Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

UNIVERSITÀ DEGLI STUDI DI MILANO

Scuola di Dottorato in Fisica, Astrofisica e Fisica Applicata
Dipartimento di Fisica

Corso di Dottorato in Fisica, Astrofisica e Fisica Applicata
Ciclo XXXVIII

Tensor networks for open quantum systems

Settore Scientifico Disciplinare FIS/03

Supervisore
Prof. Dario Tamascelli

Co-supervisore
Dott. Matteo Rossi

Tesi di dottorato di
Davide Ferracin

iD 0009-0007-2980-4677
Matricola **R13684**

Coordinatore
Prof. Aniello Mennella

Tesi redatta con il contributo finanziario dell'Unione Europea – Next Generation EU e di Algorithmiq OY.
CUP: **G43C22002320004**

Anno Accademico 2024–2025

Final examination date

27 January 2026

Thesis referees

Dr. Alex W. Chin
Institut des NanoSciences de Paris, Sorbonne Université (Paris, France)

Prof. Stephen R. Clark
H.H. Wills Physics Laboratory, School of Physics, University of Bristol (Bristol, UK)

Final examination committee

Dr. Alex W. Chin
Institut des NanoSciences de Paris, Sorbonne Université (Paris, France)

Prof. Stephen R. Clark
H.H. Wills Physics Laboratory, School of Physics, University of Bristol (Bristol, UK)

Prof. Mauro Paternostro
Dipartimento di Fisica e Chimica, Università di Palermo (Palermo, Italy)

Contents

Introduction

Is it possible to efficiently simulate a quantum physical system on a classical computer? The answer to this question is not straightforward, because of the exponential scaling of the computational resources (memory, CPU time) needed: the state of a single quantum system that lives in a d-dimensional Hilbert space is described by d complex numbers; however, the dimension of the Hilbert space of n such systems interacting with each other is d^n. Even just representing the states of a quantum system on a classical computer therefore quickly becomes dauntingly difficult as the number of components increases, let alone simulating its behaviour over time.

A variety of sophisticated numerical methods have been developed over time to tackle the challenge of simulating quantum systems: the dynamical mean-field theory (DMFT) [4], quantum Monte Carlo methods [5], approaches based on neural networks [6, 7], classical phase-space representations [8], and numerical linked-cluster expansions [9], to name a few, are representative of methods that, building upon specific assumptions, try to mitigate the aforementioned 'curse of dimensionality'. In this thesis we will focus on another simulation tool, tensor networks (TNs), which are a mathematical representation of quantum many-body states based on their entanglement structure [10]. Whenever one has a Hilbert space with a tensor-product structure, TNs naturally arise as a way to replace the original full tensor of coefficients (in a given basis) of the state by a network of interconnected tensors. As we will discuss later, in many relevant cases this approach leads to a compact description of quantum states, requiring only a number of coefficients that scales polynomially, as $O(\text{poly}(n))$, instead of exponentially, as $O(d^n)$, in the number of degrees of freedom. This method can be applied for example to a chain of n spins modelling the magnetic properties of a material, or an ensemble of n photonic modes manipulated by quantum optical devices in boson sampling experiments, or a quantum computer with n qubits.

By an appropriate truncation of the individual tensors of the network, TNs can also approximate quantum states in a controlled way, speeding up classical simulations and reducing their memory footprint while sacrificing some of the accuracy. The efficiency of TNs and their ability to represent general quantum systems are well-known advantages, which is why they are employed in a wide range of cases. A fundamental example of TN, the matrix-product state (MPS), was introduced in the 1990s as a mathematical method of representing finitely correlated many-body quantum states [11], from the AKLT state [12] to Heisenberg spin chains [13, 14].

The first successful practical TN algorithm, the density-matrix renormalisation group (DMRG), was developed in 1992 by Steven White, in the context of strongly correlated systems [15]. The DMRG is an iterative, variational algorithm devised to find the low energy properties (states, eigenenergies, and so on) of a one-dimensional many-body system. It works by systematically truncating the Hilbert space, identifying the most relevant states that describe a target state—most often the ground state—in order to find an efficient description of it [16]. It is a successful generalisation of Wilson's numerical renormalisation-group procedure, and has quickly become the method of choice in the numerical studies of static, dynamic, and thermodynamic properties of one-dimensional quantum systems, as well as in computational quantum chemistry. It was later recognised that in the thermodynamic limit, when the renormalization procedure converges to a fixed point, the DMRG ground state can be written in a matrix-product form [17]; later, the DMRG algorithm was studied as a variational optimisation procedure over the class of MPSs [14]. The reason why DMRG performs well, and in turn, why MPSs became a relevant tool in quantum many-body physics, is the small amount of entanglement in ground states of gapped Hamiltonians with short-range interactions. These states obey an *area law* of entanglement [18], according to which the entanglement entropy of the reduced state of a subregion grows like the area of the boundary, and not like its volume (which is expected from a quantum state picked at random). It is precisely this favourable scaling of the entanglement that specifies how well an MPS is able to approximate a given state [11, 16].

Later on, in the 2000s, the theory of TNs was linked to the field of quantum information through the work of Guifré Vidal [19], who treated them as proper mathematical tools for describing quantum many-body states and entanglement. The potential of MPSs then began to be systematically explored by Ignacio Cirac, Frank Verstraete and Vidal [20–22], and several other TN structures were introduced to describe different situations, even beyond one-dimensional systems, including projected entangled pair states (PEPS) [23] and multiscale entanglement renormalization ansatz (MERA) [24]. The efficiency of MPSs, and TNs in general, and their versatility are well-known advantages, which is why they are employed in an increasingly broad range of applications [10]. Beyond the fields of condensed matter physics and quantum information, TNs are employed in quantum gravity [25, 26], (classical) statistical physics [27], machine learning [28, 29], solving partial differential equations [30, 31], and many others.

In this thesis we will explore the practical application of TNs, more specifically MPSs, in two specific scenarios:

- the simulation of the dynamics of a quantum system immersed in a thermal fermionic environment;

- a full-TN simulation of the unitary and noisy behaviour of a qubit-based quantum computer.

These topics involve open quantum systems (OQSs), which introduce another level of difficulty to the simulation: they are systems that exchange energy, matter or more in general information through their interaction with their surroundings. This interaction leads to the emergence of dissipative and noisy processes. While these processes are typically seen as detrimental in the context of quantum information processing, since they reduce quantum resources such as coherence and entanglement, in other cases they can also be beneficial for instance by facilitating the transport of energy in the system. In any case, when modelling a realistic system, the effects of its environment must be included, in some form, in the description [32]. The environment consists of a large, sometimes infinite, number of degrees of freedom, therefore an explicit description of the collective state of the system and its environment is usually intractable. The goal of the theory of OQSs is instead to provide an effective description that focuses only of the system's degrees of freedom, possibly complemented by a reduced and manageable representation of the environment. Numerical simulations become important in this context as they make it possible to explore the behaviour of open systems beyond the typical assumptions of a structureless, Markovian environment.

Many algorithms have been developed to face the numerical simulation of OQSs. Numerous physical environments are well modelled by an infinite collection of harmonic oscillators, as they can describe vibrational and thermal excitation of what surrounds the system. As a result, bosonic environments are a common choice for the environment of an OQS. The non-trivial structure of these baths can induce interesting phenomena in the open system: for example, it can generate non-equilibrium processes that sustain electronic coherences in biological systems [33]. It is therefore important to develop numerical simulation tools that can treat these non-trivial effects.

The first part of this thesis will focus in particular on an extension of the time-evolving density operator with orthogonal polynomials algorithm (TEDOPA) [34], a *chain-mapping algorithm* that unitarily transforms a continuous environment into a countably infinite chain of modes. This transformation allows continuous environments to be translated into a discrete structure which can be easily written as an MPS, and which can therefore be efficiently numerically simulated. However, the size of the resulting chain grows linearly with the physical duration of the time evolution, and this can be troublesome when long-time dynamics are considered. In order to address this shortcoming, the Markovian closure (MC) method, developed in Ref. [35] for bosonic environments, truncates the chain-mapped environment to a finite size, in such a way that the dynamics of the open system is approximately preserved. In many relevant cases, fermionic environments play a central role as well. They appear, for example, when studying the behaviour of magnetic impurities in metals [36, 37] or electron transport in semiconductor quantum dots [38]. For this kind of environments, several techniques [39, 40] already exist that replace the continuous environment with a finite number of auxiliary fermionic modes, whose parameters are carefully tuned in order to mimic the effects of the original environment on the system. In this thesis, we extend the MC method to fermionic environments, providing—as in the bosonic case—a generic numerical approximation method that yields a quadratic speed-up of the wall-clock simulation time with respect to the physical simulation time, compared to standard TEDOPA. This method does not require ad-hoc fitting procedures and relies only on mild assumptions on the structure of the environment. Moreover, it also preserves the information on the first modes of the TEDOPA chain, that are responsible for the non-Markovian behaviour of the environment.

In the second part of the thesis, we turn our attention to the realm of quantum computing, and in particular to the numerical simulation of noisy quantum circuits by means of TNs. An important part of the research in quantum computing is dedicated to *quantum advantage* experiments, that seek to determine a quantum device from which an observable quantity (possibly of practical interest) can be extracted whose computation lies beyond the reach of exact classical simulation and heuristic methods [41]. Several experiments recently emerged as viable candidates for realising practical quantum advantage, such as [42] and [43]. On the hardware side, the engineering of quantum devices keeps improving with ever decreasing error rates [44], with the ultimate goal of constructing a fault-tolerant machine [45–47] where errors during computation can be efficiently corrected.Most currently (commonly) available quantum devices, however, such as those built from superconducting qubits [48, 49], trapped ions [50, 51], or nuclear spins in silicon [52, 53], are still of a relatively small size and their operations involve a significant amount of noise: they are thus termed *noisy intermediate-scale quantum* (NISQ) devices [54]. While fault tolerance requires thousands of qubits to encode and manipulate information in an error-resilient way, these smaller, noisy quantum processors are nowadays being used to perform increasingly complex quantum simulations [48, 55, 56]. These devices inevitably entangle with the surrounding environment during its execution, and as a consequence various kinds of errors appear affecting their results. The natural framework to investigate the behaviour of noisy quantum circuits is then the one of open quantum systems.

In this setting, quantum computation on noisy devices is typically supported by classical post-processing techniques, exploiting models that are usually referred to as *hybrid quantum-classical algorithms* [57]. Several *error-mitigation* schemes are being researched: they are a class of methods that allow extracting meaningful physical quantities that approximate the ideal, noiseless results from the noisy output.

In most circumstances, these quantum devices are not able by themselves to produce meaningful answers, as the signal is usually drowned by noise, and so they are unsuited for most commercial or scientific applications. Together with error-mitigation schemes, NISQ devices at pre-fault-tolerant circuit sizes may still enable useful quantum applications. Moreover, as the research rapidly evolves, understanding how noise affects the behaviour of these computers is key to advancing their development. We will concern ourselves with the tensor-network error mitigation (TEM) method introduced in Ref. [58], which defines a TN map that approximately reconstructs the ideal final state of the circuit, and thus allows estimating the ideal (if noise were not present) measured value of observables. With this method at hand, we study if and how a particular class of quantum circuits, called dual-unitary (DU) circuits [59], constitute the right compromise between integrability and complexity to stand as suitable benchmark systems for NISQ quantum devices, if we combine a capable quantum computer, an expressive theoretical model to describe noise, and an efficient error-mitigation technique. With DU circuits, the behaviour of particular correlation functions of the kicked Ising model [60] can be simulated, and the hybrid quantum-classical approach accurately reproduces such behaviour, successfully extending the reach of the quantum processor. At the same time, we see that while a naive full-state MPS simulation of the device is too complicated, TN simulations can still provide meaningful insight, for example on the performance of the noise model, and on the structure of the entanglement generated by the circuit.

Both of these scenarios involve systems that stretch the domain of applicability of TNs, which are usually meant for moderately entangled systems. The desirable behaviour is that the state of the system satisfies the area law, in that the entanglement grows no larger than $\log n$ where n is the number of constituents of the system. In the systems we are going to study, however, even when the entanglement growth is nicely bounded as a function of n, it is not so as a function of time. Even worse, quantum computers are usually specifically designed to create entanglement between its basic components, since it is one of the primary resources that make quantum computing advantageous; in fact, we do not expect that TNs perform well in simulations of deep and highly-entangled quantum circuits [61]. For this reason we must back up the use of TNs with a careful study of the convergence of the simulation method.

Thesis structure

This work is organised as follows.

- Chapter 1 provides an overview of TN definitions and algorithms, their advantages and use cases, focusing on time-evolution methods. We also cover, in Section 1.4, how to model mixed states and their evolution as MPSs.
- In Chapter 2, after a brief introduction on the basics of the theory of OQSs, we review some theoretical and numerical methods that are commonly used to reduce their complexity, with a special attention to numerical simulations.
- Chapter 3 introduces the MC approximation for bosonic environments and extends it to the case of fermionic environments. We determine the approximation in this new scenario and we study its accuracy as well as its numerical performance. The derivation of the fermionic MC and the accompanying numerical study have been published in Ref. [1]; here we complement the results by providing, by means of the superfermion formalism introduced in Section 1.4.2, a way to solve the equations of motion in specific cases.
- Finally, Chapter 4 presents the results of Ref. [2], focusing specifically on the TN methods used in the experiment. First, the theoretical framework of DU circuits is outlined, together with the specifics of the experiment. Next, the TEM algorithm is described along with the results obtained with it. This is followed by an analysis of the classical simulation with MPSs of the quantum circuits and their noise model, analysing the convergence of the simulations and whether they can be an useful companion to the results of the hybrid quantum-classical computation.

The materials presented in Chapters 3 and 4, unless stated otherwise, are either original or adapted from the author's published articles [1] and [2].

Tensor networks 1

Tensor networks (TNs) are a technique for working with tensors which have many indices: they are a representation of large, high-order tensors as products of many low-order tensors. The naive memory and computing costs of working with a dense tensor having N indices (an order-N tensor) usually scale exponentially with N. When a dense tensor is converted into a tensor network and all of the tensors in the network are low-order, the tensor network can be an efficient way to perform common important operations such as summing two high-order tensors or computing their product. These operations can still be efficient even when the high-order tensor implicitly represented by the network has hundreds or thousands of indices.

Large tensors with many indices appear naturally as the coefficients of a many-body quantum state: if $|\psi\rangle$ is a pure state in a tensor-product Hilbert space $\mathscr{H}_1 \otimes \mathscr{H}_2 \otimes \cdots \otimes \mathscr{H}_n$, given a basis $\{|e_i^{(k)}\rangle\}_{i=1}^{d_k}$ for each $\mathscr{H}_k$ we can write

$$|\psi\rangle = \sum_{i_1=1}^{d_1} \sum_{i_2=1}^{d_2} \cdots \sum_{i_n=1}^{d_n} \psi_{i_1,i_2,\ldots,i_n} |e_{i_1}^{(1)}, e_{i_2}^{(2)}, \ldots, e_{i_n}^{(n)}\rangle \tag{1.1}$$

where $|e_{i_1}^{(1)}, \ldots, e_{i_n}^{(n)}\rangle \coloneqq |e_{i_1}^{(1)}\rangle \otimes |e_{i_2}^{(2)}\rangle \otimes \cdots \otimes |e_{i_n}^{(n)}\rangle$. When dealing with TNs, we will always consider finite-dimensional Hilbert spaces, so we will never need to worry about convergence of sums, boundedness of operators and other mathematical issues which are typical of infinite-dimensional Hilbert spaces. Dealing with the coefficient tensor $\psi_{i_1,i_2,\ldots,i_n}$ directly quickly becomes expensive as n grows: indeed the exponential growth of the number of coefficients makes the representation of the tensor very expensive even when $d_i = 2$, if n is large enough. In some cases, we can use a tensor network to represent the tensor of coefficients $\psi_{i_1,i_2,\dots,i_n}$ in a more efficient manner.

1.1 Matrix-product states

The most common example of a TN is the matrix-product state (MPS) representation. Consider a Hilbert space $\mathcal{H} = \mathcal{H}_1 \otimes \mathcal{H}_2 \otimes \cdots \otimes \mathcal{H}_n$ with a local basis $\{e_i^{(k)}\}_{i=1}^{d_k}$ for each space $\mathcal{H}_k$, and a pure state $|\psi\rangle \in \mathcal{H}$ written as in Eq. (1.1). The goal of the MPS representation is to write the coefficient tensor $\psi_{i_1,i_2,\dots,i_n}$ as a TN that separates the physical degrees of freedom of each 'site' from 1 to n, so that each entry of the original tensor can be rewritten as the product of n matrices, each one corresponding to one particular site. The algorithm to obtain an MPS from $|\psi\rangle$ consists in a series of reshapings and singular-value decompositions (SVDs) on subsequent bipartitions of the tensor-product Hilbert space, as illustrated for example in Ref. [62].

First of all, we reshape the tensor of coefficients of $|\psi\rangle$ into a matrix by grouping the indices from the second to the last in a single one:

$$a_{i_1,(i_2,\dots,i_n)} := \psi_{i_1,i_2,\dots,i_n}. \tag{1.2}$$

We apply a SVD to this matrix and we obtain, for some $r_1 \in \mathbb{N}$ with $r_1 \le d_1$ and $r_1 \le d_2 \cdots d_n$,

- a $d_1 \times r_1$ matrix U satisfying $U^\dagger U = I_{r_1}$,
- a $r_1 \times d_2 \cdots d_n$ matrix V, satisfying $V^\dagger V = I_{d_2 \cdots d_n}$,
- and a list of (positive) singular values $\lambda_1, \lambda_2, \dots, \lambda_{r_1}$,

such that the previous matrix can be written as

$$a_{i_1,(i_2,\dots,i_n)} = \sum_{j=1}^{r_1} U_{i_1,j_1} \lambda_{j_1} V_{j_1,(i_2,\dots,i_n)}. \tag{1.3}$$

We reshape now λ and V into a new matrix $a'_{(j_1,i_2),(i_3,\dots,i_n)} := \lambda_{j_1} V_{j_1,(i_2,i_3,\dots,i_n)}$, and at the same time we reshape U into a set of row vectors $\Gamma^{(i_1)}$, for $i_1 \in \{1,\dots,d_1\}$, with components $\Gamma^{(i_1)}_{j_1} = U_{i_1,j_1}$. After this first step of the decomposition, the tensor of coefficients of $|\psi\rangle$ becomes

$$\psi_{i_1,i_2\dots,i_n} = \sum_{j_1=1}^{r_1} \Gamma^{(i_1)}_{j_1} a'_{(j_1,i_2),(i_3,\dots,i_n)}. \tag{1.4}$$

Now, we carry out a new SVD on the matrix a', decomposing it into

$$a'_{(j_1,i_2),(i_3,\dots,i_n)} = \sum_{j_2=1}^{r_2} U'_{(j_1,i_2),j_2} \lambda'_{j_2} V'_{j_2,(i_3,\dots,i_n)}; \tag{1.5}$$

this time, we take the matrix U' and the singular values λ from Eq. (1.3), and we define a set of matrices $\varGamma'^{(i_2)}$, for $i_2 \in \{1,\dots,d_2\}$, whose components are $\varGamma'^{(i_2)}_{j_1,j_2} = \lambda^{-1}_{j_1} U'_{(j_1,i_2),j_2}$, so that

$$\psi_{i_1,i_2\dots,i_n} = \sum_{j_1=1}^{r_1} \sum_{j_2=1}^{r_2} \varGamma^{(i_1)}_{j_1} \lambda_{j_1} \varGamma'^{(i_2)}_{j_1,j_2} \lambda'_{j_2} V'_{j_2,(i_3,\dots,i_n)}. \tag{1.6}$$

We continue this way, with $a''_{(j_2,i_3),(i_4,\dots,i_n)} := \lambda'_{j_2} V'_{j_2,(i_3,i_4,\dots,i_n)}$ and so on, each time repeating the decomposition on the a matrix. The algorithm ends when the whole tensor of coefficients is decomposed as the matrix product

$$\psi_{i_1,i_2,\dots,i_{n-1},i_n} = \sum_{j_1=1}^{r_1} \sum_{j_2=1}^{r_2} \cdots \sum_{j_n=1}^{r_n} \varGamma^{(1,i_1)}_{j_1} \lambda^{(1)}_{j_1} \varGamma^{(2,i_2)}_{j_1,j_2} \lambda^{(2)}_{j_2} \cdots \lambda^{(n-1)}_{j_{n-1}} \varGamma^{(n,i_n)}_{j_{n-1}} \tag{1.7}$$

or, in a more compact way, by treating the singular values $\lambda^{(k)}_{j_k}$ as part of an $r_k \times r_k$ diagonal matrix $\varLambda^{(k)}_{ij} = \delta_{ij} \lambda^{(k)}_i$,

$$\psi_{i_1,i_2,\dots,i_n} = \varGamma^{(1,i_1)} \varLambda^{(1)} \varGamma^{(2,i_2)} \varLambda^{(2)} \varGamma^{(3,i_3)} \cdots \varLambda^{(n-1)} \varGamma^{(n,i_n)}. \tag{1.8}$$

We have completed the decomposition of the original ψ tensor in a matrix-product state form, namely in a way that is called *Vidal form*, introduced in Ref. [21]. In Eq. (1.8) we have a $\varGamma$ tensor for each component of the tensor-product Hilbert space, which we will call *site tensor*: they are order-three tensors except for the first and the last ones, which are of order two. Between each site tensor we find a *bond tensor*, each of which is a diagonal, positive square matrix. Internal indices, linking each site tensor with the adjacent bond tensor(s), are called *bond* or *virtual indices*; their dimension r_i, called *bond dimension* is a very important parameter of the MPS which greatly determines the complexity of the calculations involving the state. The remaining, external indices are called *physical* indices.

Expressions like Eq. (1.8) can seem very complicated, but fortunately TNs can be expressed in a graphical notation that makes their structure easily understandable. In this graphical notation, a tensor is represented as a shape (of an arbitrary kind, usually

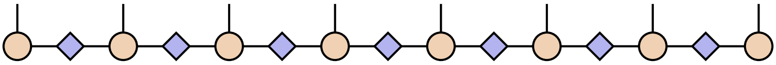

Figure 1.1. An MPS with $n = 8$ sites in Vidal form. Yellow discs represent site tensors, with the dangling leg representing their physical index; blue diamonds represent bond tensors.

simple geometric shapes) with a leg sticking out of it for each index.

$$A_{i_1,i_2,i_3} B_{j_1,j_2} =$$

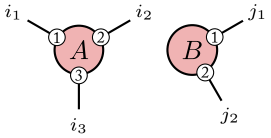

. (1.9)

The operation of *contracting* indices, that is, summing over repeated indices in a formula, is represented by simply connecting the respective legs of the tensors they belong to. It does not matter, in the graphical notation, where the individual tensors are placed: what is really important is the topology of the drawing, meaning how many legs each shape has and how the tensors are connected with each other.

$$\sum_k A_{i_1,k,i_3} B_{k,j_2} =$$

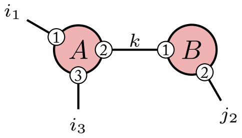

. (1.10)

Another example is the following, where we take a two-index tensor, in other words a matrix, and we contract its indices together, obtaining its trace:

$$M_{i_1,i_2} =$$

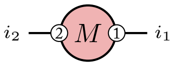

$$\Longrightarrow \quad \operatorname{tr} M = \sum_k M_{k,k} =$$

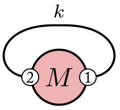

. (1.11)

In these two examples we labelled every leg in order to make it clear which index of the tensor they referred to; usually, however, which leg represents which index is understandable from the context, so we omit almost all labels. With these instructions, the MPS in Eq. (1.8) can be represented as in Fig. 1.1, where site tensors are represented by discs and bond tensors by diamonds.

The internal indices are appropriately contracted as in Eq. (1.8), while external indices remain dangling, as they are not contracted with anything. If we were to represent the initial coefficient tensor $\psi_{i_1,\ldots,i_n}$ as a TN, it would just be a single blob with

n legs: there is no particular structure beyond this. In its MPS form in Fig. 1.1 the external legs are still there, but each of them is now associated to a different tensor, showing how we decomposed the big original tensor in a sequence of smaller parts.

1.1.1 Efficiency of matrix-product states

Now we know how to write a generic pure state as an MPS: what is the point of this construction? Let us assume for the sake of simplicity that all Hilbert spaces $\mathcal{H}_k$ have the same dimension d and that $r_1 = \cdots = r_n = \chi$.[1] Then the MPS, if $n \geq 3$, is made of $2d$ vectors ($\Gamma^{(1,i_1)}$ and $\Gamma^{(n,i_n)}$ for all $i_1, i_n \in \{1, \ldots, d\}$) with χ elements, another n vectors with χ elements (the singular values) and $(n-2)d$ matrices with χ^2 elements, so we have $2d\chi + n\chi + (n-2)d\chi^2$ parameters overall. The original coefficient tensor, instead, had d^n components; the dependence on n in an MPS is apparently linear instead of exponential. Indeed the trick is that the dimension χ actually depends on n, and this dependence can be critical: since each r_k counts the amount of singular values of the SVD in the $(1 \cdots k)(k+1 \cdots n)$ bipartition of the system, also known as the Schmidt rank or number, the dimension χ is directly tied to the maximum entropy of entanglement contained in the state when we look at this particular bipartition: this entropy is given by the von Neumann entropy

$$S_{1\cdots k} = -\sum_{i=1}^{r_k} (\lambda_i^{(k)})^2 \log(\lambda_i^{(k)})^2 \tag{1.12}$$

and, given r_k, the maximal entanglement entropy we can obtain is when the singular values are all equal, $\lambda_i^{(k)} = 1/\sqrt{r_k}$, for which $S_{1\cdots k} = \log r_k$. The entanglement entropy is thus bounded by $\log r_k$. In fact the number r_k of singular values, in the worst-case scenario, increases as $d^{\max\{k,n-k\}}$, so χ may depend *exponentially* on n. It seems thus that we have only shifted the complexity of the state representation elsewhere, hiding it inside χ. But here comes the key observation that makes MPSs useful: since χ is related to the amount of entanglement in the state, it means that weakly entangled states, i.e. states for which χ does *not* increase exponentially in n, but at most polynomially, can be efficiently represented as an MPS whose dimensions scale polynomially in n, which is a great improvement over d^n!

A particularly important class of states that fall into this description is states satisfying the *area law*: in these states, contrary to expectations, the entropy of a subregion

---

[1] After slightly modifying the formula in Eq. (1.8) to allow singular values to be also equal to zero, we can take χ to be the maximum of all r_k's and pad the tensors with zeroes where needed. Orthonormal vectors in the Γ matrices corresponding to null singular values will be arbitrary.

is not extensive, scaling linearly with its volume, but it scales linearly with the area of the boundary of the subregion instead [18]. This behaviour is common for ground states of one-dimensional quantum many-body systems, where an area law holds for gapped Hamiltonians with short-range interactions and away from critical points [63]. In contrast, area laws may be violated if the interactions may be long ranged, and there exist critical lattice models for which a logarithmic divergence is encountered in the entropy. MPSs satisfy an area law by construction, since the entropy of an MPS whose bond dimension does not exceed χ is bounded by $\log \chi$; the question is whether the physical state it should represent satisfies an area law, too, in which case the MPS can describe it faithfully. However, even when the entanglement entropy is bounded by an area law, this fact does not automatically make the MPS representation convenient if we are dealing with time evolution problems. For example, there are tensor-product initial states of one-dimensional spin chains, with short-ranged Hamiltonians, whose entropy $S(t)$ satisfies at any time t an area law, but at the same time $S(t) \geq c_0 t + c_1$ for some $c_0, c_1 > 0$ [64, 65]: this fact is responsible for the hardness of computing the time evolution of many-body systems using some TN algorithms.

This section explained the relationship between dense state vectors, as in the tensor $\psi_{i_1,i_2,\dots,i_n}$, and MPSs. We want to use MPSs for systems with many sites, where it is often not possible to write down the exponentially large state vector in the first place. In that case, we simply start from an MPS description in terms of smaller tensors, and then operate on that MPS with a toolset of algorithms that allow us to transform the state without ever leaving the MPS description. In fact, contracting all the tensors back to a dense vector is usually outright impossible, and keeping the MPS form is a necessity. Luckily, we can obtain all relevant information without having to reconstruct the full state vector: there exist several techniques to manipulate quantum states as MPSs so as to carry out common tasks such as algebraic operations (sums, products), time evolution, sampling from states, measurements, partial traces, and so on [62, 66, 67].

1.1.2 Compression of matrix-product states

Even when there are many singular values, some of them may be less relevant than others. For example, suppose that $\lambda_j^{(k)} \propto \mathrm{e}^{-cj}$: then truncating away the 'tail' would amount to a small approximation of the state. This can be made precise as follows: for a given MPS given by a set of Γ, Λ tensors as in Eq. (1.8), by fixing k and summing all the matrices before and after k we can write the state, for some m, as

$$|\psi\rangle = \sum_{i_1=1}^{d} \cdots \sum_{i_n=1}^{d} \Gamma^{(1,i_1)} \Lambda^{(1)} \cdots \Gamma^{(n,i_n)} |e_{i_1}^{(1)}, \ldots, e_{i_n}^{(n)}\rangle =$$

$$= \sum_{\alpha=1}^{m} \lambda_\alpha^{(k)} |v_{\mathrm{L},\alpha}\rangle \otimes |v_{\mathrm{R},\alpha}\rangle \quad (1.13)$$

where $|v_{\mathrm{L},\alpha}\rangle$ and $|v_{\mathrm{R},\alpha}\rangle$ define a set of orthonormal vectors for $\mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_k$ and $\mathcal{H}_{k+1} \otimes \cdots \otimes \mathcal{H}_n$, respectively. By truncating the sum at $m' < m$ we obtain a new state $|\psi_{m'}\rangle$ such that

$$\left\| |\psi_{m'}\rangle - |\psi\rangle \right\|^2 = \sum_{\alpha=m'+1}^{m} (\lambda_\alpha^{(k)})^2 \quad (1.14)$$

which shows that if the singular values decrease fast enough, then an appropriate truncation of the singular values produces a controlled approximation of the state. This truncation is actually optimal, in the sense that the truncated state is the closest in the 2-norm, among all MPSs with bond dimension m', to the original one [20].[2] To carry out this truncation during the MPS construction we can act on each subsequent SVD decomposition, e.g. at Eqs. (1.3) and (1.5) by not keeping all singular values in the sums, but discarding the smallest ones. The loss of precision caused by the truncated decomposition will be small if the entanglement is weak.

1.1.3 Other MPS forms

There is no unique way to represent a state as an MPS: other forms, also called *gauges*, are commonly used. If we choose a site k, and contract all $\Lambda^{(j)}$'s with $\Gamma^{(j,i_j)}$ for $j < k$ and with $\Gamma^{(j+1,i_{j+1})}$ for $j > k$, as depicted in Fig. 1.2, we obtain another decomposition of the following form:

$$\psi_{i_1,i_2,\ldots,i_n} = A^{(1,i_1)} \cdots A^{(k-1,i_{k-1})} \Gamma^{(k,i_k)} B^{(k+1,i_{k+1})} \cdots B^{(n,i_n)}, \quad (1.15)$$

where $A^{(j,i_j)} := \Gamma^{(j,i_j)} \Lambda^{(j)}$ and $B^{(j,i_j)} := \Lambda^{(j-1)} \Gamma^{(j,i_j)}$.

This more compact form is called the *(mixed) canonical* form, and its component satisfy the nice properties

$$\sum_{i_j=1}^{d} A^{(j,i_j)\dagger} A^{(j,i_j)} = I, \qquad \sum_{i_j=1}^{d} B^{(j,i_j)} B^{(j,i_j)\dagger} = I. \quad (1.16)$$

---

[2] The optimality, however, does not necessarily hold anymore when multiple sequential SVD truncations are applied.

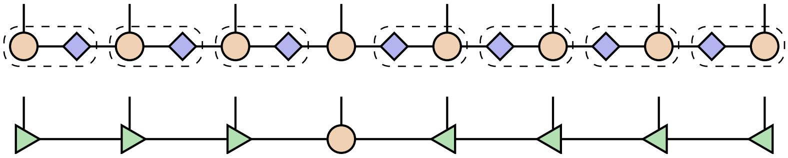

Figure 1.2. We obtain the canonical form of an MPS by merging each bond tensor with the site tensor on the left or on the right.

Tensors satisfying one of these two properties are called *left-* or *right-orthogonal*, respectively. The Γ tensor left over from the original form is called *orthogonality centre* (or orthocentre, more shortly). If the orthocentre is at the last site, all the other tensors are left-orthogonal and the MPS is termed *left-canonical*; if the orthocentre is at the first site then the MPS is *right-canonical*. A left- or right-orthogonal tensor is usually represented by a triangle pointing towards right or left, respectively, as in Fig. 1.2 (they point towards the orthogonality centre).

Another interesting gauge is the *inverse canonical* gauge introduced in Ref. [68]. In contrast to the Vidal form, where in the MPS each bond tensor, which is a diagonal matrix containing the Schmidt decomposition weights at that bond, is an orthogonality centre, in this gauge each *site* tensor is an orthogonality centre. Each bond tensor is a diagonal matrix containing the inverse of the Schmidt weights.

Yet other ways of grouping the tensors of the Vidal form exist, and of course in the end they all represent the same pure state in the Hilbert space: it is just a matter of convenience which form is the right one, a choice that is sometimes dictated by the specific algorithms used to transform the state. We will come back to this in Section 1.3.

1.1.4 Notes on fermionic systems

There is an additional step we need to consider when we want to represent fermionic states as MPSs, due to the antisymmetric structure of the Fock space they live in. In this case, the modes (i.e. the canonical creation and annihilation operators) representing different degrees of freedom anticommute and consequently, for example, an observable A which *physically* acts only on the first mode cannot be written as $A = A_1 \otimes \mathbb{1}_2 \otimes \cdots$, and similarly the global Hilbert space is not the tensor product of the Hilbert spaces relative to the individual modes.

The numerical simulation of fermions on a (classical) computer typically requires, therefore, 'translating' the anticommuting structure of fermionic Fock spaces into the structure of a ‘standard’ multi-mode Hilbert space in which different systems are combined through the tensor product, as it corresponds to standard matrices and vectors used in computer programming. There exist several of such translation maps; famous examples are the Jordan-Wigner [69] and the Bravyi-Kitaev [70] mappings. We shall use the former, which is a bijection that maps the algebra of creation and annihilation operators of N fermions, $\{c_k^{(\dagger)}\}_{k=1}^N$ in $\mathcal{F}_-(\mathbb{C}^N)$, into the algebra of N two-level systems, in $(\mathbb{C}^2)^{\otimes N}$, as

$$c_k = \sigma_z^{\otimes k-1} \otimes \sigma_- \otimes \mathbb{1}^{\otimes N-k}, \qquad c_k^\dagger = \sigma_z^{\otimes k-1} \otimes \sigma_+ \otimes \mathbb{1}^{\otimes N-k}, \tag{1.17}$$

where $\sigma_+ := \frac{1}{2}(\sigma_x + \mathrm{i}\sigma_y)$. The anticommutativity of operators on different modes is taken care of by the strings of σ_z. This transformation maps a fermion chain with nearest-neighbour interactions into a nearest-neighbour spin chain,

$$\begin{aligned}
&\sum_{k=1}^{N} \varepsilon_k c_k^\dagger c_k + \sum_{i=1}^{N-1} \lambda_k (c_k^\dagger c_{k+1} + c_{k+1}^\dagger c_k) = \\
&\quad = \sum_{k=1}^{N} \varepsilon_k \mathbb{1}^{\otimes k-1} \otimes \sigma_+\sigma_- \otimes \mathbb{1}^{\otimes N-k} + \sum_{k=1}^{N-1} \lambda_k \mathbb{1}^{\otimes k-1} \otimes T \otimes \mathbb{1}^{\otimes N-k-1},
\end{aligned} \tag{1.18}$$

with $T = \sigma_+ \otimes \sigma_- + \sigma_- \otimes \sigma_+$, but in general it may introduce artificial long-range couplings, that can make the simulation more demanding.

At the same time, fermionic systems satisfy *superselection rules* [71, 72], that are fundamental restrictions on the set of physically allowed quantum states. The fermionic modes we consider in a quantum theory should in fact be interpreted, as the special theory of relativity states, as particles of half-integer spin. This establishes a connection between the mode operators and the transformation properties under spatial rotations of the corresponding fermionic particles [73]. In particular, the state of a fermionic particle after a 2π rotation (about any axis) gains a minus sign. Consequently, states containing an odd numbers of fermions also gain a minus sign after rotating all the particles, while states with an even number of fermions remain unchanged. Given the parity operator $(-1)^N$, where N is the (total) number operator, we can then identify two classes of states: the ‘even’ and ‘odd’ states, which constitute the eigenspaces of $(-1)^N$. In our language, the state $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, a superposition of an even and an odd state, would be transformed, after a full rotation, into $\frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$. The two states are orthogonal and the fact that a 2π rotation, that is, the identity, can transform one into the other does not make sense from a physical standpoint, since the two states should be indistinguishable. A superselection rule, for this reason, imposes that *no superposition of even and odd states* may occur in physical states, and the fermionic Fock space can therefore be written as the direct sum of the two eigenspaces of the parity operator.

1.2 Operations on matrix-product states

Matrix-product states are useful not only because they can be an efficient way to encode a many-body state. By associating a separate tensor to each space $\mathcal{H}_k$, local operations that involve one or few adjacent sites can be efficiently carried out on the MPS by modifying only the relevant tensors. For example, take a linear map T that acts non trivially on the space $\mathcal{H}_k$, i.e. $T = \mathbb{1}^{\otimes k-1} \otimes t \otimes \mathbb{1}^{\otimes n-k}$, then

$$\begin{aligned} T|\psi\rangle &= \sum_{i_1=1}^{d} \cdots \sum_{i_n=1}^{d} \Gamma^{(1,i_1)}\Lambda^{(1)} \cdots \Lambda^{(n-1)}\Gamma^{(n,i_n)} T|e^{(1)}_{i_1}, \dots, e^{(n)}_{i_n}\rangle = \\ &= \sum_{i_1=1}^{d} \cdots \sum_{i_n=1}^{d} \Gamma^{(1,i_1)}\Lambda^{(1)} \cdots \Lambda^{(k-1)}\Gamma^{(k,i_k)}\Lambda^{(k)} \cdots \\ &\qquad \cdots \Lambda^{(n-1)}\Gamma^{(n,i_n)}|e^{(1)}_{i_1}\rangle \otimes \cdots \otimes t|e^{(k)}_{i_k}\rangle \otimes \cdots \otimes |e^{(n)}_{i_n}\rangle = \\ &= \sum_{i_1=1}^{d} \cdots \sum_{i_n=1}^{d} \sum_{l=1}^{d} t_{l,i_k} \Gamma^{(1,i_1)}\Lambda^{(1)} \cdots \Lambda^{(k-1)}\Gamma^{(k,i_k)}\Lambda^{(k)} \cdots \\ &\qquad\qquad \Lambda^{(n-1)}\Gamma^{(n,i_n)}|e^{(1)}_{i_1}, \dots, e^{(k)}_{l}, \dots, e^{(n)}_{i_n}\rangle \end{aligned} \tag{1.19}$$

with $t_{l,i_k} := \langle e^{(k)}_l | t | e^{(k)}_{i_k}\rangle$; this means that at the MPS level the only change is to replace $\Gamma^{(k,i_k)}$ with

$$\Gamma'^{(k,l)} = \sum_{l=1}^{d} t_{i_k,l}\Gamma^{(k,i_k)}. \tag{1.20}$$

Similarly, a linear operator acting non-trivially only on $\mathcal{H}_k \otimes \mathcal{H}_{k+1}$ acts on the MPS by modifying only the $\Gamma^{(k,i_k)}$, $\Lambda^{(k)}$ and $\Gamma^{(k+1,i_{i+k})}$ matrices, as represented in Fig. 1.3. When we act on an MPS on more than one site, after we multiplied the relevant tensors of the MPS by the tensor representing the operator we are left with a new tensor spanning more than one site, as in the dashed square in Fig. 1.3: in order to restore the original factorised form we also need to decompose the new tensor to obtain smaller blocks of the required shape. This is usually done with a series of SVDs. During this decomposition we have also the opportunity to approximate the result,

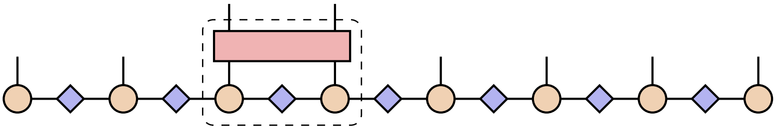

Figure 1.3. A two-site operator acting on an MPS with $n = 8$ sites, in Vidal form.

instead of keeping the exact resulting tensors: once we reveal the singular values from the SVD, we can truncate the less relevant ones, usually in two ways:

- by imposing a cutoff ε such that all singular values below ε get discarded,
- by imposing a maximum bond dimension χ, i.e. keeping at most the χ largest singular values.

