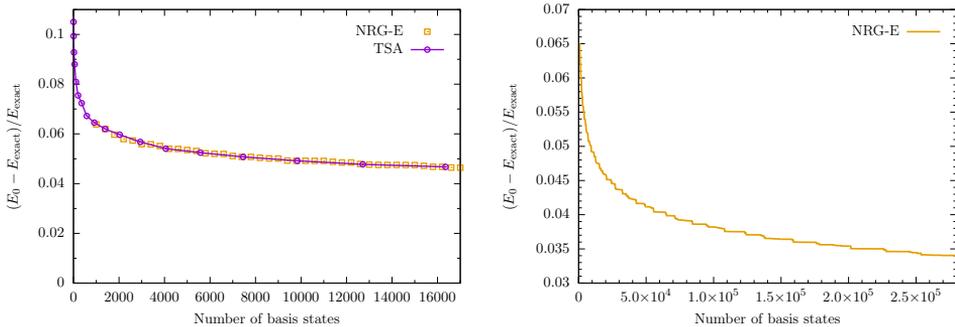
Figure 4.2: Comparison between the truncated spectrum approach (TSA) and its extension with the numerical renormalization group (NRG-E) for the convergence of the approximate ground state energy E_0 with the number of basis states. We consider the quench where $c_i = 20$ and $c_f = 10$ for $N = 10 = L$. The parameters for the numerical renormalization group extension are $N_s = 600$ and $\Delta N_s = 200$, where the first step containing 800 states corresponds to an energy cutoff of $\Lambda \approx 50$. We plot half the NRG-E steps and see excellent agreement between its results and the results from full diagonalization as obtained using the truncated spectrum approach (left panel). However, as intended, the numerical renormalization group extension is capable of considering larger computational bases (right panel).

In the left panel we see that the results between the full truncated spectrum approach agree with those obtained via the numerical renormalization group extension despite the modest size of the truncated Hamiltonian in this case. In the right panel we see that this extension enables us to deal with $280,000$ states, which is far more than we would be able to deal with if we stuck to diagonalizing the full truncated Hamiltonian. For the $280,000$ states conisdered here, the energy of the ground state has converged to below 3.5%, which corresponds to about 11% of the Fermi energy. For reference, this is significantly smaller than the level spacing $E_1 - E_0$ for the parameters we consider.

4.2.3 ORDERING BY AN ALTERNATIVE METRIC

Although the numerical renormalization group extension allows us to consider large computational bases, slow convergence of the ground state energy with the size of the computational basis can still prevent us from reaching our desired accuracy. For example, if we wish to reach a precision of 2% for the parameters considered in Fig. 4.2, we would likely have to exhaust a computational basis of over a million states. However, a closer look at the right panel in Fig. 4.2 indicates that perhaps a smaller computational basis could achieve similar convergence. To see why, note that the convergence of the energy appears to follow a step-like structure. This means that in many steps the new states of the computational basis states we introduce do not affect the convergence of the ground state energy in a meaningful way. Therefore, if we could eliminate these unimportant contributions and instead focus on those strongly affecting convergence, we could obtain the same convergence with fewer states. The question we have to ask ourselves is then if we can replace the energy as the ordering metric in order to prioritize the important states.

One alternative ordering that can be considered is based on previous works [79–82, 46] where one takes an ordering metric based on the matrix elements of the perturbing operator with respect to the lowest energy eigenstates of the initial Hamiltonian, i.e.

$$\left| \langle \lambda | g_2(0) | \tilde{E}_j \rangle \right|, \quad j = 0, 1, 2. \tag{4.8}$$

where $|\tilde{E}_j\rangle$ are the lowest energy eigenstates of $H(c_i)$. The idea behind this ordering metric is that it allows us to find the states in the computational basis which hybridize most strongly with the low energy states we are interested in. However, we cannot evaluate Eq. (4.8) directly since we do not know the off-diagonal matrix elements of the g_2 operator between eigenstates at different values of the interaction strength. We first have to construct an expansion of the lowest energy eigenstates in terms of the computational basis. This can
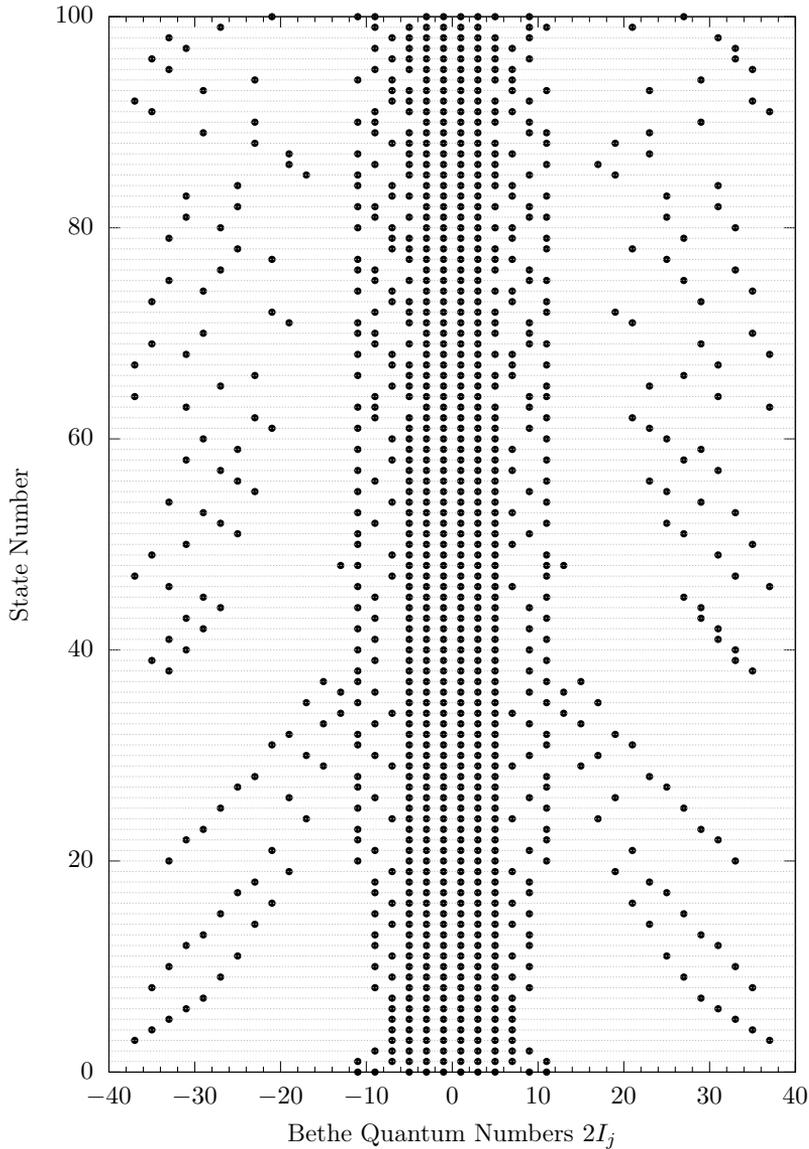
Figure 4.3: The quantum number configurations corresponding to the 100 highest weight computational basis states according to the matrix element metric as defined in Eq. (4.8) for the interaction quench from $c_i = 20$ to $c_f = 10$ at $N = 10 = L$. The sample of highest weight states shown here came from a total basis of $273,358$ states, which is the same computational basis used in 4.2. Note that some of the high weigh states have high energy excitations.

be done by using the truncated spectrum approach for a small subset of the computational basis we are interested in and then using the weight defined in 4.8 to reorder the rest of the computational basis.

To gain insight into the effect of this reordering of the computational basis based on the metric defined in Eq. (4.8), we visualize the quantum number configurations of the first one hundred states after reordering in Fig. 4.3. We see that this matrix element ordering is significantly different to the ordering by energy as many of the high weight states have pairs of highly excited quantum numbers compared to the ground state configuration. One can even think of grouping different configurations together in families of states where the same quantum numbers move out further and further, think e.g. contributions 4-8,10,12,. . . which can be characterized as being in the ground state configuration where the two outermost quantum numbers have become excited and are moving away to infinity. Interestingly, the higher energy states of this family have higher weights according to the matrix element metric, the complete opposite way they would be ordered according to the energy metric. Furthermore, we see that the ground state of the final Hamiltonian is the third highest weight state according to this metric.

1. Convergence of the ground state energy

   The reordering shown in Fig. 4.3 is not what we a priori would have expected, but proves its merit by leading to massive improvements in convergence as shown in Fig. 4.4. Here we see that after a mere seven steps of the renormalization group procedure, corresponding to 2000 states, the matrix element ordered procedure obtains convergence better than energy ordered routine does after considering $100,000$ states. Furthermore, we see that the procedure with the reordered basis quickly plateaus indicating that the computational basis states considered in those steps do not contribute appreciably to the convergence of the ground state energy. Although approximately the same threshold is reached, the alternative ordering proves very promising. Therefore we consider only such alternative metrics based on matrix elements from now on.

   One of the limiting factors at this stage, is that the convergence we can achieve is limited by the quality of the computational basis we generate. Furthermore, we have thus far generated bases based on an energy cutoff after which we computed weights for the computational basis states according to Eq. (4.8) which can also be costly computationally. A better approach would be to leverage the Hilbert space exploration algorithms we considered in Chapter 3 to generate a more optimal basis that is generated in the right order so that no additional reordering is required. We consider this approach in Sec. 4.2.4, but before then let us investigate the results we have generated thus far further.
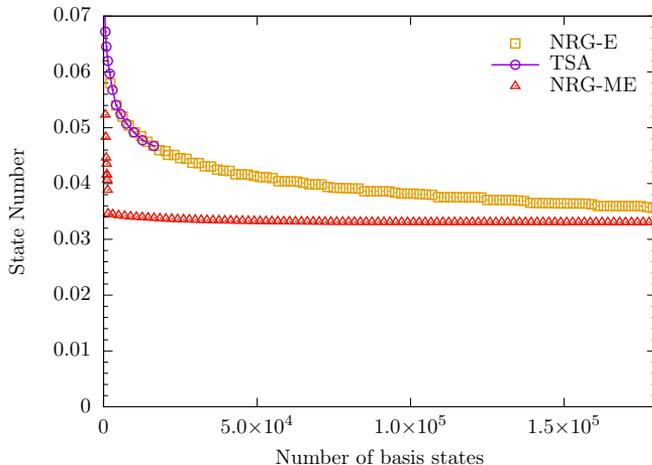
Figure 4.4: A comparison of the convergence of the ground state energy for the truncated spectrum approach extended by the numerical renormalization group where we either order the computational basis states by their energy (NRG-E) or the matrix element metric defined in Eq. (4.8) (NRG-ME). Again the parameters we consider are $c_i = 20$, $c_f = 10$, $N = 10 = L$ and $N_s = 600$, $\Delta N_s = 200$ for a total of about 280,000 basis states. The results comparing the regular truncated spectrum approach and the energy ordered renormalization group extension (NRG-E) were previously presented in Fig. 4.1. For the new results using the reordered basis (NRG-ME) however, we see a massive improvement in convergence of the ground state energy (which is also observed for other low-lying excited states).
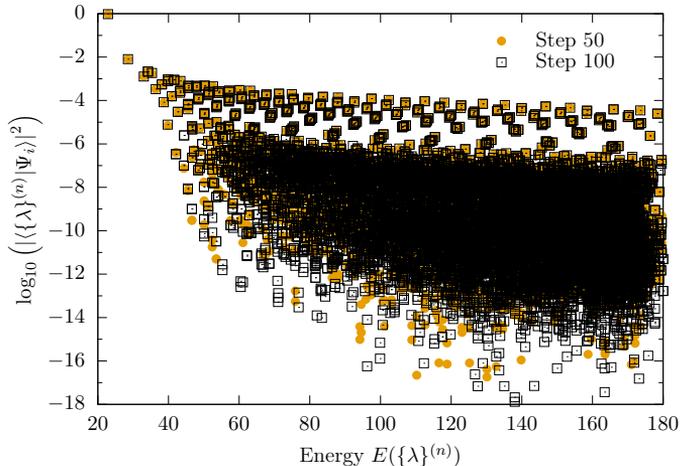
Figure 4.5: The overlaps of the computational basis states, i.e. eigenstates of the Lieb-Liniger model at $H(10)$ with the approximate ground state at $H(20)$. We obtain the overlaps by constructing the approximate ground state using the renormalization group extension of the truncated spectrum approach on the basis reorderd according to Eq. (4.8). We show the overlaps after step 50 and step 100, showing that the large overlaps are well-converged, which is not surprising given the convergence pattern observed in Fig. 4.4.

2. The overlaps: Convergence and structure

Thus far we have considered the convergence of the ground state energy, which gives us a quantitative way of measuring convergence since we know the exact ground state energy from the Bethe ansatz solution of the model. While useful, this convergence of the energy is not the only relevant quantity whose convergence we are interested in. For example, for the time evolved wave function, which forms the basis for computing the time evolution of observables, the overlaps are the central objects. Of course, our formulation of the quench in terms of a truncated Hamiltonian was precisely chosen to obtain these overlaps, so let us consider their convergence in by plotting the log of the square of the overlap of a computational basis state versus its energy for two well-separated steps in the procedure as shown in Fig. 4.5. We see that the large overlaps have converged well, which is not the case for the smaller overlaps, which are likely also affected by floating point errors. Fortunately, these small overlaps are not of physical importance, as we argued at the start of Sec. 4.2.

One subtlety that should be addressed, is that although the large over-

laps in Fig. 4.5 have converged with respect to the numer of steps of the renormalization group procedure, we would obtain different overlaps if we chose a different computational basis. After all, the diagonalization of the Hamiltonian returns normalized states and there are further strongly contributing states with energies above the effective energy cutoff $\Lambda \approx 180$ we have considered thus far. Regardless, this does not necessarily mean that physical quantities are not well-converged, as we will discuss shortly. Besides the convergence of the overlaps, Fig. 4.5 shows that the energy of a state is not a good indicator of the size of its overlap. In particular, there appear to be rare sets of states, whose energies are large, whose overlaps dominate those of the majority of states.

Having confirmed that the overlaps converge with the number of steps of the renormalization group procedure for a fixed computational basis, let us take a closer look at the structure of the high overlap states. For the quench from $c_f = 20$ to $c_f = 10$, the quantum number configurations corresponding to the 100 highest overlap states are shown in Fig. 4.6, starting from the highest overlap states on the bottom. We see that the largest overlap state is the ground state of the final Hamiltonian, and the next most important states are those with two particle-hole pairs of increasingly larger energies. The energies can be increased either by the outermost particles moving away further from the Fermi sea, or by holes moving further into the Fermi sea. This structure can be understood in conjunction with Fig. 4.5 where we see that as the energy increases, the maximal overlap decreases.

If we compare the prediction for the highest overlap states presented in Fig. 4.3, we see many of the same states, albeit ordered in a different way. This reordering is immaterial as long as the states are included in the same step of the renormalization group procedure. However, it also indicates that the ordering metric used to produce this prediction is not ideal. In the next section we introduce an alternative metric, which aims to predict the highest overlap states more accurately and in a way that does not require us to first compute approximate low-energy eigenstates.

4.2.4 Introducing a preferential scanning routine

Thus far we have shown that reordering the computational based on information about the perturbing operator can lead to massive improvements in convergence when compared to a traditional energy ordering. Currently our procedure is limited by the fact that we have to first generate a computational basis for which we compute some weighing function which requires us to first produce some approximate low-energy eigenstates using the truncated spec-
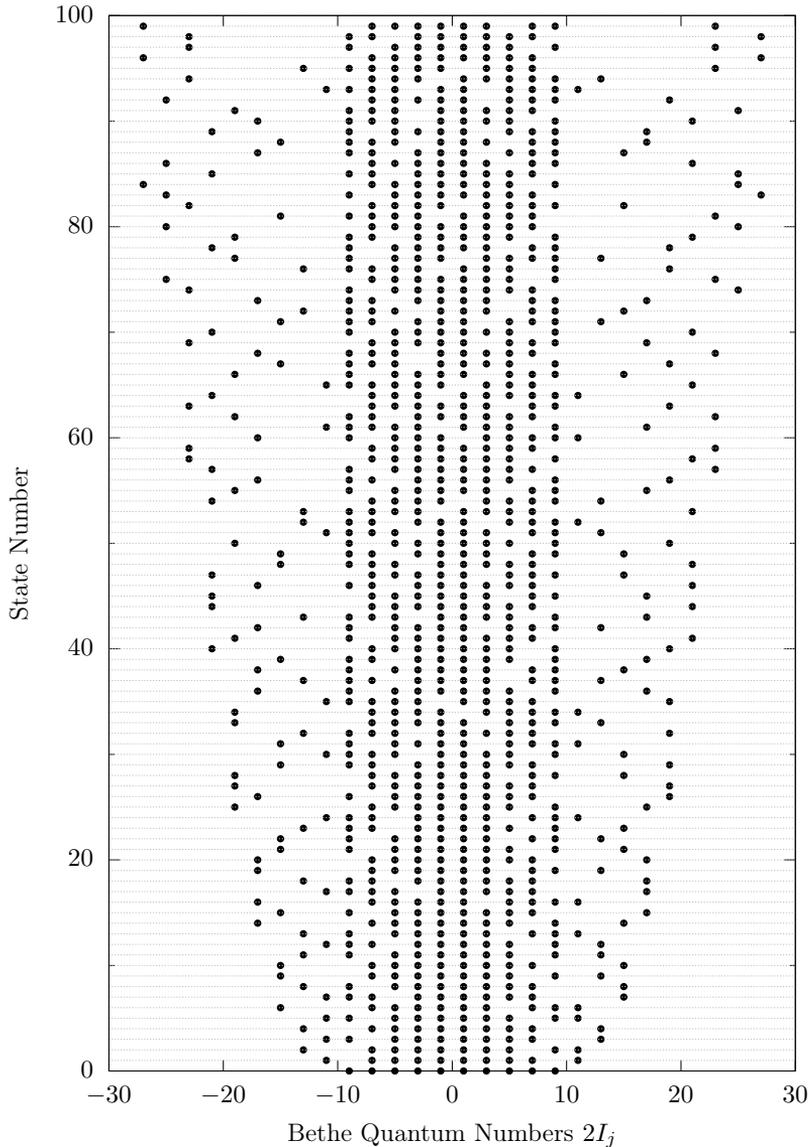
Figure 4.6: The quantum number configurations corresponding to the 100 highest overlap states for the interaction quench from $c_i = 20$ to $c_f = 10$ at $N = 10 = L$. The overlaps are computed using the numerical renormalization group procedure with the reordered basis according to 4.8 using $N_s = 600$ and $\Delta N_s = 200$. The largest overlap states are the ones at the bottom.

trum approach. If we wish to achieve better convergence, we have to eliminate the implicit energy cutoff that comes with this approach, since we saw clearly that high energy states can have large overlaps. For example, Fig. 4.5 shows that there are computational basis states with energies between 100 and 1000 which have square overlaps as large as 10^{-4}.

Given the extensive discussion on Hilbert space exploration algorithms in Chapter 3, a natural question to ask is if we can come up with a weighing function for the eigenstates of the final Hamiltonian that accurately predicts the importance of a state. We could use the metric in Eq. (4.8), but it raises the question which states we use to first compute the approximate low-energy eigenstates and we already saw that it did not reproduce the optimal ordering of computational basis states. Instead, we look to perturbation theory for a better weighing function.

For the results obtained thus far, we saw that the ground state of the final Hamiltonian has the largest overlap with the ground state of the initial Hamiltonian, which we are trying to find an expansion for. The second order perturbative expansion for the ground state of the initial Hamiltonian is then given by

$$|\Psi_0\rangle = |0\rangle + (c_i - c_f)L \sum_{m \neq 0} \frac{\langle m|g_2(0)|0\rangle}{E_0 - E_m}|m\rangle. \tag{4.9}$$

Here the integers label eigenstates of the final Hamiltonian, with 0 labelling the ground state. Given this expansion, a natural weight for a given eigenstate presents itself, namely

$$w\left(|m\rangle\right) = \left|\frac{\langle m|g_2(0)|0\rangle}{E_m - E_0 + \epsilon}\right|, \tag{4.10}$$

which is the estimate of the overlap from second order perturbation theory up to a numerical factor ϵ introduced to avoid a divergence for the case where $n = 0$ (we take $\epsilon = 0.1$). Given this weighing function, we can very efficiently generate a close to optimal computational basis, as shown in Sec. 3.6.2.

Let us consider the 400 highest weight states of the computational basis generated using the weighing function introduced in Eq. (4.10) for the quench from $c_i = 20$ to $c_f = 10$ for ten particles at unit density as shown in Fig. 4.7. We see that the ordering here is significantly different to the ordering we obtained by using the matrix element metric from Eq. (4.8), and instead much more in line with the numerical results for the states with the highest overlaps shown in Fig. 4.6. Furthermore, it is clear from Fig. 4.7 that the computational basis generated in this way is not only ordered differently to the basis obtained by considering an energy cutoff. After all, already in the first 400 states we see states with highly excited quantum numbers, corresponding
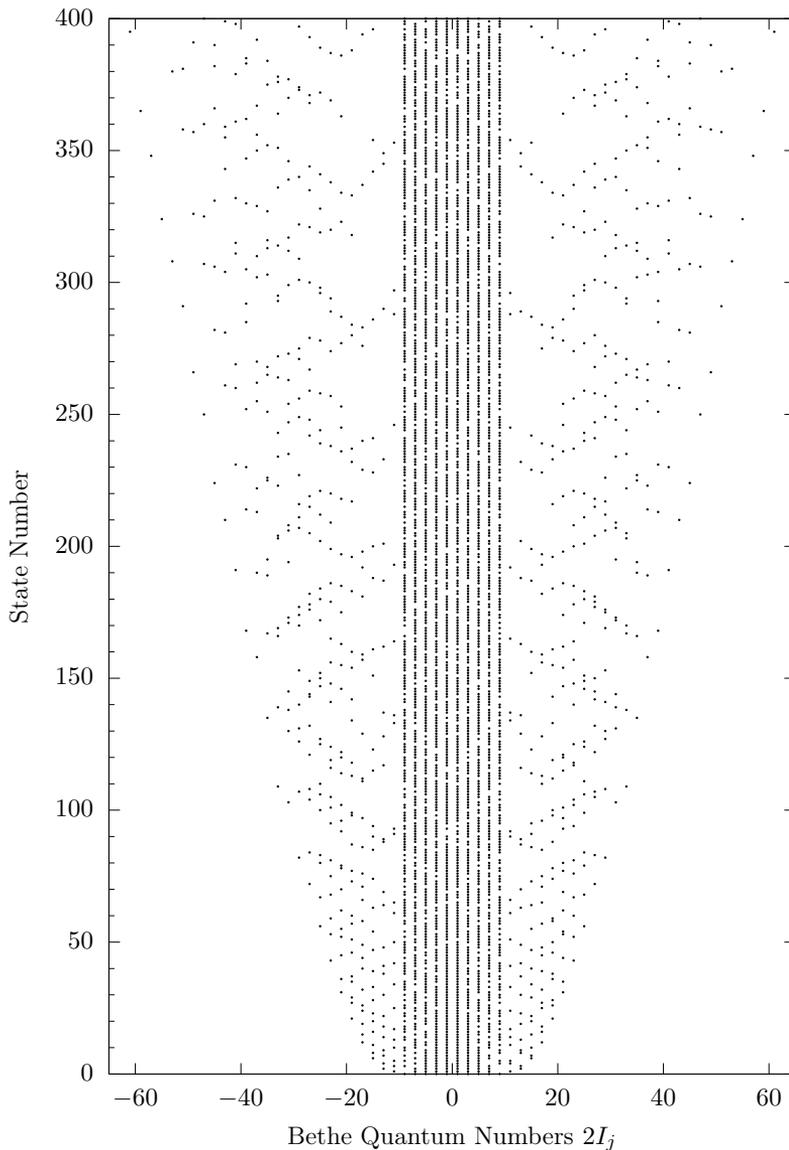
Figure 4.7: The quantum number configurations corresponding to the 400 highest highest weight states for the interaction quench from $c_i = 20$ to $c_f = 10$ at $N = 10 = L$ according to Eq. (4.10). The quantum number number configurations and their ordering are similar to those of the highest overlap states in Fig. 4.6, i.e. the preferential scanning routine generates an efficient basis.

to high energy eigenstates that would be excluded when considering an energy cutoff. Furthermore, having gotten rid of the energy cutoff opens the way to achieving better convergence as we discuss in the following sections.