This procedure provides a dynamic compression method that allows us to simultaneously update the MPS and at the same time keeping its size under control.

1.2.1 Matrix product operators

The procedure used to decompose a pure state into an MPS can be adapted to decompose a generic *operator* acting on the tensor-product Hilbert space. This time we want to write the tensor of the matrix elements of a linear operator T,

$$\langle e_{i_1}^{(1)}, \dots, e_{i_n}^{(n)} | T | e_{i'_1}^{(1)}, \dots, e_{i'_n}^{(n)} \rangle \tag{1.21}$$

as a product of smaller tensors. The decomposition is basically the same as the one we already seen for a state, if not for the fact that each tensor block has a *pair of physical indices* instead of just one (see Ref. [62] for a more detailed treatment). Then Eq. (1.21) becomes

$$W^{(1,i_1,i'_1)} \dots W^{(n,i_n,i'_n)}, \tag{1.22}$$

and is called a matrix-product operator (MPO).[3] Graphically, we represent an MPO as in Fig. 1.4: since each factor in Eq. (1.22) has two physical indices, each block in the TN will have *two* dangling legs.

---

[3] We will not care about the different gauges we could use to write an MPO, as they are not as meaningful as in the MPS case. The only important thing is that we are able to decompose a many-body operator into a TN with a separate tensor for each degree of freedom.

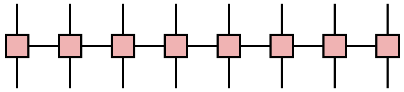

Figure 1.4. Graphical representation of an MPO.

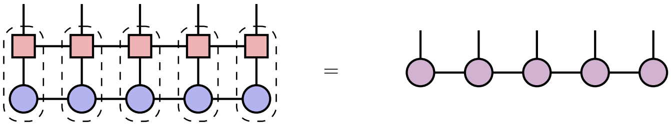

Figure 1.5. An MPO applied to an MPS: the tensors of the MPO and the MPS relative to the same site get contracted together. The bond indices in this operation are not contracted, but they are combined in a new index in the tensor product. The result is a new bond index whose dimension is the product of the dimensions of the two initial indices.

As operators can be applied to pure states to obtain other pure states, an MPO can be applied to an MPS. The result is again an MPS: given an operator T in the MPO form of Eq. (1.22) and a pure state $|\psi\rangle$ written as

$$|\psi\rangle = \sum_{i_1=1}^{d} \cdots \sum_{i_n=1}^{d} M^{(1,i_1)} \cdots M^{(n,i_n)} |e^{(1)}_{i_1}, \ldots, e^{(n)}_{i_n}\rangle, \tag{1.23}$$

the result is a new MPS

$$T|\psi\rangle = \sum_{i_1=1}^{d_1} \cdots \sum_{i_n=1}^{d_n} N^{(1,i_1)} \cdots N^{(n,i_n)} |e^{(1)}_{i_1}, \ldots, e^{(n)}_{i_n}\rangle, \tag{1.24}$$

where

$$N^{(k,i_k)}_{(a_{k-1},b_{k-1}),(a_k,b_k)} := \sum_{i'_k=1}^{d} W^{(k,i_k,i'_k)}_{a_{k-1},a_k} M^{(k,i'_k)}_{b_{k-1},b_k}. \tag{1.25}$$

Figure 1.5 depicts this operation in TN graphical notation.

It is an MPO whose ranks are the product of the ranks of the factors: the matrices W and M get multiplied contracting their physical indices, but on the bond indices we are performing a tensor product. As a consequence, multiplying MPOs and MPSs usually results in a significant (multiplicative) increase in the size of the tensors; for this reason, after the multiplication the resulting MPS must usually be compressed. Other

MPO-MPS multiplication algorithms exist which are more efficient than this apply-then-truncate method: for example the 'zip-up' and the 'density matrix' algorithms, explained in Ref. [74] perform multiplication and compression at the same time, thus avoiding the need to construct an intermediate large MPS.

1.2.2 Contractions and inner products

The calculation of expectation values of local operators is an important example that shows the power of the MPS representation: starting from the MPS form of a pure state $|\psi\rangle$

$$|\psi\rangle = \sum_{i_1=1}^{d} \cdots \sum_{i_n=1}^{d} M^{(1,i_1)} \dots M^{(n,i_n)} |e^{(1)}_{i_1}, \dots, e^{(n)}_{i_n}\rangle \tag{1.26}$$

and an MPO

$$A = \sum_{i_1=1}^{d} \sum_{i'_1=1}^{d} \cdots \sum_{i_n=1}^{d} \sum_{i'_n=1}^{d} W^{(1,i_1,i'_1)} \dots W^{(n,i_n,i'_n)} |e^{(1)}_{i_1}\rangle\langle e^{(1)}_{i'_1}| \otimes \cdots \otimes |e^{(n)}_{i_n}\rangle\langle e^{(n)}_{i'_n}| \tag{1.27}$$

the expectation value $\langle\psi|A|\psi\rangle$ is given by

$$\begin{aligned} \langle\psi|A|\psi\rangle &= \sum_{\mathbf{i},\mathbf{i}'} \overline{M^{(1,i_1)}} \dots \overline{M^{(n,i_n)}} W^{(1,i_1,i'_1)} \dots W^{(n,i_n,i'_n)} M^{(1,i'_1)} \dots M^{(n,i'_n)} = \\ &= \sum_{\mathbf{i},\mathbf{i}'} \sum_{\mathbf{a},\mathbf{a}',\mathbf{a}''} \overline{M^{(1,i_1)}_{a_1}} W^{(1,i_1,i'_1)}_{a'_1} M^{(1,i'_1)}_{a''_1} \overline{M^{(2,i_2)}_{a_1,a_2}} W^{(2,i_2,i'_2)}_{a'_1,a'_2} M^{(2,i'_2)}_{a''_1,a''_2} \cdots \\ &\qquad \overline{M^{(n,i_n)}_{a_{n-1}}} W^{(n,i_n,i'_n)}_{a'_{n-1}} M^{(n,i'_n)}_{a''_{n-1}} \end{aligned} \tag{1.28}$$

The mathematical structure of this complex expression is made much clearer by the graphical notation for TNs. A graphical example of this calculation can be found in the upper left of Figure 1.6.

Note that the TN has no dangling external indices, corresponding to the fact that the result is a scalar, i.e. a rank-zero tensor. While the order in which we multiply the tensors here does not of course affect the final result, it does affect the computational cost of the whole calculation: in fact, let us consider the worst and best cases, assuming a uniform local dimension d and bond dimensions χ and ξ throughout the MPS and the MPO, respectively. The contraction of a tensor has a computational cost equal to the product of the dimensions of all its indices, both the internal (that are summed over) and the external ones.

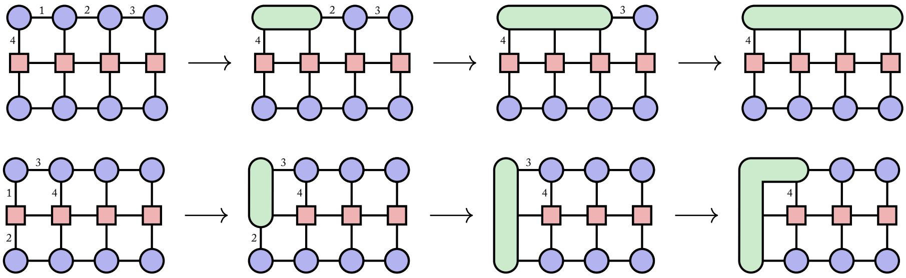

Figure 1.6. Two contraction paths (first steps) for computing the expectation value of an MPO on an MPS. The first leads to a calculation whose complexity grows exponentially in the number of sites of the MPS, the other does not.

- If we start the contraction from the top row (as illustrated in the top figure in Fig. 1.6), then the first operation would cost $O(d^2\chi)$, the second $O(d^3\chi)$, up until $O(d^n\chi)$ when the full upper MPS is contracted. This calculation is massively inefficient, and practically impossible: we are essentially computing the dense state vector from the MPS, which is exactly what we want to avoid by using the MPS representation.

- If we contract 'vertically' instead, i.e. first all the tensors belonging to the first site, then the second, and so on, the calculations are lighter. Following the order written in the bottom part of Fig. 1.6, the first contraction is $O(d^2\chi\xi)$, the second is $O(d\chi^2\xi)$, the third is $O(d\chi^3\xi)$, and so on: the complexity never grows exponentially in n. The final result is the same, but this method is smarter and, most importantly, actually feasible.

For simpler operators which act on a few number of neighbouring sites, the MPS structure allows for an even greater simplification. Consider the case of a single-site operator, and write the state as a mixed-canonical MPS with the orthocentre on the same site as the operator, as illustrated in Fig. 1.7. This case can be obtained by the previous, more general one, by replacing all but one tensor of the MPO with identities. Consider the first contraction, from the left: contracting the physical indices i_1 and i_1' we get

$$\sum_{i_1=1}^{d}\sum_{i_1'=1}^{d} \overline{M_{a_1}^{(1,i_1)}}\, W_{a_1'}^{(1,i_1,i_1')}\, M_{a_1''}^{(1,i_1')} = \sum_{i_1=1}^{d}\sum_{i_1'=1}^{d} \overline{M_{a_1}^{(1,i_1)}}\, \delta_{i_1,i_1'}\, M_{a_1''}^{(1,i_1')} =$$

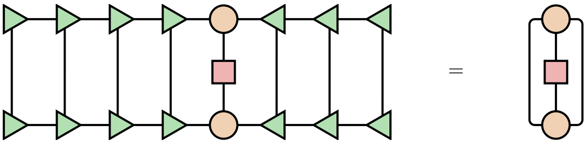

Figure 1.7. Tensor network of the expectation value of a one-site operator acting on an MPS with $n = 8$ sites. Due to the orthogonality properties of the individual tensor, the network to the left simplifies to the one to the right.

$$= \sum_{i_1=1}^{d} \overline{M^{(1,i_1)}_{a_1}} M^{(1,i_1)}_{a''_1} \tag{1.29}$$

which simplifies to δ_{a_1,a''_1} because of the left-orthogonality of the $M^{(1)}$ tensor (see Eq. (1.16)). At the next step, the same thing happens:

$$\begin{aligned}
&\sum_{i_1,i_2} \sum_{i'_1,i'_2} \sum_{a_1,a'_1,a''_1} \overline{M^{(1,i_1)}_{a_1}} W^{(1,i_1,i'_1)}_{a'_1} M^{(1,i'_1)}_{a''_1} \overline{M^{(2,i_2)}_{a_1,a_2}} W^{(2,i_2,i'_2)}_{a'_1,a'_2} M^{(2,i'_2)}_{a''_1,a''_2} = \\
&= \sum_{i_2,i'_2} \sum_{a_1,a''_1} \delta_{a_1,a''_1} \overline{M^{(2,i_2)}_{a_1,a_2}} \delta_{i_2,i'_2} M^{(2,i'_2)}_{a''_1,a''_2} = \\
&= \sum_{i_2} \sum_{a_1} \overline{M^{(2,i_2)}_{a_1,a_2}} M^{(2,i_2)}_{a_1,a''_2} = \\
&= \sum_{i_2} \sum_{a_1} (M^{(2,i_2)})^{\dagger}_{a_2,a_1} M^{(2,i_2)}_{a_1,a''_2} = \\
&= \delta_{a_2,a''_2}
\end{aligned} \tag{1.30}$$

again, and so on until we get to the site where the operator acts. The left-orthogonality property of the tensors is such that we do not need to compute their contraction as in these last equations, we already know that they will give an identity over the bond indices. The same of course happens on the right side of the operator. Consequently, the contraction for the expectation value, however many sites the MPS is defined on, reduces automatically to the contraction of just three tensors, as depicted in Fig. 1.7.

These are classic examples that illustrate how the *contracting order* in a TN is fundamentally important: it determines the number of operations needed to compute the result, and this number can greatly vary between different orders. For a general tensor network, finding the best order is a difficult task: the optimal contraction ordering problem is known to be NP-hard [75]. In special cases, the geometry of the

TNs or the even the value of the tensors (as in the previous example with left- and right-orthogonal tensors) make it possible to always find an efficient contraction order. This is the case for MPSs, for which there exist provably efficient contraction schemes for all common operations, such as the computation of expectation values we have just seen. For this reason the MPS structure is used to encode quantum systems even when their physical structure does not resemble a one-dimensional system: the great number of available algorithms and always knowing what is the best contraction order makes them a powerful tool also where they are not an obvious choice.

1.3 Time evolution of matrix-product states

There exist several algorithms to compute the time evolution of a state as an MPS, allowing us to solve e.g. the Schrödinger equation without leaving the TN formulation. Ref. [66] reviews and compares time-evolution algorithms commonly used for one-dimensional many-body states encoded in MPSs. In this section we will focus on the time-evolving block decimation (TEBD) and the time-dependent variational principle (TDVP) algorithms.

1.3.1 The Suzuki-Trotter decomposition

Typically the Hamiltonian generating the time evolution contains interaction terms which do not commute with one another: even in a very simple example, a nearest-neighbour chain

$$H = \sum_{i=1}^{n} \omega_i a_i^\dagger a_i + \sum_{i=1}^{n-1} g_i (a_i^\dagger a_{i+1} + a_{i+1}^\dagger a_i), \tag{1.31}$$

individual terms in the sums do not commute with one another, in general. Since writing the full time-evolution operator $\exp(-\mathrm{i}tH)$ as an MPO would be impossible but for very small systems, one usually wishes to express it in terms of the exponentials of smaller parts, such as the individual terms in the Hamiltonian. Unfortunately transforming the exponential of a sum in a product of 'smaller' exponentials is not straightforward because, since the addends of the sum do not commute, the correct formula to use is

$$\begin{aligned} &\exp(X+Y) = \\ &= \exp(X)\exp(Y)\exp\big(-\tfrac{1}{2}[X,Y]\big)\exp\big(\tfrac{1}{6}(2[Y,[X,Y]]+[X,[X,Y]])\big)\cdots \end{aligned} \tag{1.32}$$

known as Zassenhaus' formula [76]. For nearest-neighbour Hamiltonians such as the one in Eq. (1.31) we can identify two particular non-commuting terms whose sum is the Hamiltonian. If we write H as a sum of two-site operators, $H = \sum_{i=1}^{n-1} h_{i,i+1}$, then all the *odd* terms, i.e. $h_{i,i+1}$ operators with odd i, commute with one another; at the same time, *even* terms, where i is even, do the same.[4] It is natural then to split the Hamiltonian as

$$H = H_{\text{odd}} + H_{\text{even}} = \sum_{i=1}^{\lfloor n/2 \rfloor} h_{2i-1,2i} + \sum_{i=1}^{\lfloor n/2 \rfloor} h_{2i,2i+1} \tag{1.33}$$

so that Eq. (1.32) could be used e.g. with $X = H_{\text{odd}}$ and $Y = H_{\text{even}}$. The Suzuki-Trotter decomposition [77, 78] allows us to approximate the calculation of $\exp(X + Y)$ with a finite number of products knowing the order of the neglected terms. For example, the first two orders of approximation are

$$\begin{aligned} \exp(tX + tY) &= \exp(tX)\exp(tY) + O(t^2), \\ \exp(tX + tY) &= \exp(\tfrac{1}{2}tX)\exp(tY)\exp(\tfrac{1}{2}tX) + O(t^3). \end{aligned} \tag{1.34}$$

We usually want to apply this decomposition to a time-evolution operator $U(t) = \exp(-\mathrm{i}tH)$: once we have factorised it in terms of $U_{\text{odd}}(t) = \exp(-\mathrm{i}tH_{\text{odd}})$ and $U_{\text{even}}(t) = \exp(-\mathrm{i}tH_{\text{even}})$, we can further decompose each of these two as the products of the single $\exp(-\mathrm{i}th_{i,i+1})$, as their addends all commute. To improve the accuracy of this decomposition, it is beneficial to reduce t as much as possible: this is usually done by 'slicing' the evolution interval $[0, t]$ in N smaller intervals so that $U(t) = U(t/N)^N$, so that we need to apply the operator $U(t/N)$ to the state N times; since t/N is small, we can apply the decomposition to each $U(t/N)$, with a reduced error, at each step of the evolution.

1.3.2 The time-evolving block-decimation algorithm

Introduced in Ref. [22], the TEBD algorithm is a numerical integration scheme for nearest-neighbour Hamiltonians, which is based on MPSs and heavily uses the Suzuki-Trotter decomposition. The algorithm computes the evolved state $|\psi_t\rangle = U(t)|\psi_0\rangle$ with the following method.

---

[4] Two-site operators are associated naturally to a unique $h_{i,i+1}$, while we are free to associate a single-site operator e.g. on site i to $h_{i,i+1}$ or to $h_{i-1,i}$, unless of course it lies on the edge of the system. It is also possible to split it in both terms, i.e. assigning one half of it to $h_{i,i+1}$ and the other half to $h_{i-1,i}$; this is the most common choice.

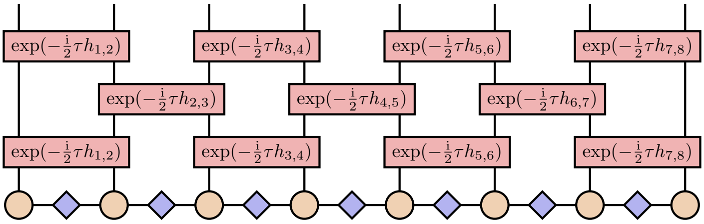

Figure 1.8. A step $|\psi_t\rangle \mapsto \exp(-i\tau H)|\psi_t\rangle$, with $H = \sum_{k=1}^{n-1} h_{k,k+1}$, of the TEBD algorithm for $n = 8$ sites, with a second-order Suzuki-Trotter approximation. The tensors within each odd or even layer, which commute with one another, are depicted on the same row to indicate that they can be applied in parallel.

1. Write the evolution operator $U(t)$ as a product of N smaller evolution steps $U(t/N)$.
2. Write the initial state $|\psi_0\rangle$ in MPS form.
3. Decompose each step $U(t/N)$ in a sequence of U_{odd} and U_{even} according to the Suzuki-Trotter decomposition, at a chosen approximation order.
4. Factorise U_{odd}, and similarly U_{even}, exactly as (with $\tau = t/N$)

$$U_{\text{odd}}(\tau) = \bigotimes_{i=1}^{\lfloor n/2 \rfloor} \exp(-i\tau h_{2i-1,2i}). \tag{1.35}$$

5. Calculate each individual $\exp(-i\tau h_{i,i+1})$; this can be done exactly, as the size of these operators is small.
6. For each time step, apply U_{odd} and U_{even} by applying each $\exp(-i\tau h_{i,i+1})$ to the state in the appropriate order to recreate the chosen approximation of U.

For example, choosing the second decomposition in Eq. (1.34) would mean evolving the state as

$$|\psi_{t+\tau}\rangle = \exp\left(-\tfrac{i}{2}\tau H_{\text{odd}}\right) \exp\left(-i\tau H_{\text{even}}\right) \exp\left(-\tfrac{i}{2}\tau H_{\text{odd}}\right)|\psi_t\rangle, \tag{1.36}$$

as depicted in Fig. 1.8.

From Fig. 1.8 it is even more clear that within an odd or even layer of operators the tensors act on different blocks of the MPS: this not only means that the order in which they are applied does not matter, but also that they can be applied together *in parallel* to the MPS.

Error sources The TEBD algorithm suffers from two errors, both of which can be controlled and estimated straightforwardly. The first one comes from the Suzuki-Trotter approximation: for a m-th order decomposition $U(\tau) \approx U_m(\tau)$ the error is $O(\tau^{m+1})$. By dividing the full simulation interval $[0, t]$ in N slices we obtain

$$\begin{aligned} U(t) = &\left[U_m\left(\frac{t}{N}\right) + O\left(\frac{t^{m+1}}{N^{m+1}}\right)\right]^N = \\ &= U_m\left(\frac{t}{N}\right)^N + O\left(\frac{t^{m+1}}{N^m}\right) = U_m(\tau)^N + O(t\tau^m) \end{aligned} \tag{1.37}$$

which shows that the error decreases as N and m are increased; at the same time it depends linearly on the total time $t = N\tau$ of the simulation. It can be rigorously bounded as shown in Suzuki's works [77, 78]. The second type of error is due to the compression of the MPS after each layer of odd or even operators is applied, which is inevitable if we want the state to have a realistically manageable size. This truncation error is described by Eq. (1.14), and it accumulates additively as the evolution goes on, step after step [22]. By recording the sum of the discarded singular values at each step we can track and control the the truncation error as the time evolution goes on. This approximation preserves the unitarity of the time evolution as each operator applied to the MPS is still unitary, however expectation values of observables may differ from their exact value if the time step is too large.

When dealing with systems of sizes where we do not have the means to compare to an exact result, we cannot know with certainty that our simulation results make sense. In order to be sure that the result is not affected too much by the errors due to the approximations, a standard strategy is to run the same simulation by gradually increasing the bond dimension and checking that some relevant quantities (usually the expectation values of selected operators) saturate and converge to extrapolated values. This is a heuristic method and the extrapolated value is not meant to be used as an estimate of the true expectation value, but the technique can be an useful way to check whether the simulation has reasonably converged, or a least to determine if it has not.

The TEBD algorithm can be extended to the case of Hamiltonians with non-nearest-neighbour interactions by using *swap operators*, that exchange the physical index of two sites, allowing us to move them to adjacent positions [79]. A generic two-site operator $T_{i,j}$, acting on the (non necessarily adjacent) sites i and j, can then be written as the composition

$$S_{i,i+1}S_{i+1,i+2}\cdots S_{j-2,j-1}T_{j-1,j}S_{j-2,j-1}\cdots S_{i+1,i+2}S_{i,i+1} \tag{1.38}$$

where $S_{i,j}$ is an operator that swaps i and j, and $T_{j-1,j}$ is the same two-body term in the Hamiltonian as $T_{i,j}$, but but acts instead on adjacent sites. With swap operators we

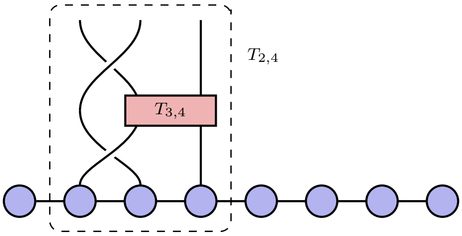

Figure 1.9. Application of a non-nearest-neighbour operator to an MPS with the help of a swap operation. The T operator on sites 2 and 4 can be written as the product of the T operator on sites 3 and 4 and two swap operators.

can write arbitrary two-site operators as a product of nearest-neighbour ones, whose exponential is easier to compute. In general, by reordering the sites within the MPS using these operators, we can change how entanglement is distributed in the state. For example, if we take two strongly correlated sites and move them far apart, we increase the range of correlations in the MPS: this may affect the ability of the MPS to (efficiently) represent the state, even if in principle we are just changing how the sites are ordered. It should be noted, however, that in an expression such as Eq. (1.38) the swap operations are immediately inverted after the application of the actual 'physical' operator $T_{j-1,j}$, in order to restore the original site numbering. For this reason, in this case, the (overall) swap operation does not affect the amount of entanglement in the system, which is instead dictated by the topology of the interactions, regardless of how we apply the operators to the MPS.

1.3.3 The time-dependent variational-principle algorithm

The time-dependent variational principle is an algorithm based on finding the state that best approximates, over time, the solution to Schrödinger's equation within a chosen manifold. Once we choose a manifold $\mathcal{M}$ within the Hilbert space of states, we want to find a time-dependent approximation $|\varphi_t\rangle \in \mathcal{M}$ that best approximates the true solution $|\psi_t\rangle$ of the differential equation [80]. This approximation is obtained from the minimisation of

$$\big\|\partial_t|\varphi_t\rangle + \mathrm{i}H|\varphi_t\rangle\big\|. \tag{1.39}$$

within $\mathcal{M}$. In particular the manifold is chosen to be the set of MPSs of fixed (both local and bond) dimensions. This minimisation leads to the following differential equation on $\mathcal{M}$:

$$\partial_t|\varphi_t\rangle = -\mathrm{i}P_{\varphi_t}\big(H|\varphi_t\rangle\big), \tag{1.40}$$

where P_{φ_t} is the orthogonal projection onto the tangent space of $\mathcal{M}$ at $|\varphi_t\rangle$. This is the differential equation we need to solve, starting from an initial condition $|\varphi_0\rangle \in \mathcal{M}$: by design, the solution will never leave $\mathcal{M}$ and so it can be described as an MPS of constant bond dimensions.

To solve Eq. (1.40) we first need to characterise the tangent space, in order to be able to write the projector. This is a highly non-trivial task. In short, MPS are invariant under the set of transformations

$$A_i \mapsto A_i^g = g_{i-1}^{-1} A_i g_i \tag{1.41}$$

where $g_0 = g_n = \mathbb{1}$ and the other g_is are invertible $\chi \times \chi$ complex matrices.[5] If we were to define tangent vectors at the point $|\varphi_t\rangle \in \mathcal{M}$ in the naive way as the derivatives of each component, we would find a set of vectors spanning the tangent space but which is not linearly independent, because of this 'gauge invariance'. Ref. [81] contains a detailed review of the geometry of MPS manifolds, and what conditions must be imposed in order to have a unique characterisation of the tangent space.[6] In the end, following Ref. [82], the tangent vectors are parameterised as follows: if $|\varphi_A\rangle \in \mathcal{M}$ is determined by the tuple $A = (A^{(1)}, \ldots, A^{(n)})$ of tensors, then a tangent vector to $|\varphi_A\rangle$ is an object $|v_A(B)\rangle$, where $B = (B^{(1)}, \ldots, B^{(n)})$ is a tuple of tensors of the same shape of A, defined as

$$\begin{split}|v_A(B)\rangle = \sum_{k=1}^{n} \sum_{i_1=1}^{d} \cdots \sum_{i_n=1}^{d} A_{\mathrm{L}}^{(1,i_1)} \cdots A_{\mathrm{L}}^{(k-1,i_{k-1})} B^{(k,i_k)} A_{\mathrm{R}}^{(k+1,i_{k+1})} \cdots \\ A_{\mathrm{R}}^{(n,i_n)} |e_{i_1}^{(1)}, \ldots, e_{i_k}^{(k)}, \ldots, e_{i_n}^{(n)}\rangle.\end{split} \tag{1.42}$$

In this expression, the k-th term is in the mixed canonical form with the orthocentre on the k-th site: the subscripts A_{L} and A_{R} denote the fact that the tensors are left- and right-orthogonal, respectively. With this parameterisation, the projector in Eq. (1.40) is

$$P_{\varphi_A} = \sum_{k=1}^{n} \underbrace{P_{\mathrm{L}}^{[1\cdots k-1]} \otimes \mathbb{1} \otimes P_{\mathrm{R}}^{[k+1\cdots n]}}_{P_1^{(k)}} - \sum_{k=1}^{n} \underbrace{P_{\mathrm{L}}^{[1\cdots k]} \otimes P_{\mathrm{R}}^{[k+1\cdots n]}}_{P_0^{(k)}} \tag{1.43}$$

---

[5] After all, we have already seen that there are different but equivalent forms (Vidal, canonical) in which an MPS can be written; Eq. (1.41) is merely a generalisation of this fact.

[6] With some mild restrictions on the MPS and on the action of the group G of gauge transformations, the MPS manifold can be seen as a *principal G-bundle*. Finding a unique parameterisation for every tangent vector becomes then equivalent to specifying a *horizontal subspace*, for which there is no unique prescription.

in which the individual left and right projectors are

$$P_{\mathrm{L}}^{[1\cdots k]} := \sum_{\alpha=1}^{M} |w_{\mathrm{L},\alpha}^{[1\cdots k]}\rangle\langle w_{\mathrm{L},\alpha}^{[1\cdots k]}|, \quad P_{\mathrm{R}}^{[1\cdots k]} := \sum_{\alpha=1}^{M} |w_{\mathrm{R},\alpha}^{[k\cdots N]}\rangle\langle w_{\mathrm{R},\alpha}^{[k\cdots N]}| \tag{1.44}$$

where $|w_{\mathrm{L},\alpha}^{[1\cdots n]}\rangle$ comes from the Schmidt decomposition of $|\varphi_A\rangle$, i.e. from the left-orthogonal matrices, and so on.

Integration scheme In order to solve Eq. (1.40) numerically, we perform a Suzuki-Trotter approximation to first order of the projection

$$P_{\varphi_A} = P_1^{(1)} - P_0^{(1)} + P_1^{(2)} - P_0^{(2)} + \cdots + P_1^{(n-1)} - P_0^{(n-1)} + P_1^{(n)} \tag{1.45}$$

by integrating each term separately, in this order, obtaining a sequence of $2n-1$ differential equations: if we evolve the state from t_0 to t_1, and $|\varphi_{A,0}\rangle$ is the initial condition, they are

$$\begin{aligned}
\mathrm{i}\partial_t|\varphi_A(t)\rangle_1 &= P_1^{(1)}H|\varphi_A(t)\rangle_1, & |\varphi_A(t_0)\rangle_1 &= |\varphi_{A,0}\rangle \\
\mathrm{i}\partial_t|\varphi_A(t)\rangle_2 &= P_0^{(1)}H|\varphi_A(t)\rangle_2, & |\varphi_A(t_0)\rangle_2 &= |\varphi_A(t_1)\rangle_1 \\
\mathrm{i}\partial_t|\varphi_A(t)\rangle_3 &= P_1^{(2)}H|\varphi_A(t)\rangle_3, & |\varphi_A(t_0)\rangle_3 &= |\varphi_A(t_1)\rangle_2 \\
\mathrm{i}\partial_t|\varphi_A(t)\rangle_4 &= P_0^{(2)}H|\varphi_A(t)\rangle_4, & |\varphi_A(t_0)\rangle_4 &= |\varphi_A(t_1)\rangle_3 \\
&\vdots \\
\mathrm{i}\partial_t|\varphi_A(t)\rangle_{2n-3} &= P_1^{(n-1)}H|\varphi_A(t)\rangle_{2n-3}, & |\varphi_A(t_0)\rangle_{2n-3} &= |\varphi_A(t_1)\rangle_{2n-4} \\
\mathrm{i}\partial_t|\varphi_A(t)\rangle_{2n-2} &= P_0^{(n-1)}H|\varphi_A(t)\rangle_{2n-2}, & |\varphi_A(t_0)\rangle_{2n-2} &= |\varphi_A(t_1)\rangle_{2n-3} \\
\mathrm{i}\partial_t|\varphi_A(t)\rangle_{2n-1} &= P_1^{(n)}H|\varphi_A(t)\rangle_{2n-1}, & |\varphi_A(t_0)\rangle_{2n-1} &= |\varphi_A(t_1)\rangle_{2n-2}.
\end{aligned} \tag{1.46}$$

The solution of the last equation is the final approximation for the solution of Eq. (1.40). This is a first-order integrator, meaning that if τ is the time step for the numerical integration then the error is $O(\tau^2)$. Composing it with the adjoint gives a second-order integrator, which means that if we solve Eq. (1.46) with a time step $\tau/2$ and then, starting from the final state, solve the equations again but in reverse order, we obtain an integration scheme with an $O(\tau^3)$ error. This is the scheme which is commonly used in practice: each integration step is composed of a first left-to-right sweep, which starts from a left-canonical MPS and ends with the MPS in the right-canonical form, and a second right-to-left sweep at the end of which the MPS is left-canonical again.

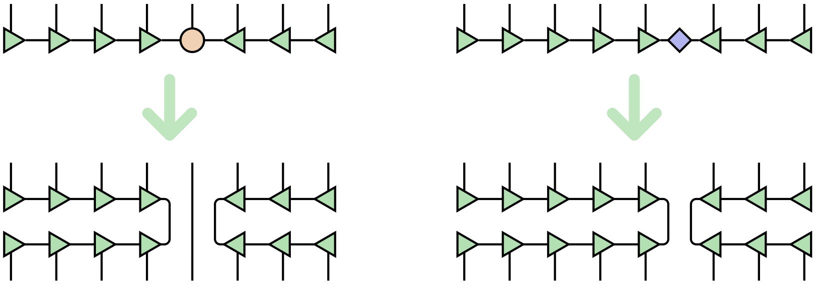

Figure 1.10. Projectors used in the TDVP algorithm. On the left, from an MPS with the orthocentre on its fifth site we obtain $P_1^{(5)}$. On the right, after further decomposing the orthocentre as a product of a left-orthogonal tensor and a bond tensor, we can compute $P_0^{(5)}$.

Each one of the split differential equations in Eq. (1.46) can be solved in closed form. Furthermore, the whole scheme can be implemented very efficiently. The graphical notation helps us navigate this procedure: Fig. 1.10, for example, shows how the projectors $P_1^{(k)}$ and $P_0^{(k)}$ can be obtained from the individual tensors of the MPS, while Fig. 1.11 shows the tensor networks involved in the time evolution of the blocks of the MPS.

We can see that in the whole TN only the orthocentre of the original MPS is modified. On the k-th site, starting from an MPS in the mixed-canonical form (note that the tensors on sites $i < k$ have already been updated to $t + \tau$)

$$A_{\mathrm{L}}^{(1,i_1)}(t+\tau)\cdots A_{\mathrm{L}}^{(k-1,i_{k-1})}(t+\tau)A^{(k,i_k)}(t)A_{\mathrm{R}}^{(n-1,i_{n-1})}(t)\cdots A_{\mathrm{R}}^{(n,i_n)}(t) \quad (1.47)$$

in the required mixed-canonical form, we obtain the solution to the differential equation involving $P_1^{(k)}$ simply by computing

$$A^{(n)}(t+\tau) = \exp(-\mathrm{i}\tau H_{\text{eff}}^{[n]})A^{(n)}(t) \qquad (1.48)$$

where $H_{\text{eff}}^{[n]}$ is the TN in Fig. 1.11 (c). Despite the size of this TN, the solution of Eq. (1.48) can be computed efficiently by caching partial contractions to the left and to the right of the gaps: for example, when moving on to site $n+1$ the contractions on the first n sites from the left are exactly the same, and do not need to be recomputed.

The zero-site projection $P_0^{(k)}$ is integrated in an analogous way; first we decompose $A^{(n)}(t+\tau)$ as $A_{\mathrm{L}}^{(n)}(t+\tau)C^{(n)}(t+\tau)$, obtaining the final left-orthogonal tensor on site n and revealing the bond tensor $C^{(n)}(t+\tau)$ between the site that was just updated

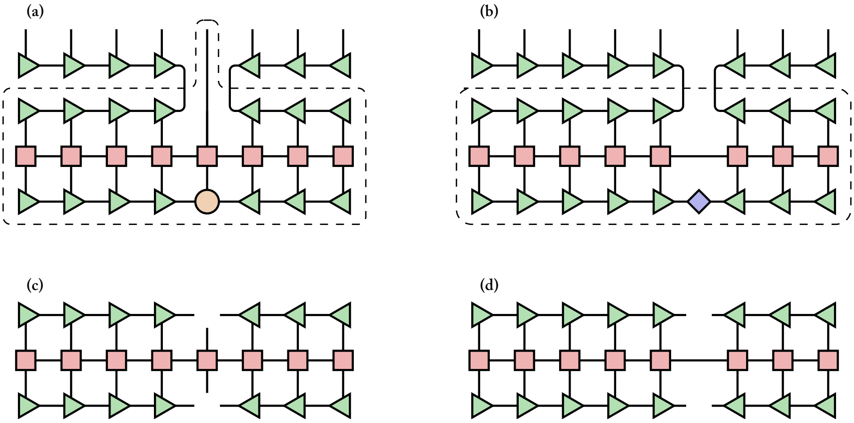

Figure 1.11. (a) TN representing $P_5^{(1)}(H|\psi\rangle)$. The highlighted tensors, when contracted, yield $H_{\text{eff}}^{[5]}|\psi\rangle$, a three-index tensor of the same shape of the orthocentre of the initial MPS. (b) TN representing $P_5^{(0)}(H|\psi\rangle)$. (c) TN representing $H_{\text{eff}}^{[5]}$, the TN we need to exponentiate in order to compute Eq. (1.48). (d) TN representing $K_{\text{eff}}^{[5]}$, the TN we need to exponentiate in order to compute Eq. (1.49).

and the next one. Then, we evolve $C^{(n)}(t+\tau)$ backwards due to the minus sign before the zero-site projector in Eq. (1.43), obtaining

$$C^{(n)}(\tau) = \exp(\mathrm{i}\tau K_{\text{eff}}^{[n]})C^{(n)}(t+\tau) \tag{1.49}$$

where K_{eff} is the TN in Fig. 1.11 (d). Finally, we move the orthocentre to site $n+1$ with $A^{(n+1)} = C^{(n)}A_{\text{R}}^{(n+1)}$ and proceed to the next step.

Error sources Two main sources of error affect the TDVP algorithm. The first one is due to the approximation which turns Eq. (1.40) into the sequence of differential equations in Eq. (1.46), which in the most common case is $O(\tau^3)$. This discretisation error corresponds, however, to an error with respect to the already projected equation, not to the original Schrödinger equation. The other source of error is in fact due to the projection itself, and this error can be bound (see Ref. [82]) by $\|(\mathbb{1} - P_{\varphi_t})H|\varphi_t\rangle\|$.

As noted in Ref. [83], since in the one-site TDVP algorithm the MPS bond dimension is fixed at the beginning, if it is too low then the MPS manifold might not contain the necessary degrees of freedom to represent the time evolution in a satisfactory way.

In some cases, especially with long-range interactions, the state can get stuck in a low-entanglement manifold and may even not evolve at all from the initial value. A way to fix the projection error is to manually expand the manifold before the time evolution, by allocating an MPS which is larger than what is necessary for the initial state, but large enough to accommodate a sufficient representation of the evolved state. This enlargement can be performed by padding the initial MPS with zeroes, or, as suggested in Ref. [83], through a subspace expansion by global Krylov vectors.However, as far as the projection error alone is considered, regardless of the chosen bond dimensions, the resulting evolution will conserve the norm exactly, together with all the quantities conserved by the Hamiltonian of the original problem.As already mentioned for TEBD, in order to judge whether the algorithm has converged we run the same simulation multiple times, each time starting with a larger MPS, until the results converge.

Variants of the algorithm There exist variants of this algorithm that target its main limitations.

- In a two-site variant of the algorithm, presented in Ref. [82] together with the original scheme, a block of *two* successive sites is updated simultaneously. During the SVD of the updated two-site block, a new bond dimension can then be chosen *dynamically*, according to the singular values.
- An *adaptive* variant, introduced in Ref. [84], evaluates the projection error in order to dynamically optimise the bond dimensions before each time step, resulting in an MPS that can dynamically restructure itself as the complexity of the evolution grows over time.
- As TDVP is an intrinsically serial algorithm, it cannot be trivially parallelised as TEBD, where different parts of the MPS can be updated simultaneously. A parallel version of two-site TDVP, based on the inverse canonical gauge, is introduced in Ref. [85]. The update algorithm is broken up into multiple blocks which are then updated simultaneously; depending on the system, this parallel splitting may lead to significant speed-ups, at the cost of some additional errors and instabilities, that can however be worked around.

In conclusion, the TDVP algorithm allows simulating the time evolution using an algorithm that can deal with arbitrary Hamiltonians, even long-range ones (without resorting to swap operations). With respect to TEBD, it is more complex and more costly to use, and cannot be easily parallelised. At the same time, while TEBD has a local approach to the approximations, in which the time evolution operator is decomposed in a product of smaller exponentials that do not ‘see’ the part of the MPS they

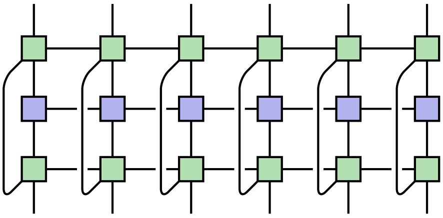

Figure 1.12. TN representation of a quantum channel $\rho \mapsto \sum_k C_k \rho C_k^\dagger$, with each C_k represented as an MPO, applied to another MPO. The additional contracted legs, which are not present in the composition of three MPOs, represent the sum over k.

do not act on, TDVP has a global approach, since the evolution operators in Eq. (1.48) and Eq. (1.49), despite their small size, contain information about all sites of the state MPS and of the Hamiltonian MPO. For this reason, TDVP can work with a greater time step than TEBD while having the same error, which partly alleviates the overhead of the former with respect to the latter.