The procedure of doing a preferential Hilbert space scan based on the perturbative expansion for the overlap is what we call the high overlap states truncation scheme (HOSTS).

4.2.5 Checking convergence within the high overlap states truncation scheme

With the high overlap states truncation scheme, we are freed of the energy cutoff that was limiting the convergence we were able to obtain previously. For example, in Fig. 4.4 we were able to find a decent ordering of the computational basis which sped up convergence compared to the energy ordering, but the energy cutoff was limiting the total convergence we could reach (about 3.5%). The preferential scanning routine introduced in the previous section is not limited in this way, allowing us to reach convergence to much less than one 1% whilst at the same time significantly decreasing the computational cost.

In the following sections, we study in more detail the convergence of the energy but also the convergence of the overlaps and local expectation values. For all these quantities, a question we should ask is if we can extrapolate these results to obtain the exact result. If this extrapolation can be done is unclear since there exists for example no scaling law for the energy as a function of the size of the computational basis. However, since our scheme is centred around the weight of eigenstates based on matrix elements of the perturbing operator it is a natural candidate to be used in an extrapolation scheme. For example, we can plot the convergence of the quantity we are interested in versus the lowest weight state included in the computation for a given level of convergence. We will show that this leads to a reasonable extrapolation scheme.

1. Convergence of the energy

    To substantiate our claim that the high overlap state truncation scheme can achieve better convergence, see Fig. 4.8(a). Here we show the convergence of the energy versus the number of states of the computational basis included in the calculation. For the results shown here, a computational basis is generated by running ABACUS for 30 seconds and reordering the states according to Eq. 4.10 since ABACUS does not always generate states in the perfect order. This gives us an ordered basis consisting of $220,743$ states which we use for the numerical renormalization group procedure with $N_s = 540$ and $\Delta N_s = 160$. Even for basis
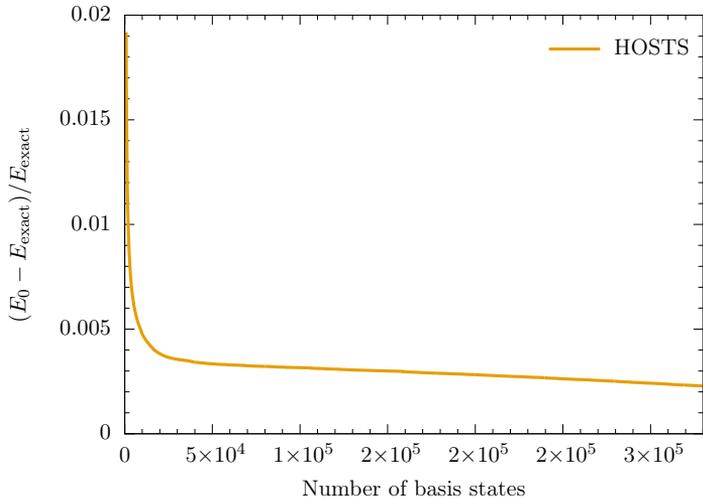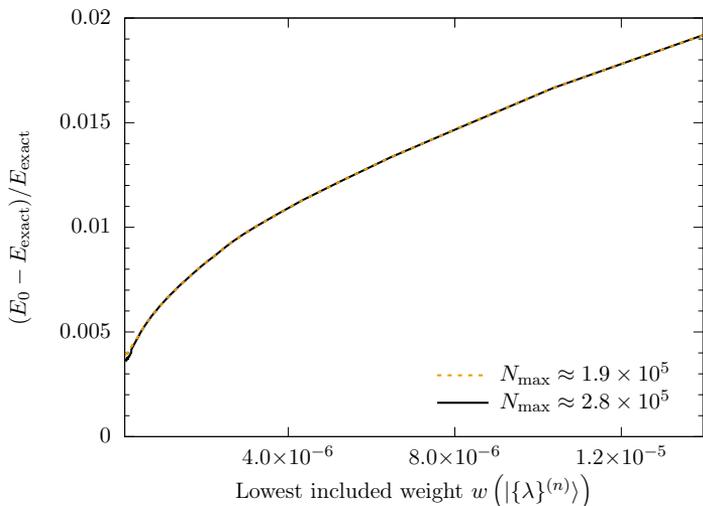
Figure 4.8: The convergence of the ground state energy for the quench from $c_i = 20$ to $c_f = 10$ using the high overlap states truncation scheme. (a) The ground state energy versus the number of basis states. (b) The ground state energy versus the lowest weight included in a given step of the renormalization group procedure for two different total basis sizes. In both cases the computational bases were generated using ABACUS and the numerical renormalization group was performed using $N_s = 540$ and $\Delta N_s = 160$. The ground state energy is converged to well below 1% for a couple of thousand basis states, which represents a huge improvement compared to the energy ordered procedure we started with, see Fig. 4.4.
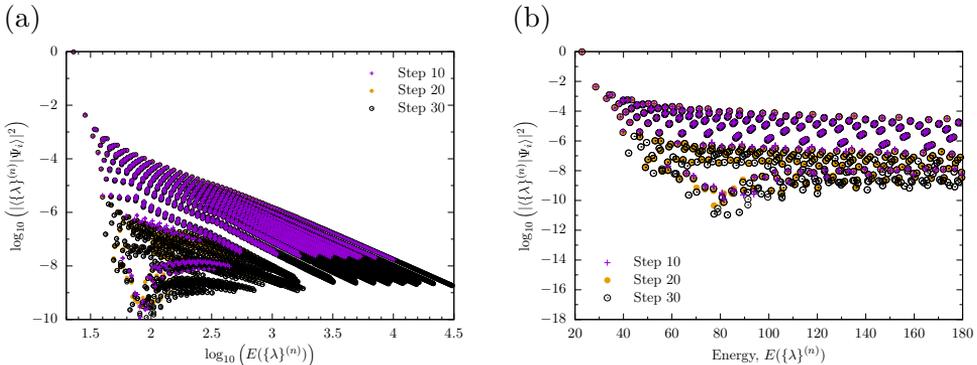
Figure 4.9: The overlaps of the computational basis states at different steps of the numerical renormalization group procedure using the high overlap states truncation scheme for the quench from $c_i = 20$ to $c_f = 10$ for $N = 10 = L$ using $N_s = 540$ and $\Delta N_s = 160$. In (a) we show all of the overlaps whereas in (b) we focus on the low energy region of the plot to enable an easier comparison to Fig. 4.5.

sizes of a couple of thousand states, which would be accessible with exact diagonalization, we see that this results in excellent convergence of below 0.5%. To indicate the improvement this represents, recall that for the energy ordered renormalization grup procedure the ground state energy was only converged up to over 3% after 200,000 states.

In Fig. 4.8(b) we consider the extrapolation scheme discussed in Sec. 4.2.5 for which we plot the convergence of the ground state energy versus the lowest weight included per step. Comparing the results for two different total basis sizes shows that there is some dependence on the total basis size for small values of the lowest included weight, which means that for the algorithm used to generate these plots not all states above a given value had been generated.

2. Convergence of the overlaps

Having seen how the high overlap states truncation scheme is able to achieve excellent convergence, let us return to the convergence of the overlaps. The convergence of the overlaps for the quench from $c_i = 20$ to $c_f = 10$ is shown in Fig. 4.9 for three different steps of the renormalization group procedure. From studying this figure, we can draw a couple of conclusions.

First of all, we see the preferential scanning routine indeed generates the high overlap states. Note that this is not a trivial statement since we

only used the second order perturbation theory formula for the overlaps to select states whereas here we consider the numerical overlaps determined using a non-perturbative method. Secondly, we again see that only the energy of a state is not a good indicator of its overlap as the maximal overlap as a function of the energy is a slowly decreasing function. For example, by just the thirteenth step of the renormalization group procedure we are probing states with energies of order $10,000$, which is enormous compared to the exact ground state energy given by $26.9684027\ldots$. Finally, there appear to be families of states, as alluded to previously in Sec 4.2.3, whose overlaps at high energies could in principle be extrapolated.

Let us compare the results obtained in Fig. 4.9 using the preferential state generation to the overlaps computed using the original basis obtained by choosing an energy in Fig. 4.5. We see that for similar bases sizes, preferential scanning probes far higher energies and also that it on average yields larger overlaps, showing that the method is successful.

3. Convergence of local expectation values

Thus far we have considered only the convergence of the energy of the initial state and the convergence of the overlaps with the number of steps of the renormalization group procedure. One can however wonder if this is sufficient to guarantee convergence for other properties of the initial state. In this section we focus on local properties of the initial state we constructed. In particular, we compare the convergence of expectation values of local operators for the constructed state to the known exact values. Knowledge of the exact values comes again from the fact that $H(c_i)$ is integrable. This verification will allow us that we are not merely reproducing a state of the same energy, but also capturing more of its structure.

The first local observable we consider is the diagonal matrix element of the density operator. Since the final Hamiltonian we consider conserves particle number and system is translationally invariant, the diagonal matrix elements of the density operator are equal to N/L for all computational basis states. Since the approximate eigenstate of the initial Hamiltonian we construct is a superposition of these states, the same holds true for the expansion.

Another set of local expectation values we consider are those of the operators appearing in the Hamiltonian. Since the energy converges, so should these local expectation values, but we can also explicitly verify this. To do so, let us consider the local expectation values of the operators that make up the Hamiltonian, namely the kinetic part $\partial_x \Psi^\dagger(0)\partial_x\Psi(0)$
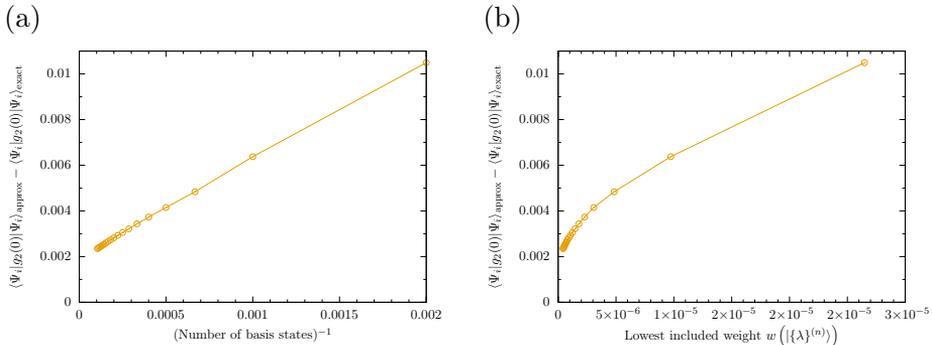
Figure 4.10: The difference between the exact value of the $g_2(0)$ expectation value in the ground state and the expecation value of the approximate eigenstate we construct using HOSTS for the quench from $c_i = 20$ to $c_f = 10$ at $N = 10 = L$. In (a) we show convergence with the number of states whereas in (b) we show convergence with the lowest weight, as given by Eq. (4.10), included in a given step.

and the interacting part $g_2(0) = \Psi^\dagger(0)\Psi^\dagger(0)\Psi(0)\Psi(0)$. Their convergence is shown in Fig. 4.10 and 4.11, where we compare the expctation values to the exact values for the ground state of $H(c_i)$. We see that the expectation values of the $g_2(0)$ operator have not converged as well as those of the kinetic energy. This is not inconsistent with the convergence of the energy since for large values of the interaction strength the kinetic energy term dominates the interaction energy[2] Overall, we see that we are correctly reproducing the expecation values of the local operators appearing in the Hamiltonian for the constructed states. Furthermore, we obtain reasonable results after extrapolating the values of the local operators as a function of the lowest included weight.

4.3 Non-equilibrium dynamics from the high overlap states truncation scheme

Now that we have successfully developed a method to numerically determine the overlaps for the ground state undergoing a quench in the interaction strength, we can compute the non-equilibrium dynamics following the quench. As mentioned in Sec. 4.2.1, determining the non-equilibrium dynamics given the overlaps is in principle straightforward using Eq. (4.2) truncated to the

---

[2]This is point is exaggerated in the limit where c goes to infinity and we end up in the Tonks-Girardeau limit in which the system can be mapped to one of non-interacting fermions.
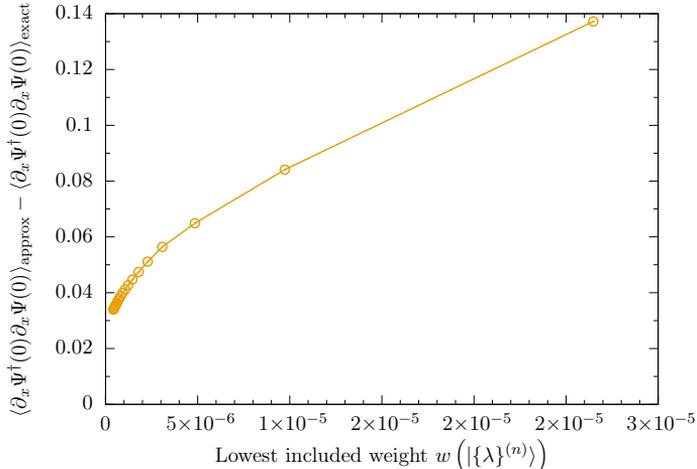
Figure 4.11: The convergence of the expectation values for the approximate ground state of $\partial_x \Psi^\dagger(0)\partial_x \Psi(0)$ minus the exact value. We consider again the quench from $c_i = 20$ to $c_f = 10$ at $N = 10 = L$. Here the exact value of $\langle 0|\partial_x \Psi^\dagger(0)\partial_x \Psi(0)|0\rangle$ is given by 2.22032.

first N_{tot} terms giving as in Eq. (4.5). In this section we use this expansion to first determine the return amplitude and the fidelity before moving on to the time evolution of $g_2(t)$.

4.3.1 THE RETURN AMPLITUDE AND THE FIDELITY

The first out-of-equilibrium quantity we consider is called the return amplitude, which is defined as

$$\langle\Psi_0|\Psi_0(t)\rangle \approx \sum_{n=0}^{N_{\text{tot}}} e^{-iE_\lambda t}\,|\langle\lambda|\Psi_0\rangle|^2\,. \tag{4.11}$$

The interest in this quantity arises from the potential non-analyticity of

$$f(t) = -\lim_{L\to\infty}\frac{1}{L}\log\langle\Psi_0|\Psi_0(t)\rangle \tag{4.12}$$

The fidelity is another quantity of interest which is given by

$$\mathcal{F}(t) = |\langle\Psi_0|\Psi_0(t)\rangle|^2, \tag{4.13}$$

i.e. it is the absolute value squared of the return amplitude.

We will use the results for the return amplitude and the fidelity to investigate differences in convergence between out-of-equilibrium quantities and the
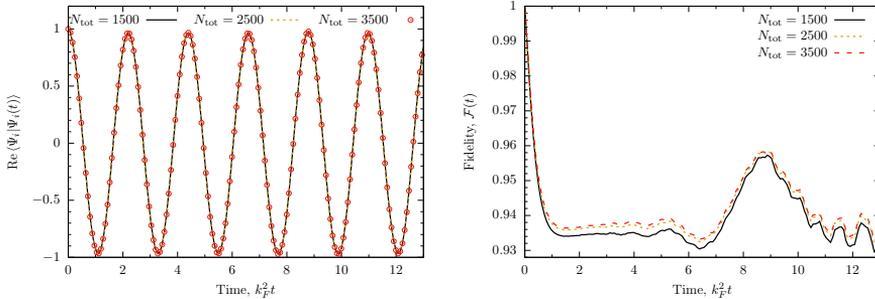
Figure 4.12: The real part of the return amplitude as defined in Eq. (4.11) (left) and the fidelity as defined in Eq. (4.13) (right) for the quenched ground state from $c_i = 20$ to $c_f = 10$ at $N = 10 = L$. We show the aforementioned quantities for three different sizes of the computational basis. Furthermore, $k_F = \pi(N-1)/L$ is the Fermi momentum as discussed in Sec. 2.3. Both the return amplitude and the fidelity are seen to converge rapidly with the size of the computational basis as compared to the energy as we investigated in Fig. 4.8.

energy and local expectation values we have considered thus far. It is generally unclear if convergence in one case implies convergence in the other. What we do know is that the energy is an unbounded operator which makes that the terms in Eq. (4.5) are unbounded, whereas each term appearing in Eq. 4.11 is bounded by one. Therefore it is perhaps unsurprising that we observe quicker convergence for the return amplitude and the fidelity, which show excellent convergence for (very) small computational basis, as shown in Fig. 4.12.

It should be noted that computing the return amplitude and the fidelity is easy since it only involves a single summation over N_{tot} overlaps which we have already determined. This will no longer be the case when we consider the time evolution of expecation values of operators in which case we will see that we have to evaluate and sum at least N_{tot}^2 matrix elements.

4.3.2 Time evolution of local observables

Consider an operator O and suppose we are interested in the time evolution of its expectation value with respect to $|\Psi_0(t)\rangle$. In constrast to the calculation of the return amplitude and the fidelity, this involves summing over N_{tot}^2 terms
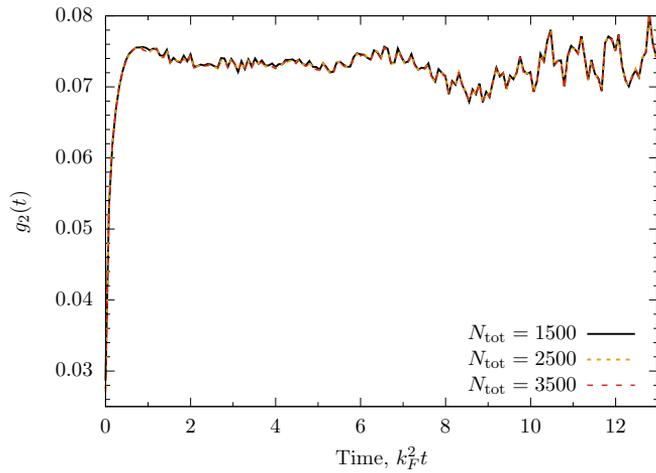
Figure 4.13: The time evolution of the expectation values of $g_2(t)$ for the ground state following a quench in the interaction strength from $c_i = 20$ to $c_f = 10$ at $N = 10 = L$. Here, $k_F = \pi(N-1)/L$ is the Fermi momentum as discussed in Sec. 2.3. We show results for three different sizes of the computational basis, which shows that the time evolution has converged even for the small basis sizes considered.

given by

$$\langle O(t)\rangle_i \equiv \langle \Psi_0(t)|O|\Psi_0(t)\rangle,$$

$$= \sum_{n,m=0}^{N_{\text{tot}}} e^{-it[E_n - E_m]}\langle \Psi_0|m\rangle\langle m|O|n\rangle\langle n|\Psi_0\rangle. \qquad (4.14)$$

Besides the fact that this involves summing over N_{tot}^2 terms rather than just N_{tot} as was the case for the return amplitude and the fidelity, it also involves computing the matrix elements $\langle m|O|n\rangle$, which in itself can be computationally difficult.

Having seen rapid convergence for the return amplitude and the fidelity for small computational bases, we can hope that the same will be true for the time evolution of more general local quantities. Here we consider the operator O to be g_2, whose matrix element expressions we discussed in Sec. 2.4.2. In Fig. 4.13 we show the time time evolution of

$$g_2(t) = \langle \Psi_0(t)|g_2(0)|\Psi(t)\rangle \qquad (4.15)$$

for the quench from $c_i = 20$ to $c_f = 10$ at $N = 10 = L$. Again we see excellent convergence for the same small computational bases.

The small computational bases required to obtain convergence for the time evolution in this section and the previous one indicate that we will be able to compute the time evolution for observables for larger particle numbers with reasonable computational resources as well. An example of this can be seen in Sec. 3.

1. Comparison to the coordinate Bethe ansatz

   In order to check the time evolution results we converge to are indeed correct, we compare to results known from the literature. Not a lot of results for quenches in the interaction strength at finite values of the interaction strength are present in the literature due to the lack of analytical expressions for the overlaps. An exception are some results regarding the time evolution of $g_2(0)$ using the coordinate Bethe ansatz presented in [49]. In this section, a quench from $c_i = 100$ to $c_f = 3.7760$ for $N = 5$ is considered and the subsequent time evolution is computed. The approach is exact, but limited to small system sizes since the coordinate Bethe ansatz approach scales extremely poorly with system size. In fact, the computation of $g_2(t)$ involves dealing with a summation over $(N!)^2$ terms.

   The extreme difference in the interaction strength considered in [49] represents a challenging quench due to the difference in energy density of
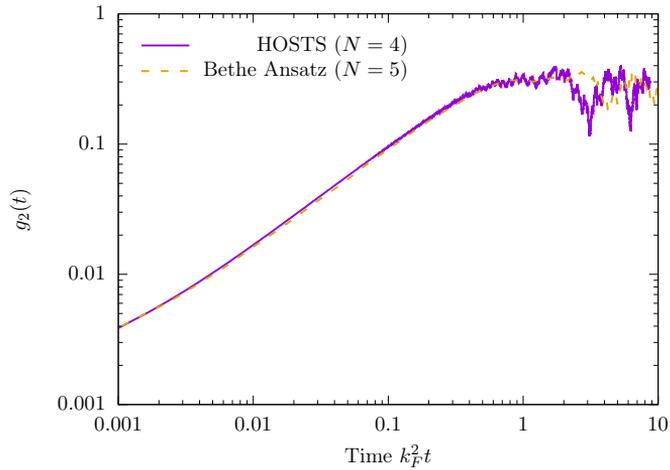
Figure 4.14: The time evolution of the expecation value of $g_2(0)$ with respect to $|\Psi_0(t)\rangle$ for the quench from $c_i = 100$ to $c_f = 3.7660$ for $N = 10 = L$. We compare to exact data obtained using the coordinate Bethe ansatz approach for $N = 5$ (dashed line) extracted from [49] to the high overlap states truncation scheme computation for $N = 4$. We discuss results from HOSTS for larger system sizes in Sec. 4.4.
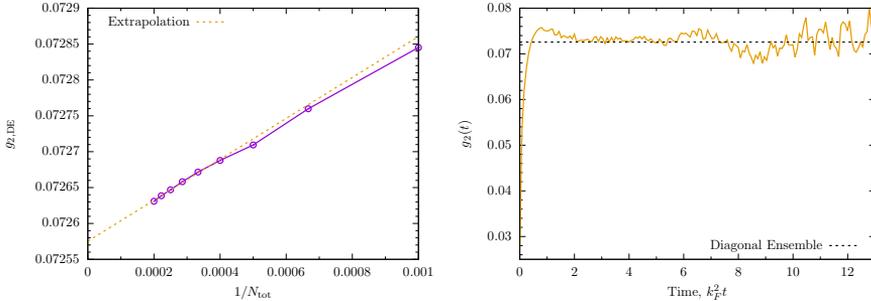
Figure 4.15: The convergence of the diagonal ensemble prediction for $g_2(0)$ as a function of the size of the computational basis (left) and its time evolution compared to the extrapolated value in the diagonal ensemble (right). Again we consider the quench from $c_i = 20$ to $c_f = 10$ for $N = 10 = L$. The extrapolated value from the diagonal ensemble is compared to the dynamics obtained using $N_{\text{tot}} = 3500$ as shown previously in Fig. 4.13.

the ground state before and after $t = 0$. This likely leads to the creation of many excitations making it a challenging scenario to test our method on. Additional numerical difficulties are introduced when considering uneven numbers of particles due to the presence of rapidities close to zero. As a result, we compare the coordinate Bethe Ansatz results for $N = 5$ to numerical data obtained for $N = 4$ (we discuss the case where $N = 6$ later). We present the time evolution data for $g_2(t)$ in Fig. 4.14. We observe excellent agreement up to the revival time which is relevant due to this being a small periodic system. More importantly, this comparison shows that we are able to capture all the physics relevant to the non-equilibrium dynamics at a significantly reduced computational cost.

2. The long time limit: the diagonal ensemble

Thus far we have shown that we can determine the out-of-equilibrium dynamics of observables at finite times, but we can also determine the long time limit of the dynamics. In this limit the problem simplifies, as is well-studied in the literature (see e.g. [47]). In our case, we have determined the overlaps, so we can directly consider the long time limit in the diagonal ensemble to be [47]

$$\langle O \rangle_{\text{DE}} = \sum_j \langle j|O|j \rangle \, |\langle j|\Psi_0 \rangle|^2 \,. \tag{4.16}$$

This is related to the long time limit of the time-averaged expectation values in the limit where the system is also large (see e.g. the appendix
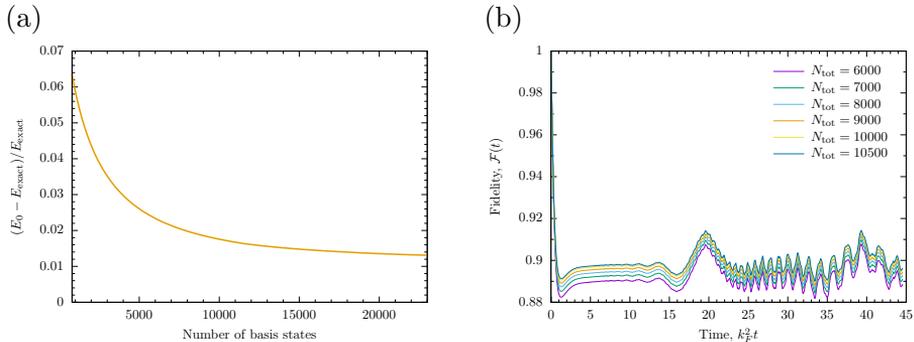
Figure 4.16: Example calculation showing (a) the energy convergence of the initial state; (b) the time evolution of the fidelity for $N = 20$ particles, for the quench $c_i = 20 \rightarrow c_f = 10$.

of [49]), i.e.

$$\lim_{T \to \infty} \frac{1}{T} \int_0^T \mathrm{d}t \langle \Psi_0(t)|O|\Psi_0(t)\rangle \to \langle O \rangle_{\mathrm{DE}}. \qquad (4.17)$$

It turns out that if O equilibrates sufficiently quickly, the long time limit is actually equal to the result in the diagonal ensemble

$$\lim_{t \to \infty} \langle \Psi(t)|O|\Psi(t)\rangle \to \langle O \rangle_{\mathrm{DE}}. \qquad (4.18)$$

as shown in for example [83]. Therefore we can use the diagonal ensemble to determine the long time limit from the overlaps.

We show results obtained for the diagonal ensemble for the g_2 operator using the high overlap states truncation scheme Fig. 4.15. In the left panel we consider the convergence of the diagonal ensemble result with growing size of the computational basis and extrapolate to find the result at $N_{\mathrm{tot}} \to \infty$. Then we compare this result obtained via extrapolation in the right panel to the dynamics of $g_2(t)$ previously shown in Fig. 4.13. We see that fitting a straight line to the convergence data fits very well and that the resulting infinite basis value agrees with the finite time dynamics up to fluctuations. The presence of fluctuations in this case are unsurprising since we are still only dealing with ten particles.

3. An example with larger numbers of particles

So far, we have examined quenches with relatively small numbers of particles, $N = 4, 6, 10$. It is worth emphasizing that even for these numbers of particles, exact calculations via the coordinate Bethe ansatz are

computationally expensive, and exact calculations with $N = 10$ corresponding to summing $\propto (N!)^2 \sim 1.3 \times 10^{13}$ terms. To even contemplate exact evaluation for $N = 20$ particles, $\propto 5.9 \times 10^{36}$ terms, seems futile. Instead the high overlap states truncation scheme, in combination with the numerical renormalization group, gives one a handle on such problems.

Here, we consider the $c_i = 20 \rightarrow c_f = 10$ quench for a larger numbers of particles ($N = 20$) as an illustrative example.

Results for the energy convergence and the time evolution of the fidelity, $\mathcal{F}(t)$, are shown in Fig. 4.16.

4.4   Introducing the Matrix Element Renormalization Group

In this section, we consider a perturbing operator $(c_i - c_f)Lg_2(0)$ whose matrix elements, with respect to some of the computational basis states, are large compared to the energy difference between these states and the unperturbed ground state. We call quenches for which the perturbing operator satisfies this property "strongly non-perturbative". In such cases relying solely on the metric (4.10) discussed in the previous section, which was motivated by leading order perturbation theory, is no longer justified. Higher order terms are expected to be relevant and, as a result, we need to modify the way we select and order states. Not only this, but we need to re-examine the assumptions behind the numerical renormalization group procedure discussed in Sec. 4.2.2, and modify these accordingly. In the final part of this section we show how the resulting procedure allows us to not only treat quenches where the initial state is a ground state, but also those where it is an excited state.

As matrix elements of the perturbing operator become large, contributions of a given computational basis state to the ground state mediated via other intermediate computational basis states can become relevant. This corresponds to the second order terms in Eq. (4.9) no longer being negligible compared to the first order terms. However, such contributions are not considered in the standard numerical group procedure as discussed in Sec. 4.2.2. As a result, these contributions are missed when states are not by chance included in the same step of the renormalization group procedure. We will see that these contributions can play an important role for strongly non-perturbative quenches, so they need to be taken into account.

An illlustration of how naively applying the algorithms developed thus far can lead to inaccurate results for strongly non-perturbative quenches is shown in Fig. 4.17. Initially the differences between the results obtained using the numerical renormalization group and full diagonalization are small, but as the number of iterations increases, the discrepancy becomes larger.
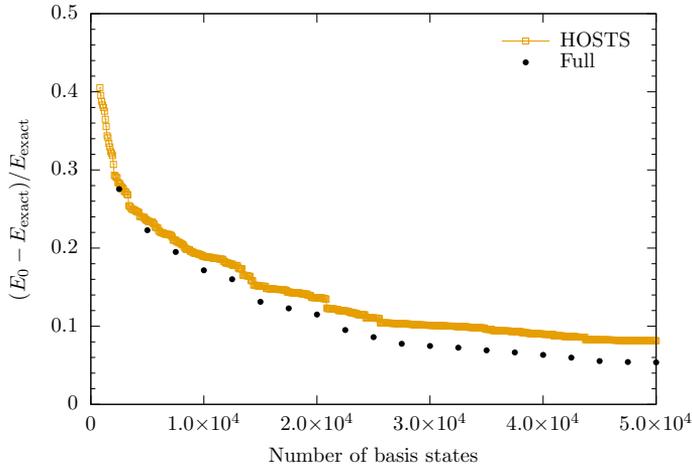
Figure 4.17: An example calculation showing that for strongly non-perturbative quenches there can be a large discrepancy between the results from full diagonalisation and the NRG-extension of HOSTS. We consider the $c_i = 100 \rightarrow c_f = 3.7660$ quench for $N = 6$ particles, cf. Fig. 4.14. $N + \Delta N_s = 800$ is fixed within each data set.

4.4.1 THE MATRIX ELEMENT RENORMALIZATION GROUP ALGORITHM FOR THE GROUND STATE

To deal with the problem for strongly non-perturbative quenches discussed in the previous section, we develop a reworking of the numerical renormalization group procedure that we refer to as the matrix element renormalization group. The main differences between the matrix element renormalization group and the conventional renormalization group algorithm are:

1. All the approximate eigenstates obtained at each step of the algorithm (from diagonalization of a truncated Hamiltonian) are kept, unlike in the conventional case where one discards ΔN_s states at each iteration.

2. When introducing new computational basis states, we select which of the previously obtained approximate eigenstates to include in the Hamiltonian using a weighing function based on the quadratic terms in the perturbation-series expansion of the wave function (4.9) instead of including the approximate eigenstates with the lowest energies.

The matrix element renormalization group takes seriously the idea that matrix elements of the perturbing operator, rather than energies, are the important quantity when operators are not strongly renormalization group relevant.

The central idea is that computational basis states $|\lambda_j\rangle$ newly included at a given step can mediate strong coupling between the approximate ground state and approximate "excited states" obtained at an earlier iteration. These "approximate states" must then be included in the truncated Hamiltonian at this diagonalization step to ensure an accurate description of the ground state. So, instead of blindly removing the high energy approximate "excited states" at each step of algorithm (as in the conventional numerical renormalization group), we keep all approximate eigenvectors, and at each iteration include the states most important for mediating the coupling between the approximate ground state and the newly added states from the computational basis.

Let $|\Omega\rangle$ be the ground state of the final Hamiltonian, then the steps of the matrix element renormalization group are as follows:

1. Generate the computational basis via preferential state generation from the ground state $|\Omega\rangle$. Order the states in the computational basis according to the metric in Eq.(4.10), to obtain $\{|\lambda_j\rangle\}$.

2. Construct a truncated Hamiltonian from the first $N_s + \Delta N_s$ computational basis states, $\{|\lambda_1\rangle, \ldots, |\lambda_{N_s + \Delta N_s}\rangle\}$ and diagonalize this Hamiltonian to obtain the first approximate eigenstates $\{|1\rangle, \ldots, |N_s + \Delta N_s\rangle\}$ with energies $\{E_1, \ldots, E_{N_s + \Delta N_s}\}$. These approximate eigenstates replace the first $N_s + \Delta N_s$ states in the computational basis, and are ordered such that $\tilde{E}_1 < \cdots < \tilde{E}_{N_s + \Delta N_s}$.

3. Define a new basis of $N_s + \Delta N_s$ eigenstates for the truncated Hamiltonian by adding the next ΔN_s states from the computational basis $\{|\lambda_j\rangle\}$ to the approximate eigenvector with the lowest energy $|1\rangle$ as well as the $N_s - 1$ approximate eigenstates $\{|i\rangle\}_{i>1}$ whose "second order weight", given by

$$w_2\left(|i\rangle\right) = \sum_j \frac{\langle i|\delta H|\lambda_j\rangle\langle\lambda_j|\delta H|1\rangle}{(\tilde{E}_1 - E_j)(\tilde{E}_1 - \tilde{E}_i)}. \tag{4.19}$$

is largest. Here $\delta H = cLg_2(0)$ is the perturbing operator (see Sec. 4.1, the sum ranges over all ΔN_s newly added computational basis states, and $E_j = E(|\lambda_j\rangle)$ are the energies of the newly added computational basis states.

4. Construct the truncated Hamiltonian in this new basis and diagonalize it to obtain $N_s + \Delta N_s$ new approximate eigenstates. These newly constructed approximate eigenstates replace the states in the computational basis used to construct the truncated Hamiltonian.

5. Return to the third step.

This process is continued, obtaining new approximate eigenstates after each cycle of steps 3 to 5, until the required convergence of the ground state energy is reached or the computational basis is exhausted.

The matrix element renormalization group has some slight disadvantages when compared to the conventional numerical renormalization group. Firstly, it is more memory intensive: a complete set of approximate eigenstates must be retained in the procedure, while in the conventional routine we only need keep track of N_s such approximate eigenstates.[3] Secondly, the matrix element renormalization group has a higher computational burden since it requires the computation of the second order weight for all approximate eigenvectors at the start of each iteration. However, we have seen that the conventional numerical renormalization group fails to produce accurate results for strongly non-perturbative quenches (see Fig.4.17) so these savings in memory and computations compared to the matrix element renormalization group are moot.

There are a couple of alternative, complementary, schemes that could be used to construct the initial state. Firstly, there exist "sweeping" improvements of the conventional numerical renormalization group (see, e.g., their discussion in [46]). If succesful however, this additional would certainly come at a higher computational cost than directly using the matrix element renormalization group. Secondly, one could invoke iterative diagonalization (via, e.g., Lanczos or Davidson) within a given truncated basis. In such a procedure, one would have to check convergence of results with basis size, but one can (in principle) deal with very large bases. How quickly such iterative diagonalization converges, with our matrix being dense, is not clear. We have not explored this avenue, but it is an interesting direction for future works.

4.4.2 Results from the matrix element renormalization group for the ground state

With the matrix element renormalization group algorithm in place, we can employ it to tackle problems that are inaccessible to the conventional numerical renormalization group.

However, first we check that the matrix element renormalization group correctly reproduces results in cases where the conventional numerical renormalization group approach works. This is a basis sanity check: can we reproduce the initial state and its dynamics in these simpler cases. Our first example is $c_i = 20 \to c_f = 10$ quench studied earlier in this chapter. We present the convergence of the energy and the overlaps in Figs. 4.18(a)–(b). In particular, Fig. 4.18(b) should be compared to Fig. 4.9 obtained previously. We see excellent agreement between the conventional numerical renormalization group

---

[3]In practice, since we need at most N_s approximate eigenvectors at any one time, these eigenvectors do not have to be stored in memory.
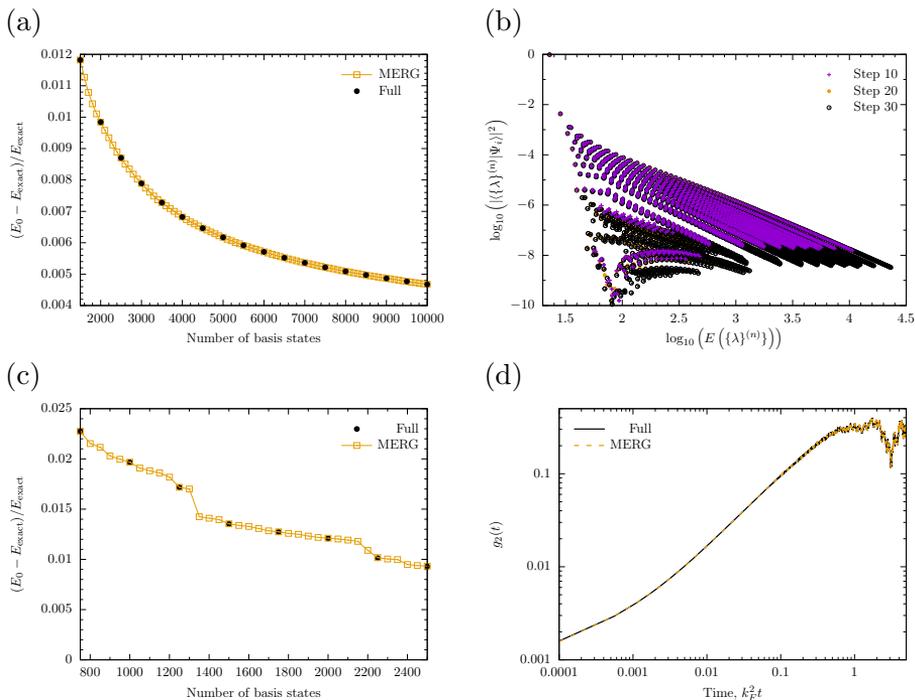
Figure 4.18: *Top row:* Matrix element renormalization group (MERG) and numerical renormalization group (NRG) for the $c_i = 20 \rightarrow c_f = 10$ quench for $N = 10$ particles, computed within the high overlap states truncation scheme. (a) The convergence of the initial state energy as a function of number of basis states; (b) the convergence of the overlaps at different steps of the MERG procedure (cf. Fig. 4.9). *Bottom row:* MERG and full diagonalization results for the $N = 4$ particle quench $c_i = 100 \rightarrow c_f = 3.7660$: (c) the convergence of the energy of the initial state with the number of basis states; (d) the time evolution of $g_2(t)$ with 2500 basis states.
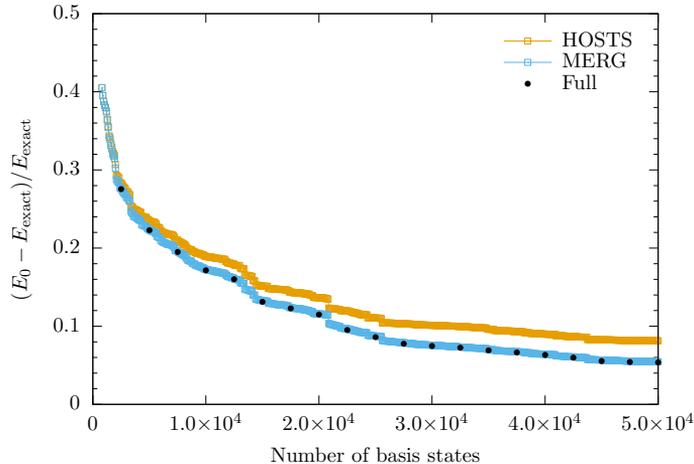
Figure 4.19: The convergence of the initial state energy E_0 for $N = 6$ particle quench $c_i = 100 \to c_f = 3.7660$ obtained with the matrix element renormalization group (MERG). MERG is performed with $N_s = 720$ and $\Delta N_s = 80$. Conventional numerical renormalization group approaches break down in this scenario (for the same $N_s, \Delta N_s$) as shown in Fig. 4.17. Full diagonalization results, for the same number of basis stats, are shown for comparison.

and the matrix element renormalization group in this scenario.

As a second check, we turn our attention to the harder quench considered in the previous section for $N = 4$ particles, $c_i = 100 \to c_f = 3.7660$. Here we check against full diagonalization of the truncated Hamiltonian (as the required number of states for excellent convergence is rather small), as shown in Figs. 4.18(c)–(d). The matrix element renormalization group gives results in excellent agreement with full diagonalization of the same basis, both in terms of energy of the initial state, Fig. 4.18(c), and the non-equilibrium dynamics of observables, Fig. 4.18(d).

With the matrix element renormalization group correctly reproducing both full diagonalization (in small bases) and conventional numerical renormalization group (in large bases) results, we examine the problematic scenario discussed in the previous section. In this strongly non-perturbative quench, the matrix element renormalization group is vital for correctly constructing the initial state. In scenarios where the conventional numerical renormalization group fails to produce results which agree with results obtained by diagonalization of the full truncated Hamiltonian, such as the one illustrated in Fig. 4.17, the matrix element renormalization group continues to produce results that agree with great accuracy as is shown in Fig. 4.19.
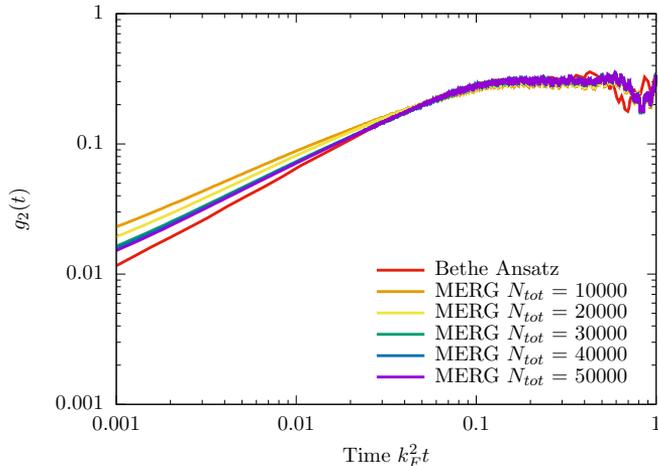
Figure 4.20: The time evolution of the local observable $g_2(0)$ following the $c_i = 100 \to c_f = 3.7660$ in the Lieb-Liniger model starting from the ground state at c_i. Exact data (dashed line) computed via the coordinate Bethe ansatz with $N = 5$ particles (from Ref. [49]) is compared to matrix element renormalization group (MERG) calculations with $N = 6$ particles.