1.4 Tensor networks for mixed states

So far we have studied the TN machinery for pure states in a Hilbert space $\mathcal{H}$, operators acting onto $\mathcal{H}$, and and how to evolve pure states over time. However, quantum states can also come in the form of density matrices, so we need to find a proper TN representation for them as well. The most obvious way to do this, since density matrices are a special type of operators on $\mathcal{H}$, is to use an MPO, but we soon run into issues when we want to describe quantum channels acting on them. Operations such as $\rho \mapsto A\rho$ or $\rho \mapsto \rho B$ can be immediately written as composition of MPOs, but it gets more difficult for generic channels. For example, if we want to write a single TN that represents a map such as $\rho \mapsto \sum_k C_k \rho C_k^\dagger$ we obtain the layout illustrated in Fig. 1.12, where each C_k is represented, in turn, by an MPO. With additional legs representing the k index in the sum, the TN layout is not as simple as an MPO-MPS product. In this section we will see two methods to solve this problem.

1.4.1 Simple vectorisation

The most direct solution is to treat matrices as vectors. Linear operators on $\mathcal{H}$ form a vector space $L(\mathcal{H})$ over $\mathbb{C}$, so we only need to choose an orthonormal basis for $L(\mathcal{H})$, and write the coefficient of mixed states in this basis as a vector. Moreover, if $\mathcal{H} = \mathcal{H}_1 \otimes \cdots \otimes \mathcal{H}_n$ then $L(\mathcal{H}) \cong L(\mathcal{H}_1) \otimes \cdots \otimes L(\mathcal{H}_n)$ so the tensor-product structure is still present at this level. Channels, in turn, become matrices acting on the coefficient vectors, which makes sums and products much easier. This way, mixed states can be written as MPSs and channels as MPOs. If the local dimension of the Hilbert space in the pure-state case is d, for a mixed-state MPS it will be d^2, which increases the overall size of the tensors, making each operation generally more expensive. Moreover, in addition to entanglement there will also be classical correlations between the components of the system, thus we expect that we will need in general larger bond dimensions to represent our states.

In order to compute expectation values, we recall that $(A, B) \mapsto \operatorname{tr}(A^\dagger B)$ is an inner product on the space of operators: since in finite-dimensional Hilbert spaces every element of $L(\mathcal{H})$ is automatically a Hilbert-Schmidt operator, this inner product always makes sense. By choosing a basis for $L(\mathcal{H})$ which is orthonormal with respect to this inner product, we can then write expectation values like $\operatorname{tr}(A\rho)$ as inner products of two MPSs: one represents ρ and the other represents $A^\dagger$. The trace of the state is given analogously as the inner product between ρ and $\mathbb{1}$, and the MPS representation of the identity is trivial. The identity is also usually already one of the elements chosen for the basis of $L(\mathcal{H})$.

For example, for a mixed state over n qubits a common choice of orthonormal basis is the tensor product of Pauli matrices, i.e. the set of the elements

$$e_{i_1,i_2,\ldots,i_n} := \sigma_{i_1} \otimes \sigma_{i_2} \otimes \cdots \otimes \sigma_{i_n} \tag{1.50}$$

with $i_k \in \{0, 1, 2, 3\}$ and $\sigma_0 = \mathbb{1}$, $\sigma_1 = \sigma_x$, $\sigma_2 = \sigma_y$, $\sigma_3 = \sigma_z$ are Pauli matrices on the individual qubits. The representation of states and channels in this basis is usually called the Pauli transfer matrix (PTM) representation, especially in the field of quantum computing.

Mixed states evolve in time not following Schrödinger's equation, but von Neumann's equation $\dot{\rho}_t = -\mathrm{i}[H, \rho_t]$ or a generalisation such as the Gorini-Kossakowski-Sudarshan-Lindblad (GKSL) equation we will encounter in Section 2.2. In any case, the time-evolution will be written as a linear differential equation

$$\begin{cases} \dot{\rho}_t = \mathcal{L}(\rho_t), \\ \rho_0 = u_0, \end{cases} \tag{1.51}$$

The time-evolution algorithms of Section 1.3 can be readily adapted to this representation, since the evolution map is linear in ρ_t. By writing u_0 and the operators in $\mathcal{L}$ as TNs we can cast this problem into the TEBD or TDVP algorithm, according to the geometry of the interactions described by $\mathcal{L}$.

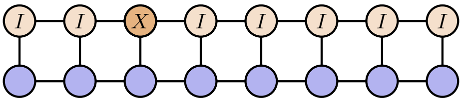

Figure 1.13. Expectation value on a mixed state computed as the Hilbert-Schmidt inner product of MPSs. Here, for example, the observable $\mathbb{1}^{\otimes 2} \otimes \sigma_x \otimes \mathbb{1}^{\otimes 5}$ acts only on site 3, and can be written as an MPS with bond dimension 1.

Some of the useful properties of MPSs do not make immediate sense anymore when they represent mixed states. For example, Eq. (1.12)—and consequently the bond dimension—does not quantify the entanglement entropy of the physical state. Other entanglement measures have been proposed for mixed states, such as the *negativity* [86], but they are not easily accessible from the MPS representation of the mixed state. In this case, Eq. (1.12) describes both classical and quantum correlations in the state, but in a way where we cannot easily tell the two contributions apart. The concept can also be extended to a generic operator, relative to a bipartition of the $L(\mathcal{H}_1) \otimes \cdots \otimes L(\mathcal{H}_n)$ space, where it is called *operator-space entanglement entropy* [87]. Just as in the pure-state case, this quantity directly affects the bond dimension of the TN representing the operator and the efficiency of relative simulation methods.

As expectation values are not computed anymore by 'sandwiching' the observable between two copies of the state like in Fig. 1.7, we cannot use the cancellation properties derived from the left- and right-orthogonality of the tensors, but we need to compute a full contraction as in Fig. 1.13 each time. In most cases, however, the observable is a trivial MPS, and the full contraction is equivalent (on most of the sites) to just picking out the local tensors associated to the basis element corresponding to the identity.

Moreover, the generator $\mathcal{L}$ of the time evolution is in general not anti-Hermitian, therefore we lose TDVP's norm-conservation property we had with pure states.[7] Unfortunately, although $\exp(t\mathcal{L})$ preserves the trace, the TDVP construction does not. Nevertheless, we can use this to our advantage as a tool to check the goodness of the simulation: if we observe that the trace of ρ_t significantly decreases as the time evolution moves forward, we can deduce that the bond dimension of the MPS is not large enough to accommodate a faithful representation of the mixed state.

In the vectorised MPS description, as well as in the original MPO one, truncating the operators as the time evolution goes on generally results in losing the positivity of the state [88]. Even if $\exp(t\mathcal{L})$ is a positive map, therefore, $\exp(t\mathcal{L})(\rho)$ may not be

---

[7] Note however that the TDVP algorithm preserves the norm *of the MPS*, and in this case the MPS norm would correspond to the square root of $\langle \rho, \rho \rangle = \mathrm{tr}(\rho^2)$, i.e. the square root of the purity.

positive semi-definite anymore once we compress it; usually the positivity assumption is simply dropped during the time evolution, which is acceptable as long as the approximation errors remain sufficiently small. This property of the state also cannot be enforced, in an MPO or an MPS such as the ones we are considering now, at the level of the individual tensors, whereas we would like to impose constraints which are *local*, instead of regarding the whole TN. Checking whether an MPO or an MPS represents a positive semi-definite operator, in fact, is computationally intractable [89].

For all of these reasons, while dealing with mixed states as MPSs with the simple vectorisation strategy we just outlined is in some cases simple and effective, in general it is technically more difficult.

A more mathematically sound approach is to employ a purification technique, representing the mixed state as a pure quantum state in a higher-dimensional Hilbert space. The larger pure state is evolved in time as any normal pure state, appropriately adapting Eq. (1.51), and the original mixed state together with all the physical quantities can then be recovered by taking the partial trace over the additional degrees of freedom.

1.4.2 The superfermion formalism

If the system is made of bosons or fermions the symmetry properties of the Fock space get lost, at best hidden, in the MPS representation. In these cases, the *superboson/superfermion formalism* is a more appropriate tool that leaves the symmetry of the space clearly visible. Introduced in [90], it has been used e.g. in [39, 91, 92] to study non-equilibrium steady states (NESSs) of strongly interacting quantum systems. Within this section we drop the braket formalism, so as to keep the notation lighter, as we will need to deal with several Hilbert spaces at once.

To begin introducing this formalism we need the concept of conjugate Hilbert space. Given a Hilbert space $\mathcal{H}$ with inner product $\langle -, - \rangle_{\mathcal{H}}$, its conjugate space $\overline{\mathcal{H}}$ is made from the same bare set of vectors of $\mathcal{H}$, but with the scalar multiplication $(\lambda, v) \mapsto \overline{\lambda} v$ and the inner product

$$\langle u, v \rangle_{\overline{\mathcal{H}}} = \overline{\langle u, v \rangle_{\mathcal{H}}}. \tag{1.52}$$

It is customary to write $\overline{v}$ for the vector v when regarded as a vector in $\overline{\mathcal{H}}$, so that $\overline{\lambda v} = \overline{\lambda}\overline{v}$. The conjugate space $\overline{\mathcal{H}}$ is isomorphic to the dual of $\mathcal{H}$, i.e. the space of bounded functionals from $\mathcal{H}$ to $\mathbb{C}$. In other words, $\overline{\mathcal{H}}$ is the space of bras to $\mathcal{H}$'s kets.

The conjugate $\overline{A}$ of a bounded operator $A \in B(\mathcal{H})$ is a bounded operator on $\overline{\mathcal{H}}$, and $\overline{A}\overline{v} = \overline{Av}$. Finally, the space $B_2(\mathcal{H})$ of Hilbert-Schmidt functions from $\mathcal{H}$ to itself is isomorphic to $\mathcal{H} \otimes \overline{\mathcal{H}}$ [93]: this is the key fact we will use to turn operators into states.

In this section, in order not to confuse the original operators with their vectorised counterparts, we explicitly denote the isomorphism between $B_2(\mathscr{H}) \equiv B_2(\mathscr{H}, \mathscr{H})$ and $\mathscr{H} \otimes \overline{\mathscr{H}}$ with V. Given an orthonormal basis $(e_n)_{n\in I}$ of $\mathscr{H}$ the products $(e_m \otimes \overline{e_n})_{m,n\in I}$ form a basis of $\mathscr{H} \otimes \overline{\mathscr{H}}$; for any state ρ on $\mathscr{H}$ we can write

$$V(\rho) = \sum_{m\in I} \sum_{n\in I} \rho_{mn} e_m \otimes \overline{e_n}. \tag{1.53}$$

If $\mathscr{H}$ has a finite dimension, then we can write this also as

$$V(\rho) = (\rho \otimes \mathbb{1}_{\overline{\mathscr{H}}}) \sum_{n\in I} e_n \otimes \overline{e_n} = (\mathbb{1}_{\mathscr{H}} \otimes \bar{\rho}) \sum_{n\in I} e_n \otimes \overline{e_n} \tag{1.54}$$

in which the sum can be considered as $V(\mathbb{1}_{\mathscr{H}})$ (the sum does not converge if $\dim \mathscr{H} = +\infty$). In this case we can also use this formula to compute the trace of operators, via the inner product on $\mathscr{H} \otimes \overline{\mathscr{H}}$:

$$\operatorname{tr}(A\rho) = \langle \mathbb{1}_{\mathscr{H}}, A\rho \rangle_{B_2(\mathscr{H})} = \langle V(\mathbb{1}_{\mathscr{H}}), (A \otimes \mathbb{1}_{\overline{\mathscr{H}}}) V(\rho) \rangle_{\mathscr{H}\otimes\overline{\mathscr{H}}}. \tag{1.55}$$

Left- and right-multiplication become

$$V(A\rho) = (A \otimes \mathbb{1}_{\overline{\mathscr{H}}}) V(\rho), \qquad V(\rho A) = (\mathbb{1}_{\mathscr{H}} \otimes \overline{A^\dagger}) V(\rho). \tag{1.56}$$

With these instructions, we can write for example von Neumann's equation in vectorised form with

$$V([H, \rho]) = (H \otimes \mathbb{1}_{\overline{\mathscr{H}}} - \mathbb{1}_{\mathscr{H}} \otimes \overline{H}) V(\rho). \tag{1.57}$$

while a simultaneous left- and right- multiplication becomes $V(A\rho B) = A \otimes \overline{B^\dagger}\, V(\rho)$.

Let us specialise to the case where $\mathscr{H}$ is a symmetric or antisymmetric Fock space. In the first case, if we consider the creation and annihilation operators $a_i^{(\dagger)}$ associated to the orthonormal basis, we get

$$\begin{aligned}
&[a_i \otimes \mathbb{1}_{\overline{\mathscr{H}}}, a_j^\dagger \otimes \mathbb{1}_{\overline{\mathscr{H}}}] = \delta_{ij}\, \mathbb{1}_{\mathscr{H}} \otimes \mathbb{1}_{\overline{\mathscr{H}}}, \\
&[\mathbb{1}_{\mathscr{H}} \otimes \overline{a_i}, \mathbb{1}_{\mathscr{H}} \otimes \overline{a_j^\dagger}] = \delta_{ij}\, \mathbb{1}_{\mathscr{H}} \otimes \mathbb{1}_{\overline{\mathscr{H}}}, \\
&[\mathbb{1}_{\mathscr{H}} \otimes \overline{a_i}, a_j^\dagger \otimes \mathbb{1}_{\overline{\mathscr{H}}}] = 0,
\end{aligned} \tag{1.58}$$

and so on, which show that $a_i^{(\dagger)} \otimes \mathbb{1}_{\overline{\mathscr{H}}}$ and $\mathbb{1}_{\mathscr{H}} \otimes \overline{a_i^{(\dagger)}}$ may be regarded as proper creation and annihilation operators on a larger symmetric Fock space, in which to each 'physical' mode we associate two 'mathematical' modes. The same construction fails for fermions, since for example

$$\{a_i \otimes \mathbb{1}_{\overline{\mathcal{H}}}, \mathbb{1}_{\mathcal{H}} \otimes \overline{a_j}\} = 2a_i \otimes \overline{a_j}$$

instead of zero, so we cannot use the new operators to describe fermionic modes.

In order to fix this issue, we tweak the definition as done in Ref. [90]. To keep the derivation simple, we consider spin-$\frac{1}{2}$ fermions and omit the subscripts indicating the spaces. In the antisymmetric N-mode Fock space $\mathcal{F}_-$ over $\mathcal{H}$, let e_0 be the vacuum state and define e_n for $n = 1, 2, \ldots, 2^N - 1$ as

$$e_n = e_{(n_N,\ldots,n_2,n_1)} := (a_N^\dagger)^{n_N} \cdots (a_2^\dagger)^{n_2}(a_1^\dagger)^{n_1}|\emptyset\rangle. \tag{1.59}$$

by associating to n the multi-index $(n_N, \ldots, n_2, n_1) \in 2^N$ made from n's digits in base 2, such that $n = \sum_{k=1}^N n_k 2^{k-1}$, e.g. $1 = (0, \ldots, 0, 1)$, $4 = (0, \ldots, 0, 1, 0, 0)$ and so on. We have, for example, $e_1 = a_1^\dagger e_0$ and $e_3 = a_2^\dagger a_1^\dagger e_0$. From a multi-index n we define its 'size' $|n| := \sum_{i=1}^N n_i$, which corresponds to the number of creation operators that need to be applied to e_0 to create e_n, i.e. the number of particles in the state. The sequence $(e_n)_{n=0}^{2^N-1}$ is an orthonormal basis for $\mathcal{F}_-$, and therefore the elements $f_{m,n} := e_m \otimes \overline{e_n}$ form an orthonormal basis of $\mathcal{F}_- \otimes \overline{\mathcal{F}_-}$. We define new creation and annihilation operators associated to this basis: for $k \in \{1, \ldots, N\}$, we set

$$A_k^\dagger f_{m,n} := a_k^\dagger e_m \otimes \overline{e_n}, \qquad \tilde{A}_k^\dagger f_{m,n} := (-1)^{|m|-|n|} e_m \otimes \overline{a_k^\dagger e_n} \tag{1.60}$$

and from the inner product

$$\begin{aligned}
\langle f_{m,n}, A_k^\dagger f_{m',n'}\rangle &= \langle e_m \otimes \overline{e_n}, a_k^\dagger e_{m'} \otimes \overline{e_{n'}}\rangle = \\
&= \langle e_m, a_k^\dagger e_{m'}\rangle \langle \overline{e_n}, \overline{e_{n'}}\rangle = \\
&= \langle a_k e_m, e_{m'}\rangle \langle \overline{e_n}, \overline{e_{n'}}\rangle = \\
&= \langle a_k e_m \otimes \overline{e_n}, e_{m'} \otimes \overline{e_{n'}}\rangle
\end{aligned} \tag{1.61}$$

from which we read $A_k f_{m,n} := a_k e_m \otimes \overline{e_n}$. Analogously,

$$\begin{aligned}
\langle f_{m,n}, \tilde{A}_k^\dagger f_{m',n'}\rangle &= (-1)^{|m'|-|n'|}\langle e_m \otimes \overline{e_n}, e_{m'} \otimes \overline{a_k^\dagger e_{n'}}\rangle = \\
&= (-1)^{|m'|-|n'|}\langle e_m, e_{m'}\rangle \langle \overline{a_k e_n}, \overline{e_{n'}}\rangle = \\
&= (-1)^{|m'|-|n'|}\langle e_m \otimes \overline{a_k e_n}, e_{m'} \otimes \overline{e_{n'}}\rangle,
\end{aligned} \tag{1.62}$$

but in order to complete the definition we need to link m' and n' to m and n. To do so, we observe that since $a_k^\dagger$ adds a particle to the state, it must be that $|m| = |m'|$

and $|n| = |n'| + 1$, otherwise the states in the inner product would have different particle numbers, and would therefore be orthogonal. Consequently, we find $\tilde{A}_k f_{m,n} = (-1)^{|m|-|n|+1} e_m \otimes \overline{a_k e_n}$.

With this definition we can rewrite the action of pairs of these creation and annihilation operators on the orthonormal basis $f_{n,m}$: if the act on the left,

$$a_k^\dagger a_l e_m \otimes \overline{e_n} = A_k^\dagger A_l f_{m,n}, \tag{1.63}$$

$$a_k a_l^\dagger e_m \otimes \overline{e_n} = A_k A_l^\dagger f_{m,n}, \tag{1.64}$$

whereas if they act on the right

$$\begin{aligned} e_m \otimes \overline{a_k^\dagger a_l e_n} &= (-1)^{|m|-|n|+1} \tilde{A}_k^\dagger (e_m \otimes \overline{a_l e_n}) = \\ &= (-1)^{|m|-|n|+1} (-1)^{|m|-|n|+1} \tilde{A}_k^\dagger \tilde{A}_l f_{m,n} = \\ &= \tilde{A}_k^\dagger \tilde{A}_l f_{m,n}, \end{aligned} \tag{1.65}$$

(when we move $a_k^\dagger$ out from the factor on the right, the 'particle count' in the exponent of -1 must take into account that a_l annihilates a particle and therefore decreases $|n|$ by 1)

$$\begin{aligned} e_m \otimes \overline{a_k a_l^\dagger e_n} &= (-1)^{|m|-|n|-2} \tilde{A}_k (e_m \otimes \overline{a_l^\dagger e_n}) = \\ &= (-1)^{|m|-|n|-2} (-1)^{|m|-|n|} \tilde{A}_k \tilde{A}_l^\dagger f_{m,n} = \\ &= \tilde{A}_k \tilde{A}_l^\dagger f_{m,n}. \end{aligned} \tag{1.66}$$

It is a little more complicated for mixed terms, since in

$$\begin{aligned} a_k e_m \otimes \overline{a_l e_n} &= (-1)^{|m|-|n|+1} A_k \tilde{A}_l f_{m,n}, \\ a_k^\dagger e_m \otimes \overline{a_l^\dagger e_n} &= (-1)^{|m|-|n|} A_k^\dagger \tilde{A}_l^\dagger f_{m,n} \end{aligned} \tag{1.67}$$

the sign does not get automatically cancelled out.[8] Anyway, from these results it is an easy exercise to check that the A and $\tilde{A}$ operators actually constitute a new set of fermionic operators, meaning that

$$\begin{gathered} \{A_k, A_l\} = \{\tilde{A}_k, \tilde{A}_l\} = 0, \\ \{A_k^\dagger, A_l^\dagger\} = \{\tilde{A}_k^\dagger, \tilde{A}_l^\dagger\} = 0, \\ \{\tilde{A}_k^\dagger, A_l^\dagger\} = \{\tilde{A}_k, A_l\} = \{\tilde{A}_k, A_l^\dagger\} = 0, \\ \{A_k, A_l^\dagger\} = \{\tilde{A}_k, \tilde{A}_l^\dagger\} = \delta_{k,l} \mathbb{1}. \end{gathered} \tag{1.68}$$

---

[8] The order of operations here does not matter: we can 'extract' first the creation and annihilation operator on e_m or the one on $\overline{e_n}$ without affecting the result.

To summarise what we have seen so far, when we employ the superfermion formalism we write states and operators in a larger space with twice the number of modes, with the same statistics as the original space.[9] To each original, 'physical' mode, we add an 'auxiliary' mode, and we have a correspondence between mixed states in the original Hilbert space and pure states in the enlarged space.

A way to get rid of the unwanted signs in Eq. (1.67) is to restrict our attention to a specific subspace of $\mathscr{F}_- \otimes \overline{\mathscr{F}_-}$, namely one of the two eigenspaces of the parity operator (in this case, the parity of $f_{m,n}$ equals $|m| + |n| \bmod 2$). Within these eigenspaces, $|n| - |m|$ is fixed and so the results of Eq. (1.67) do not depend anymore on n and m. Fortunately, physics comes to our aid here. The superselection rules we already encountered in Section 1.1.4, in fact, dictate that physical Hamiltonians must preserve the parity of the system [94]. If H is a parity-conserving Hamiltonian on $\mathscr{F}_-$, it is easy to see that $V(H) = H \otimes \mathbb{1} - \mathbb{1} \otimes \overline{H}$ preserves the parity on $\mathscr{F}_- \otimes \overline{\mathscr{F}_-}$, thus in all reasonable fermionic physical systems we can safely restrict the space to either one of the parity eigenspaces. In addition, any pure state in $\mathscr{F}_-$ corresponds to an even state in $\mathscr{F}_- \otimes \overline{\mathscr{F}_-}$:

$$\psi = \sum_n c_n e_n \qquad \Longrightarrow \qquad \psi \otimes \bar{\psi} = \sum_{n,m} c_n \overline{c_m} f_{n,m} \tag{1.69}$$

but n, m run through even (odd) numbers only if ψ is even (odd) therefore each $f_{n,m}$ term in the sum, in both cases, will be even. In the even-parity subspace, Eq. (1.67) becomes

$$\begin{aligned} a_k e_m \otimes \overline{a_l e_n} &= -A_k \tilde{A}_l f_{m,n}, \\ a_k^\dagger e_m \otimes \overline{a_l^\dagger e_n} &= A_k^\dagger \tilde{A}_l^\dagger f_{m,n}. \end{aligned} \tag{1.70}$$

Operators that preserve the parity, once written in terms of annihilation and creation operators, contain an even number of such operators. In order to determine how V transforms operators, it is sufficient to calculate its action on operators that are at most *quadratic* in a_n and $a_m^\dagger$, and the rest can be obtained through composition, i.e. $V(AB\rho) = V(A)(V(B\rho)) = V(A)(V(B)(V(\rho)))$. We report here a 'conversion table' for quadratic operators:

$$\begin{aligned} V(a_n^\dagger a_m \rho) &= A_n^\dagger A_m V(\rho), & V(a_n^\dagger \rho a_m) &= A_n^\dagger \tilde{A}_m^\dagger V(\rho), \\ V(a_n a_m^\dagger \rho) &= A_n A_m^\dagger V(\rho), & V(a_n \rho a_m^\dagger) &= \tilde{A}_m A_n V(\rho), \\ V(\rho a_n^\dagger a_m) &= \tilde{A}_m^\dagger \tilde{A}_n V(\rho), & V(\rho a_n a_m^\dagger) &= \tilde{A}_m \tilde{A}_n^\dagger V(\rho). \end{aligned} \tag{1.71}$$

---

[9] The core results that we have written so far hold for bosons too, by removing $(-1)^{|m|-|n|}$ in Eq. (1.60), albeit we need to be careful about the convergence of some infinite series. For example, the trace cannot be computed as the inner product with the element $V(\mathbb{1})$, which cannot exist since the symmetric Fock space is infinite-dimensional.

We stress that these formulae hold only in the even-parity subspace of $\mathcal{F}_-$; importantly, there can be no such expressions for $V(a_n\rho)$ or $V(a_m^\dagger\rho)$ alone, for example, since a single a_n or $a_m^\dagger$ would map even states to odd states and vice versa.

In order to compute expectation values of observables, we use Eq. (1.55): first, we need to write $V(\mathbb{1})$ as an MPS. If we call b_0 and b_1 two orthogonal basis elements of $\mathbb{C}^2$, and interpret n as the multi-index $(n_1, n_2, \dots, n_N)$ as in Eq. (1.59), we can read $V(\mathbb{1})$ as the sum of all $e_n \otimes \overline{e_n}$ elements where

$$e_n = b_{n_1} \otimes b_{n_2} \otimes \cdots \otimes b_{n_N}. \tag{1.72}$$

We observe that

$$\sum_{n=0}^{2^N-1} e_n = \sum_{n_1=0}^{1} \sum_{n_2=0}^{1} \cdots \sum_{n_N=0}^{1} b_{n_1} \otimes b_{n_2} \otimes \cdots \otimes b_{n_N} = (b_0 + b_1)^{\otimes N}, \tag{1.73}$$

so in the same fashion we derive

$$V(\mathbb{1}) = \sum_{n=0}^{2^N-1} e_n \otimes \overline{e_n} = (b_0 \otimes \overline{b_0} + b_1 \otimes \overline{b_1})^{\otimes N}. \tag{1.74}$$

We can easily write this as an MPS whose sites are pairwise entangled:

$$b_0 \otimes \overline{b_0} + b_1 \otimes \overline{b_1} = \sum_{i_1=0}^{1} \sum_{i_2=0}^{1} \sum_{a=1}^{\chi} M_a^{(1,i_1)} M_a^{(2,i_2)} b_{i_1} \otimes \overline{b_{i_2}} \tag{1.75}$$

is an MPS with bond dimension 2 and with

$$M^{(1,0)} = M^{(2,0)} = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \qquad M^{(1,1)} = M^{(2,1)} = \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \tag{1.76}$$

The effect of computing $\operatorname{tr}\rho = \langle V(\mathbb{1}), V(\rho)\rangle$ is to contract the pairs of physical and auxiliary indices of the MPS representing $V(\rho)$: we can see this in the following formula, where we consider a single physical mode.

$$\begin{aligned}\langle b_0 \otimes \overline{b_0} + b_1 \otimes \overline{b_1}, V(\rho)\rangle &= \sum_{n,m\in I} \rho_{nm} \langle b_0 \otimes \overline{b_0} + b_1 \otimes \overline{b_1}, e_n \otimes \overline{e_m}\rangle = \\ &= \sum_{n,m\in I} \rho_{nm} (\langle e_m, b_0\rangle \langle b_0, e_n\rangle + \langle e_m, b_1\rangle\langle b_1, e_n\rangle) = \\ &= \sum_{n,m\in I} \rho_{nm} \langle e_m, e_n\rangle = \sum_{n\in I} \rho_{nn},\end{aligned} \tag{1.77}$$

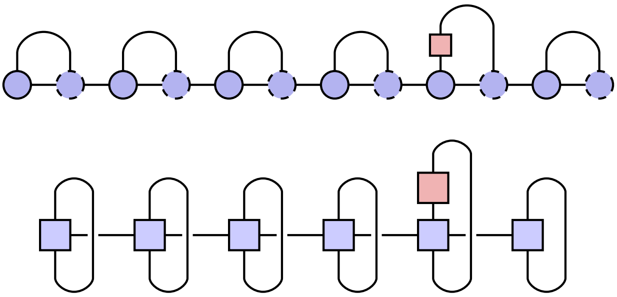

Figure 1.14. TN contraction that expresses the expectation value of a local operator on the fifth site out of six. In the superfermion formalism (*top*), solid discs represent physical modes, with the associated auxiliary mode, dashed, at their right. The trace operation contracts the free index of the tensors of each physical mode with the one of the respective auxiliary mode. In the standard MPO formalism (*bottom*), the trace contracts instead the two physical modes of each tensor.

which is the trace of the mixed state ρ. If we want to compute the expectation value $\mathrm{tr}(X\rho)$ of some operator X, first we compute $V(X\rho)$ by finding $V(X)$ in the superfermion space and applying it to $V(\rho)$, and then we compute the trace by contracting all the physical-auxiliary pairs of indices. Typically, when X represents a physical observable, it is written as a function of the creation and annihilation operators $a_k^{(\dagger)}$ of the original fermionic system, which translate directly to the creation and annihilation operators $A_k^{(\dagger)}$ of the physical modes of the superfermion system: for example, $V(a_n^\dagger a_m \rho) = A_n^\dagger A_m V(\rho)$.

With the aim of optimising the TN structure towards the efficiency of contractions, we are essentially forced to order the 'superfermionic modes' putting the auxiliary modes next to the physical mode they are associated to. This way, when we need to compute the expectation value of an observable, or the trace of a state, the contraction in Eq. (1.77) happens between neighbouring sites; this is exemplified in Fig. 1.14. Any other ordering scheme would require us to contract far-away sites, greatly increasing the complexity of the operation. There is a downside, however, in the fact that the range of interactions increases: for example, nearest-neighbour interaction between physical modes become next-to-nearest-neighbour, because we find an auxiliary mode between them.

In conclusion, with the superfermion method illustrated in this section we can write a mixed state of n fermions as a pure state of $2n$ fermions. We find at least two advantages of doing so:

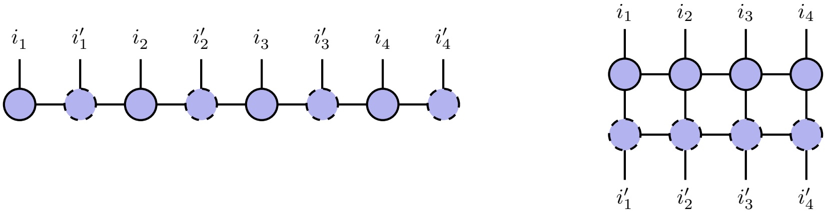

Figure 1.15. Possible alternative scheme to arrange the TN for a state in the superfermion formalism (*right*), where physical and auxiliary sites are placed on different levels, compared to the standard layout (*left*).

- the antisymmetric structure of the Fock space makes calculations easier;
- many TN software libraries have the means to describe a fermionic *pure state* as a TN, but not a mixed state.

Some remarks about MPSs we made before still hold in the superfermion formalism: we do not have automatic trace preservation from the TN time-evolution algorithms, the bond dimensions are affected by both quantum and classical correlations, and physical properties cannot always be obtained with efficient contraction techniques.

As a final note, we point out a possible way of improving this scheme. One of the drawbacks of the superfermion formalism is that the correlations spread out more than in the original state: the bond dimension needs between two consecutive sites needs to accommodate both the correlation between next-to-neighbouring physical (as well as auxiliary) modes and the correlation between a physical mode and its auxiliary one. A more clever approach would probably be to use a different TN layout in which physical and auxiliary modes are not interleaved, but on different 'levels', as in Fig. 1.15. In such a layout the correlations would be better distributed along the multiple links of the TN, possibly resulting in an overall smaller bond dimension. However, the structure of this TN is unusual. Most of the algorithms we saw in this chapter, not only to compute the time-evolution but also to perform basic operations such as expectation values or the truncated SVD, would need to be reformulated. For MPSs, on the other hand, there already exist many efficient algorithms that suit our needs. For these reasons, we will stick to the methods outlined in this section and represent mixed states as standard MPSs.

Numerical simulation of open quantum systems 2

Any realistic quantum system is coupled to another quantum system, typically much larger, which describes the world around it. We call this larger system the *environment*. The interaction between a system and its environment leads to correlations between them, which affect the dynamics of the former, usually causing decoherence and, more in general, irreversible and noisy processes. In the design of quantum computers, for example, decoherence is undesired, so they are kept as isolated as possible from the external environment. Other times, however, decoherence can play a positive role: this is the case, for example, in quantum biology systems where the environment is said to assist the biological complexes in transporting energy [95]. When controlled, dissipation can also be an essential tool for manipulating quantum information: dissipation engineering enables quantum measurement, state preparation and stabilisation [96]. In any case, it is important to find ways to incorporate the effects of the environment in the evolution of the system, in order to derive a realistic description.

2.1 Open quantum systems

When a system is considered within its environment, the two can exchange energy, particles, information, and other physical quantities. In this case, the system is called *open*, in contrast to isolated systems which are instead completely separated from their surroundings. The typical example of open quantum system (OQS) is a small quantum system, such as a spin or a quantum dot, immersed in a bosonic bath modelling the

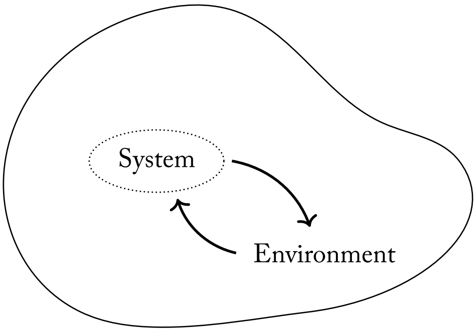

Figure 2.1. The theory of OQSs deals with systems that exchange information, energy, particles, etc. with the surrounding environment.

environment surrounding the system, which introduces, among other things, thermal noise. Typical examples are the Caldeira-Leggett model [97], and the spin-boson model [98], in which a quantum particle in one dimension or a spin, respectively, are coupled to an infinite collection of independent harmonic oscillators. These are simple yet expressive models that allow investigating decoherence and damping effects caused by the environment on the open system, and have found numerous applications, from electron transfer to quantum information processing [99].

There are instances of *fermionic* environments as well, which play a major role in several areas of condensed matter physics. The Kondo effect—which explains how the scattering of conduction electrons in a metal caused by magnetic impurities leads to a peculiar temperature dependence of the resistivity of the metal—sees the conduction band modelled as a fermionic environment [100]. Fermionic environments are also found in Anderson's impurity model, which describes the behaviour of impurities in metals in generality [36, 37, 101]. Finally, in the study of quantum transport, two fermionic environments at a certain temperature and chemical potential are connected to a quantum system that acts as a bridge to transport energy and/or particles [102–104].

In general, the set of physical states for an isolated quantum system is the set $\mathcal{S}(\mathcal{H})$ of self-adjoint, positive, unit-trace operators on a Hilbert space $\mathcal{H}$. If we write the global Hilbert space as the product of a 'system space' and an 'environment space', $\mathcal{H} = \mathcal{H}_\mathrm{S} \otimes \mathcal{H}_\mathrm{E}$, then, in principle, one can obtain the state of the open system $\rho_\mathrm{S} \in \mathcal{S}(\mathcal{H}_\mathrm{S})$ first by computing the evolution of the global state $\rho \in \mathcal{S}(\mathcal{H})$ through the von Neumann equation with a global Hamiltonian H_tot, that is $\dot{\rho}(t) = -\mathrm{i}[H_\mathrm{tot}, \rho(t)]$, then by computing the partial trace over the environment. Sometimes a 'clean' separation $\mathcal{H} = \mathcal{H}_\mathrm{S} \otimes \mathcal{H}_\mathrm{E}$ is not possible, as is the case with fermionic systems: due to the antisymmetrisation, a fermionic Fock space is not naturally equipped with a tensor-product structure with respect to the individual mode subspaces. For this reason, establishing a meaningful notion of subsystem is not trivial. Mappings such as the Jordan-Wigner transformation mappings generally do not commute with the partial trace operation [105], since they map local operators are to global spin (or qubit) operators. The superselection rule that prohibits the superposition of even and odd-parity states, that we already encountered in Section 1.1.4, allows us to define quantum channels, completely positive maps, partial traces and so on in a meaningful way by referring to the (more natural) antisymmetric way of combining systems instead of the standard tensor product [73, 94].

$$\dot{\rho}_\mathrm{S}(t) = -\mathrm{i}\,\mathrm{tr}_\mathrm{E}\big([H_\mathrm{tot}, \rho(t)]\big). \tag{2.1}$$

However, to determine the dynamics of the reduced state $\rho_\mathrm{S}(t)$ from this equation we need to compute the complete state $\rho(t)$, which in all but the simplest cases is essentially impossible. Moreover, in many situations, one is interested only in physical quantities related to the system alone, and does not need a complete description of the environment. The theory of OQSs provides a framework to approach this problem: instead of an equation describing the evolution of the entire state, it seeks to simplify the problem by searching for an equation (called *master equation*) that does not take into account the degrees of freedom of the environment, but only how they influence the state of the system.[1]

We look then for approximation techniques that, by discarding some aspects of the description, can transform Eq. (2.1) into a closed equation with a simpler form. The most important approximation in OQS theory is the assumption that the change in $\rho_\mathrm{S}(t)$ does not depend on its past evolution. Under this assumption the master equation has a simple expression: there exists a particular, time-independent operator $\mathcal{L}$ on $\mathcal{S}(\mathcal{H}_\mathrm{S})$ such that $\dot{\rho}_\mathrm{S}(t) = \mathcal{L}(\rho_\mathrm{S}(t))$. We will briefly comment in Section 2.2 on the hypotheses under which this regime, called *Markovian*, is appropriate. The basic idea behind the Markovian approximation is the clean separation of the environmental time scales from the system time scales [106]. The recent increase in the ability to control the dynamics of open systems and the properties of the environment, however, has inspired the need to understand quantum dynamics beyond the Markovian behaviour.

While there exist master equations that can capture non-Markovian phenomena [32, 106], they are typically complicated and cannot be (easily) solved. In such cases,

---

[1] In general, the 'system' is singled out as a subsystem such that all observations of interest refer to it. The physical distinction between system and environment can be blurry: in fact, we will encounter situations where the 'system' in our consideration is made of an actual, small, physical system together with a few degrees of freedom of its surrounding real-life environment.

it is therefore important to have reliable and efficient numerical tools that allow us to explore the dynamics and find key properties beyond the analytical reach. The first step in tackling the numerical simulation of OQSs is to treat the union of system and environment as an isolated system with a large number of degrees of freedom, for which a number of methods are already available in the literature. Some of them were originally developed for the simulation of strongly correlated quantum systems, such as the numerical renormalisation group in the basis of scattering states [107, 108], flow equations [109, 110], the time-dependent density-matrix renormalisation group (tDMRG) [16, 62], multilayer multiconfiguration time-dependent Hartree [111], and continuous-time quantum Monte Carlo [112–114]. In addition to these methods, many algorithms specifically designed for the simulation of the dynamics of OQSs have been proposed in recent years: among them we find the hierarchical equations of motion (HEOM) [115], the time-evolving matrix-product operators (TEMPO) [116] the dissipation-assisted matrix-product factorization (DAMPF) [117, 118], and the automated compressed environment (ACE) [119].

We are interested specifically in simulating OQSs with the TN techniques introduced in Chapter 1. These are already well suited to the simulation of systems with a large number of degrees of freedom: indeed some of the aforementioned numerical methods are already implemented as TNs. The main difficulty, in essentially all cases, is dealing with the large number of degrees of freedom of the environment affecting the reduced dynamics of the OQS. For this reason, in Section 2.3 we explore some results that reveal under which conditions two environments affect the systems in the same way: we will then be able to use those results to replace a complex environment with another one whose description (and numerical simulation) is more manageable. We conclude the chapter by studying, in Section 2.4, a particular method for simulating OQSs, the time-evolving density operator with orthogonal polynomials algorithm (TEDOPA), which reshapes a continuous environment into a discrete one, which can then be mapped into an MPS. This will be the starting point for Chapter 3, where we will put these techniques into practice.