In Fig. 4.19 we see a number of features. Firstly, we note that the agreement between the results obtained by full diagonalisation of the truncated Hamiltonian and using the matrix element renormalization group are in excellent agreement. Secondly, regardless of the procedure, the convergence of the initial state energy shows some plateaus and jumps, which implies that the metric (4.10) is not the perfect one. Understanding how to construct the most convergent metric for a given problem is an outstanding challenge, which requires further investigations. Thirdly, we see that for $N = 6$ particles the problem is very challenging: By including $50,000$ states, we still only achieve initial state energies correct to within $\sim 5.5\%$ ($\sim 10\%$ w.r.t. the Fermi energy). Whilst a better ordering metric might improve this, it still seems likely that strongly nonperturbative quenches will present a significant numerical challenge. This is further supported by Fig. 4.20, where we show the time evolution of $g_2(t)$ as compared to results from the coordinate Bethe ansatz discussed previously. We see that even a truncated wave function (4.5) with $50,000$ states included does not accurately realize the short time dynamics of observables. At longer times, once the steady state plateau is approached (and high energy modes have dephased, effectively averaging to zero), the truncated wave function does describe $g_2(t)$ well.

4.4.3 The matrix element renormalization group algorithm for excited states

In contrast to conventional renormalization group techniques, which at most target the first few excited states in addition to the ground state, the matrix element renomalization group can also be used to target more highly excited states. Furthermore, the algorithm can construct these states without the need to construct all the states of lower energy making the procedure a more efficient tool to target excited states than the methods discussed thus far. In order to do so we have to make some changes to the algorithm described in Sec. 4.4.1.

To understand why we need to change the algorithm in Sec. 4.4.1 in order to consider excited states, let us consider what happens if we replace the ground state with an excited state in the algorithm. This state will henceforth be referred to as the seed state. First of all, the preferential state generation routine leads to a different computational basis, as we now consider $|0\rangle$ in Eq.(4.10) to be an excited state. In particular, the ground state may not have a high weight according to this metric, so it may not even be included in the computational basis obtained via preferential state generation. Second of all, the algorithm retains the lowest energy approximate eigenstate $|E_1\rangle$ at every step, and selects the approximate eigenstates most relevant to this approximate eigenstate based on the second order weight. This means that the selection rules are still set to promote the convergence of the lowest energy eigenstate, rather than an excited state. Finally, it is generally unclear which eigenstate of the perturbed Hamiltonian corresponds to which of the approximate eigenstates obtained by the algorithm, as energies of different excited states can be very similar and even degenerate.

Before we discuss how we resolve these issues, note that every step of the algorithm represents a mapping between the states used to construct the truncated Hamiltonian and the approximate eigenstates obtained by diagonalization. To identify which of the approximate eigenstates a given state used in the basis for the truncated Hamiltonian is mapped to, we compute the overlaps between this state and all newly obtained approximate eigenstates. The approximate eigenstate with the largest overlap is then said to be its image provided that the RG-step is "small enough". This allows us to track the approximate eigenstate derived from the seed state throughout the procedure.

The main assumption behind our method of tracking the seed state is that the number of states added at every step of the routine is small enough, so that no single iteration wildly changes the (image of) the seed state. What step size is small enough for this assumption to hold depends on the quench and seed state under consideration.[4] However, there are some general methods

---

[4]It may happen that no step size is small enough when we consider a very strong quench

by which one can check if an appropriate step size has been chosen. Firstly, one can consider the overlaps computed at each iteration of the routine and verify that there is only one state with a significant overlap. Secondly, one can rerun the routine with a smaller step size and check that it produces the same results. With the preferential scanning routine in place, by which the most significantly states are identified and included first, the start of the routine is where the changes are most drastic and therefore the procedure is most likely to break down there. As a result, the checks proposed here need not be time-consuming.

Now that we have established how we can track the seed state, we note that we can replace the lowest energy approximate eigenstate with the image of the seed state in the second order metric used in Eq. (4.20). This change, together with the replacement of the ground state with an arbitrary seed state in the preferential scanning routine results in a routine designed to optimize the convergence of the approximate eigenstate associated to the seed state. The resulting algorithm can be summarized as follows.

Let $|\Omega\rangle$ be some eigenstate of the final Hamiltonian, which in this case can be an excited state, then the steps of the matrix element renormalization group are as follows:

1. Generate the computational basis via preferential state generation from the seed state $|\Omega\rangle$. Order the states in the computational basis according to the metric in Eq. (4.10), to obtain $\{|j\rangle\}$.

2. Construct a truncated Hamiltonian from the first $N_s + \Delta N_s$ computational basis states, $\{|\lambda_1\rangle, \ldots, |\lambda_{N_s + \Delta N_s}\rangle\}$ and diagonalize this Hamiltonian to obtain the first approximate eigenstates $\{|1\rangle, \ldots, |N_s + \Delta N_s\rangle\}$ with energies $\{\tilde{E}_1, \ldots, \tilde{E}_{N_s + \Delta N_s}\}$.

3. Compute the overlaps between $|\Omega\rangle$ and the newly acquired approximate eigenvectors $\{|1\rangle, \ldots, |N_s + \Delta N_s\rangle\}$. Then relabel the approximate eigenstates such that $|1\rangle$ refers to the approximate eigenstate with the largest overlap.

4. Take the next ΔN_s computational basis states $|\lambda_j\rangle$ and compute the "second order weight" for each of the approximate eigenstates $|i\rangle$ with $i > 1$ obtained in previous steps:

$$w_2\left(|i\rangle\right) = \sum_j \frac{\langle i|\delta H|\lambda_j\rangle\langle\lambda_j|\delta H|1\rangle}{(\tilde{E}_1 - E_j)(\tilde{E}_1 - \tilde{E}_i)}. \tag{4.20}$$

---

and/or a highly excited state.

Here $\delta H = cLg_2(0)$ is the perturbing operator (see Sec. 4.1, the sum ranges over all the newly added computational basis states, and $E_j = E(|\lambda_j\rangle)$ are the energies of the newly added computational basis states.

5. Form a truncated basis consisting of $|1\rangle$, the $N_s - 1$ states in $\{|i\rangle\}_{i>1}$ with the largest w_2-weight, and the ΔN_s computational basis states introduced in step 4, construct the truncated Hamiltonian in this basis and diagonalize it to obtain $N_s + \Delta N_s$ new approximate eigenstates $\{|1'\rangle, \ldots, |(N_s + \Delta N_s)'\rangle\}$.

6. Compute the overlaps between $|1\rangle$ and the newly acquired approximate eigenvectors, and replace $|1\rangle$ by the approximate eigenvector with the largest overlap. Replace the remaining approximate eigenvectors used to form the truncated basis in step 5 with the remaining newly obtained approximate eigenvectors.

7. Return to the fourth step.

This process is continued, obtaining new approximate eigenstates after each cycle of steps 4 to 6, until the required convergence of the eigenstate is reached or the computational basis is exhausted.

This version of the matrix element renormalization group is not more memory intensive than the routine presented for constructing ground states and it is only slightly more computationally intensive. The additonal computational cost comes from computing the overlaps at each iteration.

4.4.4 Results from the matrix element renormalization group for excited states

As mentioned in Sec. 4.4.3, one of the subtleties that arises when considering the matrix elemenent renormalization group for excited states is that, even though we know that we construct an approximate eigenstate of the perturbed Hamiltonian, we do not necessarily know a priori what eigenstate this will correspond to. For the interaction quench considered here, the most natural eigenstate of $H(c_i)$ to end up with when starting from an eigenstate of $H(c_f)$ is the eigenstate with the same quantum numbers. In this section we show some preliminary results to verify this claim, although we do note that to assert with more certainty that this claim is true, more properties of the eigenstates other than the energies would have to be considered. This is left to future works.

Consider again the quench from $c_i = 20$ to $c_f = 10$. In the following we present the results obtained from running the algorithm three times with three

different seed states, whose doubled quantum numbers are given by

$$\text{State A: } \{-9, -7, -5, -3, -1, 1, 3, 5, 7, 9\} \tag{4.21}$$

$$\text{State B: } \{-11, -7, -5, -3, -1, 1, 3, 5, 7, 11\} \tag{4.22}$$

$$\text{State C: } \{-17, -13, -9, -5, -1, 1, 5, 9, 13, 17\}. \tag{4.23}$$

The results for the energy convergence of the approximate eigenstates corresponding to these seed states obtained from the matrix element renormalization group are shown in Fig. 4.21. In order to keep track of the convergence, we again consider the percentual error of the energy only this time with respect to a different target energy of the each of the runs. The target energy E_{exact} is taken to be the energy of the eigenstate of $H(c_i)$ with quantum numbers identical to those of the seed state under consideration. We note that even though we only consider data for the energy convergence here, we have still computed the expansion of the approximate eigenstates in terms of the eigenstates of the initial basis, so could still compute the time evolution of operators if we please to do so.

The rate of convergence of the run seeded by the lowest excited state, state B, is comparable to the convergence when considering the ground state, state A. On the other hand, when considering a run seeded by a highly excited state, state C, the convergence shows characteristics reminding us of the strongly non-perturbative quenches considered in Sec. 4.4.2. Also, in order to keep track of the right approximate eigenstate throughout the procedure, we had to significantly alter the parameters characterizing the size of the renormalization group steps to $N_s = 75$ and $\Delta N_s = 25$.

The fact that considering higher energy states requires a decrease in step size is what currently limits how high the energy of the seed states may be. To overcome this limitation, we would have to reorder the computational basis so that single steps of the procedure do not change the targeted approximate eigenstate as violently. Nevertheless, even without such alterations, our current algorithm goes beyond what one could target using the conventional renormalization group techniques we discussed at the start of this paper, because in that case one would have to construct all lower energy eigenstates.

4.5 CONCLUSIONS

Even in the presence of integrability, the computation of non-equilibrium dynamics following a quantum quench remains a great challenge for theory. Well-controlled numerical approaches are vital for accessing the physics away from analytically tractable limits, including for the cases of finite-time dynamics of observables. Here we have presented a proof-of-principle investigation of finite-c to finite-c quenches in the Lieb-Liniger model using a high overlap
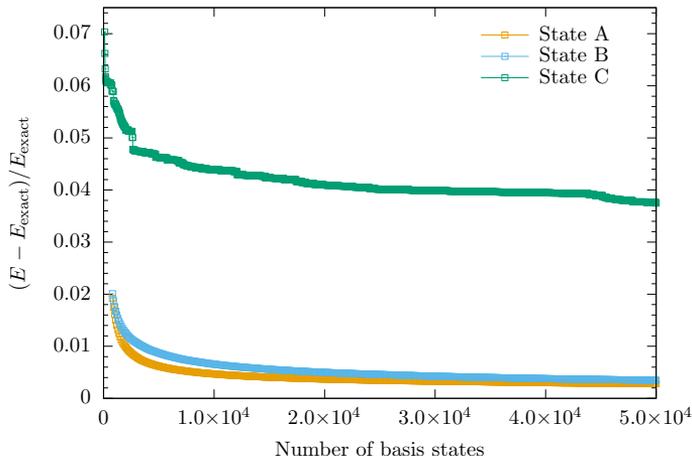
Figure 4.21: The convergence of the states A, B, and C, for the $N = 10$ particle quench $c_i = 20 \rightarrow c_f = 10$ obtained using the matrix element renormalization group (MERG). For state A, and B, MERG is performed with $N_s = 700$ and $\Delta N_s = 100$, whereas for state C, MERG is performed with $N_s = 75$ and $\Delta N_s = 25$.

states truncation scheme, in combination with full diagonalization, the numerical renormalization group, and a new matrix element renormalization group algorithm. We have worked with interacting computational basis states, which intrinsically have built-in strong correlations, and we have systematically constructed initial states in terms of high overlap states, for quenches starting from ground states as well as excited states. Using these, we have computed both real-time dynamics and the long-time limit of physical observables following a quench.

In our development of a high overlap states truncation scheme, and the matrix element renormalization group, we have highlighted the important role played by the ordering of the computational basis. Applying the conventional metric, energy of the computational basis states, we observe poor convergence of properties of the initial states. This poor convergence means applying conventional "truncated spectrum methods" (in their naive form) requires the use of unfeasibly many computational basis states. By modifying the metric, to a "matrix element" focused one that takes into account the structure of the operator coupled to the quench parameter, we achieve orders-of-magnitude improvement in the convergence of properties of the initial state with truncated Hilbert space size. This was studied in Sec. 4.1. Along the way, we were able to develop a routine that preferentially generates the states with high overlap

following a quench, and this enabled efficient convergence of the initial state energy to sub-percent precision.

This improved convergence opened the door to computing non-trivial non-equilibrium dynamics for numbers of particles far beyond the reach of brute force computations. This was discussed in Sec. 4.3. Convergence of real-time non-equilibrium dynamics of local observables with the number of computational basis states was surprisingly fast: for $N = 10$ particles $c_i = 20 \to c_f = 10$ quench, well-converged results for time evolution of $g_2(0)$ are obtained with just thousands of states (some of which are of very high in energy). The long-time limit was also shown to be efficiently accessed via the diagonal ensemble, with results agreeing with the intermediate time dynamics, as expected.

In the case of strongly non-perturbative quenches, we found that conventional numerical renormalization group improvements have to be significantly modified to achieve accurate results. This modification, the so-called matrix element renormalization group, takes seriously that the properties of the perturbing operator should govern the whole procedure. We found this modification to be necessary in the "large quench" studied previously in the literature [49], $c_i = 100 \to c_f = 3.7660$, when considering more than four particles. Our results were compared to the coordinate Bethe ansatz results of Zill *et al.* [49], and were found to be in excellent agreement. We note, however, that such strongly non-perturbative quenches remain challenging problems, with the quench projecting the initial state on to many states with sizable overlaps. This makes it tough to tackle even relatively small numbers of particles, even with our computationally efficient approach. This seems like an insurmountable problem with introducing additional approximations, beyond the scope of this work, or alternative Hilbert space ordering metrics.

Finally, we considered the construction of excited states of the perturbed Hamiltonian. We wrote down a variant of the matrix element renormalization group algorithm able to directly construct excited states of a perturbed Hamiltonian in terms of the eigenbasis of another Hamiltonian without having to construct all lower energy eigenstates. This allows one to consider more highly excited states than one normally could using the truncated spectrum approach and focusses the computational resources on this particular eigenstate rather than it being a less well-converged side-product of trying to converge the approximate eigenstate representing the ground state of the perturbed Hamiltonian.

The presented high overlap states truncation scheme, combined with full diagonalization and renormalization group improvements, can be applied to many other models and scenarios. Perhaps the most interesting is to consider the case with integrability-breaking where, provided matrix elements of the integrability-breaking terms are known, one can directly apply the same approach. This enables, for example, non-perturbative studies of prether-

malization (see, e.g., Refs. [84–87, 23]) in continuum quantum gases. Other interesting directions include extensions to other integrable continuum models, such as two-component Bose and Fermi gases or the sine-Gordon regime away from the ultra-relevant perturbation limit [55]. Finally, we would like to point out that the method developed in this paper provides, in principle, all the ingredients necessary to compute for example the time evolution of the entanglement entropy. In order to come to a tractable computation one can convert the overlaps coming from the NRG-routines to a root distribution and then use the quasi-particle picture formulas for the entanglement entropy, see e.g. [88, 89]. However, in order to ascertain the accuracy of results obtained in this way, a careful quantitative study of finite-size effects is required in order to determine if we can accurately match results in the thermodynamic and scaling limits. We leave addressing this challenge to future work.

Extending these methods to lattice models should also be possible, using strongly correlated integrable eigenstates. Such an algorithm may complement existing ones: being able to tackle longer times, but smaller systems, than the time-dependent density matrix renormalization group, but larger system sizes than exact diagonalization. It may also be interesting to implement the ideas behind the matrix element renormalization group to lattice and impurity models, invoking a Wilsonian numerical renormalization group-like picture with strongly correlated basis states. These points remain for future works.

The approach implemented within this chapter for simulating continuum one-dimensional models provides an alternative, complementary approach to continuum matrix product state methods [90, 91]. Utilizing the solvability of a proximate integrable point, time evolution is easy within our approach and can be performed to long times with high precision. This opens the door to novel, non-perturbative studies of non-equilibrium dynamics in models of relevance to cold atomic gases.

CHAPTER 5

PROPERTIES OF MATRIX ELEMENTS

The previous chapters have shown the central role played by the matrix elements to the computation of correlation functions as well as the time evolution following a quantum quench. For both problems we were interested in the rare eigenstates whose matrix elements are most important to the calculation at hand. In this chapter we study the matrix elements in the Lieb-Liniger model in more detail. In particular, we consider not only the rare eigenstates whose matrix elements with respect to some given operator are largest but also the statistics of matrix elements of the typical states. The results presented in this chapter are part of an ongoing investigation in collaboration with F.H.L. Essler, N.J. Robinson and J.-S. Caux [3]. The paper that we are preparing on the subjects discussed in this chapter also features numerical and analytical results for the infinitely repulsive limit, but since my contribution concerned the interacting limit I will restrict the discussion here to the latter.

5.1   THE ONE PARTICLE-HOLE SECTOR

The states with few particle-hole pairs are disproportionally important to the calculation of, for example, the dynamical structure factor as we saw in Chapter 3. For this calculation the single particle-hole sector is most the most important one due to the nature of the density operator as explained in Sec. 2.4. In this section we consider the matrix elements of the density operator for this important sector of atypical states with respect to the ground state as well as representative thermal states. We do so at a fixed value of the momentum so that there are at most N such states to consider rather than considering the entire infinite family of states with a single particle-hole pair we would otherwise have to deal with.

For concreteness, at any particle number N and interaction stregth c we consider the corresponding representative state $\{I_i^r\}_{i=1}^N$ describing the ground state or the best finite size approximation for the $T = 10$ thermal state at unit density. It is with respect to this reference state that we define the number of

particle-hole pairs and compute the off-diagonal matrix elements. To generate a state with a single particle-hole pair at a momentum value k, we have to increase one of the quantum numbers of $\{I_i^r\}$ by $\frac{kL}{2\pi}$. Trying this addition for every quantum number and seeing if it results in a valid state, containing no coinciding quantum numbers, gives us all states with a single particle-hole pair at the desired momentum.

The first step to understanding the matrix elements in this sector is to find the scaling with system size. To do so, we generate the states $|\mu\rangle$ with a single particle-hole pair w.r.t. the reference state $|\lambda\rangle$ and compute both their energy difference $\omega = E_\mu - E_\lambda$ as well as their off-diagonal matrix elements. By rescaling the matrix elements until the results for different system sizes collapse allows us to determine the scaling with system size. For $T = 0$ we observe that the matrix elements scale with a fractional power in L for all values of c giving

$$|\langle 0|\rho|\mu\rangle|^2 \sim \frac{g(\mu, \lambda)}{L^{\Delta(c)}} \tag{5.1}$$

as shown in Fig. 5.1a, b, c. At zero temperature, the scaling of matrix elements have previously been investigated by comparing Luttinger liquid theory to the explicitly known form factor expressions [92]. However, these results are only valid for excitations around the Fermi sea, which is no longer the case when we consider the entire single particle-hole sector with holes in the middle of the Fermi sea.

At finite temperature the scaling formula no longer involves a fractional power in system size. It however does feature an exponential factor giving

$$|\langle \lambda|\rho|\mu\rangle|^2 \sim \frac{h(\mu, \lambda)}{L^2} e^{-Lf(\mu, \lambda)}. \tag{5.2}$$

as shown in Fig. 5.1d, e, f. At first glance one might expect this exponential factor to be related to the entropy, but this is not the case. This indicates that the matrix element statistics of the single particle-hole sector deviates from the predictions we would obtain by applying the Eigenstate Thermalization Hypothesis [93]. Interestingly, the $\frac{1}{c}$ expansion can be used to accurately predict the value of the matrix elements in the one particle-hole sector away from the large interaction strength limit [94, 3].

Knowledge of the scaling of the off-diagonal matrix elements of the density operator can be used to gauge the relative importance of these states in, for example, the dynamical structure factor calculation. To see why, recall that the sumrule weight of an eigenstate $|\mu\rangle$ is given by

$$w_f(\mu) = \frac{L}{Nk^2}(E_\mu - E_\lambda)|\langle \lambda|\rho_k|\mu\rangle|^2 \tag{5.3}$$

For $T = 0$ this implies that the sum over the f sumrule contributions of the single particle-hole sector scales as $L^{-\Delta(c)}$ whereas at finite temperature it does
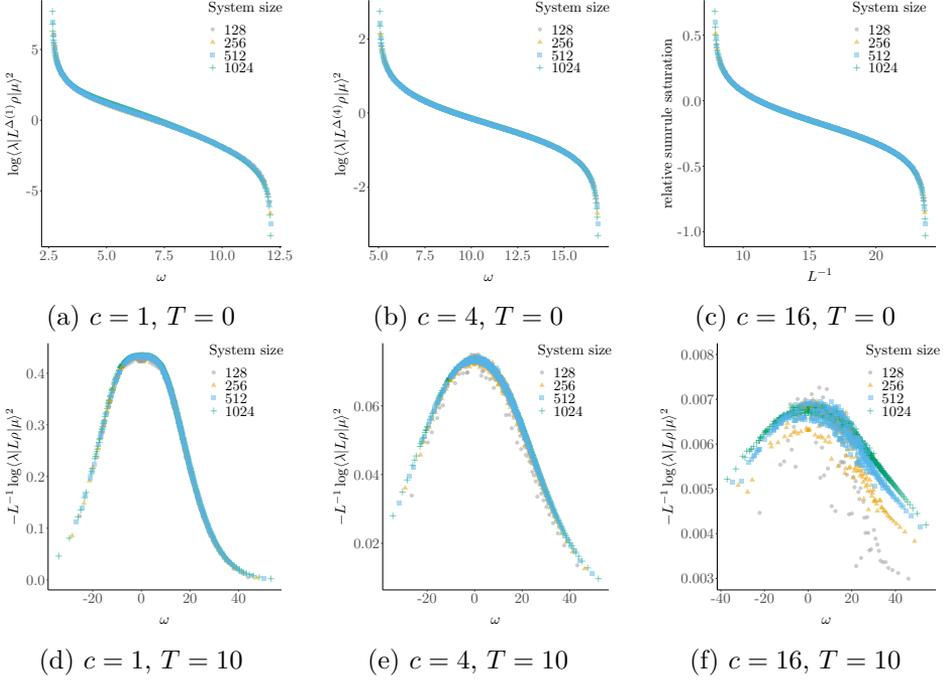
Figure 5.1: The rescaled off-diagonal density matrix elements versus the energy difference ω of the state compared to the ground state in (a)-(c) and the thermal state at $T = 10$ in (d)-(f). In both cases we consider kets at $k = \pi$ and unit density. At zero temperature, the values of delta we obtain are $\Delta(1) \approx 1.22$, $\Delta(4) \approx 1.03$, and $\Delta(16) = 1.005$.