2.2 Markovian dynamics

A major example of master equation is the Gorini-Kossakowski-Sudarshan-Lindblad (GKSL) equation [120, 121]. If the system and the environment are initially uncorrelated, i.e. the initial state is $\rho_0 = \rho_{\mathrm{S},0} \otimes \rho_{\mathrm{E},0}$, then the time evolution of the reduced system can be described by the equation

$$\rho_{\mathrm{S}}(t) = V_t\big(\rho_{\mathrm{S}}(t)\big) = \operatorname{tr}_{\mathrm{E}}\big(U(t,0)(\rho_{\mathrm{S},0} \otimes \rho_{\mathrm{E},0})U(t,0)^\dagger\big) \tag{2.2}$$

where $U(t,0)$ is the unitary operator generated by the full Hamiltonian of system and environment. The set $\{V_t\}_{t\geq 0}$, with $V_0 = \mathbb{1}$, of time-evolution operators forms a one-parameter family of convex-linear, completely positive and trace-preserving maps on $\mathcal{S}(\mathcal{H}_\mathrm{S})$, called *quantum dynamical maps*. If the characteristic time scales over which the correlations induced by the system in the environment decay are much smaller than the characteristic time scales of the system, then we can treat the environment as effectively *memoryless*. Within this approximation, the excitations caused by the system just diffuse away in the environment and never come back to influence the evolution of the system at a later time. In the end, this means that the state does not depend on information about its past. This regime is called *Markovian*, and is formalised by stating that the collection of dynamical maps $\{V_t\}_{t\geq 0}$ satisfies the semigroup property $V_t V_s = V_{t+s}$. For any such quantum dynamical semigroup there exists an operator $\mathcal{L}\colon \mathcal{S}(\mathcal{H}_\mathrm{S}) \to \mathcal{S}(\mathcal{H}_\mathrm{S})$ that generates the semigroup, in the sense that $V_t = \exp(t\mathcal{L})$. The master equation for the reduced dynamics of the system is then

$$\dot{\rho}_\mathrm{S}(t) = \mathcal{L}\big(\rho_\mathrm{S}(t)\big). \tag{2.3}$$

This equation is local in time: the right-hand side does not depend on the history of the system but only on its current state $\rho_\mathrm{S}(t)$, and $\mathcal{L}$ does not depend on time or on the state. In its most general form, the operator $\mathcal{L}$ can be written as

$$\mathcal{L}\big(\rho_\mathrm{S}(t)\big) = -\mathrm{i}[H, \rho_\mathrm{S}(t)] + \sum_{k=1}^{N} \big(L_k \rho_\mathrm{S}(t) L_k^\dagger - \tfrac{1}{2}\{\rho_\mathrm{S}(t), L_k^\dagger L_k\}\big). \tag{2.4}$$

where H, L_1, …, L_N are operators on $\mathcal{H}_\mathrm{S}$, and H in particular is self-adjoint. A complete treatment of the necessary hypotheses that lead to this master equation, as well as a microscopic derivation from the underlying unitary dynamics, which are beyond the scope of this chapter, can be found in [32, 106].

The GKSL equation can adequately capture the decoherence dynamics of many interesting systems: for example the interaction between matter and electromagnetic radiation satisfies the assumptions for Markovian dynamics in typical quantum optics situations [32]. Current advanced experimental techniques allow for the possibility of investigating regimes where the environment creates memory effects in the system dynamics, to the point where the previously mentioned separation of system and environment time scales does not hold anymore. More in general, when the underlying assumption of coupling to an essentially unchanging, memoryless environment is not fulfilled, significantly non-Markovian dynamics may arise, and system and environment can mutually influence each other's evolution on time scales which are relevant to the process under study, and cannot be easily decoupled. More complicated equations can be used to treat the evolution of the system's degrees of freedom, but they are much more complicated and cannot be (easily) solved: here, numerical simulations turn out to be an important tool that allows us to explore the behaviour of the system beyond our analytical reach. Solid-state implementations of qubits, such as quantum dots [122, 123], molecular transistors [124], unconventional superconductors [125], and molecular magnets [126], quantum thermal machines [127, 128], quantum sensing and metrology protocols [129, 130], energy-charge conversion and exciton transport in solid-state devices [131, 132], or synthetic or biological light-harvesting complexes, complex organic molecules used for light emission or solar cells [133, 134] are typical instances in which deviations from a Markovian evolution can play a significant role.

2.3 Equivalent environments

If one is interested in the evolution of the open system only, the problem of computing its reduced dynamics would be simplified if we found another environment, not necessarily directly related to the original one at the level of the overall dynamics, that induces the same reduced dynamics on the system. For a specific class of environments, called Gaussian environments, it is well known that, if their dynamics is unitary, their influence is completely determined by the one- and two-time correlation functions of the interaction operators. We start this section therefore with a brief introduction to Gaussian states and their relevance.

2.3.1 Gaussian states

Gaussian states are a class of states defined by the property that they are completely determined by their moments up to second order. Given an orthonormal basis $\{|e_i\rangle\}_{i=1}^N$ of a Hilbert space $\mathcal{H}$, let a_i be the annihilation operator associated to $|e_i\rangle$ on the (symmetric or antisymmetric) Fock space $\mathcal{F}_\pm(\mathcal{H})$. Gaussian states can be defined as those states which can be represented as ground or thermal states of quadratic Hamiltonians [135, 136], i.e. Hamiltonian operators that are polynomials of order two in the creation and annihilation operators. In this sense, a state ρ is Gaussian if there exists $\beta \in [0, +\infty]$ and such a Hamiltonian H that

$$\rho = \frac{1}{\operatorname{tr}\exp(-\beta H)} \exp(-\beta H) \tag{2.5}$$

with $\beta = +\infty$ obtained as a limit, representing the case where ρ is a ground state. All physical properties of a Gaussian state are fully determined by its low-order moments; in fact, through Wick's theorem [137] we can compute the expectation value on a

Gaussian state $\rho \in \mathcal{S}(\mathcal{F}_\pm(\mathcal{H}))$ of the product of any number of a_is and $a_i^\dagger$s as the product of terms from

$$\langle a_i \rangle_\rho = \mathrm{tr}(\rho a_i), \qquad \langle a_i a_j \rangle_\rho = \mathrm{tr}(\rho a_i a_j), \qquad \langle a_i^\dagger a_j \rangle_\rho = \mathrm{tr}(\rho a_i^\dagger a_j), \tag{2.6}$$

and their complex conjugates, for all $i, j = 1, \dots, N$.

In the context of OQSs, the Gaussian regime applies if system and environment are coupled linearly and are governed by Hamiltonians which are at most quadratic in the canonical operators; equivalently, the overall Hamiltonian of system and environment must be at most quadratic. For example, denoting by a the system modes and by b the environment ones,

$$H = \underbrace{\sum_{k=1}^{N_\mathrm{S}} \varepsilon_k a_k^\dagger a_k}_{H_\mathrm{S}} + \underbrace{\sum_{k=1}^{N_\mathrm{S}} \sum_{l=1}^{N_\mathrm{E}} g_{kl}(a_k^\dagger b_l + b_l^\dagger a_k)}_{H_\mathrm{I}} + \underbrace{\sum_{l=1}^{N_\mathrm{E}} \omega_l b_l^\dagger b_l}_{H_\mathrm{E}} \tag{2.7}$$

is Gaussian because the free Hamiltonians H_S and H_E are quadratic in the a and b operators, respectively, and the interaction H_I is linear in both as and bs. Note however that Gaussian systems can be a discrete as well as a continuous collection of degrees of freedom: what matters is that the operators are quadratic in the relevant canonical operators.

Even if Gaussian states do not exhaust the entire range of allowed dynamics, they still appear in the description of a vast number of experimental settings such as cold atom gases and ion traps [138] and optomechanical and nanomechanical oscillators [139]. On the fermionic side, Gaussian states are routinely used describe free states in condensed matter physics [140] and quantum chemistry [141].

2.3.2 Equivalence of unitary dynamics

The relevance of Gaussian states is due to the fact that in most cases of interest the Hamiltonian of the free environment is quadratic, while a linear coupling to the system arises either naturally or through a weak-coupling approximation (in which one approximates the interaction to the first order). Let us write the overall Hamiltonian as

$$H = H_\mathrm{S} + H_\mathrm{E} + \sum_{i \in I} A_{\mathrm{S},i} B_{\mathrm{E},i}. \tag{2.8}$$

The first two terms describe the free dynamics of the system and the environment, respectively, whereas the third describes the interaction between them, with $A_{\mathrm{S},i}$ acting on the system degrees of freedom and $B_{\mathrm{E},i}$ on the environment ones. We assume that the environment Hamiltonian is quadratic in the canonical creation and annihilation operators of the environment, while the $B_{\mathrm{E},i}$ operators is linear. The system Hamiltonian and the $A_{\mathrm{S},i}$ operators in the interaction terms can instead be arbitrary—the Gaussian hypothesis applies on the environment side only. This is a very general scheme that can describe typical bosonic or fermionic interactions such as an XX interaction

$$\sum_{i=1}^{N_{\mathrm{E}}} g_k(a^\dagger + a)(b_k^\dagger + b_k) \tag{2.9}$$

or an exchange interaction

$$\sum_{i=1}^{N_{\mathrm{E}}} g_k(a^\dagger b_k + a b_k^\dagger). \tag{2.10}$$

We also assume that we start from a factorised initial state $\rho_{\mathrm{S},0} \otimes \rho_{\mathrm{E},0}$. The simplicity of Gaussian environments is that their influence on the reduced system can be *entirely* described by expectation values and two-time correlation functions (TTCFs) of the coupling terms, i.e. the $B_{\mathrm{E},i}$ operators. The fundamental quantities here are the functions

$$\begin{aligned} g_i(t) &:= \mathrm{tr}_{\mathrm{E}}\big(B_{\mathrm{E},i}\,\varPhi_t^{\mathrm{E}}(\rho_{\mathrm{E},0})\big), \\ f_{i,j}(t+s,t) &:= \mathrm{tr}_{\mathrm{E}}\Big(B_{\mathrm{E},i}\,\varPhi_s^{\mathrm{E}}\big(B_{\mathrm{E},j}\,\varPhi_t^{\mathrm{E}}(\rho_{\mathrm{E},0})\big)\Big) \end{aligned} \tag{2.11}$$

where $\varPhi_t^{\mathrm{E}}$ is dynamical map of the free environment generated by H_{E}, i.e.

$$\varPhi_t^{\mathrm{E}}(\rho) = \exp(-\mathrm{i}tH_{\mathrm{E}})\rho\exp(\mathrm{i}tH_{\mathrm{E}}). \tag{2.12}$$

This means that if we swap this environment with a new one, E$'$, such that

$$H = H_{\mathrm{S}} + H_{\mathrm{E}'} + \sum_{i\in I} A_{\mathrm{S},i} B_{\mathrm{E}',i}, \tag{2.13}$$

(the system-related terms are unchanged) and

$$\begin{aligned} \mathrm{tr}_{\mathrm{E}}\big(B_{\mathrm{E},i}\,\varPhi_t^{\mathrm{E}}(\rho_{\mathrm{E},0})\big) &= \mathrm{tr}_{\mathrm{E}'}\big(B_{\mathrm{E}',i}\,\varPhi_t^{\mathrm{E}'}(\rho_{\mathrm{E}',0})\big), \\ \mathrm{tr}_{\mathrm{E}}\Big(B_{\mathrm{E},i}\,\varPhi_s^{\mathrm{E}}\big(B_{\mathrm{E},j}\,\varPhi_t^{\mathrm{E}}(\rho_{\mathrm{E},0})\big)\Big) &= \mathrm{tr}_{\mathrm{E}'}\Big(B_{\mathrm{E}',i}\,\varPhi_s^{\mathrm{E}'}\big(B_{\mathrm{E}',j}\,\varPhi_t^{\mathrm{E}'}(\rho_{\mathrm{E}',0})\big)\Big). \end{aligned} \tag{2.14}$$

for all $i, j \in I$ and for all $t, s \in \mathbb{R}$, then the reduced dynamics, i.e. $\rho_{\mathrm{S}}(t)$, will be the same in both configurations. It should be noted that the i and j indices of the $B_{\mathrm{E},i}$ operators in $g_i(t)$ and $f_{i,j}(t,s)$ are tied to the $A_{\mathrm{S},i}$ terms they appear with in the interaction Hamiltonian: it is not just the set of correlation functions and expectation values that has to coincide, their order matters as well.

Several approaches exist in the literature that allow one to map the original environment into a unitarily equivalent one, that is easier to deal with because its configuration is more suitable for the application of specific numerical tools. The TEDOPA algorithm that we will encounter in Section 2.4, for example, transforms a continuous environment in a nearest-neighbour chain of modes that can be easily simulated as a TN. Nevertheless, in these equivalent unitary models the number of degrees of freedom of the environment typically remains very large, and as a consequence the simulation of the equivalent model can still be challenging. For this reason, we explore now how an environment can be replaced not with another environment undergoing a unitary dynamics, but a *non-unitary* one.

2.3.3 Equivalence of non-unitary dynamics

With a unitary environment we are essentially treating all of its modes on an equal footing. However, in order to isolate the characteristics of the environment that most influence the dynamics of the system, we would like to devise an environment that is composed of a 'non-Markovian core' which, in turn, interacts with a Markovian environment. The former will be coupled directly, and in a coherent manner, to the OQS and will describe all the memory effects arising during the time evolution, while the latter will be effectively characterised by a GKSL equation, representing the dissipation of information out of the non-Markovian core (and away from the system). The idea behind this scheme is that computing the reduced dynamics in this new configuration, in cases where the non-Markovian core can be well described by a limited number of degrees of freedom, can be much simpler than doing so with the full unitary evolution induced by a complex environment.

A particular instance of this scheme is the concept of *pseudomodes*, where the original environment is mapped onto an alternative configuration consisting of a finite number of discrete modes, all coupled to the open system; in turn each of these discrete modes is coupled to a Markovian reservoir. To its core, the pseudomodes approach was born as a general, non-perturbative approach to study the decay of an atomic system coupled to an environment described by a non-uniform density of states [142]. The parameters (energy, damping rate, and so on) defining the pseudomodes are derived in a systematic way from the original environment, and they enable a full inclusion of non-Markovian effects by studying an enlarged OQS governed by a Lindblad master equation [143].

The equivalence between the original dynamics and the new non-unitary one rests again on the equality of the expectation values and TTCFs. It holds both for bosonic environments [144] and fermionic ones [145–148], in slightly different ways. Let us study both in more detail.

Bosonic environments The bosonic case is covered in Ref. [144]. The original configuration, that we call '1', is described by a Hamiltonian operator H_1 equal to H from Eq. (2.8). The initial state is factorised as $\rho_{\mathrm{S},0}\rho_{\mathrm{E},0}$, and both factors are Gaussian states. For this configuration the environmental expectation values and correlation functions are $g_i(t)$ and $f_{i,j}(t,s)$ from Eq. (2.11). In configuration '2' the same system S interacts with another environment R (we call it *reservoir* to reduce ambiguity) as

$$H_2 = H_\mathrm{S} + H_\mathrm{R} + \sum_{i=1}^{n} A_{\mathrm{S},i} B_{\mathrm{R},i} \tag{2.15}$$

but the reservoir modes are also subject to a Lindblad-type evolution, generated not only by H_R alone from H_2 but from

$$\mathcal{L}_\mathrm{R}(\rho) = -\mathrm{i}[H_\mathrm{R},\rho] + \sum_{j=1}^{\ell}\bigl(L_j\rho L_j^\dagger - \tfrac{1}{2}\{L_j^\dagger L_j,\rho\}\bigr) \tag{2.16}$$

where the L_1, …, L_ℓ operators act on the Hilbert space of the reservoir[2], such that the dynamical map of the free environment is $\varPhi_t^\mathrm{R} = \exp(t\mathcal{L}_\mathrm{R})$. The initial state is again factorised, as $\rho_{\mathrm{S},0}\rho_{\mathrm{R},0}$, in which $\rho_{\mathrm{R},0}$ is of course Gaussian. The environmental expectation values and correlation functions are

$$\begin{aligned} g_i'(t) &= \mathrm{tr}_\mathrm{R}\bigl(B_{\mathrm{R},i}\,\varPhi_t^\mathrm{R}(\rho_{\mathrm{R},0})\bigr), \\ f_{ij}'(t+s,t) &= \mathrm{tr}_\mathrm{R}\Bigl(B_{\mathrm{R},i}\,\varPhi_s^\mathrm{R}\bigl(B_{\mathrm{R},j}\,\varPhi_t^\mathrm{R}(\rho_{\mathrm{R},0})\bigr)\Bigr). \end{aligned} \tag{2.17}$$

Fermionic environments The fermionic case is covered in Ref. [147]. For fermions we need to start from a slightly more specific version of Eq. (2.8) as the configuration '1', and assume that the system S interacts with an environment E through the Hamiltonian operator

$$H_1 = H_\mathrm{S} + H_\mathrm{E} + \sum_{i=1}^{n}(A_{\mathrm{S},i}^\dagger B_{\mathrm{E},i} + A_{\mathrm{S},i} B_{\mathrm{E},i}^\dagger) \tag{2.18}$$

---

[2] These L_j operators can be arbitrary, as long as they determine a completely positive, trace-preserving Gaussian map. A typical example is L_j equal to a creation or annihilation operator on a reservoir mode (possibly times a positive number), in which case we obtain a quadratic Lindblad operator.

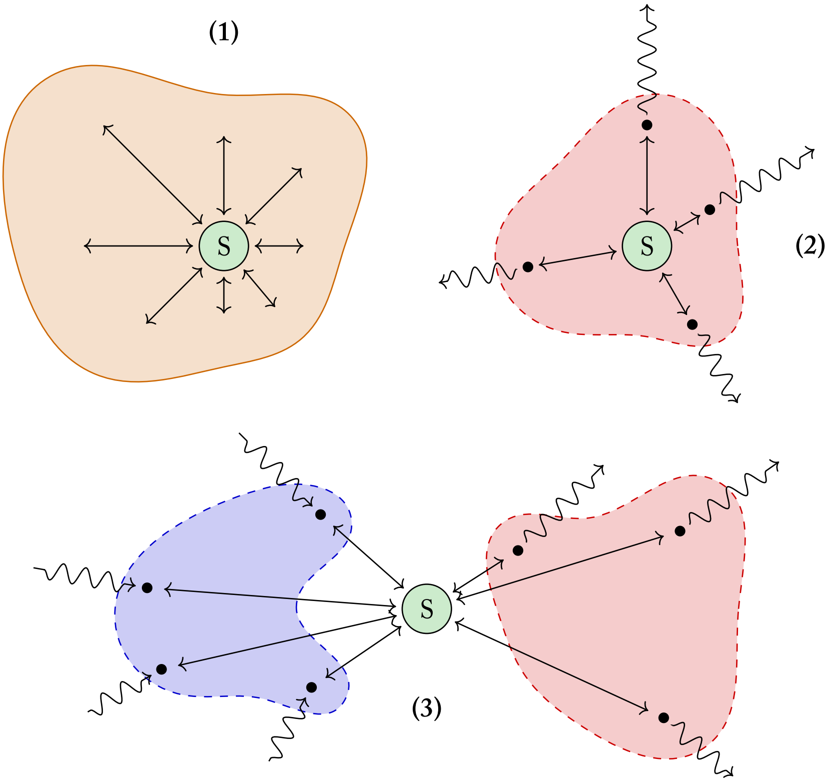

Figure 2.2. Sketches of equivalent reservoirs. 1. Original (closed) environment, with unitary dynamics. 2. Equivalent open reservoir for bosonic environments. The reservoir, representing the non-Markovian core, is open and dissipates excitations (particles, energy, information, etc.) into an external environment. 3. Equivalent open reservoirs for fermionic environments. With respect to the bosonic reservoir, here there is an additional reservoir that injects excitations.

with the usual linearity and quadraticity properties. Similarly to the bosonic case, the initial state is the product of two Gaussian states, $\rho_{\mathrm{S},0}\rho_{\mathrm{E},0}$.

In configuration '2' the open system interacts not with one but *two* reservoirs, one starting from the vacuum state and the other from the completely occupied state, acting as a sink and a source, respectively, as in Fig. 2.2.

This more complicated setting than the previous one can be explained by the fact that even at zero temperature, the ground state of a fermionic system may not be the vacuum state, depending on the chemical potential. We need then both a 'filled' and an 'empty' environment to model the modes below and above the Fermi surface, respectively. The Lindblad operators in this case have a particular form that makes the presence of two environments explicit: the generator of the free reservoir dynamics is

$$\begin{aligned}\mathscr{L}_{\mathrm{R}}(\rho) = -\mathrm{i}[H_{\mathrm{R}_0} + H_{\mathrm{R}_1}, \rho] + \sum_{i,j=1}^{\ell} \mathit{\Gamma}_{\mathrm{R}_0,ij}\big(a_{0,j}\rho a_{0,i}^{\dagger} - \tfrac{1}{2}\{a_{0,i}^{\dagger}a_{0,j}, \rho\}\big) + \\ + \sum_{i,j=1}^{\ell} \mathit{\Gamma}_{\mathrm{R}_1,ij}\big(a_{1,i}^{\dagger}\rho a_{1,j} - \tfrac{1}{2}\{a_{1,j}a_{1,i}^{\dagger}, \rho\}\big) \qquad (2.19)\end{aligned}$$

where the $a_{0,i}$s and $a_{1,i}$s are the canonical annihilation operators of the empty and filled reservoir modes, respectively. The initial state is factorised as $\rho_{\mathrm{S},0}\rho_{\mathrm{R}_1,0}\rho_{\mathrm{R}_0,0}$ where $\rho_{\mathrm{R}_1,0}$ is the completely filled state and $\rho_{\mathrm{R}_0,0}$ the vacuum state. The environmental expectation values and correlation functions are again as in Eq. (2.17), with a different dynamical map, of course, this time generated by $\mathscr{L}_{\mathrm{R}}$ in Eq. (2.19).

In both the bosonic and the fermionic case, the result is analogous to the case of unitary equivalence: if for any $t, s \geq 0$ and for any i, j we have $g_i(t) = g_i'(t)$ and $f_{i,j}(t+s, s) = f_{i,j}'(t+s, s)$ then the reduced dynamics the system is exactly the same.[3] The recipe, therefore, is the following: the original environment is replaced by a handful of pseudomodes (their number will depend on the particular situation) whose parameters such as energy, coupling constants, dissipation/absorption rates, and so on, to be determined so that the TTCFs and the expectation values coincide with those generated from the original environment. This procedure yields a way to characterise a generic dynamics, which can even deviate strongly from a Markovian behaviour, in a non-perturbative way with a smaller set of degrees of freedom. We can then compute the original unitary dynamics with a possibly simpler Lindblad-type evolution instead. In general, an exact replacement is not available, and no *finite*, damped reservoir can have the same effect, i.e. generate the same correlation functions, of an *infinite* unitary one, as explained in Ref. [149]. However, the way these equivalence theorems are proved is such that we can at least design a Lindblad-type dynamics whose parameters are such that $g_i(t) \approx g_i'(t)$ and $f_{i,j}(t+s, s) \approx f_{i,j}'(t+s, s)$, as close as possible [144]. Then the difference between the two reduced dynamics, i.e. the distance between $\rho_{\mathrm{S}}(t)$ in the original and new configurations, would be bounded by a quantity which is tied to the distance between the environment moments [150].

---

[3] The only difference (at this level) is that the time instants t and s do not take negative values, which happens because the dynamics is not reversible.

2.4 The TEDOPA chain mapping

In this section we review how the problem of a quantum system which is linearly coupled to a continuous bosonic or fermionic environment can be mapped exactly onto a one-dimensional model with nearest-neighbours interactions, as introduced in Refs. [34, 151]; the fermionic case is discussed in detail in Ref. [152]. This map, called time-evolving density operator with orthogonal polynomials algorithm (TEDOPA), is a particular instance of a *chain mapping*, a type of algorithm designed to discretise a continuous Hamiltonian operator. In our case, the procedure starts from an unspecified system S coupled to a quadratic environment

$$H_{\mathrm{E}} = \int_X g(x) f_x^\dagger f_x \,\mathrm{d}x, \tag{2.20}$$

where f_x and $f_x^\dagger$ are the annihilation and creation operators of mode $x \in X$, respectively, such that either $[f_x, f_y^\dagger] = \delta(x-y)\mathbb{1}$ for bosons or $\{f_x, f_y^\dagger\} = \delta(x-y)\mathbb{1}$ for fermions. For bosons, the (frequency) domain X is usually taken as a subset of $[0, +\infty)$. For fermions, for which the chemical potential μ is relevant, we can either make μ explicit in the Hamiltonian, by adding a $-\mu N$ term, or we can translate the frequency domain and allow negative frequencies, so that $\mu = 0$ and a negative frequency represents a mode below the chemical potential. We also assume that the interaction between system and environment is given by a linear coupling, i.e. in the creation/annihilation operators of the environment, such as

$$H_{\mathrm{I}} = \int_X A_{\mathrm{S}}(f_x + f_x^\dagger) h(x) \,\mathrm{d}x, \tag{2.21}$$

$$H_{\mathrm{I}} = \int_X (A_{\mathrm{S}}^\dagger f_x + f_x^\dagger A_{\mathrm{S}}) h(x) \,\mathrm{d}x \tag{2.22}$$

for bosons or fermions, respectively. In this configuration, the system interacts with all modes of the environment at the same time, while the environment modes do not interact with one another: this configuration is usually called *star shaped*. If the environment is in a Gaussian state, its effect on the system S is completely characterised by the *spectral density function* $J\colon X \to [0, +\infty)$ defined by the relation

$$|g'(x)| J\big(g(x)\big) = h(x)^2 \tag{2.23}$$

Under the previous hypothesis, any given spectral density J can be reproduced by setting $g(x) = cx$ and $h(x) = \sqrt{cJ(cx)}$: the spectral density enters, then, as the square of the weight of the interaction between the system and the environment mode at frequency x. In what follows we assume we are given J and define g and h this way, so that the spectral density function is the only object characterising the environment. This function defines a measure $\mathrm{d}\lambda(x) = J(x)\,\mathrm{d}x$ on X, and we can compute a set of monic polynomials $\{p_n\}_{n=0}^{+\infty}$ that are orthogonal with respect to λ. Such a set of polynomials satisfies a three-term recurrence relation

$$xp_n(x) = p_{n+1}(x) + \alpha_n p_n(x) + \beta_n p_{n-1}(x). \tag{2.24}$$

The transformation

$$f_x^\dagger = \sum_{n=0}^{+\infty} U_n(x) b_n^\dagger, \quad U_n(x) = \sqrt{J(x)}\frac{p_n(x)}{\|p_n\|_\lambda} \tag{2.25}$$

is a canonical transformation from the original f_x modes to a new, discrete set $\{b_n\}_{n=0}^{+\infty}$ with the same statistics as the original modes. After this transformation the Hamiltonian of the free environment becomes

$$H_\mathrm{E} = c\sum_{n=0}^{+\infty}\bigl(\alpha_n b_n^\dagger b_n + \sqrt{\beta_{n+1}} b_n^\dagger b_{n+1} + \sqrt{\beta_{n+1}} b_{n+1}^\dagger b_n\bigr). \tag{2.26}$$

The final form of H_I depends on the particular form of the interaction, and it is obtained by replacing $\int_X f_x h(x)\,\mathrm{d}x$ by $\|p_0\|_\lambda b_0$: for example, the operators in Eqs. (2.21) and (2.22) become, respectively,

$$H_\mathrm{I} = \|p_0\|_\lambda A_\mathrm{S}(b_0 + b_0^\dagger), \tag{2.27}$$

$$H_\mathrm{I} = \|p_0\|_\lambda (A_\mathrm{S}^\dagger b_0 + b_0^\dagger A_\mathrm{S}). \tag{2.28}$$

With the chain mapping, as illustrated in Fig. 2.3, the original continuous environment is transformed into a semi-infinite chain with nearest-neighbour interactions, with frequencies ω_n and couplings κ_n determined by the recurrence coefficients of the monic polynomials as $\omega_n = c\alpha_n$ and $\kappa_n = c\sqrt{\beta_{n+1}}$; the system, crucially, interacts with the first site of the chain only, with an interaction strength given by

$$\|p_0\|_\lambda = \left(\int_X J(x)\,\mathrm{d}x\right)^{\frac{1}{2}}. \tag{2.29}$$

For these reasons, the chain mapping is an important step towards the numerical simulation of OQSs strongly interacting with their environments: the transformed system can in fact be readily encoded into an MPS, enabling the use of powerful TN-based simulation techniques to investigate its dynamics and the properties of the ground state. Moreover, the unitarity of the map gives us the possibility of studying some properties of the environment as well, since the chain modes can be transformed back into the original environment.

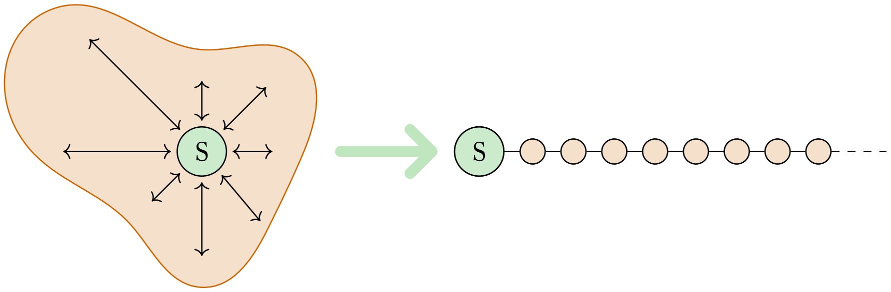

Figure 2.3. In the original configuration of the OQS, the system interacts with all the environment modes at the same time. The TEDOPA chain mapping transforms the environment into a semi-infinite chain with nearest-neighbour interactions, such that the system interacts only with the first site.

2.4.1 Thermalised TEDOPA

The TEDOPA chain mapping transforms the Hamiltonian but does not say anything in particular about the initial state of the environment. In fact, when we want to cast the problem into a TN form, for a generic initial state we are still left with the problem of how to represent it in this formalism. If the initial state is not the vacuum state—which is preserved by unitary maps—generally it does not get transformed into a simple product state of the new chain modes. When the state is not a pure state, moreover, it must be either represented as an MPO or an MPS with squared local dimensions, as seen in Section 1.4. In order to build the TN of the initial state, therefore, we must resort to other techniques such as the density-matrix renormalisation group (DMRG), when the state can be seen as the ground state of a simple Hamiltonian, or (for thermal states) imaginary-time evolution where $\exp(-\beta H)$ is interpreted as a time-evolution operator with time variable $\mathrm{i}\beta$ and the identity as initial state.

For a specific class of thermal bosonic environments, there is a way around this issue which leverages the results of Section 2.3. This technique, called T-TEDOPA, builds on TEDOPA by replacing the original environment at temperature $T > 0$ with a different environment at $T = 0$ which starts from the vacuum state [153]. This new initial state has an immediate MPS representation, thus removing the need to prepare it through other numerical methods.

The basic hypothesis is that we must be able to write the interaction Hamiltonian as a product of two self-adjoint operators, which in addition to the linearity requirements on the environment side means that

$$H_{\mathrm{I}} = A_{\mathrm{S}} \int_X (\mathrm{e}^{\mathrm{i}\varphi(x)} f_x^\dagger + \mathrm{e}^{-\mathrm{i}\varphi(x)} f_x) \sqrt{J(x)}\,\mathrm{d}x \equiv A_{\mathrm{S}} B_{\mathrm{E}} \tag{2.30}$$

with φ a real-valued function on X. The initial state is the thermal state

$$\rho_\beta = \frac{1}{\operatorname{tr}\exp(-\beta H_{\mathrm{E}})} \exp(-\beta H_{\mathrm{E}}), \qquad H_{\mathrm{E}} = \int_X x f_x^\dagger f_x \,\mathrm{d}x. \tag{2.31}$$

As there is only one term in H_{I} (compare Eq. (2.30) with Eq. (2.8)) and $\operatorname{tr}(\rho_\beta f_x) = \operatorname{tr}(\rho_\beta f_x^\dagger) = 0$, the only relevant quantity from Eq. (2.11) is the correlation function $f(t) = \operatorname{tr}(B_{\mathrm{E}}(t) B_{\mathrm{E}}(0) \rho_\beta)$, where we have already exploited the fact that the TTCF satisfies $f(t+s, s) = f(t, 0) \equiv f(t)$ because of the time-translation invariance of the dynamics. We find

$$\begin{aligned} f(t) = {} & \int_X \int_X \operatorname{tr}\Big((\mathrm{e}^{\mathrm{i}\varphi(x)} \mathrm{e}^{\mathrm{i}tx} f_x^\dagger + \mathrm{e}^{-\mathrm{i}\varphi(x)} \mathrm{e}^{-\mathrm{i}tx} f_x) \cdot \\ & \qquad \cdot (\mathrm{e}^{\mathrm{i}\varphi(x')} f_{x'}^\dagger + \mathrm{e}^{-\mathrm{i}\varphi(x')} f_{x'}) \rho_\beta \Big) \sqrt{J(x) J(x')}\,\mathrm{d}x'\,\mathrm{d}x = \\ = {} & \int_X \Big[\mathrm{e}^{\mathrm{i}tx} \operatorname{tr}(f_x^\dagger f_x \rho_\beta) + \mathrm{e}^{-\mathrm{i}tx} \big(\operatorname{tr}(f_x^\dagger f_x \rho_\beta) + 1 \big) \Big] J(x)\,\mathrm{d}x = \\ = {} & \int_X \Big[\mathrm{e}^{\mathrm{i}tx} \frac{1}{2} \Big(\coth\frac{\beta x}{2} - 1 \Big) + \mathrm{e}^{-\mathrm{i}tx} \frac{1}{2} \Big(\coth\frac{\beta x}{2} + 1 \Big) \Big] J(x)\,\mathrm{d}x. \end{aligned} \tag{2.32}$$

Now we take $X = [0, +\infty)$, as is customary with bosonic spectral densities. By splitting the two terms in separate integrals, and inverting the domain of the one that contains the $\mathrm{e}^{\mathrm{i}tx}$ factor, we get

$$\begin{aligned} f(t) = {} & \int_0^{+\infty} \mathrm{e}^{\mathrm{i}tx} \frac{1}{2} \Big(\coth\frac{\beta x}{2} - 1 \Big) J(x)\,\mathrm{d}x + \int_0^{+\infty} \mathrm{e}^{-\mathrm{i}tx} \frac{1}{2} \Big(\coth\frac{\beta x}{2} + 1 \Big) J(x)\,\mathrm{d}x = \\ = {} & \int_{-\infty}^{0} \mathrm{e}^{-\mathrm{i}tx} \frac{1}{2} \Big(-\coth\frac{\beta x}{2} - 1 \Big) J(-x)\,\mathrm{d}x + \\ & \qquad + \int_0^{+\infty} \mathrm{e}^{-\mathrm{i}tx} \frac{1}{2} \Big(\coth\frac{\beta x}{2} + 1 \Big) J(x)\,\mathrm{d}x = \\ = {} & \int_{-\infty}^{+\infty} \mathrm{e}^{-\mathrm{i}tx} \operatorname{sgn}(x) \frac{1}{2} \Big(\coth\frac{\beta x}{2} + 1 \Big) J(|x|)\,\mathrm{d}x. \end{aligned} \tag{2.33}$$

Going back to the second-to-last line in Eq. (2.32), we see that we would get the same result if we started from Eq. (2.30) with $X = (-\infty, +\infty)$, a different spectral density

$\tilde{J}_\beta \colon (-\infty, +\infty) \to [0, +\infty)$ defined as

$$\tilde{J}_\beta(x) \coloneqq \operatorname{sgn}(x) \frac{1}{2} \Big(\coth \frac{\beta x}{2} + 1 \Big) J(|x|) \tag{2.34}$$

and an initial state ρ_0 such that $\operatorname{tr}(f_x^\dagger f_x \rho_\beta) = 0$, i.e. the vacuum state. This means that the new environment E' given by

$$H_{\mathrm{E}'} = \int_{-\infty}^{+\infty} x f_x^\dagger f_x \, \mathrm{d}x, \qquad H_{\mathrm{I}'} = A_{\mathrm{S}} \int_{-\infty}^{+\infty} (\mathrm{e}^{\mathrm{i}\varphi(x)} f_x^\dagger + \mathrm{e}^{-\mathrm{i}\varphi(x)} f_x) \sqrt{\tilde{J}_\beta(x)} \, \mathrm{d}x \tag{2.35}$$

and starting from the vacuum state induces exactly the same dynamics in the open system as the original one.[4] We observe that the substitute environment does not have to be a physically realistic one: here for example the alternative one has a frequency domain that spans the whole real line, which is unrealistic. Now the TEDOPA chain mapping can be applied to the new configuration, where the dynamics is still unitary and the initial state is (exactly) represented by a trivial MPS.

2.4.2 The thermofield transformation

For fermionic environments, the T-TEDOPA method is not applicable since the typical system-environment coupling is an exchange interaction. Moreover, even the zero-temperature thermal state may contain a significant amount of correlation, since in the ground state of an ensemble of non-interacting fermions (described by Eq. (2.20)) the particles are not all in the single-mode ground state, but they are arranged in a Fermi sea. This means that with a fermionic environment, creating the TN of the initial state will be difficult in both $T = 0$ and $T > 0$ cases. To address this issue, the thermofield representation [152, 154] splits the modes above and below the chemical potential in two different environments so that (even when the temperature is not zero) the initial state is a tensor-product pure state. The transformation is essentially based on a local purification of the initial mixed state $\rho_{\beta,\mu}$ of the environment. For the following calculations, we write the frequency domain X so that it contains non-negative values only, and explicitly write the chemical potential in H_{E}. We also assume that system and environment interact through an interaction such as in Eq. (2.22). To each 'physical' mode $f_{1,x} \equiv f_x$ we associate an independent 'ancillary' mode $f_{2,x}$ so that the free

---

[4] We reintroduce the phase factors $\varphi(x)$ that were already present in Eq. (2.30), but it is clear from what we have seen so far that they have no effect at all on the reduced dynamics, since they disappear in the TTCF of Eq. (2.32). In fact, the transformation $f_x \mapsto \mathrm{e}^{\mathrm{i}\varphi(x)} f_x$ is a canonical transformation that leaves H_{E} invariant and completely eliminates the phase factors from the equations.

Hamiltonian of the extended environment is

$$H_{\mathrm{E}'} = \int_X (x-\mu)(f_{1,x}^\dagger f_{1,x} - f_{2,x}^\dagger f_{2,x})\,\mathrm{d}x, \tag{2.36}$$

essentially doubling the initial environment over negative energies. The additional environment is completely decoupled from the system and from the original environment, thus it does not change anything in the dynamics we are interested in. The next step is to combine the $f_{1,x}$ and $f_{2,x}$ modes into new fermionic modes $g_{0,x}$ and $g_{1,x}$ through the canonical transformation

$$\begin{pmatrix} g_{0,x} \\ g_{1,x} \end{pmatrix} = \begin{pmatrix} \cos\theta_{\beta,\mu}(x) & -\sin\theta_{\beta,\mu}(x) \\ \sin\theta_{\beta,\mu}(x) & \cos\theta_{\beta,\mu}(x) \end{pmatrix} \begin{pmatrix} f_{1,x} \\ f_{2,x}^\dagger \end{pmatrix}, \tag{2.37}$$

with the $\theta_{\beta,\mu}(x)$ angle to be determined by equating the TTCFs. From the original H_I

$$H_\mathrm{I} = B_\mathrm{E}^\dagger A_\mathrm{S} + A_\mathrm{S}^\dagger B_\mathrm{E}, \quad B_\mathrm{E} = \int_X f_{1,x} h(x)\,\mathrm{d}x \tag{2.38}$$

the only relevant TTCFs we need to compare are

$$\begin{aligned} c_1(t) &= \mathrm{tr}\big(B_\mathrm{E}(t)^\dagger B_\mathrm{E}(0)\rho_{\beta,\mu}\big), \\ c_2(t) &= \mathrm{tr}\big(B_\mathrm{E}(t) B_\mathrm{E}(0)^\dagger \rho_{\beta,\mu}\big) \end{aligned} \tag{2.39}$$

where $B_\mathrm{E}(t)$ is evolved according to the Hamiltonian H_E of the free environment:

$$B_\mathrm{E}(t) = \int_X f_{1,x}(t) h(x)\,\mathrm{d}x = \int_X \mathrm{e}^{-\mathrm{i}t(x-\mu)} f_{1,x} h(x)\,\mathrm{d}x. \tag{2.40}$$

We obtain

$$\begin{aligned} c_1(t) &= \int_X \mathrm{e}^{\mathrm{i}t(x-\mu)}\,\mathrm{tr}(\rho_{\beta,\mu} f_{1,x}^\dagger f_{1,x}) h(x)^2\,\mathrm{d}x, \\ c_2(t) &= \int_X \mathrm{e}^{-\mathrm{i}t(x-\mu)}\big[1 - \mathrm{tr}(\rho_{\beta,\mu} f_{1,x}^\dagger f_{1,x})\big] h(x)^2\,\mathrm{d}x. \end{aligned} \tag{2.41}$$

In terms of the newly defined modes g, the free Hamiltonian Eq. (2.36) reads

$$H_\mathrm{E}' = \int_X (x-\mu)(g_{0,x}^\dagger g_{0,x} + g_{1,x}^\dagger g_{1,x})\,\mathrm{d}x \tag{2.42}$$

while the interaction Hamiltonian becomes $H_\mathrm{I}' = (B_\mathrm{E}')^\dagger A_\mathrm{S} + A_\mathrm{S}^\dagger B_\mathrm{E}'$ with

$$B_\mathrm{E}' = \int_X (\cos\theta_{\beta,\mu}(x) g_{0,x} + \sin\theta_{\beta,\mu}(x) g_{1,x}) h(x)\,\mathrm{d}x. \tag{2.43}$$

We initialise the new environment with the '0' modes in the vacuum state and the '1' modes in the filled state, i.e. in a state $|\emptyset_0, \mathrm{F}_1\rangle$ such that $g_{0,x}|\emptyset_0, \mathrm{F}_1\rangle = 0$ and $g_{1,x}^\dagger|\emptyset_0, \mathrm{F}_1\rangle = 0$ for all $x \in X$; the TTCFs are

$$\begin{aligned}
c_1'(t) &= \langle\emptyset_0, \mathrm{F}_1|B_\mathrm{E}'(t)^\dagger B_\mathrm{E}'(0)|\emptyset_0, \mathrm{F}_1\rangle = \\
&= \int_X \mathrm{e}^{\mathrm{i}t(x-\mu)} h(x)^2 \big[(\cos\theta_{\beta,\mu}(x))^2 \underbrace{\langle\emptyset_0, \mathrm{F}_1|g_{0,x}^\dagger g_{0,x}|\emptyset_0, \mathrm{F}_1\rangle}_{0} + \\
&\qquad + (\sin\theta_{\beta,\mu}(x))^2 \underbrace{\langle\emptyset_0, \mathrm{F}_1|g_{1,x}^\dagger g_{1,x}|\emptyset_0, \mathrm{F}_1\rangle}_{1} + \\
&\qquad + \sin\theta_{\beta,\mu}(x)\cos\theta_{\beta,\mu}(x) \underbrace{\langle\emptyset_0, \mathrm{F}_1|g_{0,x}^\dagger g_{1,x} + g_{1,x}^\dagger g_{0,x}|\emptyset_0, \mathrm{F}_1\rangle}_{0}\big] \,\mathrm{d}x = \\
&= \int_X \mathrm{e}^{\mathrm{i}t(x-\mu)} (\sin\theta_{\beta,\mu}(x))^2 h(x)^2 \,\mathrm{d}x,
\end{aligned} \tag{2.44}$$

and analogously

$$c_2'(t) = \langle\emptyset_0, \mathrm{F}_1|B_\mathrm{E}'(t) B_\mathrm{E}'(0)^\dagger|\emptyset_0, \mathrm{F}_1\rangle = \int_X \mathrm{e}^{-\mathrm{i}t(x-\mu)} (\cos\theta_{\beta,\mu}(x))^2 h(x)^2 \,\mathrm{d}x, \tag{2.45}$$

so in order to have the TTCFs match, i.e. $c_1(t) = c_1'(t)$ and $c_2(t) = c_2'(t)$ for all t, it is sufficient to choose $\theta_{\beta,\mu}(x)$ so that

$$(\sin\theta_{\beta,\mu}(x))^2 = \mathrm{tr}(\rho_{\beta,\mu} f_{1,x}^\dagger f_{1,x}) = \frac{1}{\mathrm{e}^{\beta(x-\mu)}+1}, \tag{2.46}$$

just as in Ref. [154]. The interaction Hamiltonian H_I' becomes then

$$H_\mathrm{I}' = \int_X \big[h_{\beta,\mu}^{(0)}(x)(A_\mathrm{S}^\dagger g_{0,x} + g_{0,x}^\dagger A_\mathrm{S}) + h_{\beta,\mu}^{(1)}(x)(A_\mathrm{S}^\dagger g_{1,x} + g_{1,x}^\dagger A_\mathrm{S})\big] \,\mathrm{d}x, \tag{2.47}$$

with couplings $h^{(0)}(x)$ and $h^{(1)}(x)$ defined by

$$\begin{aligned}
h_{\beta,\mu}^{(0)}(x) &= \cos\theta_{\beta,\mu}(x)\, h(x), \\
h_{\beta,\mu}^{(1)}(x) &= \sin\theta_{\beta,\mu}(x)\, h(x).
\end{aligned} \tag{2.48}$$

The square of these two functions will be the spectral density functions of the new environments, the initially empty and the initially filled one, respectively. A graphical example of this transformation can be found in Fig. 2.4. Just as in T-TEDOPA, we have moved the temperature dependence into the spectral density function of the environment in a way that preserves the TTCFs of the free environment, and with them the dynamics of the reduced state of the OQS. The initial state of the new environment, both in the empty and in the filled case, can be trivially represented as an MPS: therefore, at the cost of doubling the size of the environment, we can avoid the issue of computing the MPS of the correlated initial state.

2.4.3 Environment merging

The equivalence under equality of TTCFs allows us to combine different baths as well. For example, imagine a system that interacts with *two* environments at once: by applying the thermofield transformation to the environments, we obtain two pairs of empty/full environments. This configuration is not one-dimensional anymore and so it would not fit naturally in an MPS, without introducing artificial long-range interactions, and with that a highly non-local entanglement structure which could be detrimental to the efficiency of time-evolution algorithms. To remedy this, we can perform another replacement which merges the two empty environments together, and the two filled environments together (clearly the same approach works for any number of environments). It is easy to see that if two quadratic environments[5]

$$H_{\mathrm{E}_1} = \int_{X_1} x f_{1,x}^\dagger f_{1,x} \,\mathrm{d}x, \qquad H_{\mathrm{E}_2} = \int_{X_2} x f_{2,x}^\dagger f_{2,x} \,\mathrm{d}x \tag{2.49}$$

interact in the same way with the system, for example as

$$H_{\mathrm{I}_1} = \int_{X_1} (A_{\mathrm{S}}^\dagger f_{1,x} + f_{1,x}^\dagger A_{\mathrm{S}}) h_1(x) \,\mathrm{d}x, \quad H_{\mathrm{I}_2} = \int_{X_2} (A_{\mathrm{S}}^\dagger f_{2,x} + f_{2,x}^\dagger A_{\mathrm{S}}) h_2(x) \,\mathrm{d}x, \tag{2.50}$$

they can be added together into a new environment defined by

$$H_{\mathrm{E}} = \int_{X} x f_{x}^\dagger f_{x} \,\mathrm{d}x, \quad H_{\mathrm{I}} = \int_{X} (A_{\mathrm{S}}^\dagger f_{x} + f_{x}^\dagger A_{\mathrm{S}}) h(x) \,\mathrm{d}x, \tag{2.51}$$

such that $X = X_1 \cup X_2$ and $h^2 = h_1^2 + h_2^2$, meaning that by adding the two original spectral density functions (extended as zero outside their respective domains) we obtain the one for the combined environment.

---

[5] Here it is crucial that the frequency domains X_1 and X_2 are shifted according to the chemical potential so that the latter is zero: this way the two Hamiltonian operators H_{E_1} and H_{E_2} have an identical functional form.

The procedure is graphically illustrated in Fig. 2.4: in that picture, we start on step (1) from two leads with spectral densities

$$J_\alpha(x) = \frac{2}{\pi}\sqrt{(x - x^\alpha_{\text{min}})(x^\alpha_{\text{max}} - x)}, \tag{2.52}$$

with $\alpha = \{\mathrm{L}, \mathrm{R}\}$, with different chemical potentials $\mu_\mathrm{L} < \mu_\mathrm{R}$ and inverse temperatures $\beta_\mathrm{L} < \beta_\mathrm{R}$. At step (2) the spectral densities are split in the respective components

$$\begin{aligned} J^{(0)}_\alpha(x) &= \frac{1}{\mathrm{e}^{-\beta_\alpha x} + 1} J_\alpha(x), \\ J^{(1)}_\alpha(x) &= \frac{1}{\mathrm{e}^{\beta_\alpha x} + 1} J_\alpha(x) \end{aligned} \tag{2.53}$$

and finally at step (3) we recombine $J^{(0)}_\mathrm{L}$ with $J^{(0)}_\mathrm{R}$, and $J^{(0)}_\mathrm{L}$ with $J^{(0)}_\mathrm{R}$ obtaining two new environments with spectral densities

$$J^{(0)}(x) = J^{(0)}_\mathrm{L}(x) + J^{(0)}_\mathrm{R}(x), \qquad J^{(1)}(x) = J^{(1)}_\mathrm{L}(x) + J^{(1)}_\mathrm{R}(x), \tag{2.54}$$

starting from the vacuum and the filled state, respectively.

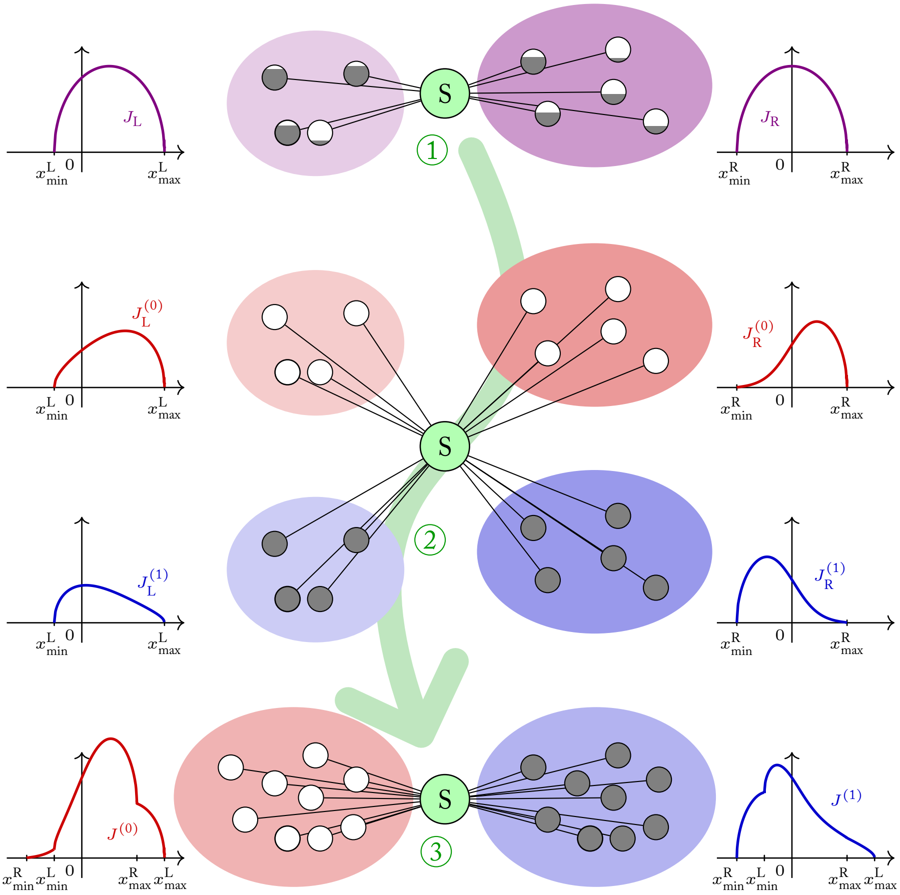

Figure 2.4. Recombination of two fermionic leads. (1) The two leads ('L' and 'R') interact with a system. (2) By using the thermofield transformation, each lead is split into its '0' and '1' components, depicted in red and blue, which start from the empty and filled state, respectively. (3) The two hot components are merged, and so the two cold ones.

Markovian closure for fermionic environments 3

One of the main limitations of the original TEDOPA formulation is its quadratic scaling $O(t^2)$ with the physical simulation time t [153]. As a matter of fact, the TEDOPA system we obtain from the chain mapping is infinite, but in order to run an actual computer simulation we would need to truncate it at a certain length. Such a truncation must be suitably chosen so as not to introduce finite size effects: in Ref. [155] it is shown that when the system is composed of a single site with free Hamiltonian $H_\text{S} = \varepsilon c^\dagger c$, then the orthogonal polynomial discretisation scheme of a spectral density function supported on $[a, b]$ with N sites reproduces basically the exact time evolution of the system Green's function up to a time

$$t_\text{max} = 2\frac{2N+1}{b-a}. \tag{3.1}$$

This result shows that the number of sites must be incremented linearly the longer we want the simulation to go on. While, more in general, the exact dependence on the simulation time is hard to define, for sufficiently large t_max, Lieb-Robinson bound techniques suggest that $N(t_\text{max})$ scales linearly in t_max [156]. Coupled with the fact that the simulation also already depends linearly on t_max since the state has to be computed at more time instants, one obtains the $O(t^2)$ dependence stated above.

As a result, long-time dynamics may prove difficult to simulate with TEDOPA. A context where long-time dynamics are typically required is the investigation of NESS properties of an OQS interacting with external environments. For example, in the dynamical mean-field theory (DMFT) approach to strongly correlated materials, the challenging part of the program is to solve the impurity model; this requires finding the impurity spectral function, which can be accurately determined only if a NESS is reached. A way around this problem has been put forward, for the case of bosonic environments, in Ref. [35] where a semi-infinite uniform chain of harmonic oscillators is mimicked by a small collection of damped harmonic modes, called Markovian closure (MC). The use of the MC allows for a quadratic reduction of the dependence of the simulation complexity on the physical simulation time, i.e. from $O(t^2)$ to $O(t)$, and makes it possible to use the algorithm to determine long-time OQS dynamics.

The analysis presented in Refs. [157, 158], in particular, reveals a steady and fast build-up of strong correlations between the environmental fermionic modes induced by the interaction with the impurity. Even by adopting state-of-the-art strategies for the mitigation of the effects of the appearance of such correlations, the determination of the OQS dynamics over long times, which is typically required to reach the NESS state, remains a computationally challenging task. The possibility of exploiting the MC construct in the fermionic setting would therefore represent a powerful asset. In this chapter we fill this gap by introducing a fermionic MC, and illustrate its application and computational impact by means of some relevant case studies.

Data presented in this chapter were obtained using the ITensor library [159] in the *Julia* programming language. All code is available at [3].

3.1 Reference model

Let us write down the physical model to which we will apply the MC in this chapter: consider a general impurity model interacting with continuous fermionic baths (leads). The derivation that follows is essentially independent of the details of the system, but for simplicity we model it here as an m-mode fermionic system with annihilation and creation operators d_λ and $d_\lambda^\dagger$, $\lambda \in \{1, \dots, m\}$, satisfying the canonical anticommutation relations (CARs) $\{d_\lambda, d_\nu^\dagger\} = \delta_{\lambda\nu}$; the system Hamiltonian H_S is an arbitrary function of these operators. The environment is composed of two leads, called left (L) and right (R) leads, i.e. two continua of (non-interacting) fermionic modes, for which we assume a linear dispersion relation; each of them is indexed by the frequency $x \in X_\alpha$ (we set the chemical potential to zero, and allow modes with negative frequency), with $\alpha \in \{\mathrm{L}, \mathrm{R}\}$; the operators $f_{\alpha,x}$ and $f_{\alpha,x}^\dagger$ obey the CARs $\{f_{\alpha,x}, f_{\alpha',x'}^\dagger\} = \delta_{\alpha\alpha'}\delta(x - x')$. The evolution of the overall system-bath complex is generated by the Hamiltonian

$$H_\mathrm{SE} = H_\mathrm{S} + H_\mathrm{E} + H_\mathrm{I}, \tag{3.2}$$

$$H_{\mathrm{E}} = H_{\mathrm{E,L}} + H_{\mathrm{E,R}} = \sum_{\alpha\in\{\mathrm{L,R}\}} \int_{X_\alpha} x f^\dagger_{\alpha,x} f_{\alpha,x}\,\mathrm{d}x, \tag{3.3}$$

$$H_{\mathrm{I}} = \sum_{\alpha\in\{\mathrm{L,R}\}} \sum_{\lambda=1}^{m} \int_{X_\alpha} h_{\lambda,\alpha}(x)(A^\dagger_\lambda f_{\alpha,x} + f^\dagger_{\alpha,x} A_\lambda)\,\mathrm{d}x. \tag{3.4}$$

The interaction between the system and the environment is specified by H_{I}, with A_λ and $f_{\alpha,\omega}$ and $f^\dagger_{\alpha,\omega}$ operators acting on mode λ of the system and on the left ($\alpha = \mathrm{L}$) or right ($\alpha = \mathrm{R}$) lead, respectively. The interaction (hybridisation) strength between the system and the fermionic mode at frequency x is given by $h_{\lambda,\alpha}(x) = \kappa_{\lambda,\alpha}\sqrt{J_\alpha(x)}$, where the functions $J_\alpha\colon X_\alpha \to [0,+\infty)$, $\alpha \in \{\mathrm{L},\mathrm{R}\}$, are spectral densities, while the coefficients $\kappa_\lambda \in \mathbb{R}$ account for possibly different interaction strengths between the environments and the different modes of the system. In simpler words, our scheme represents a central set of modes, which make up the *system*, interacting with external environments. Each system mode can interact with more than one environment at a time, and vice versa; it is essential, however, that the interaction between a given environment and different system modes differs by at most a constant, so that we can compress the environments together.

The system-bath complex of Eq. (3.2) can describe different relevant fermionic (spinless) OQSs such as the resonant-level model (RLM) and the interacting resonant-level model (IRLM), and can be straightforwardly generalised to the case of spinful fermions, as well as more than two leads. We point out that if the system Hamiltonian H_{S} comprises only quadratic terms in d_λ and $d^\dagger_\lambda$, and A_λ is linear in those operators, then the model fixed by Eq. (3.2) is Gaussian and exactly solvable, i.e. by means of exact diagonalization [160].

We assume that the overall system is initially in a factorised state

$$\rho_{\mathrm{SE}}(0) = \rho_{\mathrm{S}}(0)\,\rho_{\mathrm{E,L}}(0)\,\rho_{\mathrm{E,R}}(0), \tag{3.5}$$

for an arbitrary $\rho_{\mathrm{S}}(0)$, while $\rho_{\mathrm{E,L}}(0)$ ($\rho_{\mathrm{E,R}}(0)$) is the thermal state of the left (right) environment at a given inverse temperature β_{L} (β_{R}):

$$\rho_{\mathrm{E},\alpha}(0) = \frac{1}{\mathcal{Z}_\alpha}\exp(-\beta_\alpha H_{\mathrm{E},\alpha}), \tag{3.6}$$

where $\mathcal{Z}_\alpha = \operatorname{tr}\exp(-\beta_\alpha H_{\mathrm{E},\alpha})$ is the partition function.

With these assumptions, following the recombination procedure outlined in Section 2.4.3, we can reproduce the reduced dynamics of the systems replacing the two leads by two new separate environments, an initially empty one and an initially filled one, with zero temperature and chemical potential. Starting from the spectral density functions J_L and J_R, we obtain with this procedure the spectral densities

$$\begin{aligned} J^{(0)}(x) &= \frac{1}{\mathrm{e}^{-\beta_\mathrm{L}x}+1}J_\mathrm{L}(x) + \frac{1}{\mathrm{e}^{-\beta_\mathrm{R}x}+1}J_\mathrm{R}(x), \\ J^{(1)}(x) &= \frac{1}{\mathrm{e}^{\beta_\mathrm{L}x}+1}J_\mathrm{L}(x) + \frac{1}{\mathrm{e}^{\beta_\mathrm{R}x}+1}J_\mathrm{R}(x). \end{aligned} \tag{3.7}$$

We shall call the environment that starts from the empty state 'hot' (with spectral density $J^{(0)}$) and the environment that starts from the filled state 'cold' (with spectral density $J^{(1)}$), as they represent the upper and lower band of the original frequency domain, respectively. With the chain mapping introduced in Section 2.4 we then map the 'hot' and the 'cold' environment each onto a TEDOPA chain (as in Fig. 3.1), obtaining at last a pair of environments that we can directly represent as MPSs.

3.2 Derivation of the fermionic Markovian closure

As we already observed at the beginning of this chapter, in actual numerical simulations the chains generated by TEDOPA need to be truncated beyond a certain length, which must scale proportionally to the total physical simulation time in order not to induce finite-size effects on the dynamics. The MC mechanism, proposed in Ref. [35] for bosonic environments, provides a solution to this problem. In the following text we will introduce the basic ideas behind the MC and formulate an equivalent mechanism for fermionic environments. The key ingredient in this is the equivalence theorem seen in Section 2.3, which allows us to establish when a unitarily evolving environment induces the same reduced dynamics on an OQS as an auxiliary environment undergoing a non-unitary Lindblad evolution.

3.2.1 Asymptotic coefficients

The main starting point of the MC construction is the fact that a certain class of spectral density functions behave 'asymptotically the same'. The *Szegő class*, originally defined in Ref. [161, 1, Section 12.1], can be defined as the set of non-negative measurable functions on $[a,b]$ such that

$$\int_a^b \frac{\log f(t)}{\sqrt{(b-x)(x-a)}}\,\mathrm{d}t > -\infty. \tag{3.8}$$

If such an f is considered as a measure on $[a,b]$, as in the TEDOPA chain mapping procedure, the recurrence coefficients α_n and β_n of the monic orthogonal polynomials

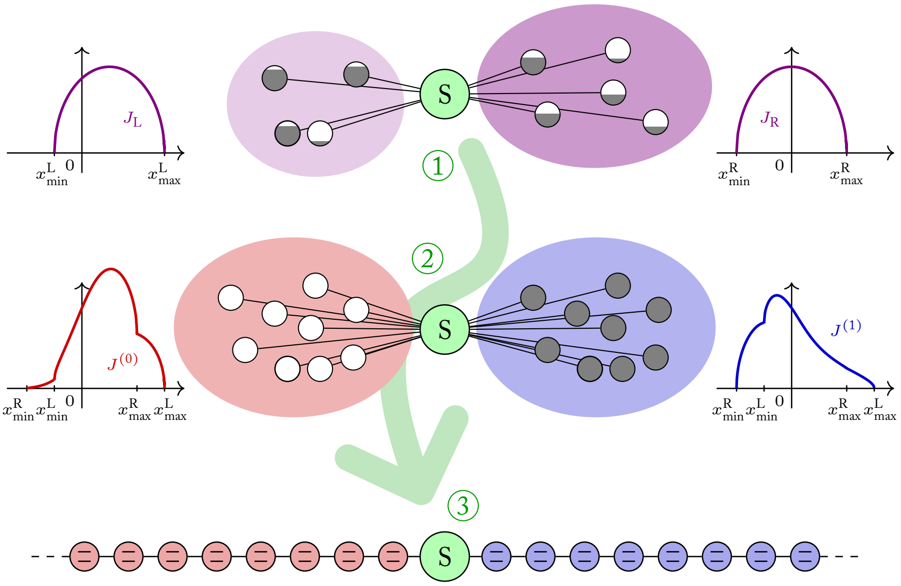

Figure 3.1. Chain mapping of two fermionic leads. (1) The two leads ('L' and 'R') interact with a system. (2) The spectral density functions of the two leads are recombined into a 'hot' environment and a 'cold' environment, as in Eq. (3.7). (3) Each of the two new environments is transformed into TEDOPA chains.

it generates show an interesting convergence property [151]: their sequences converge to limiting values

$$\lim_{n\to+\infty} \alpha_n = \frac{a+b}{2}, \qquad \lim_{n\to+\infty} \beta_n = \frac{(b-a)^2}{16} \tag{3.9}$$

that depend only on the endpoints of the support of the function. This convergence reflects on the coefficients of the Hamiltonian of the TEDOPA chain, namely the on-site energies and the coupling parameters, that converge to $\mathit{\Omega} = \lim_{n\to+\infty} \alpha_n$ and $K = \lim_{n\to+\infty} \sqrt{\beta_n}$, respectively.

The Szegő class is a quite broad set, that contains many of the spectral density functions commonly encountered in a physical setting, as shown for example in [151]: power-law, Lorentzian, flat, semicircle, exponential functions all fall into this class.

Band gaps are a clear obstacle to a function belonging to this class, since Eq. (3.8) would not be satisfied. Numerically, if there is a band gap then the chain coefficients do not converge, thus the Markovian closure cannot be used. More in general, Szegő-class functions may show a certain degree of non-analiticity: spectral density functions can even be discontinuous or diverging within the domain, as long as they satisfy Eq. (3.8). Discontinuities, for example, easily arise when combining environments as we showed Section 2.4.3.

Let us consider, without loss of generality, the Szegő class on the interval $[-1, 1]$, which we will denote by G. It easy to see that

- if f, g are in G then $f + g$ and fg are too in G,
- if f is in G and h is a bounded function, then fh is in G,
- the concatenation of two functions, i.e. the new function

  $$u(t) = \begin{cases} f(2t+1) & -1 \leq t < 0 \\ g(2t-1) & 0 \geq t \geq 1, \end{cases} \tag{3.10}$$

  is in G, too.

When combining environments as in Section 2.4.3, the only critical operation remains truncating the domain of a function $f \in G$, or equivalently multiplying it by the characteristic function of a subinterval, for example $[a, 1]$ with $a > -1$. This could be an issue if $f(x) \to 0$ as $x \to a^+$ in such a way that $\log f(x)$, combined with the denominator in Eq. (3.8), makes the integral diverge.[1] If f is positive in the interior of its domain, instead, truncation is never problematic. In conclusion, we can say that the conditions for spectral density functions, as well as the result of the merging multiple ones, to belong to the Szegő class are very light. When merging two baths, the condition $X_\mathrm{L} \cap X_\mathrm{R} \neq \emptyset$ on their supports is indeed a very mild one, since in conduction schemes both leads do typically include the system transition frequency.

Once the coefficients have converged towards their asymptotic values up to a desired tolerance $\varepsilon > 0$, meaning that we can identify N_E such that for all $n > N_\mathrm{E}$ we have $|\omega_n - \Omega| < \varepsilon$ and $|\kappa_n - K| < \varepsilon$, the part of the chain that includes the modes $c_n^{(\dagger)}$ for $n > N_\mathrm{E}$ can be approximated by a uniform chain whose coefficients ω_n and κ_n are all set equal to Ω and K respectively. The resulting chain Hamiltonian reads

---

[1] A concrete example can be found in [1, Appendix B], where we analyse the Szegő class in more detail.

$$H'_{\mathrm{E}} = \sum_{n=0}^{N_{\mathrm{E}}} \omega_n c_n^\dagger c_n + \sum_{n=0}^{N_{\mathrm{E}}} \kappa_n (c_{n+1}^\dagger c_n + c_n^\dagger c_{n+1}) + \\ + \sum_{n=N_{\mathrm{E}}+1}^{+\infty} \mathit{\Omega} c_n^\dagger c_n + \sum_{n=N_{\mathrm{E}}+1}^{+\infty} K (c_{n+1}^\dagger c_n + c_n^\dagger c_{n+1}). \quad (3.11)$$

For any assigned spectral density, it is always possible to make a suitable choice of N_{E} and to estimate the corresponding error [150]. As we will see in the next sections, moreover, in the case of spectral densities that are relevant in the fermionic setting, the approximate convergence is reached within a very small number ($N_{\mathrm{E}} < 10$) of chain sites. It helps the physical intuition to think of the uniform chain as a 'runway'. While the inhomogeneities in the couplings or in the frequencies in the part of the chain closer to the system ($n < N_{\mathrm{E}}$) scatter the excitations travelling in that region, particles/holes entering the uniform region are instead propagated towards $n \to +\infty$ without the possibility of being scattered back [162].

Such a flat TEDOPA chain can be traced back (by inverting the chain mapping) to a continuous fermionic environment, that we call *residual environment*, characterised by

$$H_{\mathrm{E}}^{\mathrm{res}} = \int_{\mathit{\Omega}-2K}^{\mathit{\Omega}+2K} x f_x^\dagger f_x \,\mathrm{d}x, \quad (3.12)$$

interacting with the N_{E}-th mode (the last one of the inhomogeneous part) through

$$H_{\mathrm{I}}^{\mathrm{res}} = \int_{\mathit{\Omega}-2K}^{\mathit{\Omega}+2K} \sqrt{J_\infty(x)} (c_{N_{\mathrm{E}}}^\dagger f_x + f_x^\dagger c_{N_{\mathrm{E}}}) \,\mathrm{d}x, \quad (3.13)$$

where $J_\infty \colon [\mathit{\Omega} - 2K, \mathit{\Omega} + 2K] \to [0, +\infty)$ is the spectral density function

$$J_\infty(x) = \frac{1}{2\pi} \sqrt{(2K - \mathit{\Omega} + x)(2K + \mathit{\Omega} - x)}. \quad (3.14)$$

also known as the 'Wigner' semicircle [157].

The convergence toward a flat residual spectrum can be seen as an embedding of the original system into an enlarged set of degrees of freedom that evolve in a Markovian way [163]. In Ref. [35] it was proven, for a bosonic environment mapped onto a single TEDOPA chain, that the residual environment having spectral density J_∞ can be replaced with an auxiliary system, the Markovian closure, consisting of a small number of damped, interacting bosonic modes subject to a non-unitary evolution. With this mechanism the semi-infinite homogeneous chain is replaced by a finite auxiliary environment acting as an absorber for the excitations travelling along the homogeneous part of the chain. The MC is moreover *universal*, in the sense that it can be applied to all chain mappings of bosonic environments whose spectral density belongs to the Szegő class.

In what follows, we will provide an analogous construction for fermionic environments. We will show how the semi-infinite homogeneous chain of fermionic modes governed by the Hamiltonian

$$\sum_{n=N_{\mathrm{E}}+1}^{+\infty} \mathit{\Omega} c_n^\dagger c_n + \sum_{n=N_{\mathrm{E}}+1}^{+\infty} K(c_{n+1}^\dagger c_n + c_n^\dagger c_{n+1}) \tag{3.15}$$

from Eq. (3.11) can be replaced by a fermionic Markovian closure (FMC), i.e. where the auxiliary modes follow the Fermi statistics, having the same structure as the one devised for bosonic systems. The resulting closure, moreover, has the same universality character of its bosonic counterpart: it can be applied to all fermionic environments with a Szegő-class spectral density.

3.2.2 Derivation

We saw in Section 2.3 how a fermionic bath can be mapped to an auxiliary system comprising two different baths, one of which is initially completely filled and the other completely empty, representing the original environmental modes below and above the chemical potential, respectively. After applying the chain mapping to both the new environments separately, we obtain two TEDOPA chains, and we will map their uniform regions to sets of damped fermionic modes with appropriate parameters. The equivalence of the reduced dynamics generated by the different environments rests, as usual, on the equivalence of the TTCFs of the interaction Hamiltonian.

This time, since we want to replace not the whole environment but only the uniform part of the chain, the role of the 'system' is taken up by the original OQS *and* the TEDOPA chain up to $n < N_{\mathrm{E}}$, while the 'environment' is composed of the chain modes with $n \geq N_{\mathrm{E}}$. We immediately observe two facts.

1. The two chains are independent, therefore the construction of the MC for one of them will not depend on the other one. This means that we can focus on the two environments separately.

2. Within each of the two chains, we can discard the original OQS and everything up to $n = N_{\mathrm{E}} - 1$, because those modes do not interact with the new 'environment': only the $n = N_{\mathrm{E}}$ mode matters.

For both the initially empty and the initially filled environments, we start by defining

$$A_{\mathrm{S}} := c_{N_{\mathrm{E}}}, \quad B_{\mathrm{E}} := \int_{\Omega-2K}^{\Omega+2K} \sqrt{J_{\infty}(x)} f_x \, \mathrm{d}x \tag{3.16}$$

so that we can write the interaction Hamiltonian in Eq. (3.13) as

$$H_{\mathrm{I}}^{\mathrm{res}} = A_{\mathrm{S}}^{\dagger} B_{\mathrm{E}} + B_{\mathrm{E}}^{\dagger} A_{\mathrm{S}}. \tag{3.17}$$

Given the initial condition of the transformed environments, the only non-vanishing TTCFs are

$$\begin{aligned} c_0(t) &= \langle B_{\mathrm{E}}(t) B_{\mathrm{E}}(0)^{\dagger} \rangle, \\ c_1(t) &= \langle B_{\mathrm{E}}(t)^{\dagger} B_{\mathrm{E}}(0) \rangle \end{aligned} \tag{3.18}$$

where the average is computed on the vacuum state $|\emptyset\rangle$ for the initially empty lead, and from the filled state $|\mathrm{F}\rangle$ for the initially filled one. In the former case, we have $B_{\mathrm{E}}|\emptyset\rangle = 0$, so $c_1(t)$ is identically zero and only

$$c_0(t) = \langle \emptyset | B_{\mathrm{E}}(t) B_{\mathrm{E}}(0)^{\dagger} | \emptyset \rangle = \int_{\Omega-2K}^{\Omega+2K} \mathrm{e}^{-\mathrm{i}tx} J_{\infty}(x) \, \mathrm{d}x. \tag{3.19}$$

remains; in the latter case, $B_{\mathrm{E}}^{\dagger}|\mathrm{F}\rangle = 0$ instead so we only have

$$c_1(t) = \langle \mathrm{F} | B_{\mathrm{E}}(t)^{\dagger} B_{\mathrm{E}}(0) | \mathrm{F} \rangle = \int_{\Omega-2K}^{\Omega+2K} \mathrm{e}^{\mathrm{i}tx} J_{\infty}(x) \, \mathrm{d}x. \tag{3.20}$$

By computing the Fourier transform of J_{∞}, we get

$$\begin{aligned} c_0(t) &= K^2 \mathrm{e}^{-\mathrm{i}\Omega t} C_{\mathrm{sc}}(2Kt), \\ c_1(t) &= K^2 \mathrm{e}^{\mathrm{i}\Omega t} C_{\mathrm{sc}}(2Kt), \end{aligned} \tag{3.21}$$

where

$$C_{\mathrm{sc}}(x) = \int_{-1}^{1} \mathrm{e}^{-\mathrm{i}xy} \frac{2}{\pi} \sqrt{1-y^2} \, \mathrm{d}y = j_0(x) + j_2(x) \tag{3.22}$$

is the transform of the unit-radius semicircle spectral density

$$J_{\mathrm{sc}}(y) = \frac{2}{\pi} \sqrt{1-y^2} \tag{3.23}$$

and j_n denotes the n-th Bessel functions of the first kind.

We have now to define a pseudomode environment that reproduces the TTCFs in Eq. (3.21). We will define new set of N_C fermionic modes $\{a_i\}_{i=1}^{N_\mathrm{C}}$. Let us start from the residual environment associated to a chain in the vacuum state: this system acts (exactly as in the bosonic case) as a perfect absorber of excitations, so we take its initial state to be the vacuum $|\emptyset\rangle\langle\emptyset|$ and describe its free dynamics as generated by the Lindblad operator $\mathcal{L}_\mathrm{R}^{(0)} = -\mathrm{i}[H_\mathrm{R}^{(0)}, -] + \mathcal{D}_\mathrm{R}^{(0)}$ (in which $[x, -]$ denotes the function $y \mapsto [x, y]$) with

$$H_\mathrm{R}^{(0)} = \sum_{i,j=1}^{N_\mathrm{C}} \Lambda_{ij}^{(0)} a_i^\dagger a_j \tag{3.24}$$

and

$$\mathcal{D}_\mathrm{R}^{(0)}(\rho) = \sum_{i,j=1}^{N_\mathrm{C}} \Gamma_{ij}^{(0)} \big(a_j \rho a_i^\dagger - \tfrac{1}{2}\{\rho, a_i^\dagger a_j\}\big) \tag{3.25}$$

for some complex matrices $\Lambda^{(0)}$ and $\Gamma^{(0)}$ to be determined; note that the vacuum is stationary with respect to the evolution operator, i.e. $\mathcal{L}_\mathrm{R}^{(0)}(|\emptyset\rangle\langle\emptyset|) = 0$. We define a new operator B_R, which plays the same role as B_E in the exchange-interaction Hamiltonian in Eq. (3.17),

$$B_\mathrm{R} \coloneqq \sum_{k=1}^{N_\mathrm{C}} \zeta_k a_k, \tag{3.26}$$

for some $\zeta \in \mathbb{C}^N$. This operator generates the new correlation functions

$$\begin{aligned} c_0'(t) &= \langle B_\mathrm{R}(t) B_\mathrm{R}(0)^\dagger \rangle, \\ c_1'(t) &= \langle B_\mathrm{R}(t)^\dagger B_\mathrm{R}(0) \rangle, \end{aligned} \tag{3.27}$$

that must match, respectively, $c_0(t)$ and $c_1(t)$ from Eq. (3.21) if we want the reduced dynamics to be the same.

Let us start by considering $c_0'(t)$: given the initial condition $a_j(0) = a_j$ for the time evolution of the annihilation operators, the only non-zero TTCF is

$$\begin{aligned} c_0'(t) = \langle\emptyset| B_\mathrm{R}(t) B_\mathrm{R}(0)^\dagger |\emptyset\rangle &= \\ = \sum_{k=1}^{N_\mathrm{C}} \sum_{l=1}^{N_\mathrm{C}} \zeta_k \overline{\zeta_l} \, \langle\emptyset| a_k(t) a_l^\dagger |\emptyset\rangle &= \\ = \sum_{k=1}^{N_\mathrm{C}} \sum_{l=1}^{N_\mathrm{C}} \zeta_k \overline{\zeta_l} \, \langle\emptyset| a_k(t) |1_l\rangle&, \end{aligned} \tag{3.28}$$

where $|1_l\rangle \coloneqq a_l^\dagger|\emptyset\rangle$ is the state in which the l-th mode (only) is occupied.