(a) $c = 1$,   (b) $c = 4$,   (c) $c = 16$,
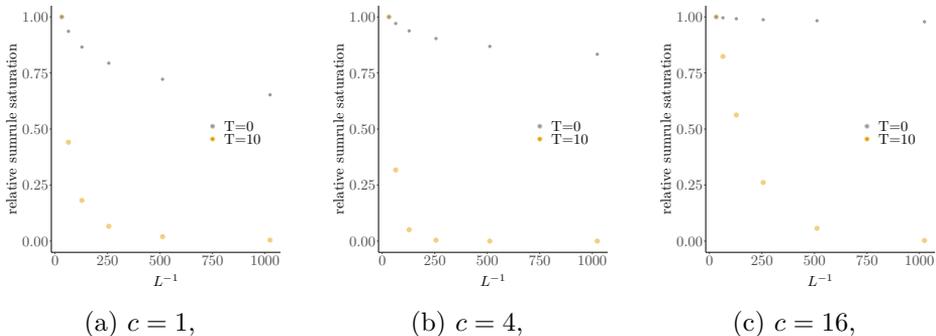
Figure 5.2: The relative importance of the 1ph sector. We see that in all cases the importance of the single particle-hole sector decreases more slowly with system size when compared to the finite temperature case. Furthermore, we see that the decrease in importance of the single particle-hole sector is slower at larger values of the interaction strength.

not imply a clear scaling of the sum due to the fact that the exponential factors in that case are state-dependent. It is however clear from the exponential suppression that at finite temperature the importance of the single particle-hole sector vanishes much more rapidly as can be seen in Fig. 5.2. One way of understand this difference is that at finite temperature there are many more multi particle-hole states where all but one particle-hole pair are soft modes. Such states carry similar weights to the pure single particle-hole states, diminishing their relative importance. Of course the importance of the single particle-hole sector in absolute terms at fixed system size depends also on the interaction strength, where its importance increases with the interaction strength.

5.2   The two particle-hole sector

In constrast to the sector of states with a single particle-hole pair with respect to the representative state, the two particle-hole sector contains an infinite number of eigenstates, even at fixed momentum. One way of eliminating this problem is by introducing an energy cutoff in addition to the momentum restriction. However, even with a cutoff in the energy, the number of states grows very rapidly with system size making it difficult to generate the entire two particle-hole sector for large system sizes as is necessary to determine the scaling relations. An alternative approach is to sample the two particle-hole sector with a cutoff rather than generate it exhaustively. This is the approach we take in this section.

We impose a cutoff on the two particle-hole sector by restricting the max-

imal and minimal values a rapidity can take. Note that this implicitly also introduces an energy cutoff, although we are still unable to generate all states up to the cutoff for all but the smallest system sizes. Therefore we choose to sample this space of states instead. Let $N = L$ be the system size under consderation and let $\{I_i^r\}$ be the quantum numbers of the state with respect to which we want to consider the two particle-hole sector. Then the steps of the sampling procedure used to generate states with two particle-hole pairs at some momentum $\Delta k = \frac{2\pi \Delta I}{L} \neq 0$ with rapidities in $[\lambda_{\min}, \lambda_{\max}]$ are as follows:

1. Determine which two quantum numbers will be changed by picking two random integers $k, l \in \{1, \ldots, N\}$ such that $k \neq l$.

2. Estimate the bounds on the quantum numbers for the randomly chosen quantum numbers such that the rapidity bounds are respected: Replace λ_l and λ_k in $\{\lambda_i^r\}_{1 \leq i \leq N}$ by the maximal and minimal rapidity. We replace the larger of the two by $\lambda_{\max}$ and the smaller rapidity by $\lambda_{\min}$. Plugging the resulting rapidities $\{\lambda_i\}_{1 \leq i \leq N}$ into the Bethe equations for $\lambda_{\min}$ and $\lambda_{\max}$ and solving for the corresponding quantum numbers gives us an estimate of the quantum number bounds $I_{\min}, I_{\max}$ corresponding to this choice of the cut-off.

3. Change the randomly selected quantum numbers by adding to I_l a random non-zero integer from $[I_{\min} - I_l, I_{\max} - I_l] \backslash \{0\}$. Since we target a fixed momentum sector, this also determines the change in I_k. If the resulting state does not satisfy the rapidity bounds, it is discarded.

The procedure outlined above generates a single eigenstate with two particle-hole pairs compared to the initial state $\{I_i^r\}_{1 \leq i \leq N}$.

Let us now consider the results we obtain when we use this procedure to generate $10,000$[1] states with respect to the ground state at unit density for different system sizes and interaction strengths. It is clear from Fig. 5.3a, b, c that the matrix element distributions collapse if we multiply the matrix elements by the square of the system size. This implies that the scaling of the matrix element satisfies

$$|\langle \lambda | \rho | \mu \rangle|^2 \sim \frac{1}{L^4} \tag{5.4}$$

where $|\mu\rangle$ and $|\lambda\rangle$ differ by two particle-hole pairs. Thus we no longer observe the fractional scaling we encountered in the zero temperature case for the single particle-hole sector.

Whereas the scaling of the matrix elements of the two particle-hole sector are the same regardless of the interaction strength at zero temperature, this

---

[1]We have studied the number of states required to obtain consistent results by generating data for multiple sampling sizes and comparing the results. We are confident that the conclusions we draw throughout this chapter are not effected by sampling sizes.

(a) $c = 1$, $T = 0$ (b) $c = 4$, $T = 0$ (c) $c = 16$, $T = 0$

(d) $c = 1$, $T = 10$ (e) $c = 4$, $T = 10$ (f) $c = 16$, $T = 10$
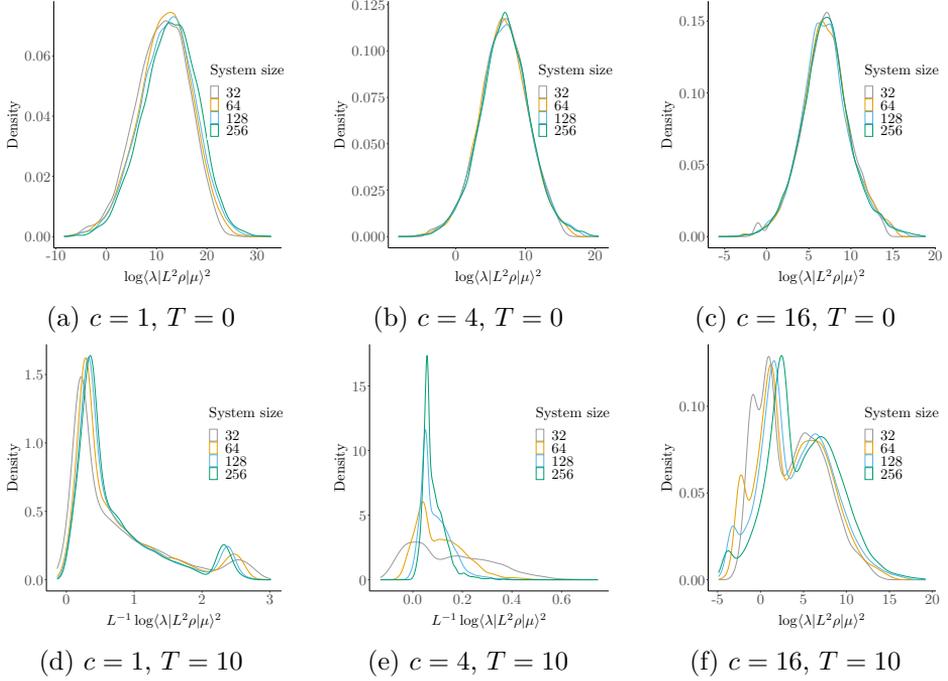
Figure 5.3: Distributions of the rescaled off-diagonal density matrix elements of the $10,000$ states $|\mu\rangle$ obtained for different system sizes via sampling of the two particle-hole sector with respect to the ground state in (a)-(c) and the thermal state at $T = 10$ in (d)-(f) both represented by $|\lambda\rangle$. In either case we consider a momentum difference of between ket and bra of $k = \pi$ and all states have unit density.
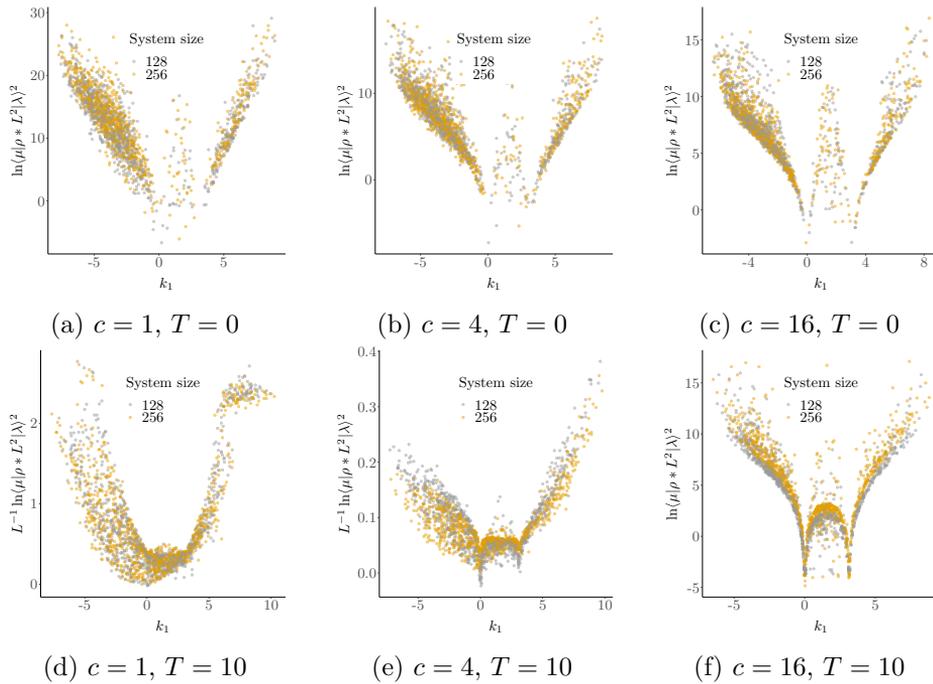
Figure 5.4: The rescaled size of the matrix element versus the momentum jump of the first particle of the corresponding state. The data displayed here is a random sampling of 10% of the data shown in Fig. 5.3.

is not the case for the finite entropy states. For example, in 5.3d we see that at $c = 1$ the distributions collapse after multiplying the matrix element by L^2 taking the logarithm of the resulting square and dividing by the system length. The corresponding scaling relation for the matrix elements is therefore

$$|\langle\lambda|\rho|\mu\rangle|^2 \sim \frac{h(\mu,\lambda)}{L^4}e^{-Lf(\mu,\lambda)}. \tag{5.5}$$

much like the one observed in Eq. (5.2). However, when we consider the distributions at $c = 16$ in Fig. 5.3f the data nearly collapses when performing the same rescaling as we do in the zero temperature case. At $c = 4$ neither scaling relation yields a collapse indicating that there is a cross-over between the weakly and strongly interacting regime.

In order to understand these differences in scaling relations, let us consider how the sectors with different numbers of particle-hole pairs are connected. We know that in all cases the f-sumrule is satisfied as described in Eq. (3.2) and Eq. (3.3), which is independent of the interaction strength. Since the number of states in a sector grows more strongly the more particle-hole pairs are allowed, the matrix elements of states with more particle-hole pairs have to be suppressed more strongly in order to preserve the sumrule. Furthermore, there is an interplay between the size of the matrix elements and the energy since both contribute to the f-sumrule, which has to remain unchanged. The fact that we see a stronger suppression of the matrix elements in the weakly interacting regime is consistent with the fact that in this case the density of states at a given energy is higher.

Considering the distributions of matrix elements can be convenient for determining the scaling relations, but from this we cannot deduce which properties of a state matter for the size of the matrix element. Note that there are two properties which distinguish the different two particle-hole states, namely which quantum numbers we turn into particle-hole pairs and how we distribute the momentum jumps amongst these particle-hole pairs. In Fig. 5.4 we plot the momentum jump of the first particle versus the size of the rescaled matrix element for 10% of the points we obtained via sampling. We see that generally as one of the two jumps is small, i.e. around $k_1 = 0$ and $k_1 = \pi$, the matrix elements are smallest. This corresponds to the so-called soft modes, where one of the particles moves only a little. Another observation we can make based on Fig. 5.4 is that for a fixed value of k_1 the spread of matrix element values is larger for smaller values of the interaction strength. This indicates that as the interaction strength decreases, which quantum numbers are moved has a bigger influence on the resulting matrix element.

What the implications of the scalings determined in this section are to the importance of the two particle-hole sector is more difficult to determine than for the single particle-hole sector. After all, the number of states in the single

particle-hole sector is finite and scales with N whereas here we are considering an infinite sector. Furthermore, the number of soft modes increases with system size whilst their matrix elements decrease. It is therefore inconclusive from these results if the sumrule contribution of the two particle-hole sector increases or decreases as system size increases. Furthermore, this may be state-dependent since there are more soft modes possible for states with a higher entropy.

5.3 Matrix elements between different macrostates

Thus far we have considered the matrix element between a reference state and the corresponding one and two particle-hole sectors. Despite their differences, such states still correspond to the same macrostate in the thermodynamic limit. In this section we consider the scaling of matrix elements when we consider microstates corresponding to different macrostates instead. We do so by considering the representative states at different temperatures. In particular, we consider the thermal states at $T = 0, 5$ and $T = 10$ at unit density for both $c \in \{1, 4, 16\}$ at different system sizes and compare the off-diagonal matrix elements of the g_2 operator[2] as shown in Fig. 5.5. We consider the g_2 operator in this section rather than the density operator because we will be comparing the representative states at different temperatures, which all have a vanishing momentum. This is due to the fact that the root distributions are symmetric, i.e. a root distribution $\rho(x)$ satisfies $\rho(x) = \rho(-x)$.

In Sec. 2.4.2 we argued that as states become increasingly different (in terms of their quantum numbers), their off-diagonal matrix elements of physically important local operators such as the density operator and the g_2 operator become smaller. Here we are looking at an extreme case of dissimilarity, the states do not even correspond to the same macrostate. Therefore we expect the matrix elements to vanish rapidly, but the question remains as to how these matrix elements approach zero in the thermodynamic limit. If these matrix elements would follow the Eigenstate Thermalization Hypothesis [93], they would be of the form

$$\langle \lambda | g_2 | \mu \rangle = f(\omega, \overline{E}) e^{-S(\overline{E})/2} R_{ij} \tag{5.6}$$

where $\omega = E_\mu - E_\lambda$ and $\overline{E} = (E_\mu + E_\lambda)^{\frac{1}{2}}$ and R_{ij} is a real or complex random number with mean zero and variance one. However, from Fig. 5.5 we see that in actuality the dominant scaling suppresses the matrix elements much more

---

[2]Due to the rapidly vanishing nature of the matrix elements of the between microstates corresponding to different macrostates arbitrary precision numerics is quickly required as the number of particles increases.

(a) $T = 0$ and $T = 5$ (b) $T = 0$ and $T = 10$ (c) $T = 5$ and $T = 10$
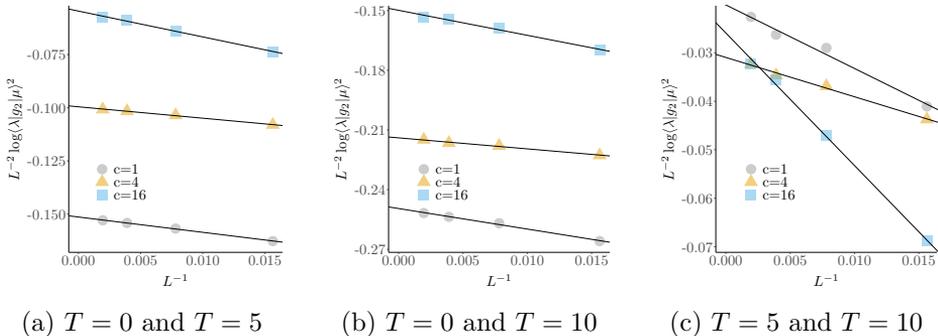
Figure 5.5: Scaling of the off-diagonal matrix elements between the finite size representative states at for $T \in \{0, 5, 10\}$ and $N = L \in \{64, 128, 256, 512\}$.

strongly giving

$$|\langle \lambda | g_2 | \mu \rangle|^2 \sim e^{-L^2 g(\lambda, \mu)}. \tag{5.7}$$

Rather than just being a numerical observation, the same scaling of the off-diagonal matrix elements can also be derived analytically for the off-diagonal matrix elements of the field operator in the $c \to \infty$ limit [3].

5.4   Sampling macrosates

In Sec. 2.5 we explained that the relation between eigenstates at finite size and root distributions is one-to-many. Until now we have used the representative state to play the role of the thermal state. In this section we investigate the surprisingly subtle problem of sampling eigenstates corresponding to a given macrostate.

5.4.1   Microcanonical Sampling

Since the thermal state represents the maximal entropy state at a given energy density, one way of sampling microstates corresponding to this macrostate is by randomly generating states at the same energy density. Provided that we do so in an unbiased way, the proportion of microstates corresponding to the thermal state approaches one as we increase the system size. In this section we consider a sampling of states at a given energy density by imposing a cutoff in quantum number space and sampling every allowed state with equal probability.

More precisely, the procedure required to generate a N particle state we consider can be formulated as follows:

1. Choose maximal and minimal quantum numbers $I_{\max}$ and $I_{\min}$.

2. Construct the set of allowed quantum numbers $\{I_1, \ldots I_M\}$ between $I_{\min}$ and $I_{\max}$ and randomly shuffle this set using for example the Fischer-Yates shuffle algorithm.

3. Pick the N first entries of the quantum numbers after the shuffle as the quantum numbers of the state.

This procedure uniformly samples the space of states respecting the quantum number bounds which we will refer to as the microcanonical sampling procedure. In order to ensure that we sample the thermal state we are interested in, we compute the energy of each state we generate in this way. If the energy of the state is within some designated window around the energy density of the thermal state we keep the state whereas we discard it otherwise. From now on we take the energy window to be $\omega \in [-1, 1]$, where ω is the difference in energy of a microstate with respect to the reference state.

In this procedure an important role is played by the cutoff chosen for the quantum numbers since this determines which part of Hilbert space we are sampling. When considering results produced using this procedure we therefore have to ensure that our results are independent of the cutoff chosen. This is the case when the overwhelming majority of eigenstates we are interested in are included in the portion of Hilbert space being sampled by our algorithm. Unfortunately, the number of states grows rapidly with an increasing cutoff making it difficult to consider big variations in the cutoff.

First we consider the density function of $100,000$ quantum numbers generated using the microcanonical sampling for a window of size 1 around the energy of the representative states of $T = 10$ thermal states at unit density for $c = 1, 4, 16$. The results of this sampling are show in 5.6a, b, c for two different choices of the cutoff. The dataset labelled $\Delta I_{\max} = 0$ refers to a quantum number cutoff equal to the largest and smallest quantum numbers of the representative states whereas the second dataset increases $I_{\max}$ by 5 and decreases the lower bound by the same amount. We can see how well the resulting normalized distributions coincide to the thermodynamic result shown in black. The data generated using the smaller cutoff, is lower in the middle and falls off more slowly until the cutoff whereas for the larger cutoff the agreement for most quantum numbers is approximately good up to the cutoff[3]. Surprisingly, the momentum distribution of the states sampled agree well despite their differences on the level of the root density as shown in Fig. 5.6d, e, f. The differences become more pronounced when considering the distributions of the third moment Q_3 as shown in Fig. 5.6g, h, i. This is to be expected as it is much more sensitive to large rapidities, where these samplings are bound to differ.

---

[3]The larger cutoff dataset presented here is the best we can achieve with the computational means available to us.
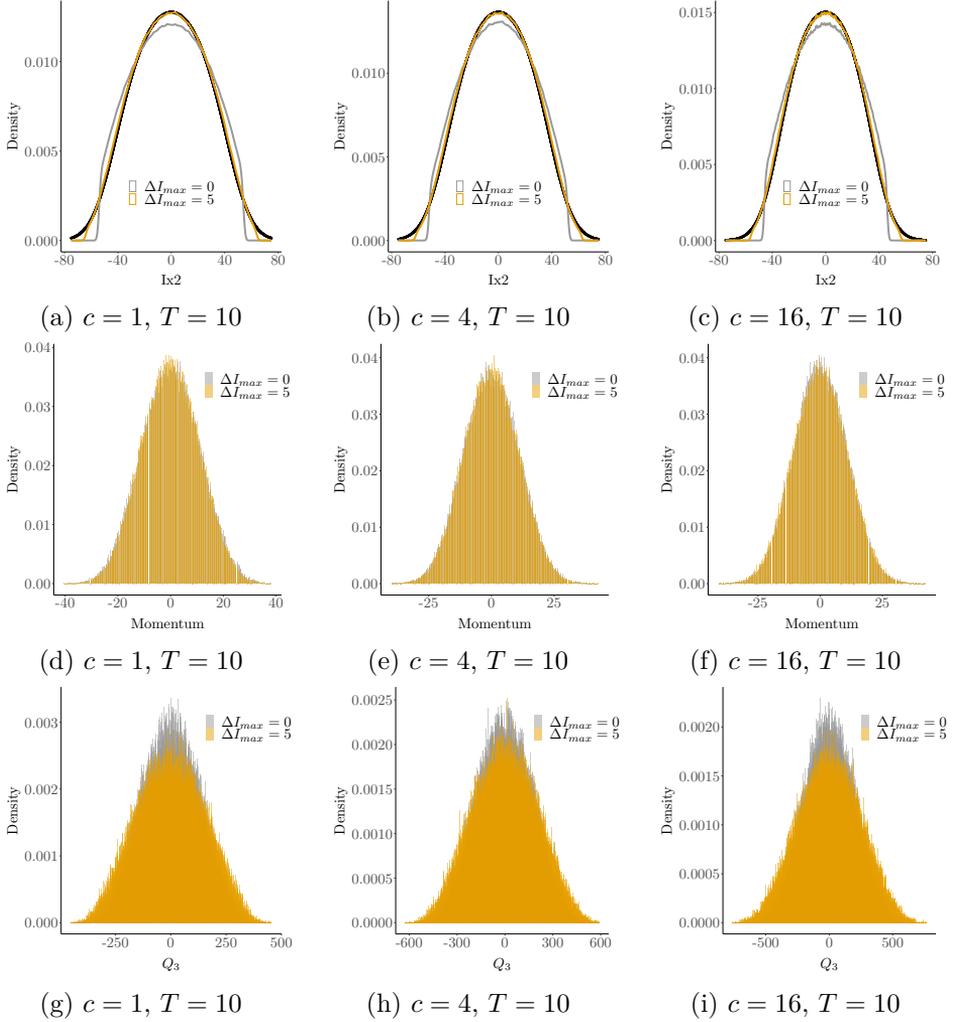
Figure 5.6: A comparison of the results obtained by sampling $100,000$ states with $\omega \in [-1, 1]$ using the microconical sampling procedure for two different values of the quantum number cutoff at $T = 10$ and unit density. For the data labelled by $\Delta I_{max} = 0$, the cutoff is equal to the maximal and minimal quantum number of the representative state, whereas for $\Delta I_{max} = 5$ they are increased/decreased by 5. In (a)-(c) we compare the root distribution in the thermodynamic limit with normalized histograms of the sampled quantum numbers. We also consider the resulting normalized momentum distributions in (d)-(f) as well as the normalized distributions of the third moment in (g)-(i).