Let $v^i(t)$ be the vector whose j-th component is $v^i_j(t) = \langle\emptyset|a_i(t)|1_j\rangle$. It satisfies

$$\frac{\mathrm{d}v^i_j(t)}{\mathrm{d}t} = \langle\emptyset|\dot{a}_i(t)|1_j\rangle = \langle\emptyset|\mathcal{L}'_{\mathrm{R}}\big(a_i(t)\big)|1_j\rangle, \tag{3.29}$$

where $\mathcal{L}'_{\mathrm{R}} = \mathrm{i}[H^{(0)}_{\mathrm{R}}, -] + \mathcal{D}^{(0)\prime}_{\mathrm{R}}$ is the adjoint Lindblad operator that appears in the equation of motion in the Heisenberg picture, and starts from the value

$$v^i_j(0) = \langle\emptyset|a_i(0)|1_j\rangle = \langle\emptyset|a_i|1_j\rangle = \delta_{ij}. \tag{3.30}$$

Within the right-hand side of the previous differential equation we find

$$\begin{aligned}
&\langle\emptyset|H^{(0)}_{\mathrm{R}}a_i(t)|1_j\rangle = 0,\\
&\langle\emptyset|a_i(t)H^{(0)}_{\mathrm{R}}|1_j\rangle = \sum_{m=1}^{N_{\mathrm{C}}} \Lambda^{(0)}_{mj}\,\langle\emptyset|a_i(t)|1_m\rangle,\\
&\langle\emptyset|\mathcal{D}^{(0)\prime}_{\mathrm{R}}\big(a_i(t)\big)|1_j\rangle = -\frac{1}{2}\sum_{m=1}^{N_{\mathrm{C}}} \Gamma^{(0)}_{mj}\,\langle\emptyset|a_i(t)|1_m\rangle :
\end{aligned} \tag{3.31}$$

the differential equation becomes then

$$\begin{cases}\dfrac{\mathrm{d}v^i(t)}{\mathrm{d}t} = \big(-i\Lambda^{(0)} - \tfrac{1}{2}\Gamma^{(0)}\big)^{\mathrm{T}} v^i(t),\\ v^i_j(0) = \delta_{ij},\end{cases} \tag{3.32}$$

which is solved by

$$v^i_j(t) = \sum_{m=1}^{N_{\mathrm{C}}} \exp(-\mathrm{i}t\Lambda^{(0)} - \tfrac{1}{2}t\Gamma^{(0)})_{mj} v^i_m(0) = \exp(-\mathrm{i}t\Lambda^{(0)} - \tfrac{1}{2}t\Gamma^{(0)})_{ij}. \tag{3.33}$$

The correlation function we are now considering is therefore

$$c'_0(t) = \sum_{k=1}^{N_{\mathrm{C}}}\sum_{l=1}^{N_{\mathrm{C}}} \zeta_k\overline{\zeta_l}\exp(-\mathrm{i}t\Lambda^{(0)} - \tfrac{1}{2}t\Gamma^{(0)})_{kl} = \langle\zeta, \exp(-\mathrm{i}t\Lambda^{(0)} - \tfrac{1}{2}t\Gamma^{(0)})^{\mathrm{T}}\zeta\rangle. \tag{3.34}$$

When the TEDOPA chain starts from the completely filled state $|\mathrm{F}\rangle\langle\mathrm{F}|$ instead, we choose a Lindblad operator with the following terms:

$$\begin{aligned} H_{\mathrm{R}}^{(1)} &= \sum_{i,j=1}^{N_{\mathrm{C}}} \Lambda_{ij}^{(1)} a_i^\dagger a_j, \\ \mathcal{D}_{\mathrm{R}}^{(1)}(\rho) &= \sum_{i,j=1}^{N_{\mathrm{C}}} \Gamma_{ij}^{(1)} \big(a_i^\dagger \rho a_j - \tfrac{1}{2}\{\rho, a_j a_i^\dagger\} \big) \end{aligned} \tag{3.35}$$

which satisfies $\mathcal{D}_{\mathrm{R}}^{(1)}(|\mathrm{F}\rangle\langle\mathrm{F}|) = 0$. The only nonzero correlation function is now

$$c_1'(t) = \langle \mathrm{F}|B_{\mathrm{R}}(t)^\dagger B_{\mathrm{R}}(0)|\mathrm{F}\rangle = \sum_{k=1}^{N_{\mathrm{C}}}\sum_{l=1}^{N_{\mathrm{C}}} \overline{\zeta_k}\zeta_l \, \langle \mathrm{F}|a_k^\dagger(t) a_l|\mathrm{F}\rangle. \tag{3.36}$$

A calculation analogous to the one we made for c_0' brings us to

$$c_1'(t) = \langle \zeta, \exp(\mathrm{i}t\Lambda^{(1)} - \tfrac{1}{2}t\Gamma^{(1)})^{\mathrm{T}}\zeta\rangle. \tag{3.37}$$

Now that we know the TTCFs of the original (c_0 and c_1) and pseudomode (c_0' and c_1') environments, we can say that the reduced dynamics of the non-uniform part of the chain is the same if and only if

$$K^2 \mathrm{e}^{-\mathrm{i}\Omega t} C_{\mathrm{sc}}(2Kt) = \langle \zeta, \exp(-\mathrm{i}t\Lambda^{(0)} - \tfrac{1}{2}t\Gamma^{(0)})^{\mathrm{T}}\zeta\rangle \tag{3.38}$$

for the initially empty environment, and

$$K^2 \mathrm{e}^{\mathrm{i}\Omega t} C_{\mathrm{sc}}(2Kt) = \langle \zeta, \exp(\mathrm{i}t\Lambda^{(1)} - \tfrac{1}{2}t\Gamma^{(1)})^{\mathrm{T}}\zeta\rangle \tag{3.39}$$

for the initially filled one, for all $t \geq 0$. Note that by taking the conjugate of Eq. (3.39) we have

$$\begin{aligned} K^2 e^{-i\Omega t} C_{\mathrm{sc}}(2Kt) &= \overline{\langle \zeta, \exp(it\Lambda^{(1)} - \tfrac{1}{2}t\Gamma^{(1)})^{\mathrm{T}}\zeta\rangle} = \\ &= \langle \zeta, \exp(-it\overline{\Lambda^{(1)}} - \tfrac{1}{2}t\overline{\Gamma^{(1)}})\zeta\rangle \end{aligned} \tag{3.40}$$

which has the same structure as Eq. (3.38), so once we solve Eq. (3.38) we can just set $\Lambda^{(1)} = \overline{\Lambda^{(0)}}$ and $\Gamma^{(1)} = \overline{\Gamma^{(0)}}$ in order to obtain a solution for Eq. (3.39). Ultimately, this means that we need to find Λ, Γ and ζ so that

$$\langle \zeta, \exp(-\mathrm{i}t\Lambda - \tfrac{1}{2}t\Gamma)^{\mathrm{T}}\,\zeta\rangle = K^2 \mathrm{e}^{-\mathrm{i}\Omega t} C_{\mathrm{sc}}(2Kt). \tag{3.41}$$

Define $\Lambda' \coloneqq \Lambda - \Omega I$ and $\zeta' \coloneqq K\zeta$: we get

$$\begin{aligned} K^2 \mathrm{e}^{-\mathrm{i}\Omega t} \langle \zeta', \exp(-\mathrm{i}t\Lambda' - \tfrac{1}{2}t\Gamma)^{\mathrm{T}}\, \zeta' \rangle &= K^2 \mathrm{e}^{-\mathrm{i}\Omega t} C_{\mathrm{sc}}(2Kt) \\ \langle \zeta', \exp(-\mathrm{i}t\Lambda' - \tfrac{1}{2}t\Gamma)^{\mathrm{T}}\, \zeta' \rangle &= C_{\mathrm{sc}}(2Kt); \end{aligned}$$

now rescale the matrices by $2K$, with $\Lambda' = 2K\Lambda''$ and $\Gamma = 2K\Gamma''$ to obtain

$$\langle \zeta', \exp(-2\mathrm{i}Kt\Lambda'' - \tfrac{1}{2}2Kt\Gamma'')^{\mathrm{T}}\, \zeta' \rangle = C_{\mathrm{sc}}(2Kt),$$

then rescale the time variable, too, with $s = 2Kt$, so that

$$\langle \zeta', \exp(-\mathrm{i}s\Lambda'' - \tfrac{1}{2}s\Gamma'')^{\mathrm{T}}\, \zeta' \rangle = C_{\mathrm{sc}}(s).$$

This means that we can first find a set of pseudomodes that solve the last equation, which is the TTCFs generated by J_{sc} in Eq. (3.23), and then rescale back the results to obtain pseudomodes that generate the original TTCFs.

As we anticipated in Section 2.3, we will not be able to find an exact solution of this last equation with a finite number of pseudomodes only. Finding the solution involves a non-trivial inversion problem (see Ref. [149]) and, in general, only an approximate solution (e.g. by means of approximated Prony methods [164]) can be found. In any case, the last equation is the same as the one determined in Ref. [35] for the bosonic case, so the solution computed there for the bosonic MC is also valid in the fermionic setting: given a tridiagonal $N_{\mathrm{C}} \times N_{\mathrm{C}}$ complex matrix

$$M = \begin{pmatrix} \alpha_1 & \beta_1 & 0 & 0 & \cdots & 0 \\ \beta_1 & \alpha_2 & \beta_2 & 0 & \cdots & 0 \\ 0 & \beta_2 & \alpha_3 & \beta_3 & \cdots & 0 \\ 0 & 0 & \beta_3 & \alpha_4 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & 0 & \cdots & \alpha_{N_{\mathrm{C}}} \end{pmatrix} \tag{3.42}$$

and a complex vector $w \in \mathbb{C}^{N_{\mathrm{C}}}$, the quantity $\|\langle w, \exp(tM)w\rangle - C_{\mathrm{sc}}(t)\|$ is minimised in the cases when N_{C} is equal to 6, 8 and 10. The coefficients α_j, β_j and w_j solving the minimisation problem are reported in Table 3.1. Given the equivalence, in the end, of the structures of pseudomodes in the bosonic and the fermionic MCs, the observations in [35] about how the expected semicircle spectral density of the residual environment is well approximated by the surrogate pseudomode environment carry over to our case.

Going back to the original TTCFs in Eq. (3.41), by setting $M = -\mathrm{i}\Lambda'' - \frac{1}{2}\Gamma''$ and $w = \zeta'$ we obtain the initial problem, of which we can then write the solution as

$$\Lambda^{(0)} = \begin{pmatrix} \omega_1 & g_1 & 0 & \cdots & 0 \\ g_1 & \omega_2 & g_2 & \cdots & 0 \\ 0 & g_2 & \omega_3 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & \omega_{N_\mathrm{C}} \end{pmatrix} \tag{3.43}$$

and

$$\Gamma^{(0)}_{ij} = \gamma_j \delta_{ij}, \tag{3.44}$$

where

$$\begin{aligned} \omega_j &= \Omega - 2K \operatorname{Im} \alpha_j, \\ g_j &= -2K \operatorname{Im} \beta_j, \\ \gamma_j &= -4K \operatorname{Re} \alpha_j, \\ \zeta_j &= K w_j. \end{aligned} \tag{3.45}$$

This way, the minimisation of $\|\langle w, \exp(tM) w\rangle - C_\mathrm{sc}(t)\|$ is then equivalent to the minimisation of

$$\left\| K^2 e^{-i\Omega t} C_\mathrm{sc}(2Kt) - \langle B_\mathrm{R}(t) B_\mathrm{R}(0)^\dagger \rangle \right\| \tag{3.46}$$

i.e. of $\|c_0(t) - c_0'(t)\|$, for all $t \geq 0$. This is the recipe for obtaining the parameters defining the pseudomodes that mimic the initially empty environment; the relative solution for the initially filled environment follows immediately by using $\overline{\alpha_j}$ and $\overline{\beta_j}$ instead of α_j and β_j, respectively.

To wrap up this derivation, we have found that the auxiliary environment defined by Eqs. (3.26), (3.43) and (3.44) is a collection of N_C pseudomodes, with nearest-neighbour interaction, undergoing local dissipation and each starting from the vacuum state and interacting with the truncation site N_E of the chain; this provides an FMC for a TEDOPA chain of fermionic modes starting from the vacuum state. The pseudomodes mimicking the initially filled environment work in the same way, except that the Lindblad operator $\mathcal{D}_1$ does not dissipate but *absorbs* particles from the external (unmodeled) environment, injecting them into the system.[2] In Fig. 3.2 we provide a

---

[2] The 'cold' closure, actually, still has a dissipative behaviour, but it does not dissipate particles, but rather *holes*. In fact, holes are the proper way to describe the excitations of the 'cold' environment, that starts from a completely filled state and is perturbed by the system. This perturbation creates holes in the initially filled environment that travel along the chain, get scattered, and eventually reach the closure, where they are dissipated away into the external environment. If we look at the particle balance, instead, the net effect is that this side of the FMC injects particles into the system.

j	$\operatorname{Re}\alpha_j$	$\operatorname{Im}\beta_j$	$\operatorname{Re}w_j$	$\operatorname{Im}w_j$
1	$-1.60\cdot10^{-2}$	$7.85\cdot10^{-1}$	$2.74\cdot10^{-5}$	$-1.11\cdot10^{-5}$
2	$-1.48\cdot10^{-10}$	$-8.13\cdot10^{-1}$	$-4.79\cdot10^{-1}$	$3.99\cdot10^{-1}$
3	$-2.18\cdot10^{0}$	$-1.08\cdot10^{0}$	$6.34\cdot10^{-6}$	$-3.53\cdot10^{-6}$
4	$-1.44\cdot10^{-11}$	$-6.75\cdot10^{-1}$	$4.82\cdot10^{-1}$	$-3.84\cdot10^{-1}$
5	$-4.79\cdot10^{-3}$	$8.05\cdot10^{-1}$	$-1.40\cdot10^{-6}$	$2.45\cdot10^{-6}$
6	$-1.57\cdot10^{-9}$		$3.83\cdot10^{-1}$	$-2.93\cdot10^{-1}$

j	$\operatorname{Re}\alpha_j$	$\operatorname{Im}\beta_j$	$\operatorname{Re}w_j$	$\operatorname{Im}w_j$
1	$-1.06\cdot10^{-9}$	$-8.88\cdot10^{-1}$	$-6.58\cdot10^{-2}$	$-2.48\cdot10^{-1}$
2	$-1.64\cdot10^{-10}$	$4.07\cdot10^{-1}$	$-1.31\cdot10^{-1}$	$3.47\cdot10^{-2}$
3	$-2.70\cdot10^{-11}$	$-9.96\cdot10^{-1}$	$-1.79\cdot10^{-1}$	$-6.75\cdot10^{-1}$
4	$-2.98\cdot10^{0}$	$-1.49\cdot10^{0}$	$1.92\cdot10^{-2}$	$-5.08\cdot10^{-3}$
5	$-1.02\cdot10^{-9}$	$-1.04\cdot10^{0}$	$9.77\cdot10^{-2}$	$3.68\cdot10^{-1}$
6	$-3.61\cdot10^{-9}$	$-4.55\cdot10^{-1}$	$-1.36\cdot10^{-1}$	$3.60\cdot10^{-2}$
7	$-3.53\cdot10^{-11}$	$8.48\cdot10^{-1}$	$-1.06\cdot10^{-1}$	$-4.01\cdot10^{-1}$
8	$-3.73\cdot10^{-11}$		$-2.91\cdot10^{-1}$	$7.73\cdot10^{-2}$

j	$\operatorname{Re}\alpha_j$	$\operatorname{Im}\beta_j$	$\operatorname{Re}w_j$	$\operatorname{Im}w_j$
1	$-3.43\cdot10^{-1}$	$1.13\cdot10^{0}$	$-1.32\cdot10^{-3}$	$4.62\cdot10^{-4}$
2	$-8.67\cdot10^{-5}$	$1.05\cdot10^{0}$	$3.32\cdot10^{-3}$	$5.49\cdot10^{-4}$
3	$-2.73\cdot10^{0}$	$-1.08\cdot10^{0}$	$-2.40\cdot10^{-3}$	$-1.48\cdot10^{-3}$
4	$-7.09\cdot10^{-1}$	$8.35\cdot10^{-1}$	$1.94\cdot10^{-2}$	$-3.55\cdot10^{-2}$
5	$-3.24\cdot10^{-6}$	$-6.04\cdot10^{-1}$	$-3.32\cdot10^{-2}$	$-1.20\cdot10^{-2}$
6	$-4.50\cdot10^{-7}$	$-5.09\cdot10^{-1}$	$1.04\cdot10^{-1}$	$-3.53\cdot10^{-1}$
7	$-2.79\cdot10^{-6}$	$6.77\cdot10^{-1}$	$1.21\cdot10^{-1}$	$2.08\cdot10^{-2}$
8	$-9.48\cdot10^{-5}$	$1.61\cdot10^{-1}$	$1.65\cdot10^{-1}$	$-8.17\cdot10^{-1}$
9	$-1.37\cdot10^{-3}$	$-9.51\cdot10^{-1}$	$-1.21\cdot10^{-1}$	$-4.45\cdot10^{-3}$
10	$-5.95\cdot10^{-6}$		$4.72\cdot10^{-2}$	$-3.67\cdot10^{-1}$

Table 3.1. Universal parameters for the MC for N_C equal to 6, 8 and 10 (from top to bottom). In all three cases $\operatorname{Im}\alpha_j$ and $\operatorname{Re}\beta_j$ are zero for all j.

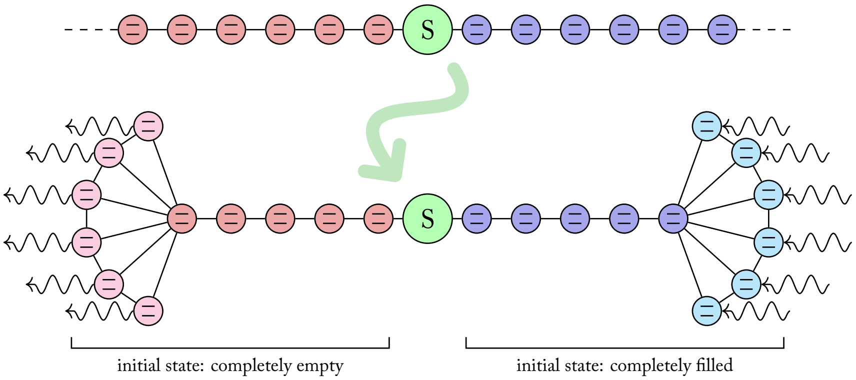

Figure 3.2. The chain-mapped system of Fig. 3.1 before and after truncating the TEDOPA chains and replacing the residual environments, on both the initially empty and initially filled side, with pseudomodes.

graphical representation of the finite-size system resulting from the use of the FMC construct. It is clear from the picture that we can expect a flow of particles/excitations from the filled environment, which will constantly absorb particles from its reservoir, through the system and into the empty environment, which will dissipate them away into its reservoir.

We can see in Eq. (3.45) that the parameters needed to define the MC modes depend on the original environments only through Ω and K, which are easily obtainable from the lower and upper bounds of the domain of the original spectral density function, therefore it is very easy to determine the desired pseudomode parameters for a chosen N_C. An analysis of the accuracy of the approximation, which holds here as well, is performed in the Supplemental Material of Ref. [35].

We remark that the semicircular density of states is not the only way the residual environment can be described. Other fitting schemes have been proposed, e.g. in Refs. [146, 165], which however target specific spectral densities. Moreover, different interaction patterns, e.g. with fully connected pseudomodes or next-nearest neighbour interactions, and alternative methods for the determination of the closure coefficients, as the one recently proposed in Refs. [166, 167], could be used to either reduce the number of modes in the closure or further reduce the approximation error. While a tailored fitting procedure might produce, by directly replacing the original environment with some pseudomodes and skipping the TEDOPA construction altogether, equivalent environments with a smaller number of modes, the construction of the pseudomodes via the FMC method is a more generic and straightforward approach which does not rely on particular features of the initial spectral densities.

The FMC method, as well as its the original bosonic version, is not the only way to embed a non-Markovian dynamics in a Markovian one. In the *reaction coordinate* approach, used for example in Refs. [163, 168, 169] a set of central, collective degrees of freedom of the environment (the reaction coordinate) is identified and extracted. The reaction coordinates and the open system form then a new 'supersystem', that is responsible for non-Markovian behaviour and short-time dynamics, and in turn interacts with a residual environment, such that this dynamics can be described by employing standard Markovian weak-coupling techniques. Just as the MC, this method allows capturing non-Markovian effects quantitatively at strong system-bath coupling and for structured baths. The chain mapping and the pseudomodes in the MC method clearly share many similarities to the reaction coordinates. Pseudomodes, however, arise from a quite abstract derivation, entirely based on the TTCF equivalence, which is very general but does not tell us how exactly they represent the residual environment besides the TTCFs. In fact, a question that remains open is whether we can use the state of the pseudomodes to infer any physical information, such as the mean energy or occupation number, about the modes of the residual environment we have replaced.

3.3 The fermionic Markovian closure as a tensor network

In rest of this chapter we will benchmark the FMC construction by testing it on a concrete model, the single-impurity Anderson model (SIAM) [101]. Generally, this model describes electrons, for which we would have two fermionic levels $\sigma \in \{\uparrow,\downarrow\}$ (representing the spin degrees of freedom) making up the system Hamiltonian

$$H_{\mathrm{S}} = \varepsilon_{\uparrow} a_{\uparrow}^{\dagger} a_{\uparrow} + \varepsilon_{\downarrow} a_{\downarrow}^{\dagger} a_{\downarrow} + U a_{\uparrow}^{\dagger} a_{\uparrow} a_{\downarrow}^{\dagger} a_{\downarrow}. \tag{3.47}$$

In the continuum limit the environment—also composed of electrons—can be described by the following Hamiltonian operators:

$$\begin{aligned} H_{\mathrm{E}} &= \sum_{\sigma\in\{\uparrow,\downarrow\}} \int_X x f_{\sigma,x}^{\dagger} f_{\sigma,x} \,\mathrm{d}x, \\ H_{\mathrm{I}} &= \sum_{\sigma\in\{\uparrow,\downarrow\}} \int_X \sqrt{J(x)} (f_{\sigma,x}^{\dagger} a_{\sigma} + a_{\sigma}^{\dagger} f_{\sigma,x}) \,\mathrm{d}x \end{aligned} \tag{3.48}$$

for some spectral density function J, starting from the thermal state at a given temperature and chemical potential. For the most part we will consider the *non-interacting* version, i.e. a particular instance of this model in which $U = 0$. In this version, the system does not couple the two spin degrees of freedom, which are therefore completely independent: we can consider then a single spin component only, and instead of electrons we will speak of *spinless fermions*, i.e. each mode will represent just a two level (empty/occupied) degree of freedom.

After the transformation we described in the previous section, we transform the initial model into a system such as in Fig. 3.2, described by the following operators. For more clarity, we denote by a the annihilation/creation operators relative to the system, with b those relative to the TEDOPA chains (what is left of them) and with c those relative to the pseudomodes. We call 0 the 'hot' environment and 1 the 'cold' one.

$$H = H_\mathrm{S} + H_{\mathrm{E},0} + H_{\mathrm{C},0} + H_{\mathrm{E},1} + H_{\mathrm{C},1} \tag{3.49}$$

where

$$\begin{aligned}
H_\mathrm{S} &= \varepsilon a^\dagger a, \\
H_{\mathrm{E},\alpha} &= \eta_\alpha (a^\dagger b_{\alpha,1} + b^\dagger_{\alpha,1} a) + \sum_{n=1}^{N_\mathrm{E}} \omega_{\alpha,n} b^\dagger_{\alpha,n} b_{\alpha,n} + \\
&\qquad + \sum_{n=1}^{N_\mathrm{E}-1} \kappa_{\alpha,n} (b^\dagger_{\alpha,n} b_{\alpha,n+1} + b^\dagger_{\alpha,n+1} b_{\alpha,n}), \\
H_{\mathrm{C},\alpha} &= \sum_{n=1}^{N_\mathrm{C}} \nu_{\alpha,n} c^\dagger_{\alpha,n} c_{\alpha,n} + \sum_{n=1}^{N_\mathrm{C}} (\zeta_{\alpha,n} b^\dagger_{\alpha,N_\mathrm{E}} c_{\alpha,n} + \overline{\zeta_{\alpha,n}} c^\dagger_{\alpha,n} b_{\alpha,N_\mathrm{E}}) + \\
&\qquad\qquad + \sum_{n=1}^{N_\mathrm{C}-1} \lambda_{\alpha,n} (c^\dagger_{\alpha,n} c_{\alpha,n+1} + c^\dagger_{\alpha,n+1} c_{\alpha,n}).
\end{aligned} \tag{3.50}$$

together with the Lindblad operators

$$\begin{aligned}
\mathcal{D}_0 &= \sum_{n=1}^{N_\mathrm{C}} \gamma_{0,n} \big(c_{\alpha,n} \rho c^\dagger_{\alpha,n} - \tfrac{1}{2} \{ c^\dagger_{\alpha,n} c_{\alpha,n}, \rho \} \big) \\
\mathcal{D}_1 &= \sum_{n=1}^{N_\mathrm{C}} \gamma_{1,n} \big(c^\dagger_{\alpha,n} \rho c_{\alpha,n} - \tfrac{1}{2} \{ c_{\alpha,n} c^\dagger_{\alpha,n}, \rho \} \big)
\end{aligned} \tag{3.51}$$

make up the master equation

$$\dot{\rho}_t = \mathcal{L}(\rho_t) = -\mathrm{i}[H, \rho_t] + \mathcal{D}_0(\rho_t) + \mathcal{D}_1(\rho_t). \tag{3.52}$$

With respect to the unitary picture, we see that there are new hurdles we must overcome if we want our simulation to be efficient, therefore we must be more careful in how we shape the TN representing the system.

1. The FMC structure is not one-dimensional, so we need to introduce long-range couplings if we want to represent it as an MPS.
2. The GKSL equation resulting from the FMC construction forces us to use mixed states. We can still use MPSs, but we need to double the size of the local Hilbert space on each mode.
3. The presence of a mixed state means also that classical correlations, in addition to quantum ones, will emerge during the time evolution, possibly increasing the bond dimension necessary for an MPS simulation.
4. The TN representation can be more efficient when some quantities (e.g. parity, number of particles) are conserved in the time evolution, by using sparse arrays for the MPS tensors. Even when the original Hamiltonian conserves such quantities, though, it is not immediately clear whether the transformation into the FMC system preserves this property.

With respect to the bosonic MC scheme, additionally, in the most general setting there are now *two* environments, with the OQS in the middle. This means that we can expect a significant amount of correlation, both quantum and classical, to build up across the two leads.

Here is how we can address these issues. The number of pseudomodes usually needed to represent the residual environments is small, and in this case it suffices to 'flatten' the pseudomodes into a linear configuration, at the cost of creating non-nearest-neighbour interactions between the sites (see Fig. 3.3).

Unfortunately, other simple TN topologies would not be viable representations. For example, a tree TN, whose time evolution is also computable with the TDVP algorithm [170], cannot (directly) represent the system we are considering, because the FMC does not have a tree structure since the pseudomodes interact with each other. A simple MPS, for which we already know efficient evolution algorithms, appears then to be the best compromise.

Once we have chosen the MPS representation, there does not appear to be a way to avoid doubling the dimension of the local Hilbert spaces, but there are several ways to do that. We will use the superfermion formalism we saw in Section 1.4, where the doubling translates to adding an auxiliary mode, of the same shape, to each physical one. We prefer this method because the auxiliary modes are fermions, too, so they

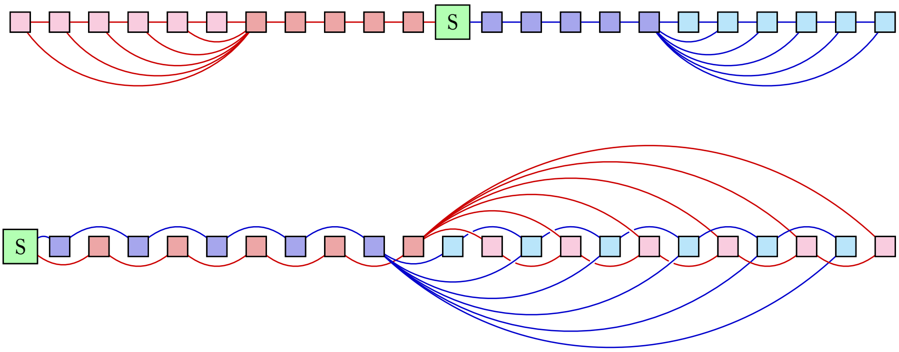

Figure 3.3. TN layout representing the application of the FMC on both environment chains. Squares denote a pair of a physical mode and the relative auxiliary mode; lines denotes coupling terms between modes. *Top*: original flattened layout. *Bottom*: interleaved layout.

seamlessly integrate with the already existing physical structure. With this formalism, however, only the parity of the system is conserved (by design) during the time evolution, but with an additional trick we can restore the conservation of the number of fermions.

With the superfermion formalism, denoting the auxiliary modes with a tilde, we transform the terms in Eq. (3.52) following Eq. (1.71). The Hamiltonian part is very simple: for the $\rho_t \mapsto H\rho_t$ term of the commutator, we just interpret the canonical creation and annihilation operators as their superfermionic counterpart, for example

$$\omega_{\alpha,n} b^\dagger_{\alpha,n} b_{\alpha,n} \rho \mapsto V(\omega_{\alpha,n} b^\dagger_{\alpha,n} b_{\alpha,n} \rho) = \omega_{\alpha,n} b^\dagger_{\alpha,n} b_{\alpha,n} V(\rho), \tag{3.53}$$

while in the $\rho_t \mapsto \rho_t H$ we replace all operators with their auxiliary counterparts, and move them to the left:

$$\omega_{\alpha,n} \rho b^\dagger_{\alpha,n} b_{\alpha,n} \mapsto V(\omega_{\alpha,n} \rho b^\dagger_{\alpha,n} b_{\alpha,n}) = \omega_{\alpha,n} \tilde{b}^\dagger_{\alpha,n} \tilde{b}_{\alpha,n} V(\rho). \tag{3.54}$$

Denoting with $[H, -]$ the map $\rho \mapsto [H, \rho]$, we get

$$[H_\mathrm{S}, -] = \varepsilon a^\dagger a - \varepsilon \tilde{a}^\dagger \tilde{a}, \tag{3.55}$$

$$
\begin{aligned}
[H_{\mathrm{E},\alpha}, -] = \eta_\alpha(a^\dagger b_{\alpha,1} + b^\dagger_{\alpha,1} a - \tilde{a}^\dagger \tilde{b}_{\alpha,1} - \tilde{b}^\dagger_{\alpha,1} \tilde{a}) + {}&\\
+ \sum_{n=1}^{N_\mathrm{E}} \omega_{\alpha,n}(b^\dagger_{\alpha,n} b_{\alpha,n} - \tilde{b}^\dagger_{\alpha,n} \tilde{b}_{\alpha,n}) + {}&\\
+ \sum_{n=1}^{N_\mathrm{E}-1} \kappa_{\alpha,n}(b^\dagger_{\alpha,n} b_{\alpha,n+1} + b^\dagger_{\alpha,n+1} b_{\alpha,n} - \tilde{b}^\dagger_{\alpha,n} \tilde{b}_{\alpha,n+1} - \tilde{b}^\dagger_{\alpha,n+1} \tilde{b}_{\alpha,n}),&
\end{aligned}
\tag{3.56}
$$

$$
\begin{aligned}
[H_{\mathrm{C},\alpha}, -] = {}& \sum_{n=1}^{N_\mathrm{C}} \nu_{\alpha,n}(c^\dagger_{\alpha,n} c_{\alpha,n} - \tilde{c}^\dagger_{\alpha,n} \tilde{c}_{\alpha,n}) + \qquad (3.57)\\
&+ \sum_{n=1}^{N_\mathrm{C}} (\zeta_{\alpha,n} \tilde{b}^\dagger_{\alpha,N_\mathrm{E}} \tilde{c}_{\alpha,n} + \overline{\zeta_{\alpha,n}} \tilde{c}^\dagger_{\alpha,n} \tilde{b}_{\alpha,N_\mathrm{E}} + \\
&- \overline{\zeta_{\alpha,n}} \tilde{b}^\dagger_{\alpha,N_\mathrm{E}} \tilde{c}_{\alpha,n} - \zeta_{\alpha,n} \tilde{c}^\dagger_{\alpha,n} \tilde{b}_{\alpha,N_\mathrm{E}}) + \\
&+ \sum_{n=1}^{N_\mathrm{C}-1} \lambda_{\alpha,n}(c^\dagger_{\alpha,n} c_{\alpha,n+1} + c^\dagger_{\alpha,n+1} c_{\alpha,n} + \\
&\qquad - \tilde{c}^\dagger_{\alpha,n} \tilde{c}_{\alpha,n+1} - \tilde{c}^\dagger_{\alpha,n+1} \tilde{c}_{\alpha,n}),
\end{aligned}
$$

$$
\mathcal{D}_0 = \sum_{n=1}^{N_\mathrm{C}} \gamma_{0,n}\bigl(-c_{0,n} \tilde{c}_{0,n} - \tfrac{1}{2} c^\dagger_{0,n} c_{0,n} - \tfrac{1}{2} \tilde{c}^\dagger_{0,n} \tilde{c}_{0,n}\bigr), \tag{3.58}
$$

$$
\mathcal{D}_1 = \sum_{n=1}^{N_\mathrm{C}} \gamma_{1,n}\bigl(c^\dagger_{1,n} \tilde{c}^\dagger_{1,n} + \tfrac{1}{2} c^\dagger_{1,n} c_{1,n} + \tfrac{1}{2} \tilde{c}^\dagger_{1,n} \tilde{c}_{1,n} - \mathbb{1}\bigr). \tag{3.59}
$$

All these operators act on the state $V(\rho)$ by multiplying it on the left. Note that the physical and auxiliary modes are not coupled by the unitary part: the terms $-c_{0,n}\tilde{c}_{0,n}$ and $c^\dagger_{1,n}\tilde{c}^\dagger_{1,n}$ in the Lindblad operators $\mathcal{D}_\alpha$ are the only ones connecting them: dissipation by the 'hot' environment is described by a pair of annihilation operators, one on the physical and one on the relative auxiliary mode, while absorption by the 'cold' environment is described by a pair of creation operators.

It is evident that exactly these two terms in Eqs. (3.58) and (3.59) do not conserve the particle number. To remedy this issue, as pointed out in Ref. [92], we 'invert' the auxiliary modes, switching particles and holes via a simple $\tilde{c}_k \leftrightarrow \tilde{c}^\dagger_k$ canonical transformation. We can see in fact that all the other terms (which do conserve the particle number) involve either two physical modes or two auxiliary modes, therefore the particle-hole inversion would change either none or both of them, preserving the number-conserving property. In the offending terms instead only one of the factor changes, and they become

$$-c_{0,n}\tilde{c}^{\dagger}_{0,n} \quad \text{and} \quad c^{\dagger}_{1,n}\tilde{c}_{1,n} \tag{3.60}$$

respectively: they now conserve the number of fermions. In the end, without introducing new symbols, the particle-hole inversion on the auxiliary modes transforms the previous operators in

$$\begin{align}
[H_{\mathrm{S}}, -] &= \varepsilon a^{\dagger} a + \varepsilon \tilde{a}^{\dagger}\tilde{a}, \tag{3.61}\\
[H_{\mathrm{E},\alpha}, -] &= \eta_{\alpha}(a^{\dagger} b_{\alpha,1} + b^{\dagger}_{\alpha,1} a + \tilde{a}^{\dagger}\tilde{b}_{\alpha,1} + \tilde{b}^{\dagger}_{\alpha,1}\tilde{a}) + \\
&\quad + \sum_{n=1}^{N_{\mathrm{E}}} \omega_{\alpha,n}(b^{\dagger}_{\alpha,n} b_{\alpha,n} + \tilde{b}^{\dagger}_{\alpha,n}\tilde{b}_{\alpha,n}) + \\
&\quad + \sum_{n=1}^{N_{\mathrm{E}}-1} \kappa_{\alpha,n}(b^{\dagger}_{\alpha,n} b_{\alpha,n+1} + b^{\dagger}_{\alpha,n+1} b_{\alpha,n} + \tilde{b}^{\dagger}_{\alpha,n}\tilde{b}_{\alpha,n+1} + \tilde{b}^{\dagger}_{\alpha,n+1}\tilde{b}_{\alpha,n}), \tag{3.62}\\
[H_{\mathrm{C},\alpha}, -] &= \sum_{n=1}^{N_{\mathrm{C}}} \nu_{\alpha,n}(c^{\dagger}_{\alpha,n} c_{\alpha,n} + \tilde{c}^{\dagger}_{\alpha,n}\tilde{c}_{\alpha,n}) + \tag{3.63}\\
&\quad + \sum_{n=1}^{N_{\mathrm{C}}}(\zeta_{\alpha,n}\tilde{b}^{\dagger}_{\alpha,N_{\mathrm{E}}}\tilde{c}_{\alpha,n} + \overline{\zeta_{\alpha,n}}\tilde{c}^{\dagger}_{\alpha,n}\tilde{b}_{\alpha,N_{\mathrm{E}}} + \\
&\quad + \overline{\zeta_{\alpha,n}}\tilde{b}^{\dagger}_{\alpha,N_{\mathrm{E}}}\tilde{c}_{\alpha,n} + \zeta_{\alpha,n}\tilde{c}^{\dagger}_{\alpha,n}\tilde{b}_{\alpha,N_{\mathrm{E}}}) + \\
&\quad + \sum_{n=1}^{N_{\mathrm{C}}-1} \lambda_{\alpha,n}(c^{\dagger}_{\alpha,n} c_{\alpha,n+1} + c^{\dagger}_{\alpha,n+1} c_{\alpha,n} + \\
&\qquad + \tilde{c}^{\dagger}_{\alpha,n}\tilde{c}_{\alpha,n+1} + \tilde{c}^{\dagger}_{\alpha,n+1}\tilde{c}_{\alpha,n}), \\
\mathscr{D}_0 &= \sum_{n=1}^{N_{\mathrm{C}}} \gamma_{0,n}\big(-c_{0,n}\tilde{c}^{\dagger}_{0,n} - \tfrac{1}{2}c^{\dagger}_{0,n} c_{0,n} - \tfrac{1}{2}\tilde{c}_{0,n}\tilde{c}^{\dagger}_{0,n}\big), \tag{3.64}\\
\mathscr{D}_1 &= \sum_{n=1}^{N_{\mathrm{C}}} \gamma_{1,n}\big(c^{\dagger}_{1,n}\tilde{c}_{1,n} + \tfrac{1}{2}c^{\dagger}_{1,n} c_{1,n} + \tfrac{1}{2}\tilde{c}_{1,n}\tilde{c}^{\dagger}_{1,n} - \mathbb{1}\big). \tag{3.65}
\end{align}$$

In the Hamiltonian part, we rearranged the terms into a normal form by taking advantage of the CARs: in this case the multiples of $\mathbb{1}$ produced by the anticommutation cancel out; we cannot do the same thing in the Lindblad operators.

The next step is to convert the antisymmetric structure of the fermion Fock space into a standard tensor-product space, in order to be able to encode states and operators into an MPS. We will use the Jordan-Wigner mapping, for which we need to choose a specific enumeration of the sites, flattening the system onto a line. The naive way of doing it would separate the 'hot' and the 'cold' environment on opposite sides of the central system. However, as shown in [157], we expect that a significant amount of correlation will build up between the two environments. This configuration would therefore lead to the appearance of long-range correlations, ultimately causing a significant increase in the bond dimensions in the middle of the MPS, which may be prohibitive for our numerical simulations. An alternative enumeration is introduced in Ref. [157]: we interleave the two environments, alternating a mode from the empty environment with one from the full one; we keep the auxiliary modes close to their associated physical modes. This is applied already at the level of the unitary dynamics, i.e. on the TEDOPA chains with no FMC yet. A similar result holds for bosonic environments as well, as noted in Ref. [171], where it is also shown how this configuration, resulting in a simple MPS, can also be more efficient than more complex TNs.

In the interleaved configuration, the open system is put at one edge of the linear chain (see Fig. 3.3). The two environments are brought 'closer' (in the TN topology) to each other, allowing us to use a lower overall bond dimension in the MPS; to do so, we pay the price of transforming the original nearest-neighbour interactions of H_{E} into coupling operators spanning *four* sites: between $c_{\alpha,n}$ and $c_{\alpha,n+1}$ we find the auxiliary mode of the former, and two modes from the opposite environment.

3.4 Exact solution with Gaussian states

In the time evolution of the MPS, with the FMC we are introducing another approximation error, so it is not always easy to discern whether the errors are coming from the MPS time-evolution algorithm, from having bond dimensions which are too low or from the FMC itself. With the exact solution at our disposal, we will be able to analyse the FMC-related error directly, without the other issues interfering. Once the FMC error is ruled out, it will be easier to know how far we need to increase the bond dimension of the MPS representation of the state, by comparing the MPS solution with the exact one. In practice, we select a few operators, let the state evolve for a certain time (ideally until it reaches its steady state) and then compare the expectation values of those operators over time, setting a threshold that the MPS evolution needs to satisfy.

If the system Hamiltonian is quadratic, then every term of the GKSL equation is quadratic, so it maps Gaussian states into Gaussian states; in this case, the evolved state

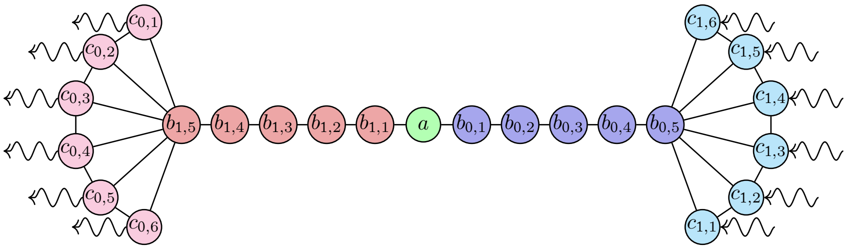

Figure 3.4. Numbering scheme for the superfermionic modes, as in Section 3.4. Each site shows the name of the annihilation operator related to it.

ρ_t is completely determined by the TTCFs such as $\operatorname{tr}(b_n^\dagger b_m \rho_t)$. We can therefore try to compute an exact solution of the equation, at least numerically: we will not obtain an analytic solution, but we will have a differential equation that can be solved *efficiently* with arbitrary numerical precision. By evolving the matrix of TTCFs directly, instead of the state MPS, it will be sufficient to perform computations on a $N \times N$ matrix (where N is the number of modes): the complexity is only polynomial in N.

Together with the superfermion formalism, we can write the time-evolution differential equation of the TTCF in a very simple way, as follows. We assume the system is composed of a single fermionic mode, which starts from the occupied state. First of all, let us use a common notation for all modes. Let $M := N_\mathrm{E} + N_\mathrm{C}$: we have in total $M + 1$ physical modes and just as many auxiliary modes. We will use the following scheme to name the modes in a uniform way:

$$\begin{array}{ccccccccccccc}
c_{0,N_\mathrm{C}} & \cdots & c_{0,1} & b_{0,N_\mathrm{E}} & \cdots & b_{0,1} & a & b_{1,1} & \cdots & b_{1,N_\mathrm{E}} & c_{1,1} & \cdots & c_{1,N_\mathrm{C}} \\
f_{-M} & \cdots & f_{-N_\mathrm{E}-1} & f_{-N_\mathrm{E}} & \cdots & f_{-1} & f_0 & f_1 & \cdots & f_{N_\mathrm{E}} & f_{N_\mathrm{E}+1} & \cdots & f_M
\end{array}$$

(the previous naming scheme is illustrated in Fig. 3.4, for reference).

We join physical and auxiliary modes in a single vector of $2(2M+1)$ creation and annihilation operators

$$\varphi = (\, f_{-M} \;\; f_{-M+1} \;\; \cdots \;\; f_M \;\; \tilde{f}^\dagger_{-M} \;\; \tilde{f}^\dagger_{-M+1} \;\; \cdots \;\; \tilde{f}^\dagger_M \;\;)^\mathrm{T}.$$

and in terms of the φ_j operators we have a new set of CARs

$$\{\varphi_j, \varphi_k\} = 0, \quad \{\varphi_j^\dagger, \varphi_k^\dagger\} = 0, \quad \{\varphi_j, \varphi_k^\dagger\} = \delta_{jk}\mathbb{1}. \tag{3.66}$$

Since the equations will be more convoluted than usual, we use in this section the superfermion notation introduced in Section 1.4.2; recall that V is the 'vectorisation map' from the original space of states $\mathcal{S}(\mathscr{H})$, containing density matrices, to $\mathscr{H} \otimes \overline{\mathscr{H}}$, containing vectors, where $\mathscr{H}$ is the antisymmetric Fock space of $2M+1$ fermionic modes.

We start by writing the vectorisation of $\mathscr{L}(\rho)$, that is $V(\mathscr{L}(\rho))$, with $f^\dagger$ and $\tilde{f}$ on the left, and f and $\tilde{f}^\dagger$ on the right in every term by making use of the CARs. We call $\mathcal{H}$ the (Hermitian) matrix containing all parameters coming from the unitary part of the evolution, i.e. the $-\mathrm{i}[H,-]$ term, and Γ_0 and Γ_1 the matrices containing the dissipation/absorption rates, equal to $\Gamma^{(0)}$ and $\Gamma^{(1)}$ from Section 3.2 (in particular $\Gamma^{(0)}$ appears in Eq. (3.44)), respectively, extended as zero on the rest of the system:

$$\begin{aligned}\Gamma_{0,nm} &:= \Theta(n \leq -N_\mathrm{E} - 1)\delta_{nm}\gamma^{(0)}_{-n-N_\mathrm{E}} = \\ &= \begin{cases} \delta_{nm}\gamma^{(0)}_{-n-N_\mathrm{E}} & -N_\mathrm{C} - N_\mathrm{E} \leq n \leq -N_\mathrm{E} - 1, \\ 0 & \text{otherwise,} \end{cases}\end{aligned} \tag{3.67}$$

and similarly $\Gamma_{1,nm} := \Theta(n \geq N_\mathrm{E} + 1)\delta_{nm}\gamma^{(1)}_{n-N_\mathrm{E}}$. We find

$$\begin{aligned} V(\mathscr{L}(\rho)) = \sum_{j,k=-M}^{M} \big[&-\mathrm{i}\mathcal{H}_{jk} f_j^\dagger f_k - \mathrm{i}\overline{\mathcal{H}_{jk}}\, \tilde{f}_k^\dagger \tilde{f}_j + \Gamma_{0,jk}(\tilde{f}_k^\dagger f_j - \tfrac{1}{2} f_k^\dagger f_j + \tfrac{1}{2}\tilde{f}_j^\dagger \tilde{f}_k) + \\ &+ \Gamma_{1,jk}(f_j^\dagger \tilde{f}_k + \tfrac{1}{2} f_j^\dagger f_k - \tfrac{1}{2}\tilde{f}_k^\dagger \tilde{f}_j)\big] V(\rho) + \\ &+ \sum_{j=-M}^{M} (\mathrm{i}\overline{\mathcal{H}_{jj}} + \tfrac{1}{2}\Gamma_{0,jj} - \tfrac{1}{2}\Gamma_{1,jj}) V(\rho) = \\ = \sum_{j,k=-M}^{M} \big[&(-\mathrm{i}\mathcal{H}_{jk} - \tfrac{1}{2}\Gamma_{0,jk} + \tfrac{1}{2}\Gamma_{1,jk}) f_j^\dagger f_k + \\ &+ (-\mathrm{i}\overline{\mathcal{H}_{kj}} + \tfrac{1}{2}\Gamma_{0,jk} - \tfrac{1}{2}\Gamma_{1,jk})\tilde{f}_j^\dagger \tilde{f}_k + \\ &+ \Gamma_{0,jk}\tilde{f}_j^\dagger f_k + \Gamma_{1,jk} f_j^\dagger \tilde{f}_k\big] V(\rho) + \operatorname{tr}\big(\mathrm{i}\overline{\mathcal{H}} - \tfrac{1}{2}(\Gamma_1 + \Gamma_0)\big) V(\rho) \end{aligned} \tag{3.68}$$

so we can write $V(\mathscr{L}(\rho))$ as $LV(\rho)$ where

$$L = \underbrace{\varphi^\dagger \begin{pmatrix} -\mathrm{i}\mathcal{H} + \tfrac{1}{2}(\Gamma_1 - \Gamma_0) & \Gamma_1 \\ \Gamma_0 & -\mathrm{i}\mathcal{H} - \tfrac{1}{2}(\Gamma_1 - \Gamma_0) \end{pmatrix}}_{K} \varphi + \operatorname{tr}\big(\mathrm{i}\mathcal{H} - \tfrac{1}{2}(\Gamma_1 + \Gamma_0)\big)\mathbb{1}. \tag{3.69}$$

Now that we have reshaped the GKSL equation in this very simple form, we compute the time evolution of expectation values of pairs of physical operators, $\operatorname{tr}(\rho_t f_j^\dagger f_k)$, from which we can extract the expectation values of all relevant physical operators, due to the Gaussianity of the evolution. In the vectorised setting the time evolution works as follows: for a generic operator X on $\mathcal{H}$, call $\tilde{X}$ its counterpart in the enlarged space $\mathcal{H} \otimes \overline{\mathcal{H}}$, such that $\tilde{X}V(\rho) = V(X\rho)$:

$$\begin{aligned} \operatorname{tr}(X\rho_t) &= \langle V(\mathbb{1}), V(X\rho_t)\rangle = \langle V(\mathbb{1}), \tilde{X}V(\rho_t)\rangle = \langle V(\mathbb{1}), \tilde{X}\exp(tL)V(\rho_0)\rangle = \\ &= \langle V(\mathbb{1}), \exp(tL)\exp(-tL)\tilde{X}\exp(tL)V(\rho_0)\rangle = \\ &= \langle \exp(tL')V(\mathbb{1}), \exp(-tL)\tilde{X}\exp(tL)V(\rho_0)\rangle \end{aligned} \tag{3.70}$$

Now $\exp(tL')V(\mathbb{1}) = V(\mathbb{1})$, where L' is the adjoint of L, since $L'V(\mathbb{1}) = 0$, therefore

$$\operatorname{tr}(X\rho_t) = \langle V(\mathbb{1}), \exp(-tL)\tilde{X}\exp(tL)V(\rho_0)\rangle \tag{3.71}$$

from which we can see that $\exp(-tL)\tilde{X}\exp(tL)$ is the vectorisation of the time-evolved operator X_t in the original formulation. Now, we are interested in the value of

$$\begin{aligned} \operatorname{tr}(f_j^\dagger f_k\rho_t) &= \langle V(\mathbb{1}), V(f_j^\dagger f_k\rho_t)\rangle = \\ &= \langle V(\mathbb{1}), f_j^\dagger f_k V(\rho_t)\rangle = \langle V(\mathbb{1}), \varphi_{j+M+1}^\dagger \varphi_{k+M+1} V(\rho_t)\rangle \end{aligned} \tag{3.72}$$

for $1 \le n, m \le 2M+1$. We want therefore to compute the evolution over time of the matrix

$$\mathcal{B}_{nm}(t) := \langle V(\mathbb{1}), \varphi_n^\dagger \varphi_m V(\rho_t)\rangle. \tag{3.73}$$

The time evolution of a single mode operator is computed from the differential equation

$$\partial_t\varphi_j(t) = \partial_t\big(\exp(-tL)\varphi_j\exp(tL)\big) = [\varphi_j(t), \varphi^\dagger K\varphi] = \sum_{l=-2M}^{2M} K_{jl}\varphi_l(t) \tag{3.74}$$

from which we see that the vector of operators φ evolves as $\varphi(t) = \exp(tK)\varphi$. Consequently,

$$\begin{aligned} \mathcal{B}_{nm}(t) &= \langle V(\mathbb{1}), \varphi_n^\dagger\varphi_m V(\rho_t)\rangle = \\ &= \langle V(\mathbb{1}), \exp(-tL)\varphi_n^\dagger\exp(tL)\exp(-tL)\varphi_m\exp(tL)V(\rho_0)\rangle = \\ &= \sum_{j,l=-2M}^{2M} \langle V(\mathbb{1}), \exp(-tK)_{jn}\varphi_j^\dagger \exp(tK)_{ml}\varphi_l V(\rho_0)\rangle \end{aligned} \tag{3.75}$$

or, more compactly,

$$\mathcal{B}(t) = \exp(-tK^{\mathrm{T}})\mathcal{B}(0)\exp(tK^{\mathrm{T}}). \tag{3.76}$$

We can write $\mathcal{B}(t)$ as the block matrix

$$\mathcal{B}(t) = \begin{pmatrix} B^{(1)} & B^{(2)} \\ B^{(3)} & B^{(4)} \end{pmatrix} \tag{3.77}$$

where, for $n, m \in \{1, 2, \dots, 2M+1\}$,

$$\begin{aligned} B^{(1)}_{nm} &= \langle V(\mathbb{1}), f^{\dagger}_{n-M-1} f_{m-M-1} V(\rho_t)\rangle = \mathrm{tr}(\rho_t f^{\dagger}_{n-M-1} f_{m-M-1}), \\ B^{(2)}_{nm} &= \langle V(\mathbb{1}), f^{\dagger}_{n-M-1} \tilde{f}^{\dagger}_{m-M-1} V(\rho_t)\rangle = \mathrm{tr}(\rho_t f_{m-M-1} f^{\dagger}_{n-M-1}), \\ B^{(3)}_{nm} &= \langle V(\mathbb{1}), \tilde{f}_{n-M-1} f_{m-M-1} V(\rho_t)\rangle = \mathrm{tr}(\rho_t f^{\dagger}_{n-M-1} f_{m-M-1}), \\ B^{(4)}_{nm} &= \langle V(\mathbb{1}), \tilde{f}_{n-M-1} \tilde{f}^{\dagger}_{m-M-1} V(\rho_t)\rangle = \mathrm{tr}(\rho_t f_{m-M-1} f^{\dagger}_{n-M-1}). \end{aligned} \tag{3.78}$$

By defining the matrix $\mathcal{R}(t)$ as $\mathcal{R}_{jk}(t) \coloneqq \mathrm{tr}(\rho_t f^{\dagger}_j f_k)$, which is Hermitian, we can then write

$$\mathcal{B}(t) = \begin{pmatrix} \mathcal{R}(t) & I - \mathcal{R}(t) \\ \mathcal{R}(t) & I - \mathcal{R}(t) \end{pmatrix}. \tag{3.79}$$

At the beginning of the time evolution we have $\mathcal{R}_{jk}(0) = \mathrm{tr}(\rho_0 f^{\dagger}_j f_k) = 0$ unless $j = k > 0$ where it equals one (the modes with $j = k > 0$ correspond to the 'cold' environment), and $j = k = 0$ where we have the initial population of the open system at the centre.

3.4.1 Solution

We compute here a numerical solution to Eq. (3.76). The exponential character of the evolution, and the fact that the generator is not anti-Hermitian, are such that as t grows the components of the $\exp(\pm tK^{\mathrm{T}})$ matrix become either really small or really large, thus we first need to fine-tune the actual calculation in order to avoid numerical issues. One precaution to avoid this is to evolve the state not by computing Eq. (3.76) directly at any t, but evolving the $\mathcal{B}$ matrix step by step, i.e. once by partitioning the $[0, t_{\mathrm{f}}]$ interval into a sequence $(0, t_1, t_2, \dots, t_n \equiv t_{\mathrm{f}})$ and compute a sequence of matrices $(\mathcal{B}(t_i))_{i=0}^{n}$ with

$$\mathcal{B}(t_i) = \exp\big(-(t_i - t_{i-1})K^{\mathrm{T}}\big)\,\mathcal{B}(t_{i-1})\,\exp\big((t_i - t_{i-1})K^{\mathrm{T}}\big). \tag{3.80}$$

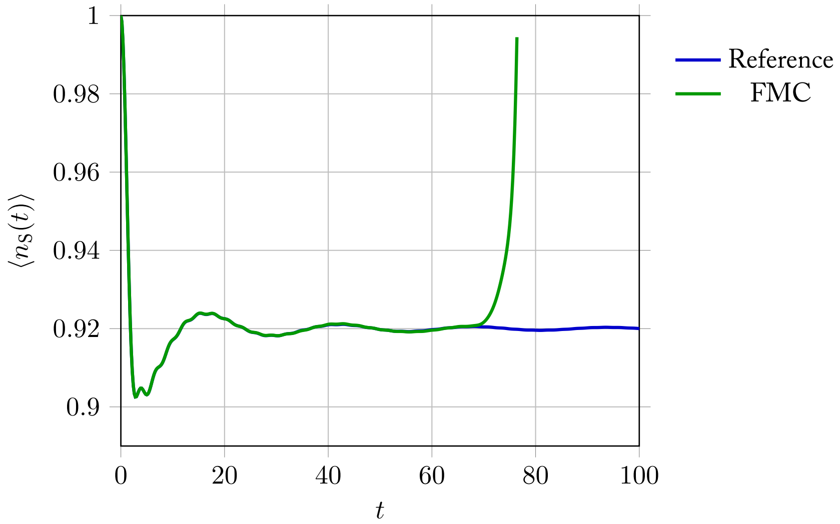

Figure 3.5. Expectation value of the population of the open system from the solution of Eq. (3.76) (FMC) compared with a reference solution, computed by integrating the equation obtained with a truncated TEDOPA.

If the $t_i - t_{i-1}$ steps are small enough, then the matrix elements do not diverge. If the time step is uniform this is also a much faster method since we only need to exponentiate the generator twice.

We compute a concrete solution to the following setting, derived from Refs. [157, 158]: a non-interacting SIAM with the impurity, starting from the occupied state of energy $\varepsilon = -\pi/8$, in contact with an environment described by a semicircle spectral density

$$J(x) = \frac{1}{10\pi}\sqrt{x(2-x)} \tag{3.81}$$

on the domain $[0,2]$, with temperature $T = 0.4$ and chemical potential $\mu = 0.2$. As a figure of merit, we plot the occupation number of the open system, $\langle n_S(t) \rangle = \langle V(\mathbb{1}), f_0^\dagger f_0 V(\rho_t) \rangle$, comparing it with a numerical solution of the differential equation of the TEDOPA system (without the FMC). Figure 3.5 shows the result of calculations performed with a uniform time step of 0.2. Unfortunately, even with a step-by-step evolution the solution clearly diverges from the correct path.

This issue can be traced back to the fact that, due to the finite precision of floating-

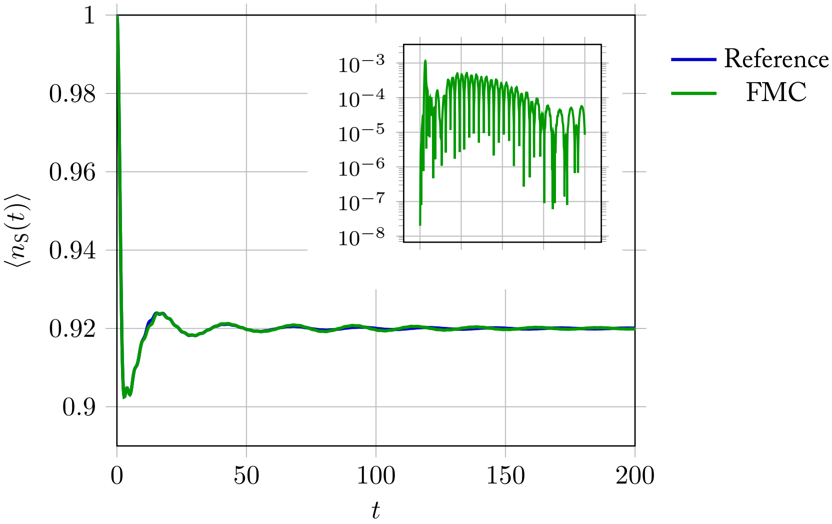

Figure 3.6. Expectation value of the population of the open system from the solution of Eq. (3.76), obtained enforcing the form in Eq. (3.79) after each time step. The inset plot shows the absolute difference between the FMC plot and a reference solution computed by integrating the equations obtained with a truncated TEDOPA.

point calculations, and the very different orders of magnitude appearing in K and in its spectrum, the correlation matrix deviates from the form illustrated in Eq. (3.79) as the evolution proceeds. These deviations accumulate exponentially over time, causing the appearance of spurious populations in the unphysical modes, and resulting in a dramatic divergence from the expected result as in Fig. 3.5. A simple solution to this problem is to enforce, after each time step, that the correlation matrix $\mathcal{B}(t)$ has the form of Eq. (3.79), with $\mathcal{R}(t)$ a Hermitian matrix. The result is shown in Fig. 3.6.

3.5 Numerical analysis of the MPS solution

We put the FMC technique to the test in a true MPS setting, as promised in Section 3.3, with the TN described in that chapter, as in Fig. 3.3. Since the FMC gives rise to long-range interactions once we flatten the system onto an MPS, we can either use TEBD with swap operators or TDVP, that can deal with arbitrary interaction geometries. We choose the TDVP algorithm, namely the one-site version, imposing the conservation of particle number and using sparse tensors to represent the state and the observables, for increased efficiency. As anticipated in Section 1.3.3, because we have long-range interactions, we need to artificially increase the bond dimensions of the MPS at the beginning of the evolution. The initial bond dimension will not change throughout the time evolution, since the one-site TDVP algorithm keeps it constant. Therefore, in order to judge whether the chosen bond dimension is sufficient, we run different simulations with an increasing initial bond dimension, until the results we are interested in have satisfactorily converged.[3]

3.5.1 Correlation functions

Since the FMC method is entirely built on the fact that the environment correlation functions completely determine the dynamics of the open system, we can quickly check whether the approximation is working by making sure that, in fact, the open system 'sees' the expected correlation functions, at least to a sufficient degree of approximation. To this purpose we choose for the SIAM environment a non-trivial example, merging the two semicircle spectral densities

$$\begin{aligned} J_{\mathrm{L}}(x) &= \tfrac{1}{2\pi}\sqrt{x(2-x)}, \quad T_{\mathrm{L}} = 0.2, \quad \mu_{\mathrm{L}} = 1, \\ J_{\mathrm{R}}(x) &= \tfrac{1}{4\pi}\sqrt{x(2-x)}, \quad T_{\mathrm{R}} = 1, \qquad \mu_{\mathrm{R}} = \tfrac{1}{4}, \end{aligned} \tag{3.82}$$

both defined on the domain $[0,2]$. Through the procedure detailed in the previous chapters we derive the two equivalent environments

$$\begin{aligned} J_0(x) &= \sum_{\alpha\in\{\mathrm{L,R}\}} \frac{1}{\mathrm{e}^{-x/T_\alpha}+1} J_\alpha(x+\mu_\alpha), \\ J_1(x) &= \sum_{\alpha\in\{\mathrm{L,R}\}} \frac{1}{\mathrm{e}^{x/T_\alpha}+1} J_\alpha(x+\mu_\alpha), \end{aligned} \tag{3.83}$$

and reshape the respective environments with the chain mapping: the top half of Fig. 3.7 shows how the chain coefficients converge towards their asymptotic values. We apply the FMC leaving $N_{\mathrm{E}} = 13$ or $N_{\mathrm{E}} = 20$ chain sites and adding $N_{\mathrm{C}} = 6$ pseudomodes on each side: we choose these numbers so that for $n > N_{\mathrm{E}}$ the distance of the n-th coefficients from their asymptotic value is less than 10^{-2} for $N_{\mathrm{E}} = 13$, and

---

[3] The adaptive TDVP scheme of [84] could be used to dynamically adjust the bond dimensions, but in our numerical simulations we observed that the bond dimensions saturate very quickly during the initial transient regime of the evolution. The advantage of a flexible bond dimension would therefore be confined only to a small initial window, and for this reason we stuck with the standard TDVP algorithm.

less than 5×10^{-3} for $N_{\mathrm{E}} = 20$. We then reproduced the spectral densities by calculating the inverse Fourier transform of the correlation function, which for the initially empty environment gives the following function

$$\frac{1}{\pi} \operatorname{Re} \int_0^{+\infty} \mathrm{e}^{\mathrm{i}xt} \eta_0^2 \operatorname{tr}\big(b_{0,1} \varPhi_t(b_{0,1}^\dagger |\emptyset_0, \mathrm{F}_1\rangle\langle\emptyset_0, \mathrm{F}_1|)\big) \, \mathrm{d}t \tag{3.84}$$

of the frequency x, where $b_{0,1}$ is the annihilation operator of the first mode of the chain derived from the initially empty environment (see Eq. (3.50)), and $\varPhi_t$ the time-evolution map, i.e. $\exp(t\mathcal{L})$ where $\mathcal{L}$ is the Lindbladian operator fixed by $H_{\mathrm{E},0}$, $H_{\mathrm{C},0}$ and $\mathcal{D}_0$ from Eqs. (3.50) and (3.51). An analogous formula holds for the initially filled environment. We run the simulation up until $t = 400$; this truncation, that we cannot avoid, introduces artefacts in the Fourier transform. We smooth them out by multiplying the correlation function by an exponentially decaying factor $\exp(-gt^2)$ with g such that the resulting product is 10^{-15} at $t = 400$. This does not alter the correlation function significantly since it is already exponentially decaying. We can see in Fig. 3.7 that the expected and simulated spectral densities are in very good agreement: the absolute difference is at worst 1×10^{-2}, at the non-differentiable points. Since we can compute the correlation function only for a finite amount of physical time, we expect, after all, that this procedure is not able to describe correctly the spectral density function in a neighbourhood of its non-differentiable points.

3.5.2 Accuracy of MPS simulations

We now perform a full MPS simulation of the same configuration as in Section 3.4, considering the non-interacting SIAM with $\varepsilon = -\pi/8$ and the environment characterised by the spectral density of Eq. (3.81).

In Figs. 3.8 and 3.9 we compare our approach to standard TEDOPA, i.e. without the Markovian closure, and to the exact solution with Gaussian states developed in Section 3.4, with the solution obtained by integrating the ODE of the TEDOPA system taken as our reference result. We show the time evolution expectation value of the population of the impurity site. In both figures different closure sizes (for N_{C} equal to 6, 8 and 10) attached to the truncation point of the TEDOPA chain are considered.

We observe that the MPS formulation of the FMC in Fig. 3.8 deviates from the reference solution by 1–2 orders of magnitude more than the exact solution obtained with the Gaussian state formalism: the discrepancy however is present (albeit smaller) in the TEDOPA MPS simulation, too, which strongly suggests that the MPS representation of the FMC is not inherently worse than the one of the TEDOPA system.

The number N_{C} of pseudomodes does not have a strong effect on the accuracy of the simulations: counter-intuitively, the use of a larger number of pseudomodes does

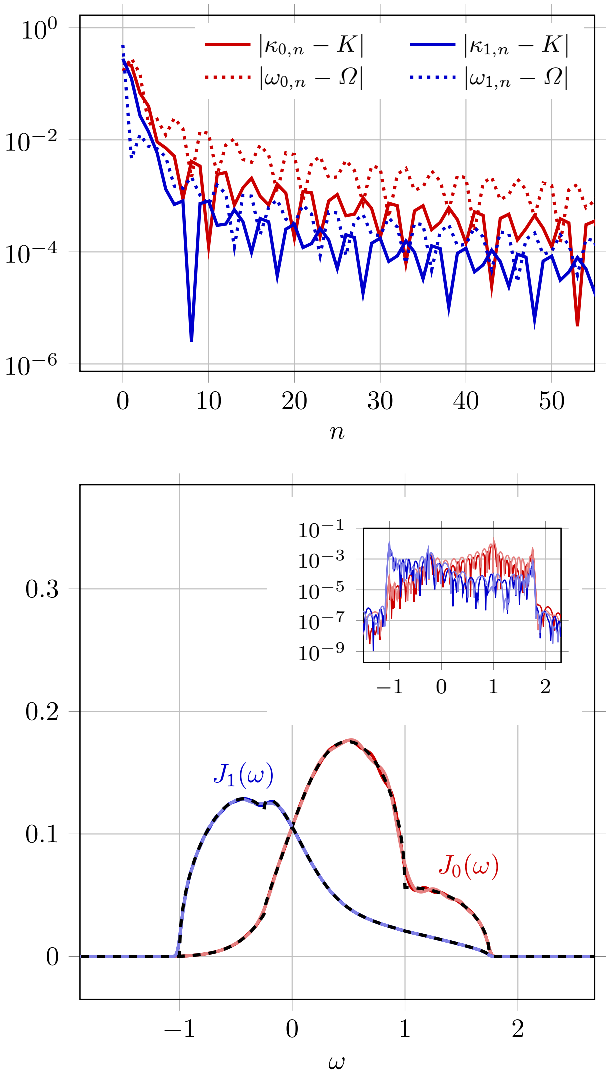

Figure 3.7. *Top*: convergence of chain coefficients of the equivalent 'hot' (0) and 'cold' (1) environments obtained from the spectral densities in Eq. (3.82). *Bottom*: comparison of the simulated (solid lines) correlation functions of the equivalent environments, computed with Eq. (3.84), against the expected ones (dashed lines). The inset plot shows the absolute error between simulated and expected functions. Plots with darker colours represent simulations with $N_\mathrm{E} = 20$, lighter ones with $N_\mathrm{E} = 13$.

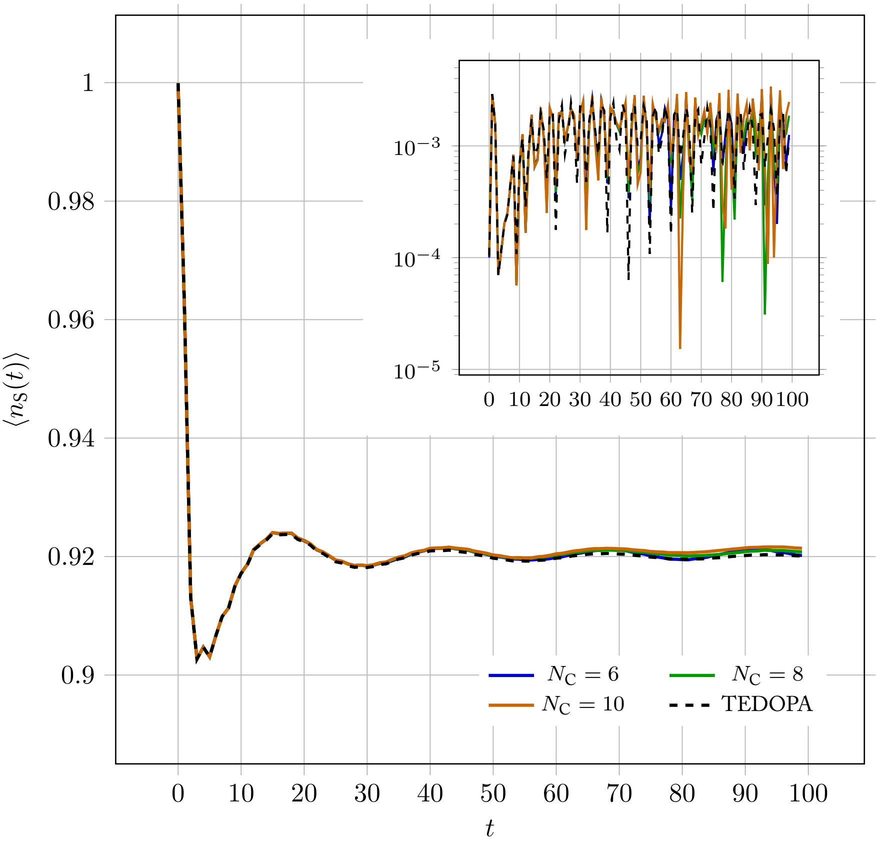

Figure 3.8. Population of the open system, starting from a filled state, comparing the FMC and standard TEDOPA methods on a spinless SIAM simulation with $\Omega = [0, 2]$, $\mu = 0.2$ and $T = 0.4$, $N_E = 6$, $\varepsilon = -\pi/8$. Inset: absolute error between MPS simulations and the exact TEDOPA solution of Fig. 3.6, taken as a reference curve.

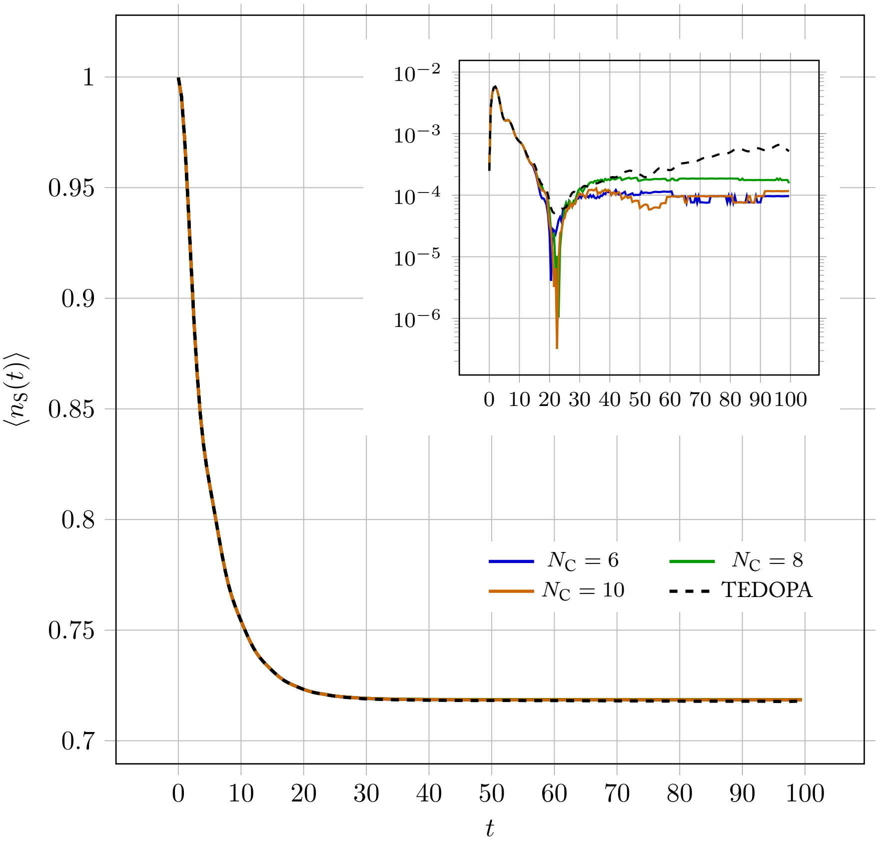

Figure 3.9. Population of the open system, starting from a filled state, comparing the FMC and standard TEDOPA methods on a spinless SIAM simulation with $\Omega = [0, 2]$, $\mu = 1$, $T = 0.4$ and $N_E = 6$, $\varepsilon = -\pi/8$. Inset: absolute error between MPS simulations and the exact TEDOPA solution, in the same way as in Fig. 3.8.

not necessarily improve the quality of the results. As observed in Ref. [35], this can be related to the fact that the $N_\mathrm{C} = 8$ closure better approximates the residual spectral density near the border of the domain than the $N_\mathrm{C} = 10$ one. The quality of the approximation of the semicircle spectral density at the edges of the support impacts on low- and high-momentum components of the wave-packet travelling along the chain. The effect of the deviations from the ideal spectral density in these regions, therefore, strongly depends on the excitation dynamics on the chain, which is, in turn, determined by the chain coefficients. Unbalanced environments, namely when the support of the spectral density $J_{\beta,\mu}(\omega)$ is not symmetric with respect to the chemical potential, will lead to different system-chain coupling strength η_α (see Eq. (3.50)) for the vacuum and filled chains. This asymmetry is responsible for the generation of slowly travelling packets, which, due to the chain momentum/energy dispersion relation, sample the frequency region where the approximation of the asymptotic spectral density is less accurate. If the chemical potential sits at the middle point of the domain of the spectral density, as in the case considered in Fig. 3.9, the situation is more favourable: the system-chain couplings η_α have the same magnitude and the FMC performs better, without a significant difference between the three $N_\mathrm{C} = 6, 8, 10$ cases.

The accuracy does not depend also on N_E, i.e. how many sites we leave on the original chains before we attach the FMC; if we increase this parameter, the errors due to the FMC start to appear at a later time (as expected), but then they increase until they reach the same value. It is clear from the plots in Figs. 3.10 and 3.11 that the chain coefficients are very close to their asymptotic values already from $n = 5$, which motivates the choices for N_E in the simulations shown in the main text. Increasing N_E would increase the numerical costs in general, since the system would be bigger, and possibly the bond dimension too, so we do not see any practical advantage in choosing bigger values for this parameter.

The FMC can be extended straightforwardly to the spinful case of the SIAM, i.e. the interacting model, just by using spinful fermions as modes instead of spinless ones; the simulations become inevitably much more costly than the ones we have already seen, since the local dimension increases from $2^2 = 4$ to $4^2 = 16$. In this case, as discussed in Ref. [152], a simulation in the Heisenberg picture, which in our vectorised setting involves minimal modifications to the time-evolution algorithm, allows us to use an overall lower bond dimension, at the price of being able to simulate the evolution of only a specific observable. In Fig. 3.12 we compare the results in the spinful case obtained by a Heisenberg-picture simulation of the FMC system, with Eq. (3.81) as spectral density, $\varepsilon_\uparrow = \varepsilon_\downarrow = -\pi/8$ and $U = -\pi/4$, $T = 0.4$ and $\mu = 1$. This time, because of the quartic term $U a_\uparrow^\dagger a_\uparrow a_\downarrow^\dagger a_\downarrow$ describing the interactions between the two spin degrees of freedom, the model is not Gaussian, therefore we have no exact

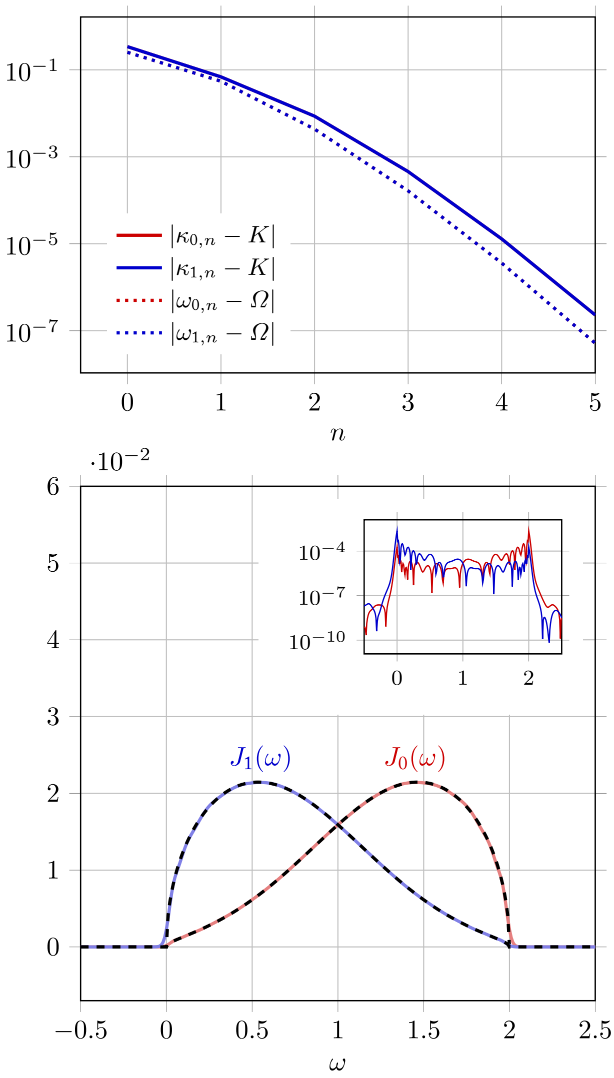

Figure 3.10. *Top*: absolute distance between chain coefficients of the equivalent environments in Fig. 3.9. The blue and red lines coincide, since $\mu = 1$ lies in the middle of the support, consequently $J_0(x) = J_1(\mu - x)$, which leads to $\Omega = 1$, $\omega_{0,n} - \Omega = -(\omega_{1,n} - \Omega)$ and $\kappa_{0,n} = \kappa_{1,n}$ for all n. *Bottom*: comparison of the simulated (solid lines) correlation functions of the equivalent environments, computed with Eq. (3.84), against the expected ones (dashed lines). The inset plot shows the absolute error between simulated and expected functions.

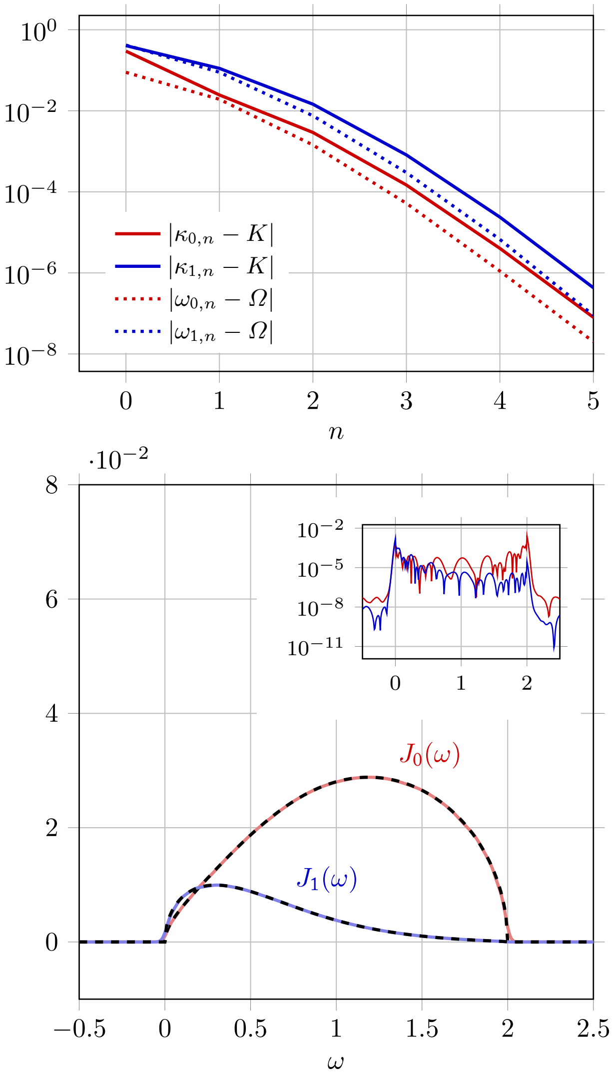

Figure 3.11. *Top*: absolute distance between chain coefficients of the equivalent environments of the simulations in Fig. 3.8. *Bottom*: comparison of the simulated (solid lines) correlation functions of the equivalent environments, computed with Eq. (3.84), against the expected ones (dashed lines). The inset plot shows the absolute error between simulated and expected functions.