Although this method is easy to formulate, implement, and known to generate unbiased samples of the macrostate[4], it also has significant downsides. The most important limitation is that the computational cost of running the algorithm scales exponentially with system size. This makes it impossible to use it to properly consider scaling with system size or generate results not troubled by finite size effects. Furthermore, the cutoff also artificially restricts the portion of Hilbert space we are sampling and increasing it is exponentially hard. Finally, it can only be used to sample the entropically most likely states, i.e. the thermal states rather than general root distributions.

5.4.2 Box Sampling

The most commonly used method of sampling a macrostate is based on the derivation of the entropy in Sec. 2.5. The idea is to divide the allowed quantum numbers at finite size into boxes and to compute the number of quantum numbers a given box has to hold from the root distribution. Then shuffling the quantum numbers around in their respective boxes does not change the box occupation number and therefore generates microstates corresponding to the macrostate we used to determine the box occupation numbers. The problem with this approach, commonly called box sampling, at the sizes we are able to access is that the results depend strongly on our choice of boxes.

Practically, the box sampling procedure we employ works as follows. Let $N = L$ be the system size under consideration and let $\{I_i^r\}$ be the quantum numbers of the representative state. Then the steps of the box sampling procedure we employ are as follows:

1. Choose the number of boxes N_b.

2. Compute the minimal box size Δ such that N_b boxes can fit all the quantum numbers of the representative state, i.e. choose the minimal integer Δ such that $I_1^r + \Delta * N_b \geq I_N^r$.

3. Center the boxes w.r.t. the state, i.e. choose the starting point of the first box I_s such that $|I_s - I_1^r|^2 + |I_s + N_b * \Delta - I_N^r|^2$ is minimal.

4. Determine the box occupation numbers, i.e. the number of quantum numbers of the representative in a given box $\{O_i\}_{i \leq N_b}$, which ranges from $I_s + \Delta * (i - 1)$ to $I_s + \Delta * i$.

5. Generate a set of quantum numbers $\{I_i\}$ by randomly choosing O_i unique quantum numbers for each box.

---

[4]Provided that we consider large enough system sizes.

The procedure outlined above generates a single microstate characterised by quantum numbers $\{I_i\}$.

An illustration of the results from such a sampling procedure are shown in Fig. 5.7. Here the aforementioned dependence on the number of boxes that one chooses is visible. Furthermore, Fig. 5.7d - i show that the resulting momentum and third moment distributions are much more narrow compared to the results from microcanonical sampling. The latter can be understood from the fact that the portion of Hilbert space sampled by the box sampling routine is smaller than that sampled by our microcanonical sampling routine. One can expect this discrepancy between the two methods of sampling to diminish as we increase both the system size as well as the number of boxes. This procedure is regularly followed to obtain the thermodynamic limit by letting the system size go to infinity first and then the number of boxes. However, it turns out that, despite its appeal, in this approach the effects due to the finite number of boxes are not well-controlled. After an extensive study of these finite size effects, where we considered systems with up to hundreds of particles, we have therefore decided to look for alternative sampling methods not troubled by such problems.

5.4.3 RANDOM WEIGHTED SAMPLING

Having observed the problem with the complexity of the microcanonical sampling in Sec. 5.4.1 and the problem with the ambiguities of box sampling in Sec. 5.4.2, the question arises if we can sample the root distribution in a less ambiguous yet efficient way. One approach that comes to mind is to consider all possible integers with a very large cutoff and associate to them probabilities based on the value of the root distribution at this point forming a discrete probability distribution. In order to generate a state from this distribution we have to pick N distinct integers, which can be achieved by removing a quantum number from the discrete probability distribution after it has been picked and rescaling the remaining probabilities.

This gives rise to the following procedure. Let $N = L$ be the system size under consideration and let $\{I_i^r\}$ be the quantum numbers of the representative state. Then the steps of the box sampling procedure which we call random weighted sampling (RWS) we employ are as follows:

1. Choose a cutoff $I_{\max} = -I_{\min}$ and generate the set of allowed integers between the maximal and minimal values $\{I_1, \ldots, I_M\}$.

2. Associate a weight $\rho(I_1)$ to each of the integers and normalize by the sum to obtain a discrete probability distribution.

3. Pick a quantum number from the set of allowed quantum numbers weighted by their probabilities and afterwards remove it from the set
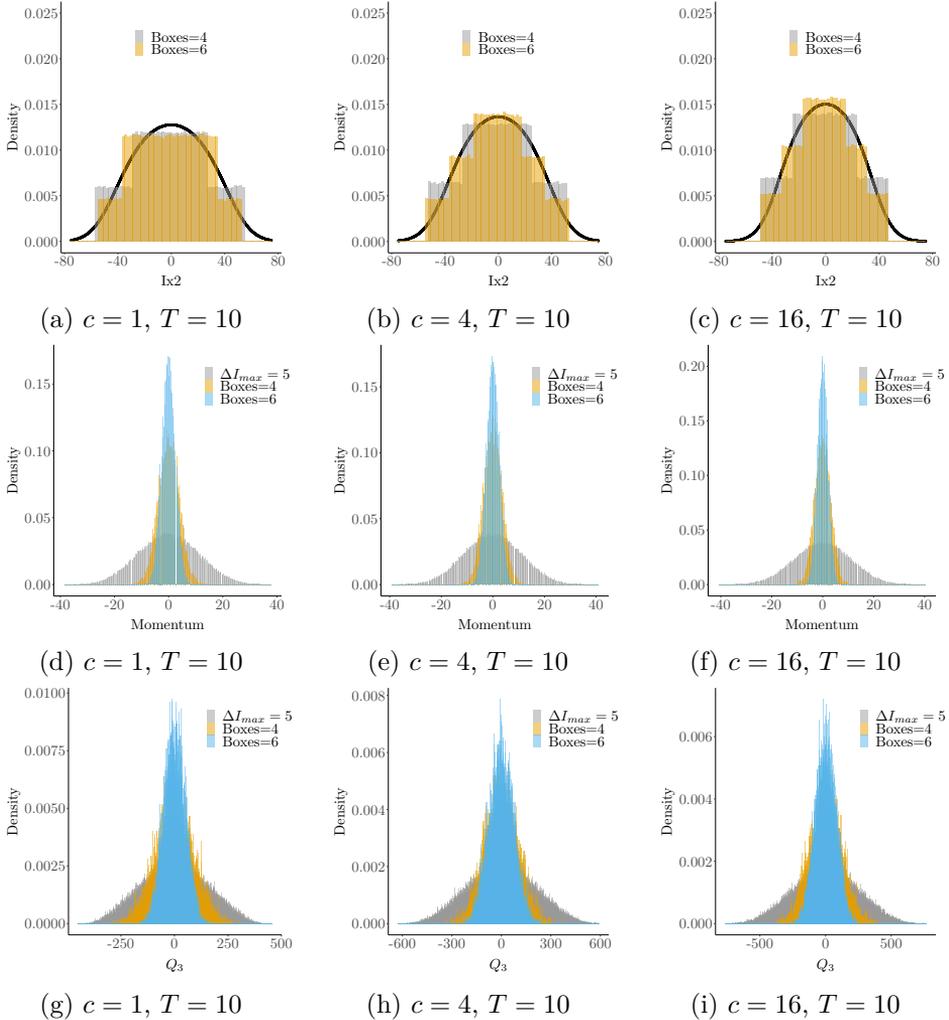
Figure 5.7: A comparison of the results obtained by sampling $50,000$ states with using the box sampling procedure where the number of boxes is equal to 4 and 6 at $T = 10$ and unit density. In (a)-(c) we compare the root distribution in the thermodynamic limit with normalized histograms of the sampled quantum numbers. We also consider the resulting normalized momentum distributions in (d)-(f) as well as the normalized distributions of the third moment in (g)-(i). In the latter case we also compare to the results obtained using the microcanical sampling procedure previously shown in Fig. 5.6.

and rescale to retain normalization of the discrete probability distribution.

4. Repeat the previous step until N integers have been chosen.

If the energy density of the resulting state is within some designated window around the energy density of the thermal state we keep the state whereas we discard it otherwise. We note that even though there is a cutoff in quantum number space due to the infinite dimensionality of the Hilbert space, it does not carry the same implications as the cutoff in the microcanonical sampling procedure. For example, it can be taken such that $\rho(I_{\max})$ is tiny so that we can capture the root density with a minimal computational burden. Also, we can again only keep states in some designated energy window to allow for a better comparison with the other methods we considered thus far. One of the downsides of this algorithm is that it requires us to construct discrete probability distributions at each step. This numerical overhead can be avoided by implementing a more efficient yet equivalent method as discussed in [95].

Despite its appeal, there is a subtlety associated to sampling in this way. Namely, the resulting integers are not distributed according to the root distribution. This is caused by the rescaling step after we have removed an integer from the discrete probability distribution. To see this in detail, let us consider a simple example where we have three possible integers $\{I_1, I_2, I_3\}$ which we sample with probabilities $\{p_1, p_2, p_3\}$. Sampling two states from this, we can obtain six different results with the following probabilities given by

$$(I_1, I_2) \text{ with probability } p_1 \tfrac{p_2}{1-p_1} \qquad (I_1, I_3) \text{ with probability } p_1 \tfrac{p_3}{1-p_1} \qquad (5.8)$$

$$(I_2, I_1) \text{ with probability } p_2 \tfrac{p_1}{1-p_2} \qquad (I_2, I_3) \text{ with probability } p_2 \tfrac{p_3}{1-p_2} \qquad (5.9)$$

$$(I_3, I_1) \text{ with probability } p_3 \tfrac{p_1}{1-p_3} \qquad (I_3, I_2) \text{ with probability } p_3 \tfrac{p_2}{1-p_3} \qquad (5.10)$$

From this, we can calculate the probability of drawing I_1 is equal to

$$p_1 \left(\tfrac{p_2+p_3}{1-p_1} + \tfrac{1}{1-p_2} + \tfrac{1}{1-p_3} \right) = p_1 \left(1 + \tfrac{p_2}{1-p_2} + \tfrac{p_3}{1-p_3} \right) \qquad (5.11)$$

rather than p_1.

One of the ways to make this algorithm work is by determining the initial discrete probability distribution not based on the root distribution, but instead sample a different distribution $\tilde{P}$ specifically chosen such that sampling it gives us quantum numbers distributed according to the root distribution. This can be done numerically by starting from a random $\tilde{P}$, sampling it a bunch of times and turn the results into a discrete probability distribution. Where the resulting discrete probability distribution differs from the discrete probability distribution describing the root density add the difference to $\tilde{P}$ taking care

not to accidentally create negative probabilities. The resulting probability distribution $\tilde{P}$, which we call the prior, is then what we use in the second step of the random weighted sampling algorithm to determine the discrete probability distribution.

In Fig. 5.8 we compare the results obtained using random weighted sampling with an appropriately determined prior to the root distribution and the results from microcanonical sampling. We see in Fig. 5.8a b c that the random weighted sampling algorithm is able to capture well the root distributions apart from some minor deviations at the outermost tails. However, when considering the momentum and Q_3 distributions in Fig. 5.8d, e, f, g, h, i we see that the difference between this sampling approach and the microcanonical sampling start to differ as the interaction strength is decreased. Before drawing a definite conclusion about the cause of these differences, let us consider another alternative sampling method.

5.4.4 DIRECT RANDOM SAMPLING

The random weighted sampling algorithm introduced in the previous section seemed capable of sampling the root distribution well, but did not fully agree with the microcanonical sampling for the observables we considered at smaller values of the interaction strength. In this section we propose another way of sampling the root distribution and compare the results to the random weighted sampling algorithm. Rather than determining a prior discrete probability distribution and sampling it without replacement, we instead do the following.

1. Choose a cutoff $\mathrm{I}_{\max} = -\mathrm{I}_{\min}$ and generate the set of allowed integers between the maximal and minimal values $\{I_1, \ldots, I_M\}$.

2. For every $I_l \in \{I_1, \ldots, I_M\}$ draw a random number u_l from $(0, 1)$ and compute their weight $w(I_l) = \rho(I_l) - u_l$.

3. Choose the N quantum numbers for which the weight is largest to construct the resulting state.

Again we keep only states with energy densities within some designated window and discard them otherwise. This procedure we call direct random sampling (DRS).

In Fig. 5.9 we see that the direct random sampling approach, like the random weighted sampling routine, captures the root density correctly up to some minor mismatch in the outermost tails. Also, the momentum and Q_3 distributions of both approaches agree. However, this also means that neither method agrees with the microcanonical sampling at $c = 1$ and $c = 4$ for the parameters considered. In particular, the root distributions generated by DRS and RWS do a better job at capturing the tails of the thermodynamic root distribution,
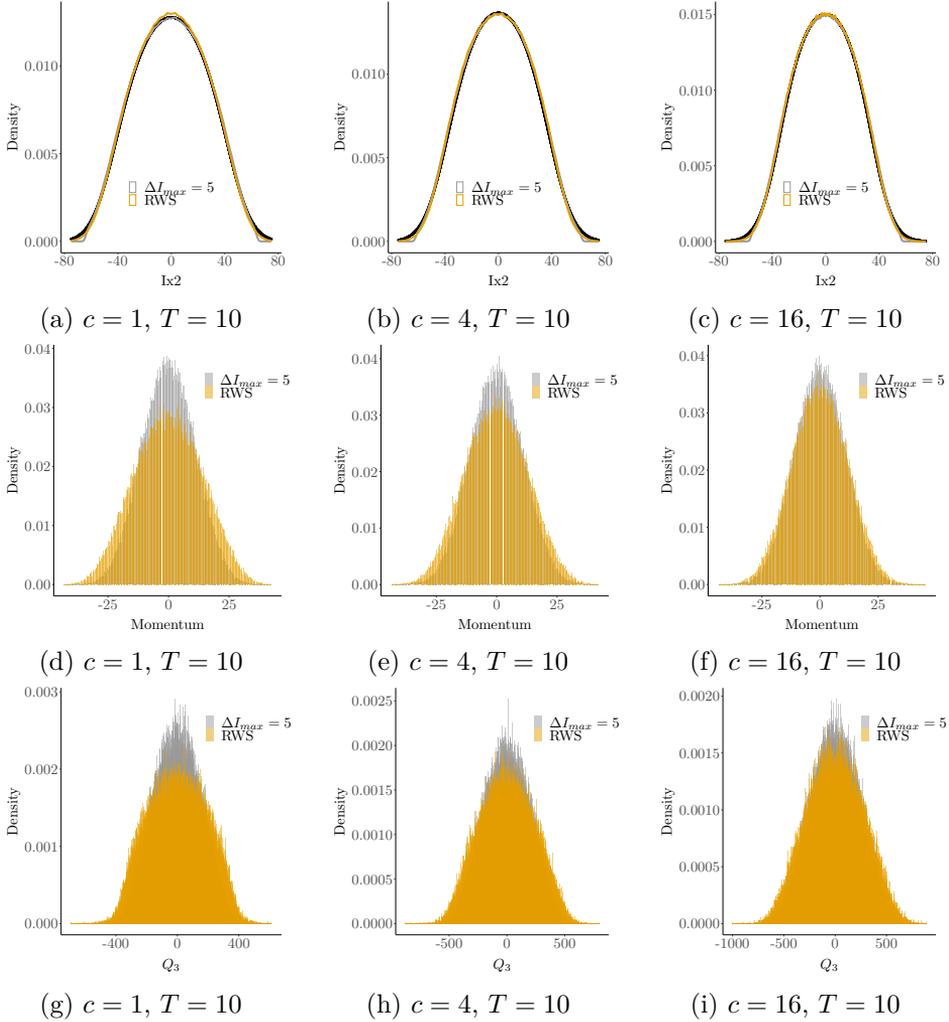
Figure 5.8: A comparison of the results obtained by sampling $100,000$ states with $\omega \in [-1, 1]$ using the microconical sampling procedure and the random weighted sampling routine at $T = 10$ and unit density. In (a)-(c) we compare the root distribution in the thermodynamic limit with normalized histograms of the sampled quantum numbers. We also consider the resulting normalized momentum distributions in (d)-(f) as well as the normalized distributions of the third moment in (g)-(i).

(a) $c = 1$, $T = 10$ (b) $c = 4$, $T = 10$ (c) $c = 16$, $T = 10$

(d) $c = 1$, $T = 10$ (e) $c = 4$, $T = 10$ (f) $c = 16$, $T = 10$

(g) $c = 1$, $T = 10$ (h) $c = 4$, $T = 10$ (i) $c = 16$, $T = 10$
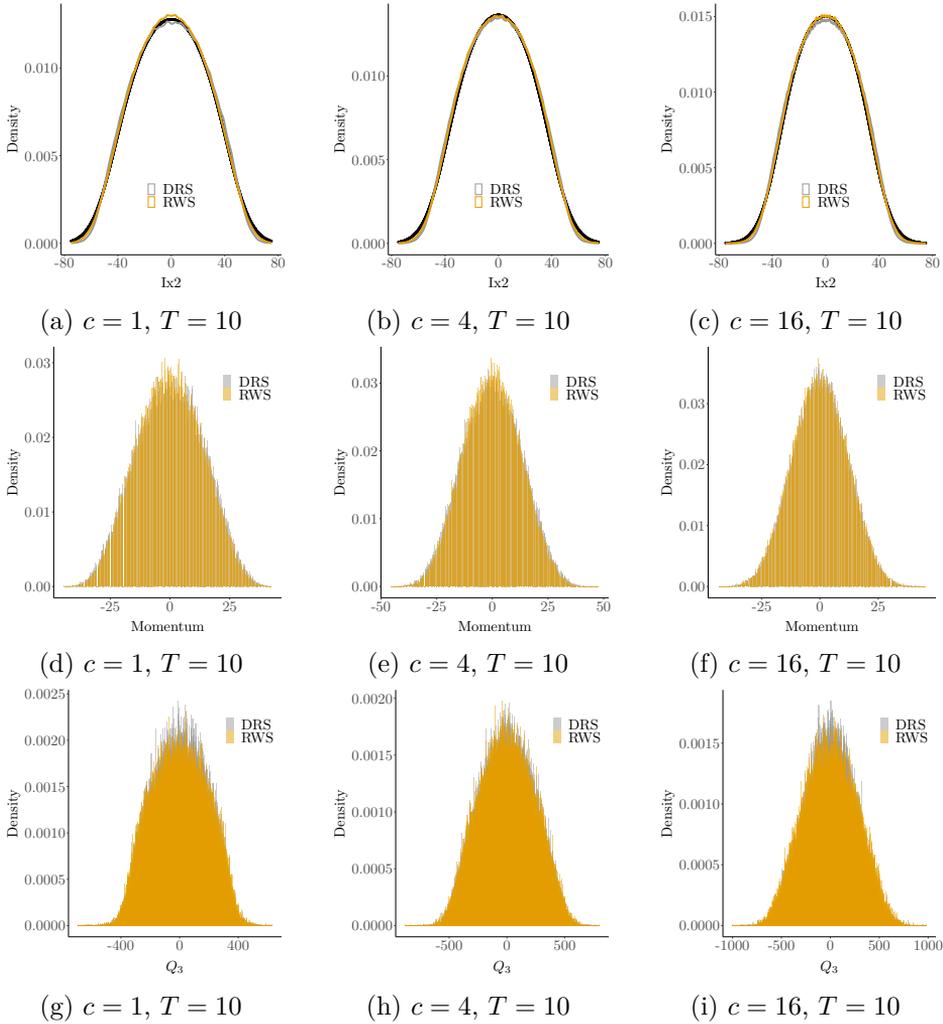
Figure 5.9: A comparison of the results obtained by sampling $100,000$ states with $\omega \in [-1, 1]$ using the direct random sampling procedure and the random weighted sampling routine at $T = 10$ and unit density. In (a)-(c) we compare the root distribution in the thermodynamic limit with normalized histograms of the sampled quantum numbers. We also consider the resulting normalized momentum distributions in (d)-(f) as well as the normalized distributions of the third moment in (g)-(i).

Also, they generate broader momentum and Q_3 distributions compared to the microcanonical sampling approach. If it were not for the limitations imposed by the cutoff, the microcanonical would be the most unrestrictive form of sampling and would therefore yield the widest momentum and Q_3 distributions. The fact that RWS and DRS generate broader distributions at $c = 1$ and $c = 4$ therefore indicates that the mismatch is due to the cutoff that is necessarily imposed when implementing the microcanonical sampling routine.

Given these considerations, we conclude that DRS and RWS are capable of generating microstates corresponding to a given root distribution in an unambiguous way. Furthermore, unlike the microcanonical sampling approach, they are able of correctly capturing the tails of the distribution and they are much more computationally efficient, allowing us to use them in order to sample states with hundreds of particles. The data for all of the figures in the remainder of this chapter have been generated using the random weighted sampling routine.

5.5 Matrix element statistics

Having developed an efficient and unbiased method of sampling the microstates of a given macrostate in the previous sections, we use this section to illustrate how these methods can be used to study the statistics of matrix elements.

5.5.1 Diagonal matrix elements

Let us start by considering the diagonal matrix elements of the g_2 operator. Note that we do not consider the diagonal matrix elements of the density operator as they are all equal to the particle density. The different microstates corresponding to a given macrostate should agree on the diagonal matrix element in the thermodynamic limit if the latter is to make sense. We verify this by sampling microstates from the thermal state at $T = 10$ and unit density and computing the corresponding diagonal matrix elements as shown in Fig. Fig. 5.10. As the system size increases, the variance of the sampled matrix elements decreases according to

$$\sigma \langle \lambda | g_2 | \lambda \rangle \propto aL^{-\frac{1}{2}} + bL^{-1} + \dots \tag{5.12}$$

as shown in Fig. 5.10b. Furthermore, means of the distributions of sampled diagonal matrix elements can be nicely extrapolated to their value in the thermodynamic limit as shown in Fig. 5.10a.
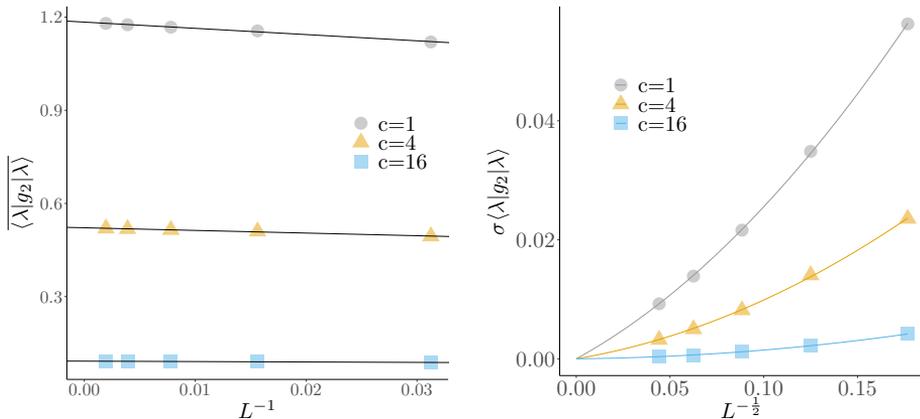
(a) The mean of the diagonal g_2 matrix element for the sampled data.

(b) The variance of the diagonal g_2 matrix element for the sampled data.

Figure 5.10: The means and variances of the diagonal g_2 matrix elements for states sampled according to the eandom weighted sampling algorithm from the thermal state at $T = 10$ at unit density and with $\omega \in [-1, 1]$. Each point is calculated using a sampling size of $50,000$ states.

5.5.2 OFF-DIAGONAL MATRIX ELEMENTS

Now that we have verified that the diagonal matrix elements behave as expected, the question arises if we can say anything about the off-diagonal matrix elements between typical states and the representative state. To investigate this, we consider typical states at a given momentum, which we sample by generating states using the random weighted sampling procedure. Having generated such microstates, we can compute the off-diagonal matrix elements with respect to the representative state of the macrostate.

Unlike in previous cases, we are considering matrix elements between states which differ by extensively many particle-hole pairs. This introduces some numerical dificulties due to the smallness of the resulting matrix elements and the complicated formulas required to compute them, see Sec. 2.4.2. In order to ensure that our numerical evaluation of these formulas is correct, we use arbitrary precision numerics to compute the matrix elements. Furthermore, we can verify our calculations by not only computing the matrix element but also its transpose. Although mathematically equivalent, the two problems are numerically quite distinct (we observe a mismatch between the matrix element and its transpose when insufficient precision is used to do the calculations). The fact that our routines produce agreeing results gives us confidence in the validity of the resulting matrix element distributions.

The matrix element distributions obtained by sampling are shown in Fig. 5.11. We see that after a rescaling and a shift, the matrix element distributions for the different system sizes collapse. This means that the matrix elements scale like

$$\langle \lambda|O|\mu\rangle^2 \propto e^{-a_O(c)L\log(L)-b_O(c)L} \tag{5.13}$$

where a_O and b_O are coefficients which depend on the system size as well as the operator under consideration. This collapse is more complete for the weakly interacting regime, whereas for larger values of the interacting strength we are still plagued by some finite size effects. Remarkably, a Fréchet distribution appears to describe the distributions well for all values of the interaction sts well for both operators and all values of the interaction strength. The fact that we can fit a Fréchet distribution in all these cases[5] raises the question if this fact is universal to a larger class of systems to which the Lieb-Liniger model belongs.

5.6 Conclusions

The central role played by the matrix elements in the calculation of correlation functions makes them a worthy subject of study. For integrable models the computation of correlation functions is generally dominated by a select set of atypical states with few particle-hole pairs. In this chapter we started by studying such atypical states in the form of the single and two particle-hole sectors in order to further quantify some of the commonly known intuitions regarding their importance.

For the single particle-hole sector we determined the scaling relations and found a distinct difference between the zero temperature and finite entropy case, where in the zero temperature case we observed a power law decay with system size in constrast to an exponential scaling at finite temperature. Moving on to the two particle-hole sector we saw a shift from different power law scalings depending on the interaction strength for the single particle-hole sector to a uniform power law decay with system size. Furthermore, for the finite temperature case we found a crossover between an exponential scaling of the matrix elements in the weakly interacting limit to a power law scaling in the strongly interacting regime.

Having considered the atypical states, we moved on to the matrix elements between representative states corresponding to different macrostates. One might expect the matrix elements between such typical states to have properties similar to typical state in non-integrable models. If that were the case, the matrix elements between typical states would be suppressed entropically

---

[5]A Fréchet distribution also excellently fits the distribution of the off-diagonal matrix elements of the field operator in the infinitely repulsive limit [[cite:fhl$_e ssler_b e_n odate$]].

(a) $c = 1$, $T = 10$ (b) $c = 4$, $T = 10$ (c) $c = 16$, $T = 10$

(d) $c = 1$, $T = 10$ (e) $c = 4$, $T = 10$ (f) $c = 16$, $T = 10$
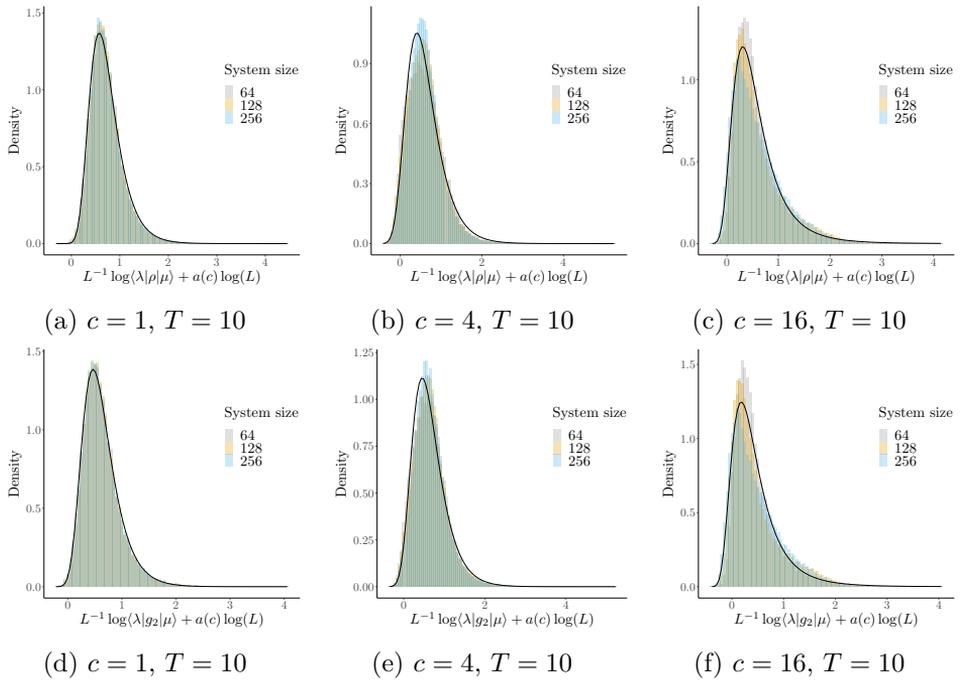
Figure 5.11: The rescaled off-diagonal matrix elements of the density operator in (a)-(c) and the g_2 operator in (d)-(f) between randomly generated microstates with $\omega \in [-1, 1]$ corresponding to the $T = 10$ thermal state at unit density and the its representative state.

according to the Eigenstate Thermalization Hypothesis. We discovered however, that the matrix elements between macrostates of the g_2 operator are suppressed much more strongly than entropically.

In order to study the typical states in more detail, the need arose for unbiased and efficient sampling algorithms generating microstates corresponding to a given macrostate. Starting with a simple microcanonical sampling, which requires the introduction of a cutoff in quantum number space, we discovered that it was numerically difficult to property capture the tails of the root distribution with this approach. This fact, in combination with the poor scaling of the numerical cost with system size, made us conclude that it is an unsuitable method to study the scaling relations of matrix elements.

Moving on, we considered the box sampling approach. Surprisingly, for this approach we encountered encountered a strong dependence on the box number, which persisted even for large system sizes. Furthermore, the portion of Hilbert space sampled by box sampling in the finite case is too restrictive and the method does not capture the root distribution well. As such, this method was also deemed unsuitable. To overcome these limitations we developed two independent approaches which we dubbed *random weighted sampling* and *direct random sampling*. These approaches reproduce the targeted root distributions well and are numerically efficient, making them suitable for studying the matrix element statistics of typical states.

Using these new sampling methods we were finally able to consider the matrix elements of typical states. We started by considering the diagonal matrix elements and verified that the standard deviations of the resulting distributions vanish in the thermodynamic limit. Following this, we considered the off-diagonal matrix elements of both the density and the g_2 operator between a fixed representative state and sampled typical states. We discovered that, for all operators and interaction strengths we considered, we can fit a Fréchet distribution to the distributions of off-diagonal matrix elements between the representative state and the typical microstates. This surprising result raises the question if this result is particular to the Lieb-Liniger model or simply represents an example from a larger class of systems sharing the same statistical properties.

Summary

Throughout this thesis we considered different types of problems, but all were centred around the same model, the Lieb-Liniger model. Therefore we began this thesis by introducing the Lieb-Liniger model, which describes a one-dimensional gas of Bosons, in Chapter 2. We started this chapter by reviewing the coordinate Bethe Ansatz, which yields the spectrum and a convenient description of eigenstates in terms of sets of quantum numbers. After a brief discussion of the Algebraic Bethe Ansatz, we summarized some of the major results obtained using this approach, namely the matrix element expressions for physically important operators such as the density operator. Finally, we considered the thermodynamic limit of the formalism whose derivation gives rise to a correspondence between eigenstates at finite size, called microstates, with root distributions in the thermodynamic limit, called macrostates.

In order to go from calculating individual microstates to correlation functions generally requires the numerical evaluation of summations over eigenstates in Hilbert space. Since there are infinitely many eigenstates in Hilbert space, such summations have to be truncated in order to evaluate them numerically. The question is if and how we can obtain accurate results in the presence of such a truncation. In an effort to obtain the best possible results, we devoted Chapter 3 to the study of Hilbert space exploration/scanning algorithms aimed at finding the eigenstates most important for a given calculation. Starting from a basic example of a scanning algorithm, we made stepwise improvements until we ended up with an approach which outperforms previously available algorithms for finite temperature calculations.

Having discussed Hilbert space exploration algorithms, we turn to the problem of computing the time evolution following a quench in the interaction strength. We rephrased the problem so that it is equivalent to diagonalizing a perturbed Hamiltonian in terms of a basis of Bethe states. In this formulation the problem can be approximately solved numerically by truncating the Hamiltonian and diagonalizing it. The resulting problem, of choosing which eigenstates to include in the calculation of the truncated Hamiltonian, bears clear similarity to the problem of choosing which eigenstates to prioritise in the calculation of a correlation function. By choosing an appropriate weighing

metric for eigenstates we can therefore use the same scanning algorithms to most accurately capture the time evolution following the quench for a given size of the truncated Hamiltonian.

Even with an optimally chosen basis, however, limitations associated to the maximal size of the truncated basis we can consider remain. These limitations can be overcome by introducing an iterative renormalization-group based scheme we call a *High Overlap State Truncation Scheme*. With this approach we are able to consider truncated bases of hundreds of thousands of eigenstates. However, applying this approach to a quench from the strongly interacting regime to a regime at intermediate interaction strength, reveals a discrepancy between the results obtained using the iterative scheme and the full diagonalization of the corresponding truncated Hamiltonian. We resolve this issue by modifying the iterative scheme resulting in what we call the *Matrix Element Renormalization Group* approach.

In Chapter 5 we switched gears and studied the more fundamental properties of matrix elements in the Lieb-Liniger model. We determined scaling relations for the one and two particle-hole sectors, which further quantified existing intuitions regarding their diminishing importance as we increase system size or consider finite entropy states. Moving on, we considered matrix elements between microstates corresponding to different macrostates, which revealed a surprising deviation from the type of scaling behaviour expected the Eigenstate Thermalization Hypothesis.

In order to study the statistics of matrix elements of typical states required a mechanism of sampling microstates corresponding to a given macrostate. We considered different existing methods and concluded that they are limited to small system sizes or generate biased data. To resolve this we developed two independent alternatives whose results agree beautifully. These algorithms allowed us to study the diagonal and off-diagonal matrix elements for system sizes up to hundreds of particles. In the off-diagonal case this resulted in distributions following a Fréchet distribution, which has not yet been encountered in the study of matrix element statistics.

Samenvatting

In dit proefschrift behandelen we verschillende vraagstukken waarbij de gemene deler het onderliggende model is. Daarom beginnen we in hoofdstuk 2 van dit proefschrift met het introduceren van dit model, het Lieb-Liniger model, dat een één-dimensionaal gas van bosonen beschrijft. Allereerst introduceren we de Bethe Ansatz die ons in staat stelt de eigenwaardes van het model te bepalen en tevens een beschrijving van eigentoestanden aan de hand van kwantumgetallen biedt. Na een korte herhaling van de Algebraïsche Bethe Ansatz behandelen we enkele van de belangrijkste resultaten die door middel van deze methode afgeleid zijn. Voor het Lieb-Liniger model zijn dit de matrixelementen van belangrijke fysische operatoren zoals bijvoorbeeld de dichtheidsoperator. Als laatste leiden we in dit hoofdstuk de thermodynamische limiet van het model af. Hieruit komt ook een correspondentie voort tussen eigentoestanden in het eindige geval en eigenfuncties in de thermodynamische limiet, die we vanaf nu macrotestanden noemen.

Om de stap te zetten van het berekenen van individuele eigentoestanden naar het berekenen van correlatiefuncties moeten we een (numerieke) sommatie doen over de eigentoestanden. Gezien er oneindig veel van zulke eigentoestanden bestaan is het doorgaans noodzakelijk om dergelijke sommaties te benaderen door middel van een eindige sommatie. De vraag is of en hoe we ondanks deze benadering accurate resultaten kunnen verkrijgen. In een poging om tot het best mogelijke resultaat te komen wijden we hoofdstuk 3 aan het ontwikkelen van algoritmes die de belangrijkste eigentoestanden voor een gegeven sommatie proberen te vinden. Beginnend met een zo simpel mogelijk algoritme brengen we stapsgewijs verbeteringen aan. Dit leidt uiteindelijk tot een algoritme dat beter presteert dan bestaande methodes voor gevallen waar de temperatuur eindig is.

Na deze discussie over zoekalgoritmes voor eigentoestanden wenden we ons tot het bepalen van de tijdsevolutie na een plotselinge verandering in de sterkte van de interacties tussen de deeltjes in het Lieb-Liniger model. Dit probleem herformuleren we zodat het opgelost kan worden door het diagonaliseren van een verstoorde Hamiltoniaan opgesteld in een basis van Bethe toestanden. Wederom is een benadering noodzakelijk om het probleem numeriek op te

kunnen lossen omdat ook de verstoorde Hamiltoniaan oneindig groot is. Dit doen we door een eindige basis van Bethe toestanden te kiezen en de resulterende afgekapte Hamiltoniaan te beschouwen in plaats van de originele. Om deze benadering zo goed mogelijk te laten zijn moeten we wederom de juiste Bethe toestanden weten te kiezen. Door een passende wegingsfunctie te kiezen kunnen we de algoritmes uit het vorige hoofdstuk gebruiken om een zo nauwkeurig mogelijke benadering van de tijdsevolutie te vinden voor een gegeven grootte van de afgekapte Hamiltoniaan.

Zelfs met een optimaal gekozen verzameling van Bethe toestanden blijft de grootte van de verstoorde Hamiltoniaan die we kunnen diagonaliseren een beperking, wat zijn weerslag heeft op de validiteit van de benadering. Deze beperking kan weggenomen worden door een iteratieve procedure te introduceren die gebaseerd op de numerieke renormalizatiegroep. Met deze methodiek kunnen we bases bestaande uit honderdduizenden eigentoestanden beschouwen wat, in combinatie met onze zoekalgoritmes, uitstekende resultaten oplevert. Als we echter een zeer drastische verandering van de interacties tussen de deeltjes beschouwen, treden er problemen op met de iteratieve aanpak. Om ook in dit geval de meest accuraat mogelijke resultaten te behalen maken we verdere aanpassingen aan de iteratieve aanpak. Dat resulteert in een nieuw algoritme dat nog steeds grote bases kan beschouwen, maar waarvoor ook in alle gevallen de iteratieve procedure accurate resultaten oplevert.

In hoofdstuk 5 bestuderen we de fundamentele eigenschappen van de matrixelementen van eigentoestanden. We bepalen onder andere de schalingsrelaties voor de sectoren van eigentoestanden die het belangrijkst zijn voor de meeste correlatiefuncties. Deze relaties kwantificeren veel van de bestaande intuities over hoe het belang van deze sectoren verandert als we grotere systemen of toestanden met een eindige entropie beschouwen. Verder bestuderen we de matrixelementen tussen eigentoestanden behorend tot verschillende macrotoestanden. De schalingsrelatie die hieruit voortkomt toont een verrassende afwijking van het gedrag dat verwacht zou worden aan de hand van de ETH[6].

Het bestuderen van de statistiche eigenschappen van matrixelementen van typische eigentoestanden vereist een algoritme dat op willekeurige wijze eigentoestanden genereert die behoren tot een gegeven macrotoestand. We beschouwen allereerst verschillende bestaande methodes en concluderen dat deze niet volledig onbevooroordeeld zijn of alleen gebruikt kunnen worden voor kleine systemen. Daarom ontwikkelen we alternatieve methodes die ons in staat stellen de matrixelementen te bestuderen voor systeemgrootten van wel honderden deeltjes. De daarmee behaalde resultaten tonen een niet eerder in deze context geobserveerde statistiche distributie, de Fréchet distributie.

---

[6]Eigenstate Thermalization Hypothesis

Appendix A

Solving the Bethe Ansatz Equations numerically

The logarithmic Bethe Ansatz Equations represent a system of non-linear coupled equations for the rapidities. In order to solve this system for a given set of quantum numbers the Newton-Rahpson algorithm can be used. In this appendix we review the algorithm first for the one-dimensional case and then for the multi-dimensional case after which we discuss how the method can be used to solve the Bethe Ansatz Equations.

A.1 The Newton-Raphson algorithm

A.1.1 The one-dimensional case

Consider a function $f : \mathbb{R} \to \mathbb{R}$ and suppose we are interested in finding one of its roots x^*. Starting from some point x_0 we can linearise the function f around this given initial point provided that $f'(x_0) \neq 0$, giving

$$f(x_0 + \epsilon_0) = f(x_0) + \epsilon_0 f'(x_0) + O(\epsilon_0^2). \tag{A.1}$$

Neglecting the higher order terms, we can find the value of ϵ_0 for which the linearisation intersects with the real axis, which is given by

$$\epsilon_0 = -\frac{f(x_0)}{f'(x_0)} \tag{A.2}$$

Taking this zero of the linearization, provided that $f'(x_0)$ is non-zero, as the new starting point and repeating the linearization step gives rise to a sequence of points $\{x_0, x_1, x_2, \dots\}$, where the x_i are given by

$$x_i = x_{i-1} - \frac{f(x_{i-1})}{f'(x_{i-1})}. \tag{A.3}$$

A sufficient condition for this procedure to converge is that $f(x)$ has a root and is convex provided that $f'(x_0) = \neq 0$.

A.1.2 The N-dimensional case

Let us now instead consider a function $g : \mathbb{R}^N \to \mathbb{R}^N$. We can generalise the procedure described in the previous section as follows, first consider the Taylor expansion of g around some $\mathbf{x}_0$, i.e.

$$g(\mathbf{x}_0 + \epsilon_0) = g(\mathbf{x}_0) + Jg(\mathbf{x}_0) \cdot \epsilon_0 + O(\epsilon_0^2) \tag{A.4}$$

where Jg is the jacobian of g. Then we can again find a root of the linearisation by solving for ϵ_0 which gives

$$\epsilon_0 = -Jg^{-1}(\mathbf{x}_0)g(\mathbf{x}_0) \tag{A.5}$$

This procedure again gives rise sequence $\{\mathbf{x}_0, \mathbf{x}_1, \mathbf{x}_2, \dots\}$, which converges if g has a root and the Jacobian is positive-definite in addition to the requirement that $Jg(\mathbf{x}_0) \neq \mathbf{0}$.

A.2 Solving the Bethe Ansatz Equations

In order to use the Newton-Raphson algorithm to solve the logarithmic Bethe equations, we have to know the corresponding Jacobian. Consider the action given by

$$S(\{\lambda\}) = \frac{L}{2} \sum_j \left\{ \lambda_j^2 + \sum_{l=j+1}^{N} \Phi(\lambda_j - \lambda_l) - 2\pi I_j \lambda_j \right\}. \tag{A.6}$$

which is called the Yang-Yang action where

$$\Phi(z) = \left[2z \, \mathrm{atan}\frac{z}{c} - c \ln\left(1 + \frac{z^2}{c^2}\right) \right] \tag{A.7}$$

The extremum conditions are the logarithmic Bethe equations for the Lieb-Liniger model, i.e.

$$\frac{\partial}{\partial \lambda_j} S(\{\lambda\}) = \lambda_j L + 2 \sum_l \mathrm{atan}\left(\frac{\lambda_j - \lambda_l}{c}\right) - 2\pi I_j. \tag{A.8}$$

The Jacobian of the logarithmic Bethe equations is therefore given by the Hessian of the Yang-Yang matrix called the Gaudin matrix whose entries are given by

$$\frac{\partial^2}{\partial \lambda_j \partial \lambda_l} S(\{\lambda\}) = \delta_{jl} \left\{ L + \sum_{m=1}^{N} \frac{2c}{(\lambda_j - \lambda_m)^2 + c^2} \right\} - \frac{2c}{(\lambda_j - \lambda_l)^2 + c^2}. \tag{A.9}$$

This is again the Gaudin matrix, whose determinant describes the norm of an eigenstate. Since the Gaudin matrix is positive definite, using the Newton-Raphson method can be used to solve the logarithmic Bethe Ansatz Equations.

ACKNOWLEDGEMENTS

The first person I want to thank is my supervisor J-S, without whom I would not have written this thesis. When I applied to a position in your group it was not because I had read your papers, but because I watched you talk enthusiastically about Go on YouTube and read about your crusade against the big time publishers. That was the type of person I wanted to have as my supervisor. Although your crusade sometimes meant we did not have as much time together as we would have otherwise, you taught me a lot and always did so in a kind, patient and supporting manner.

Another person without whom I am not sure I would have been able to write this thesis is Neil. What started as a simple collaboration on a project turned into you becoming one of the guiding forces throughout my PhD. Even when you left academia you continued supporting and mentoring me, and I am forever grateful for it. I strongly admire you and hope to be more like you one day.

I am also deeply endebted to my paranymphs, Rebekka and Sergio. Both have contributed greatly to my emotional welfare and my mental sanity during the sometimes not so stressfree times over the last four years.

Sergio, on my first day you were asked by J-S to show me around and get me settled because he was travelling. However, I do not think you ever stopped being my PhD-guide. You were always there for me when I needed some words of encouragement or just a big hug. I am very happy to have you with me as one of my paranymphs.

Rebekka, you are such a warm person that makes the entire institute a happier place. I am sorry for stealing your job initially, but I am very happy you joined us regardless. Thank you for supporting me as one of my paranymphs.

Thanks also to all of my other serious and not-so-serious fellow gang members, Alvise, Vincenzo, Sasha, Enej, Alberto, Daniel and Yuan. Although our numbers have whittled over the years, I am grateful for the time we spent together.

I also owe my thanks to all of the wonderful colleagues with whom I shared our sunlight-deprived office cave. Yinin, Lieuwe, Patrick, Erik, Ward, and Boris, I wish you all brighter offices in the future. In addition I would like to

thank all of my other fellow PhD-students and postdocs for fun chats at the coffee machine, Friday drinks and so much more. Carla, Francesca, Karina, Jans, Schelto, Corentin, Jiri, Juraj, Bart, Liam, Juan-Diego, and Rodrigo, you are all amazing people.

Coming to Amsterdam I immediately appreciated the friendly atmosphere at the intstitute. All the faculty and staff deserve a thank you for making it so. I want to especially thank Jasper for giving me an enjoyable teaching experience and for all the fun chats we had.

Another group of people that, sometimes unknowingly, supported me throughout my PhD are my friends. Nick, Sebas, Nien, Eveline, EPL and my German crew, I hope to see more of you now that I have emerged from this ordeal as a civil servant. Max and Brin deserve special thanks for also putting up with me during our intense two-year co-habitation project prior to the PhD, and our lovely hiking holiday last year that did me a world of good. I love you to bits.

Thanks also to my cat Snorrie, despite being the fun-police that did not let me work from home as well as I would have liked.

Lieve papa, mama, Jan, Carina en Cornelieke, ik ben jullie zo dankbaar voor alle steun en al het lieve advies dat ik van jullie gehad heb de afgelopen jaren. Hoewel ik niet zo goed ben in het vragen om hulp is het erg fijn om te weten dat jullie altijd voor me klaar zouden staan als ik dat wel zou doen. Wat ik hier ook probeer te schrijven, het dekt de lading niet hoe belangrijk jullie voor me zijn. Ik hou van jullie.

Finally I owe my thanks to Marieke. You always support me without question and make my life so much more enjoyable. I am grateful for all the times you told me what I needed to hear instead of what I wanted to hear. If it were not for you this defence would surely have taken place at a later point in time.

Bibliography

[1] Albertus J. J. M. De Klerk and Jean-Sébastien Caux. Improved Hilbert space exploration algorithms for finite temperature calculations. *SciPost Physics Core*, 6(2):039, May 2023.

[2] Neil Robinson, Albertus de Klerk, and Jean-Sébastien Caux. On computing non-equilibrium dynamics following a quench. *SciPost Phys.*, 11(6):104, December 2021.

[3] F.H.L. Essler, A.J.J.M. de Klerk, N.J. Robinson, and J.-S. Caux. To be published.

[4] Albertus J. J. M. de Klerk, Roland I. van der Veen, Jan Willem Dalhuisen, and Dirk Bouwmeester. Knotted optical vortices in exact solutions to Maxwell's equations. *Phys. Rev. A*, 95(5):053820, May 2017.

[5] Hans Albrecht Bethe. Zur Theorie der Metalle. *Z. Phys.*, 75(205):205–226, June 1931.

[6] Elliott H. Lieb and Werner Liniger. Exact Analysis of an Interacting Bose Gas. I. The General Solution and the Ground State. *Phys. Rev.*, 130(4):1605–1616, May 1963.

[7] Elliott H. Lieb. Exact Analysis of an Interacting Bose Gas. II. The Excitation Spectrum. *Phys. Rev.*, 130(4):1616–1624, May 1963.

[8] Elliott H. Lieb. Residual Entropy of Square Ice. *Phys. Rev.*, 162(1):162–172, October 1967.

[9] Elliott H. Lieb. Exact Solution of the Problem of the Entropy of Two-Dimensional Ice. *Phys. Rev. Lett.*, 18(17):692–694, April 1967.

[10] Rodney J Baxter. Partition function of the Eight-Vertex lattice model. *Annals of Physics*, 70(1):193–228, March 1972.

[11] C. N. Yang. Some Exact Results for the Many-Body Problem in one Dimension with Repulsive Delta-Function Interaction. *Phys. Rev. Lett.*, 19(23):1312–1315, December 1967.

[12] Jean-Sébastien Caux and Jorn Mossel. Remarks on the notion of quantum integrability. *J. Stat. Mech.*, 2011(02):P02023, February 2011.

[13] Anatoli Polkovnikov, Krishnendu Sengupta, Alessandro Silva, and Mukund Vengalattore. *Colloquium* : Nonequilibrium dynamics of closed interacting quantum systems. *Rev. Mod. Phys.*, 83(3):863–883, August 2011.

[14] Tim Langen, Remi Geiger, and Jörg Schmiedmayer. Ultracold Atoms Out of Equilibrium. *Annu. Rev. Condens. Matter Phys.*, 6(1):201–217, March 2015.

[15] Toshiya Kinoshita, Trevor Wenger, and David S. Weiss. A quantum Newton's cradle. *Nature*, 440(7086):900–903, April 2006.

[16] Fabian H L Essler and Maurizio Fagotti. Quench dynamics and relaxation in isolated integrable quantum spin chains. *J. Stat. Mech.*, 2016(6):064002, June 2016.

[17] Pasquale Calabrese and John Cardy. Quantum quenches in $1 + 1$ dimensional conformal field theories. *J. Stat. Mech.*, 2016(6):064003, June 2016.

[18] M. A. Cazalilla, R. Citro, T. Giamarchi, E. Orignac, and M. Rigol. One dimensional bosons: From condensed matter systems to ultracold gases. *Rev. Mod. Phys.*, 83(4):1405–1466, December 2011.

[19] Denis Bernard and Benjamin Doyon. Conformal field theory out of equilibrium: a review. *J. Stat. Mech.*, 2016(6):064005, June 2016.

[20] Jean-Sébastien Caux. The Quench Action. *J. Stat. Mech.*, 2016(6):064006, June 2016.

[21] Lev Vidmar and Marcos Rigol. Generalized Gibbs ensemble in integrable lattice models. *J. Stat. Mech.*, 2016(6):064007, June 2016.

[22] E. Ilievski, J. De Nardis, B. Wouters, J.-S. Caux, F.H.L. Essler, and T. Prosen. Complete Generalized Gibbs Ensembles in an Interacting Theory. *Phys. Rev. Lett.*, 115(15):157201, October 2015.

[23] Tim Langen, Thomas Gasenzer, and Jörg Schmiedmayer. Prethermalization and universal dynamics in near-integrable quantum systems. *J. Stat. Mech.*, 2016(6):064009, June 2016.

[24] Romain Vasseur and Joel E Moore. Nonequilibrium quantum dynamics and transport: from integrability to many-body localization. *J. Stat. Mech.*, 2016(6):064010, June 2016.

[25] Andrea De Luca and Giuseppe Mussardo. Equilibration properties of classical integrable field theories. *J. Stat. Mech.*, 2016(6):064011, June 2016.

[26] V. E. Korepin, N. M. Bogoliubov, and A. G. Izergin. *Quantum Inverse Scattering Method and Correlation Functions*. Cambridge University Press, March 1997.

[27] Jean-Sébastien Caux and Pasquale Calabrese. Dynamical density-density correlations in the one-dimensional Bose gas. *Phys. Rev. A*, 74(3):031605, September 2006.

[28] Pasquale Calabrese and Jean-Sébastien Caux. Dynamics of the attractive 1D Bose gas: analytical treatment from integrability. *J. Stat. Mech.*, 2007(08):P08032–P08032, August 2007.

[29] Miosz Panfil and Jean-Sébastien Caux. Finite-temperature correlations in the Lieb-Liniger one-dimensional Bose gas. *Phys. Rev. A*, 89(3):033605, March 2014.

[30] N. A. Slavnov. Calculation of scalar products of wave functions and form factors in the framework of the alcebraic Bethe ansatz. *Theor Math Phys*, 79(2):502–508, May 1989.

[31] N. A. Slavnov. Nonequal-time current correlation function in a one-dimensional Bose gas. *Theor Math Phys*, 82(3):273–282, March 1990.

[32] T. Kojima, V. E. Korepin, and N. A. Slavnov. Determinant Representation for Dynamical Correlation Functions of the Quantum Nonlinear Schrödinger Equation. *Communications in Mathematical Physics*, 188(3):657–689, October 1997.

[33] Balázs Pozsgay. Local correlations in the 1D Bose gas from a scaling limit of the XXZ chain. *J. Stat. Mech.*, 2011(11):P11017, November 2011.

[34] Lorenzo Piroli and Pasquale Calabrese. Exact formulas for the form factors of local operators in the LiebLiniger model. *J. Phys. A: Math. Theor.*, 48(45):454002, November 2015.

[35] C. N. Yang and C. P. Yang. Thermodynamics of a OneDimensional System of Bosons with Repulsive DeltaFunction Interaction. *Journal of Mathematical Physics*, 10(7):1115–1122, July 1969.

[36] M. Girardeau. Relationship between Systems of Impenetrable Bosons and Fermions in One Dimension. *Journal of Mathematical Physics*, 1(6):516–523, November 1960.

[37] Michel Gaudin. *The Bethe Wavefunction*. Cambridge University Press, March 2014. Google-Books-ID: 4XLtAgAAQBAJ.

[38] E. M. Lifshitz. *Statistical Physics, vol. 5*. Butterworth-Heinemann, 1980.

[39] Jean-Sébastien Caux and Pasquale Calabrese. Dynamical density-density correlations in the one-dimensional Bose gas. *Phys. Rev. A*, 74(3):031605, September 2006.

[40] Jean-Sébastien Caux. Correlation functions of integrable models: A description of the ABACUS algorithm. *Journal of Mathematical Physics*, 50(9):095214, September 2009.

[41] V. P. Yurov and Al. B. Zamolodchikov. Truncated comformal space approach to scaling lee-yang model. *Int. J. Mod. Phys. A*, 05(16):3221–3245, August 1990. Publisher: World Scientific Publishing Co.

[42] V.p. Yurov and Al.b. Zamolodchikov. Truncated-fermionic-space approach to the critical 2d ising model with magnetic field. *Int. J. Mod. Phys. A*, 06(25):4557–4578, October 1991. Publisher: World Scientific Publishing Co.

[43] Robert M. Konik and Yury Adamov. Numerical Renormalization Group for Continuum One-Dimensional Systems. *Phys. Rev. Lett.*, 98(14):147205, April 2007.

[44] Jacopo De Nardis, Bram Wouters, Michael Brockmann, and Jean-Sébastien Caux. Solution for an interaction quench in the Lieb-Liniger Bose gas. *Phys. Rev. A*, 89(3):033601, March 2014.

[45] J De Nardis and J-S Caux. Analytical expression for a post-quench time evolution of the one-body density matrix of one-dimensional hard-core bosons. *J. Stat. Mech.*, 2014(12):P12012, December 2014.

[46] Andrew J A James, Robert M Konik, Philippe Lecheminant, Neil J Robinson, and Alexei M Tsvelik. Non-perturbative methodologies for low-dimensional strongly-correlated systems: From non-Abelian bosonization to truncated spectrum methods. *Rep. Prog. Phys.*, 81(4):046002, April 2018.

[47] Marcos Rigol, Vanja Dunjko, and Maxim Olshanii. Thermalization and its mechanism for generic isolated quantum systems. *Nature*, 452(7189):854–858, April 2008.

[48] Jan C. Zill, Tod M. Wright, Karén V. Kheruntsyan, Thomas Gasenzer, and Matthew J. Davis. Relaxation dynamics of the Lieb-Liniger gas following an interaction quench: A coordinate Bethe-ansatz analysis. *Phys. Rev. A*, 91(2):023611, February 2015.

[49] Jan C Zill, Tod M Wright, Karén V Kheruntsyan, Thomas Gasenzer, and Matthew J Davis. A coordinate Bethe ansatz approach to the calculation of equilibrium and nonequilibrium correlations of the one-dimensional Bose gas. *New J. Phys.*, 18(4):045010, April 2016.

[50] T. Rakovszky, M. Mestyán, M. Collura, M. Kormos, and G. Takács. Hamiltonian truncation approach to quenches in the Ising field theory. *Nuclear Physics B*, 911:805–845, October 2016.

[51] Marton Kormos, Mario Collura, Gabor Takács, and Pasquale Calabrese. Real-time confinement following a quantum quench to a non-integrable model. *Nature Phys*, 13(3):246–249, March 2017. Number: 3 Publisher: Nature Publishing Group.

[52] Kristóf Hódsági, Márton Kormos, and Gábor Takács. Quench dynamics of the Ising field theory in a magnetic field. *SciPost Phys.*, 5(3):027, September 2018.

[53] Neil J. Robinson, Andrew J. A. James, and Robert M. Konik. Signatures of rare states and thermalization in a theory with confinement. *Phys. Rev. B*, 99(19):195108, May 2019.

[54] D. X. Horváth and G. Takács. Overlaps after quantum quenches in the sine-Gordon model. *Physics Letters B*, 771:539–545, August 2017.

[55] I. Kukuljan, S. Sotiriadis, and G. Takacs. Correlation Functions of the Quantum Sine-Gordon Model in and out of Equilibrium. *Phys. Rev. Lett.*, 121(11):110402, September 2018.

[56] D. X. Horváth, M. Kormos, and G. Takács. Overlap singularity and time evolution in integrable quantum field theory. *J. High Energ. Phys.*, 2018(8):170, August 2018.

[57] Jean-Sébastien Caux and Fabian H. L. Essler. Time Evolution of Local Observables After Quenching to an Integrable Model. *Phys. Rev. Lett.*, 110(25):257203, June 2013.

[58] Jean-Sébastien Caux. The Quench Action. *J. Stat. Mech.*, 2016(6):064006, June 2016.

[59] Balázs Pozsgay. Overlaps between eigenstates of the XXZ spin-1/2 chain and a class of simple product states. *J. Stat. Mech.*, 2014(6):P06011, June 2014.

[60] M Brockmann. Overlaps of q -raised Néel states with XXZ Bethe states and their relation to the LiebLiniger Bose gas. *J. Stat. Mech.*, 2014(5):P05006, May 2014.

[61] M Brockmann, J De Nardis, B Wouters, and J-S Caux. A Gaudin-like determinant for overlaps of Néel and XXZ Bethe states. *J. Phys. A: Math. Theor.*, 47(14):145003, April 2014.

[62] M Brockmann, J De Nardis, B Wouters, and J-S Caux. Néel-XXZ state overlaps: odd particle numbers and LiebLiniger scaling limit. *J. Phys. A: Math. Theor.*, 47(34):345003, August 2014.

[63] Lorenzo Piroli and Pasquale Calabrese. Recursive formulas for the overlaps between Bethe states and product states in XXZ Heisenberg chains. *J. Phys. A: Math. Theor.*, 47(38):385003, September 2014.

[64] Marius de Leeuw, Charlotte Kristjansen, and Konstantin Zarembo. One-point functions in defect CFT and integrability. *J. High Energ. Phys.*, 2015(8):98, August 2015.

[65] O Foda and K Zarembo. Overlaps of partial Néel states and Bethe states. *J. Stat. Mech.*, 2016(2):023107, February 2016.

[66] Isak Buhl-Mortensen, Marius de Leeuw, Charlotte Kristjansen, and Konstantin Zarembo. One-point functions in AdS/dCFT from matrix product states. *J. High Energ. Phys.*, 2016(2):52, February 2016.

[67] Marius de Leeuw, Charlotte Kristjansen, and Stefano Mori. AdS/dCFT one-point functions of the SU(3) sector. *Physics Letters B*, 763:197–202, 2016.

[68] Bruno Bertini, Elena Tartaglia, and Pasquale Calabrese. Quantum quench in the infinitely repulsive Hubbard model: the stationary state. *J. Stat. Mech.*, 2017(10):103107, October 2017.

[69] Lorenzo Piroli, Balázs Pozsgay, and Eric Vernier. What is an integrable quench? *Nuclear Physics B*, 925:362–402, 2017.

[70] Lorenzo Piroli and Pasquale Calabrese. Exact dynamics following an interaction quench in a one-dimensional anyonic gas. *Phys. Rev. A*, 96(2):023611, August 2017.

[71] Marius de Leeuw, Charlotte Kristjansen, and Georgios Linardopoulos. One-point functions of non-protected operators in the *SO* (5) symmetric D3D7 dCFT. *J. Phys. A: Math. Theor.*, 50(25):254001, June 2017.

[72] Marius de Leeuw, Tamás Gombor, Charlotte Kristjansen, Georgios Linardopoulos, and Balázs Pozsgay. Spin chain overlaps and the twisted Yangian. *J. High Energ. Phys.*, 2020(1):176, January 2020.

[73] Georgios Linardopoulos. Solving holographic defects. In *Proceedings of Corfu Summer Institute 2019 "School and Workshops on Elementary Particle Physics and Gravity" PoS(CORFU2019)*, page 141, Corfù, Greece, August 2020. Sissa Medialab.

[74] Kristóf Hódsági, Márton Kormos, and Gábor Takács. Quench dynamics of the Ising field theory in a magnetic field. *SciPost Phys.*, 5(3):027, September 2018.

[75] Andrew J.A. James, Robert M. Konik, and Neil J. Robinson. Nonthermal States Arising from Confinement in One and Two Dimensions. *Phys. Rev. Lett.*, 122(13):130603, April 2019.

[76] Overlaps after quantum quenches in the sine-Gordon model | Elsevier Enhanced Reader.

[77] Kenneth G. Wilson. The renormalization group: Critical phenomena and the Kondo problem. *Rev. Mod. Phys.*, 47(4):773–840, October 1975.

[78] G P Brandino, R M Konik, and G Mussardo. Energy level distribution of perturbed conformal field theories. *J. Stat. Mech.*, 2010(07):P07013, July 2010.

[79] Robert M. Konik. Exciton Hierarchies in Gapped Carbon Nanotubes. *Phys. Rev. Lett.*, 106(13):136805, March 2011.

[80] Jean-Sébastien Caux and Robert M. Konik. Constructing the Generalized Gibbs Ensemble after a Quantum Quench. *Phys. Rev. Lett.*, 109(17):175301, October 2012.

[81] G.P. Brandino, J.-S. Caux, and R.M. Konik. Glimmers of a Quantum KAM Theorem: Insights from Quantum Quenches in One-Dimensional Bose Gases. *Phys. Rev. X*, 5(4):041043, December 2015.

[82] Robert M. Konik, Matthew Y. Sfeir, and James A. Misewich. Predicting excitonic gaps of semiconducting single-walled carbon nanotubes from a field theoretic analysis. *Phys. Rev. B*, 91(7):075417, February 2015.

[83] Lorenzo Piroli, Eric Vernier, Pasquale Calabrese, and Marcos Rigol. Correlations and diagonal entropy after quantum quenches in XXZ chains. *Phys. Rev. B*, 95(5):054308, February 2017.

[84] Michael Moeckel and Stefan Kehrein. Interaction Quench in the Hubbard Model. *Phys. Rev. Lett.*, 100(17):175702, May 2008.

[85] Matteo Marcuzzi, Jamir Marino, Andrea Gambassi, and Alessandro Silva. Prethermalization in a Nonintegrable Quantum Spin Chain after a Quench. *Phys. Rev. Lett.*, 111(19):197203, November 2013.

[86] F. H. L. Essler, S. Kehrein, S. R. Manmana, and N. J. Robinson. Quench dynamics in a model with tuneable integrability breaking. *Phys. Rev. B*, 89(16):165104, April 2014.

[87] Bruno Bertini, Fabian H.L. Essler, Stefan Groha, and Neil J. Robinson. Prethermalization and Thermalization in Models with Weak Integrability Breaking. *Phys. Rev. Lett.*, 115(18):180601, October 2015.

[88] Pasquale Calabrese and John Cardy. Evolution of entanglement entropy in one-dimensional systems. *J. Stat. Mech.*, 2005(04):P04010, April 2005.

[89] Vincenzo Alba and Pasquale Calabrese. Entanglement dynamics after quantum quenches in generic integrable systems. *SciPost Phys.*, 4(3):017, March 2018.

[90] F. Verstraete and J. I. Cirac. Continuous Matrix Product States for Quantum Fields. *Phys. Rev. Lett.*, 104(19):190405, May 2010.

[91] Damian Draxler, Jutho Haegeman, Tobias J. Osborne, Vid Stojevic, Laurens Vanderstraeten, and Frank Verstraete. Particles, Holes, and Solitons: A Matrix Product State Approach. *Phys. Rev. Lett.*, 111(2):020402, July 2013.

[92] Aditya Shashi, Miosz Panfil, Jean-Sébastien Caux, and Adilet Imambekov. Exact prefactors in static and dynamic correlation functions of one-dimensional quantum integrable models: Applications to the Calogero-Sutherland, Lieb-Liniger, and X X Z models. *Phys. Rev. B*, 85(15):155136, April 2012.

[93] Luca D'Alessio, Yariv Kafri, Anatoli Polkovnikov, and Marcos Rigol. From quantum chaos and eigenstate thermalization to statistical mechanics and thermodynamics. *Advances in Physics*, 65(3):239–362, May 2016.

[94] Etienne Granet and Fabian Essler. A systematic $1/c$-expansion of form factor sums for dynamical correlations in the Lieb-Liniger model. *SciPost Phys.*, 9(6):082, December 2020.

[95] Pavlos S. Efraimidis and Paul G. Spirakis. Weighted random sampling with a reservoir. *Information Processing Letters*, 97(5):181–185, March 2006.