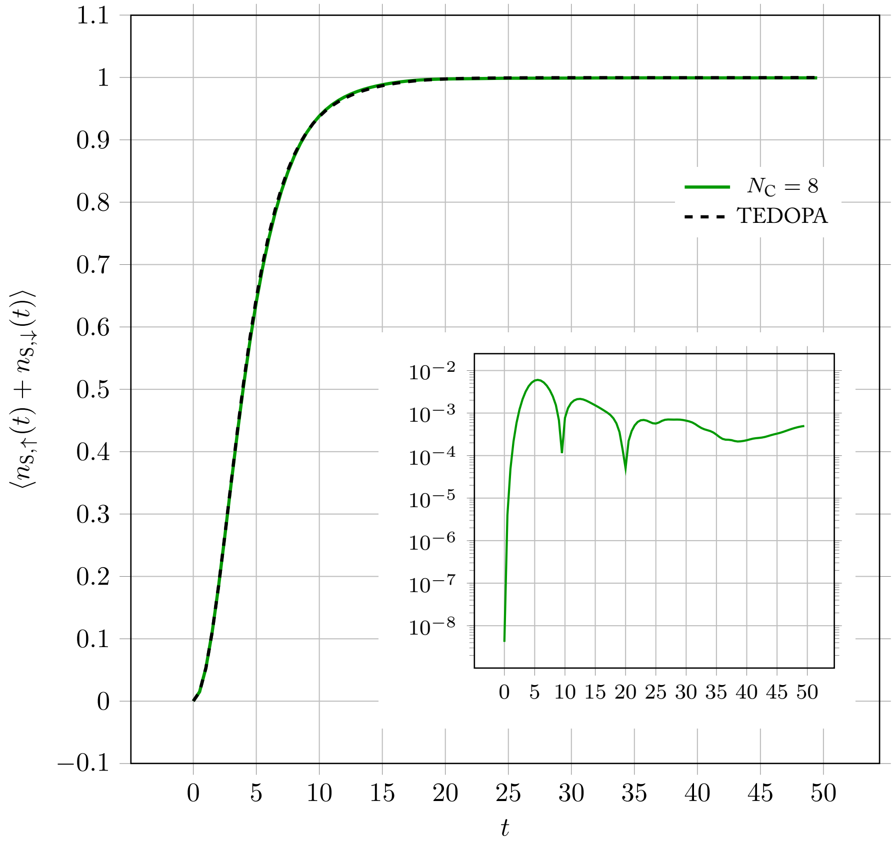

Figure 3.12. Expectation value of the total population $\langle n_{S,\uparrow}(t) + n_{S,\downarrow}(t)\rangle$ of the system site in the spinful SIAM: we show an FMC simulation in the Heisenberg picture, with J as in Eq. (3.81), $\varepsilon_\uparrow = \varepsilon_\downarrow = -\pi/8$, $U = -2\varepsilon_\uparrow$, $T = 0.4$ and $\mu = 1$, and a standard TEDOPA one taken as a reference. Inset: absolute error between FMC and TEDOPA simulations.

solution at our disposal: we compare the FMC results with a simple TEDOPA-only simulation. We observe that after a reasonably long initial transient period the FMC has introduced an error of the order of 10^{-3}, just as in the spinless case. These results show that the FMC method provides an accurate representation, even in a TN form, of the influence of the environment on the dynamics of the OQS.

3.5.3 Computational cost

For a given chain of length L, the complexity of a TEDOPA simulation scales as $O(Lt_{\max}(d\chi)^3)$ [153], where d is the local dimension, χ the bond dimension of the MPS, and $t_{\max}$ the total physical simulation time. In a uniform TEDOPA chain, the excitations propagate at rate that equals twice the coupling constant between the sites. In a typical TEDOPA simulation the asymptotic value of the coupling constant, K, is a good estimate of the actual propagation speed if the coefficient, site by site, converges relatively quickly towards this value. The chain should therefore contain at least $2Kt_{\max}$ sites if we want to avoid that the excitations 'bounce back', creating artefacts due to the finite size; with $L = Kt_{\max}$ we get that the TEDOPA complexity is $O(Kt_{\max}^2 d^3\chi^3)$.

In FMC simulations, the size of the system (i.e. its length L) is fixed, but d gets squared and the required bond dimension χ' is generally higher than χ, in order to account for classical (in addition to quantum) correlations and the presence of longer-range interactions: overall we find $O(t_{\max}d^6\chi'^3)$. Whether the FMC is an advantage over standard TEDOPA depends on how much χ' needs to be higher than χ.

Given these scaling properties, it is clear that the FMC can be an advantage over a standard TEDOPA simulation only if we want to study long-time dynamics, when the $t_{\max}^2$-dependence of the latter can overcome the simpler $t_{\max}$-dependence of the former. We will show an example of this in this section. This tells us that the two methods are really meant to be used in different situations:

- a standard TEDOPA simulation is more efficient, even with very long chains, if the entanglement is low, effectively making use of the favourable scaling properties of matrix-product states;

- the FMC method works best with very long simulations or if the whole system develops moderate levels of entanglement, in which case the determination of the evolution of long chains would require very long computational times.

It is however important to stress that, if the open system starts from a mixed state or if there are Lindblad terms acting on the open system (e.g. dephasing terms), so that both TEDOPA and FMC simulation methods have to deal with density matrices from the start, and thus with a d^6 dependence on the local dimension, then the FMC provides a clear major speed-up, which is not limited to long-time simulations, due to its naturally finite size.

The possibility to reach longer physical times can be beneficial e.g. for the computation of spectral functions: usually, extrapolation of the retarded or advanced Green functions by means of some ingenious method, such as linear prediction [172], is needed. The determination of the retarded Green's function typically requires the preparation of some equilibrium state through an adiabatic evolution from a simple initial state. This evolution must be sufficiently slow so as to reach the correct equilibrium state, so that one can easily end up with very long TEDOPA chains in order to accommodate a high t_{max}. We try out this scheme on the non-interacting model, following the technique presented in Ref. [157]: system and environment are initialised in a product state, then they are evolved using a time-dependent Hamiltonian where the system-environment interaction term is slowly ramped up from zero to its actual value.

We use the spectral density J as in Eq. (3.81) and $\varepsilon = -\pi/8$. At the beginning of the evolution, our open system is in the empty state and the environment in its thermal equilibrium state with $T = 0.4$ and $\mu = 1$; the composite system is slowly evolved under the Hamiltonian $H(t) = H_{\mathrm{E}} + r(t)(H_{\mathrm{I}} + H_{\mathrm{S}})$ where $r(t) = \min\{1, t/\tau\}$, $\tau = 20$, until $t = 100$. Figure 3.13 displays the results of this procedure: we take a standard TEDOPA simulation as a reference. We found that a bond dimension equal to 200 is sufficient for an FMC structure with $N_{\mathrm{E}} = 6$ and $N_{\mathrm{C}} = 6$ to obtain a state ρ_{MC} such that, if ρ_{TEDOPA} is the state obtained with a standard TEDOPA evolution, $|\mathrm{tr}(A(\rho_{\mathrm{MC}} - \rho_{\mathrm{TEDOPA}}))| < 10^{-3}$ for local and non-local observables A acting on the common sites between the two systems, i.e. on the open system and on the first N_{E} sites on each environment chain. We show at last in Fig. 3.14 the difference in computational time between the standard TEDOPA and FMC simulations used to compute these equilibrium states. The plot illustrates the wall-clock time required for each time step in the relaxation phase (i.e. when $t > \tau$) of the adiabatic evolution towards an equilibrium state, considering environments of various lengths. In this situation, entanglement is slowly spreading along the chain, requiring more and more time for each sweep of the TDVP algorithm in the normal TEDOPA cases. This can be explained by observing that while the operations on the tensors, especially the matrix decompositions, admit a general dependence on the *dimensions* of the involved tensors, they can also depend on the actual *values* of the coefficients of the tensors. For example, the SVD algorithm found in commonly used linear algebra libraries terminates sooner if fewer singular values are not zero (or above the cutoff we have set to discard them): the operations can then, in practice, perform better than what the reported scaling suggests. In the case of Fig. 3.14, in particular, the tensor operations are more efficient in the regions of the TEDOPA chain far from the central system, and for this reason the time required for each step is not constant in time—even though it increases if the chain is longer—but generally increases as the evolution proceeds (and eventually settles on a final value). With the FMC, instead, the entanglement increases and spreads rapidly at the beginning (which here already happened in the ramp-up phase of the adiabatic evolution, not shown) and reaches very soon the maximum quantity allowed by the limited bond dimension. After this point, the computational time

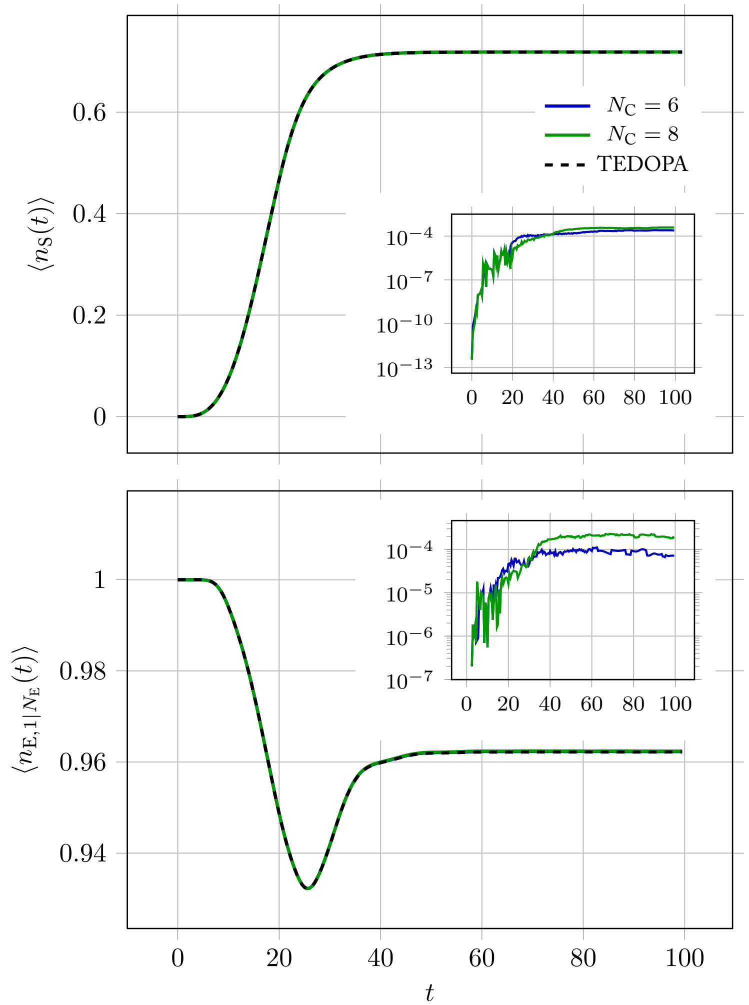

Figure 3.13. Adiabatic evolution of non-interacting SIAM with $T = 0.4$, $\mu = 1$, $\varepsilon = -\frac{\pi}{8}$ and J given by Eq. (3.81). Population of the open system (top frame) and of the last site before the FMC in the chain of the initially filled environment (bottom frame). The composite system is evolved from a product state towards a state where the impurity is in equilibrium with the environment. Insets: absolute error between the FMC and TEDOPA simulations.

remains constant as the evolution goes on.

This example clearly shows that the FMC can reduce the time complexity of the simulation and is more efficient than standard TEDOPA when long chains and moderate levels of entanglement are involved.

3.5.4 Technical improvements

The comparison of the time evolution obtained with normal TEDOPA and with the FMC, considering several systems and environments, has proven to be a satisfactory match as far as the expectation values of operators and the environment correlation functions were concerned.

There are of course some improvements that can be studied in order to enhance the performance of the FMC method. While the quality of the results obtained by means of FMC is already more than satisfying, we could look for a better parameterisation of the closure, e.g. by means of the techniques introduced in Refs. [166, 167], in order to improve the accuracy of the fitting of the asymptotic spectral density, as well as fermion-to-spin representations leading to a decrease of the entanglement in the TN. Another improvement comes from using a different factorisation technique for the re-orthogonalisation of the MPS during a TDVP sweep, where the cost of a traditional SVD decomposition can become very demanding (it scales as $(d\chi)^3$ where d is the local dimension and χ the bond dimension) especially in the spinful case where $d = 16$. The *reduced-rank randomised* SVD [173] has a more favourable scaling with respect to the standard SVD, and might be useful to speed up the calculations. In situations where the memory of the environment saturates on time scales which are shorter than relaxation, the *dynamical map extraction* technique shown in [92, 174] can be used to compute equilibrium TTCFs by extrapolating the initial transient dynamics beyond the memory time of the environment, up to the stationary regime. We finally observe that, for very long simulation times that exceed the memory time of the environment, the transfer tensors formalism [175] can be used to further enhance efficiency.

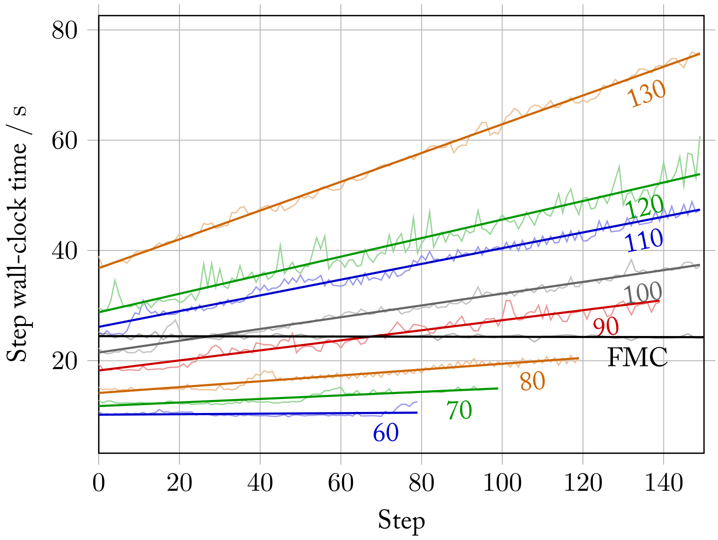

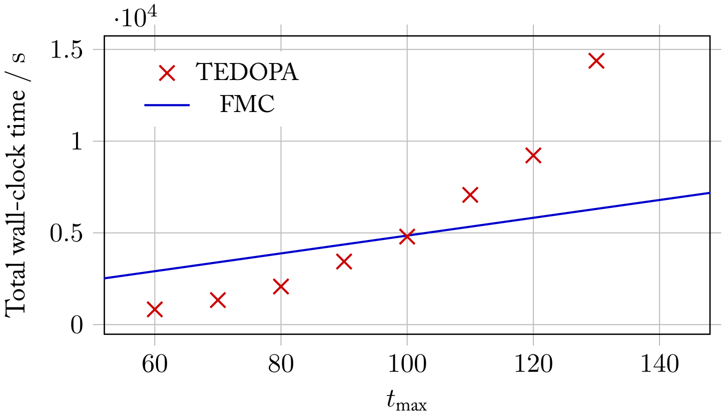

Figure 3.14. *Top*: wall-clock time required with standard TEDOPA for each step of the relaxation phase of the system described in Fig. 3.13, for different $t_{\max}$ (shown alongside the curves). While the duration of individual steps (lighter curves) can vary due to situational factors, we observe a clear increasing trend (straight, darker lines), and the increase is steeper when the chain is longer. *Bottom*: wall-clock time required for the whole evolution in the top figure, showing the how the FMC outperforms standard TEDOPA over long times.

Quantum simulation of a many-body chaotic system 4

In recent years considerable effort has been devoted to the design of quantum computers with an increasing number of qubits and capable of executing circuits with the best possible accuracy. The information handled by quantum computers is extremely fragile, and building devices that can shield it from noise, or correct errors that emerge during the computation, is a very difficult engineering challenge. For example, qubits can interact with each other in uncontrolled ways, leading to differences between the expected and the actual calculations. The qubits may also interact with the external environment, causing decoherence phenomena to such an extent that the quantum behaviour of the system is partially or completely lost during the circuit execution. Some amount of physical error is inevitable, but if the physical error rate is kept below a certain threshold then, through the application of error-correction schemes, the logical error rate can be suppressed to arbitrarily low levels, making the computer *fault tolerant*—essentially correcting errors faster than they appear [45–47]. While research on designing and building better and more resilient quantum computers continues, and several new advanced error-correction schemes are being devised, quantum computers available today are noisy and with a relatively small size: they are commonly called *noisy intermediate-scale quantum* (NISQ) devices [54].

Among the various possible application of quantum computers, we find their use as simulators, as already theorised by Richard Feynman [176]. By virtue of being quantum systems themselves, quantum computers are well-suited to simulate the behaviour of *other* quantum systems such as complex molecules or materials. The parameters of the physical quantum platform can be tuned to match the behaviour of the system being studied, whose dynamics can be investigated by measuring the state on the platform. In particular, quantum circuits with local unitaries have emerged as a rich playground for the exploration of many-body quantum dynamics of discrete-time systems, which can be implemented exactly, without being subject to discretisation errors due e.g. to the Suzuki-Trotter approximation. This *digitalisation* of the dynamics of quantum systems can be applied to investigate, for example, the behaviour of measurement-induced phase transitions [177], ergodic phases of many-body systems [178], phase transitions in random unitary circuit models [179].

In this chapter we will analyse a particular quantum-computing experiment which involves *dual-unitary (DU)* circuits [59], a class of quantum circuits with specific unitarity properties which allows the exact computation of some properties of the system that would typically be very hard to calculate. DU circuits have recently attracted attention because they are an analytically tractable model of *quantum chaos*, which is an interesting open problem in many-body physics. A number of measures have been proposed to characterise whether a quantum system can be said to be chaotic: for example, from the point of view of spectral statistics, a chaotic system possesses an energy spectrum with correlations described by random matrix theory [60]; from the point of view of dynamical complexity, instead, it is a system where local-operator entanglement grows linearly over time [180, 181]. In DU circuits, many of these measures can be computed analytically, and in most cases, these systems turn out to be maximally chaotic [60]. This makes DU systems interesting because we have a chaotic system that can be studied analytically without it necessarily being an integrable model. Since they can be studied classically, DU systems have also found some use as a testing ground for benchmarking quantum processors, as in Refs. [56, 182]. The intrinsic locality of the interaction, moreover, makes them particularly well suited to run on current quantum processors, where error-mitigation techniques can be used to recover the ideal (noiseless) value of physical quantities by processing the output of the quantum device on a classical computer [183–185]. However, in cases where we cannot have an exact analytical formula for the quantity of interest, verifying the results at non-trivial scales is a highly complex task. One of the possible methods to benchmark these devices is to run a classical simulation of the calculations and to compare the results with the ones obtained on quantum hardware. An exact description of the computer's state at any time during the execution of a circuit is usually possible only for a very small number of qubits, because of the exponential scaling of the Hilbert space dimensions with the number of qubits; the most common approach is therefore to build an approximate representation of the state, for example by means of TNs.

The drawing of a set of one-, two- or three-site local operators applied sequentially to an MPS has a stark resemblance to a logic circuit where a sequence of gates is applied to the input state: it is not surprising then that TNs are widely used to describe gate-based quantum computers. Quantum objects such as state vectors, mixed states, operators and channels can be encoded into TNs, so that the execution of a quantum circuit can be simulated on a classical computer. Since one of the main features that makes quantum computers powerful is their ability to entangle qubits and perform calculations on them, and since the MPS representation becomes more expensive as the state entanglement increases, we expect classical simulation with MPSs to become unfeasible as we increase the number of qubits. Nevertheless, in the case of NISQ devices, TNs can still contribute to the simulation in different ways:

- by performing a simulation of the noisy output, which can validate the theoretical model describing the noise and error sources;
- by using them to perform *error mitigation* on the results, instead of a full simulation.

Both cases, as we will see, lie more readily within the reach of TN tools. In the experiment presented in this chapter TNs will be employed not only for a straight simulation of the execution of the quantum circuit, but also to perform benchmark simulations for the noise model of the computer, as well as to apply a particular error-mitigation scheme aimed at recovering an approximation of the ideal state from the noise-affected output.

4.1 Dual-unitary circuits

Given a two-mode operator $U\colon \mathcal{H}\otimes\mathcal{H} \to \mathcal{H}\otimes\mathcal{H}$, and an orthonormal basis $\{|e_i\rangle\}_{i\in\mathcal{I}}$ on $\mathcal{H}$, from the components

$$U = \sum_{i,j,k,l\in\mathcal{I}} U_{i,j,k,l} |e_k, e_l\rangle\langle e_i, e_j| \tag{4.1}$$

one defines a *dual operator* by

$$U_{\mathrm{D}} := \sum_{i,j,k,l\in\mathcal{I}} U_{i,j,k,l} |e_j, e_l\rangle\langle e_i, e_k| \tag{4.2}$$

through a shuffling of two input/output subsystems of U (we have swapped j and k in the bra/ket indices). If both U and U_{D} are unitary operators, then U is called *dual unitary* (DU). When viewed as a two-qubit gate, a DU operator gate can be seen as a gate that is unitary when viewed as a propagator along the ‘spatial’ direction as well as the ‘temporal’ direction.

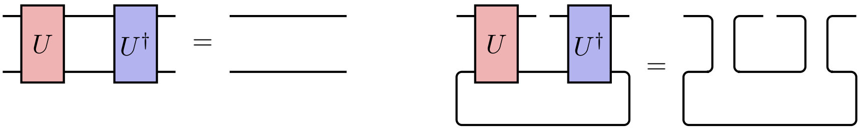

Figure 4.1. Cancellation rules for unitary and dual-unitary gates. The left half shows the usual rule for unitary gates: a two-qubit unitary gate U composed with its adjoint $U^\dagger$ returns $\mathbb{1} \otimes \mathbb{1}$, the identity on both qubits. If U is also dual unitary, then the cancellation rule on the right applies as well, yielding another way to simplify gates in a circuit.

DU circuits, which will be the primary focus of the simulations in this chapter, consist of N qubits (labelled with $n \in \{0, 1, \dots, N-1\}$) evolving by a brickwork pattern of DU gates $U_{n,n+1}$ [59, 180, 186]. The brickwork is composed of 'even' and 'odd' layers of DU gates,

$$U_{\text{even}} = \bigotimes_{j=0}^{\lfloor \frac{N}{2} \rfloor - 1} U_{2j,2j+1} \quad \text{and} \quad U_{\text{odd}} = \bigotimes_{j=1}^{\lfloor \frac{N}{2} \rfloor} U_{2j-1,2j} \tag{4.3}$$

respectively, repeated one after the other.

A periodic gate structure such as this one can be associated to a class of interesting physical systems, *periodically driven systems*, which are widely used as one of the simplest settings that can display chaos, a famous example being the kicked rotor [187]. Recent motivation to study these systems comes from trying to understand many-body physics as well as from quantum information [188]. Floquet's theorem [189] states that for a periodically driven system we can write the evolution operator $U(t', t)$ as

$$U(t', t) = \exp(-\mathrm{i}K_{t'}) \exp(-\mathrm{i}(t' - t)H) \exp(\mathrm{i}K_t) \tag{4.4}$$

where K_t is a time-dependent, periodic Hermitian operator, and H is a time-independent Hamiltonian. These systems are sometimes easy to understand: for example in the kicked-rotor case the propagator is written as the product of two other propagators, $U = V_1 V_2$ where V_1 is the kinetic term and V_2 the potential; still, it can result in rich physics even when the constituents V_1 and V_2 alone (provided they do not commute) have simple non-chaotic dynamics. Continuing this splitting idea one can write U in terms of individual two-qubit gates as in quantum computation, resulting in a quantum circuit. Such formulation is also extremely handy for modern numerical simulation methods based on MPSs.

4.1.1 The kicked Ising model

A natural example of a chaotic model with local interactions, on a finite Hilbert space, in the class of DU systems is the *kicked Ising model* [190, 191]. It describes a classical Ising model subject to a longitudinal magnetic field,

$$H_{\mathrm{I}} = J \sum_{n=0}^{N-2} Z_n Z_{n+1} + h \sum_{n=0}^{N-1} Z_n, \tag{4.5}$$

which is periodically 'kicked' by a transverse field

$$H_{\mathrm{K}} = b \sum_{n=0}^{N-1} X_n, \tag{4.6}$$

where X_n, Y_n, Z_n are the local Pauli operators on qubit n.

We will see now how we can represent the time evolution of the kicked Ising model with a quantum circuit made of two-qubit DU gates. The kicked Ising model is described by the time-dependent Hamiltonian:

$$H_{\mathrm{KI}}(t) = H_{\mathrm{I}} + \sum_{m\in\mathbb{Z}} \delta(t-m) H_{\mathrm{K}}, \tag{4.7}$$

that generates the time-evolution operator over one period (which in this case is equal to 1) [192]

$$U_{\mathrm{KI}} = \mathrm{e}^{-\mathrm{i}H_{\mathrm{K}}} \mathrm{e}^{-\mathrm{i}H_{\mathrm{I}}}, \tag{4.8}$$

which is the form of Floquet's theorem in Eq. (4.4), with $t' = 1$ and $t \to 0^+$, so that only one kick is included in a single period. We can therefore view the temporal evolution of this system as a sequence of discrete steps, at each of which we apply U_{KI} to the state. This means that we can *exactly* represent the propagator U_{KI} with a sequence of quantum gates. This is a very convenient feature of the kicked Ising model, in that the continuous time evolution does not require a Suzuki-Trotter approximation in order to be digitally simulated.

This Floquet evolution can be implemented through a brickwork circuit where the two-qubit building blocks are specified by the model parameters h, J, and b. The propagator can be written as

$$U_{\mathrm{KI}} = \exp\Big(-\mathrm{i}b\sum_n X_n\Big)\exp\Big(-\mathrm{i}J\sum_n Z_nZ_{n+1}\Big)\exp\Big(-\mathrm{i}h\sum_n Z_n\Big) = \tag{4.9}$$

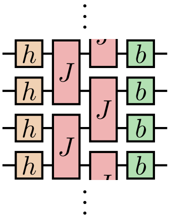

where we have introduced the gates

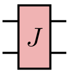

$$= \exp(-\mathrm{i}JZ\otimes Z),$$

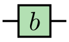

$$= \exp(-\mathrm{i}bX),$$

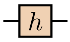

$$= \exp(-\mathrm{i}hZ). \tag{4.10}$$

Now, consider the two-qubit gate

$$\begin{aligned} U_{n,n+1} = &\exp(-\mathrm{i}hZ_n)\exp(-\mathrm{i}JZ_nZ_{n+1})\exp(-\mathrm{i}b(X_n+X_{n+1}))\,\cdot \\ &\cdot\exp(-\mathrm{i}JZ_nZ_n+1)\exp(-\mathrm{i}hZ_n) = \end{aligned} \tag{4.11}$$

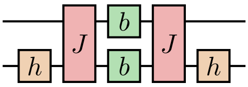

A brickwork circuit made up of this gate consists of an even and an odd layer of gates as in Eq. (4.3), repeated periodically. Two time steps of the brickwork are given by

$$U = U_{\mathrm{odd}}U_{\mathrm{even}} = \tag{4.12}$$

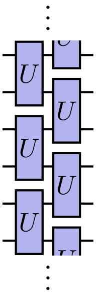

$$=$$

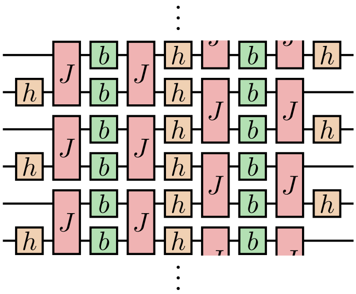

Define the unitary operator

$$\Sigma = \prod_{n=1}^{\frac{N-1}{2}} \exp(-\mathrm{i} J Z_{2n+1} Z_{2n}) \prod_{n=1}^{\frac{N-1}{2}} \exp(-\mathrm{i} h Z_{2n+1}) = \tag{4.13}$$

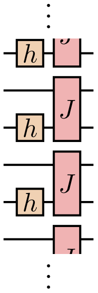

and consider the circuit given by

$$\Sigma^\dagger U \Sigma = \tag{4.14}$$

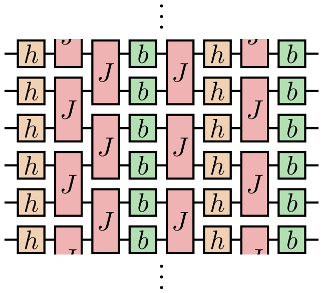

Using the fact that the h-gates and the J-gates commute, and comparing Eq. (4.14) with the circuit diagram for U_{KI} in Eq. (4.9), we can see that $\Sigma^\dagger U \Sigma$ is equal to two periods of the Floquet unitary operator of the kicked Ising model:

$$\Sigma^\dagger U \Sigma = U_{\mathrm{KI}}^2. \tag{4.15}$$

This shows that the Floquet unitary for the kicked Ising model can be related to a brickwork circuit with a fixed two-qubit gate given by Eq. (4.11).

For $|J| = |b| = \pi/4$ the $U_{n,n+1}$ gates are dual-unitary, for any choice of h [193]. If $h = 0$ the model is integrable, as it can be mapped to free fermions via a Jordan-Wigner transformation: the corresponding brickwork circuit is composed of Clifford gates, and circuits consisting of Clifford gates only can be efficiently simulated with a classical computer, as stated by the Gottesman–Knill theorem [194]. In general, for other choices of h, the model is not integrable, yet analytical solutions exist for the time evolution of certain correlation functions.

4.2 Experiment on the quantum computer

In this section we comment briefly on how the experiment was run on the quantum computer and on the results, in order to establish the context in which the TN simulations were performed.

4.2.1 Setup

We assume that the number of qubits N is odd, and we initialise the quantum computer in the state

$$|\psi_0\rangle = |+\rangle \otimes |\psi_{\text{Bell}}\rangle^{\otimes \frac{N-1}{2}}. \quad |\psi_{\text{Bell}}\rangle = \frac{1}{\sqrt{2}}(|0\rangle \otimes |0\rangle + |1\rangle \otimes |1\rangle) \tag{4.16}$$

that prepares the 0-th qubit in the state $|+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$ and all other qubits in a product of Bell pairs $|\psi_{\text{Bell}}\rangle$ (for all qubit pairs $i, i+1$, with $i \in \{1, 3, \dots, N-2\}$). We want to measure the function $t \mapsto \langle\psi_0|X_n(t)|\psi_0\rangle$, where $X_n(t)$ is the Pauli X operator on the n-th qubit evolved with U_{KI} in the Heisenberg picture. Figure 4.2 (step 1) depicts, for example, the circuit representing $2^{\frac{N-1}{2}}\langle\psi_0|X_n(t)|\psi_0\rangle$ with $n = t = \frac{N-1}{2}$. Bell pairs are represented as simple arcs connecting the qubits, since applying a two-qubit gate to $|\psi_{\text{Bell}}\rangle$ is equivalent, except for the $\sqrt{2}$ normalisation factor, to connecting (i.e. contracting) the input legs of the gate:

$$\sqrt{2}\,(\langle e_i| \otimes \langle e_j|) A \otimes B |\psi_{\text{Bell}}\rangle = \sum_{l=0}^{1} \langle e_i|A|e_k\rangle\langle e_j|B|e_k\rangle. \tag{4.17}$$

If the system is at one of its DU points, i.e. if $|J| = |b| = \pi/4$, then using the DU cancellation rules of Fig. 4.1 we can simplify the circuit representing $\langle\psi_0|X_t(t)|\psi_0\rangle$ as in Fig. 4.2. The result is

$$\langle\psi_0|X_t(t)|\psi_0\rangle = \langle +|\mathcal{T}_U^t(X)|+\rangle \tag{4.18}$$

where $\mathcal{T}_U \colon L(\mathcal{H}) \to L(\mathcal{H})$ is the map

$$A \mapsto \frac{1}{2}\,\mathrm{tr}_2\big(U^\dagger(\mathbb{1} \otimes A)U\big) = \frac{1}{2} \tag{4.19}$$

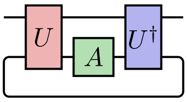

in which tr_2 is the partial trace on the second factor. This map, with our DU gate

$$U = \exp(-\mathrm{i}hZ \otimes \mathbb{1}) \exp(-\mathrm{i}JZ \otimes Z) \exp(-\mathrm{i}b(X \otimes \mathbb{1} + \mathbb{1} \otimes X)) \cdot \\ \cdot \exp(-\mathrm{i}JZ \otimes Z) \exp(-\mathrm{i}hZ \otimes \mathbb{1}) \quad (4.20)$$

and $|b| = |J| = \pi/4$, is such that

$$\begin{aligned} \mathcal{T}_U(\mathbb{1}) &= \mathbb{1}, \\ \mathcal{T}_U(X) &= \cos(2h)X + \sin(2h)Y, \\ \mathcal{T}_U(Y) &= 0, \\ \mathcal{T}_U(Z) &= 0, \end{aligned} \quad (4.21)$$

therefore $\langle\psi_0|X_t(t)|\psi_0\rangle = \cos(2h)^t$. With a similar reasoning we can prove that $\langle\psi_0|X_n(t)|\psi_0\rangle = 0$ if $n \neq t$.

Noting that $|+\rangle\langle+| = \frac{1}{2}(\mathbb{1} + X)$ and that $\mathrm{tr}(\mathcal{T}_U(X)) = 0$, we can also write

$$\langle\psi_0|X_t(t)|\psi_0\rangle = \frac{1}{2}\,\mathrm{tr}\big(X\mathcal{T}_U^t(X)\big). \quad (4.22)$$

From Fig. 4.2 we observe that the final expression we arrive at is identical to the one given in [59] where an infinite-temperature initial state $\rho_\infty = 2^{-N}\mathbb{1}$ is used in place of our Bell pairs.[1] We have therefore an intuitive physical interpretation for our expectation value: it is equivalent to the infinite-temperature auto-correlation function

$$C_n(t) = \mathrm{tr}\big(\rho_\infty X_0(0) X_n(t)\big), \quad (4.23)$$

Let us briefly sum up what we have derived in this section. We have seen how in a brickwork circuit of DU gates limits information spread not only within a light-cone shape in the 'temporal' direction, that is for $t < n$, but also in the 'spatial' direction for $n < t$. As a result, non-zero correlations are only found on the 'diagonal' of the circuit where $t = n$. The expectation value we are interested in can be computed exactly: for odd N and $t \leq (N-1)/2$,

$$\langle\psi_0|X_n(t)|\psi_0\rangle = C_n(t) = \begin{cases} \cos(2h)^t & \text{if } n = t \\ 0 & \text{otherwise,} \end{cases} \quad (4.24)$$

On the light cone, the auto-correlation function is constant at the integrable Clifford point $h = 0$, but otherwise decays exponentially in time. When the circuit is not DU, the equivalence with the infinite-temperature model and Eq. (4.24) do not hold anymore. We expect anyway that when the parameters are tuned slightly away from the DU points, the spreading of information will still be, for the most part, confined near the diagonal of the circuit.

---

[1] This equivalence holds as long as $N \geq 2t + 1$, which will always be the case in this experiment.

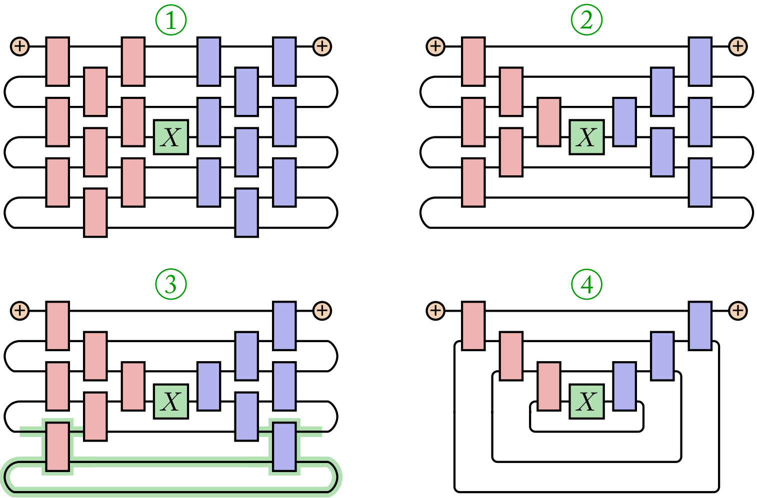

Figure 4.2. (1) Calculation of $2^{\frac{N+1}{2}} \langle\psi_0|X_t(t)|\psi_0\rangle$ in a DU circuit with $N = 7$ qubits, and $t = 3$, where each red block represents a DU gate U and blue blocks represent $U^\dagger$. (2) Unitary gates cancel with their adjoints, leaving only gates within the light cone defined by the $|+\rangle$ state on the first qubit and the X_t gate in the middle. (3) The DU property of the gates allows us to further simplify the expression. The first pair of gates we can simplify, as in Fig. 4.1, is highlighted in green. (4) By using the DU cancellation rules we can then remove all gates which lie off the diagonals joining the $|+\rangle$ state and the X_t operator. We can also join the two ⊕ into the $|+\rangle\langle+|$ operator, thereby exactly replicating the diagrams in [59].

4.2.2 Measurement

The most general kind of measurement of a quantum system is described by a positive operator-valued measure (POVM), that is a set $\{\Pi_k\}_{k=1}^M \subseteq L(\mathcal{H})$ of positive semi-definite operators that sum to the identity.[2] An element Π_k of the POVM, also called *effect*, is associated with the measurement outcome k, such that the probability of obtaining it when performing a measurement on the state ρ is given by $\mathrm{tr}(\rho\Pi_k)$. We call a POVM *informationally complete* if, in addition, its effects span $L(\mathcal{H})$. To an informationally complete (IC) POVM one can associate a dual frame, that is a set

---

[2] We consider here only the discrete case, which is enough for our experiment as we only deal with finite-dimensional Hilbert spaces.

$\{D_k\}_{k=1}^M$ of Hermitian operators such that

$$A = \sum_{k=1}^{M} \mathrm{tr}(D_k A) \mathit{\Pi}_k \tag{4.25}$$

for any operator A. This means that the expectation value of an observable A on a state ρ can be written as

$$\mathrm{tr}(\rho A) = \sum_{k=1}^{M} \mathrm{tr}(D_k A)\, \mathrm{tr}(\rho \mathit{\Pi}_k) = \mathrm{tr}\left(\sum_{k=1}^{M} \mathrm{tr}(\rho \mathit{\Pi}_k) D_k A \right) \tag{4.26}$$

and we can write, therefore,

$$\rho = \sum_{k=1}^{M} \mathrm{tr}(\rho \mathit{\Pi}_k) D_k \tag{4.27}$$

meaning that we can reconstruct the state from the average of the dual operators over the probabilities $\mathrm{tr}(\rho \mathit{\Pi}_k)$. Equation (4.27) is a generic decomposition of the state in terms of the effects and their duals, and if we actually used such a formula to reconstruct ρ from the measurements the computational cost would scale exponentially with the number of qubits, as in a full state tomography procedure [195]. However, we are not interested in knowing the final state of the quantum device, but rather in measuring Eq. (4.24), and we can optimise the measurement procedure towards this goal. We consider as our IC POVM a set of projective single-qubit measurements in the eigenbases of the Pauli operators X, Y, Z. On each qubit, for each circuit run, we choose one of the effects according to a set probability p_X, p_Y, p_Z, namely $p_X = p_Y = p_Z = \frac{1}{3}$ for all qubits but the n-th qubit (the one associated with the observable X_n of interest), for which $p_X = \frac{8}{10}$ and $p_Y = p_Z = \frac{1}{10}$.The full-system POVM is then determined by the tensor products of the elements of the single-qubit POVMs effects, and the associated dual operators.

$$\begin{aligned}
\mathit{\Pi}_{Z,+} &= p_Z |0\rangle\langle 0|, & \mathit{\Pi}_{Z,-} &= p_Z |1\rangle\langle 1|, \\
\mathit{\Pi}_{X,+} &= p_X |+\rangle\langle +|, & \mathit{\Pi}_{X,-} &= p_X |-\rangle\langle -|, \\
\mathit{\Pi}_{Y,+} &= p_Y |\mathrm{i}\rangle\langle \mathrm{i}|, & \mathit{\Pi}_{Y,-} &= p_Y |-\mathrm{i}\rangle\langle -\mathrm{i}|
\end{aligned} \tag{4.28}$$

with $|\pm\rangle \coloneqq \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle)$ and $|\pm\mathrm{i}\rangle \coloneqq \frac{1}{\sqrt{2}}(|0\rangle \pm \mathrm{i}|1\rangle)$, with respective dual operators

$$D_{\alpha,\pm} = \frac{1}{2}\left(\mathbb{1} \pm \frac{1}{p_\alpha} P_\alpha \right), \tag{4.29}$$

for $\alpha \in \{X, Y, Z\}$, where the P_α are Pauli matrices. Suppose we performed a series of measurements yielding a sequence $\mathcal{S}$ of results, each of which is a list of elements $(\alpha, \sigma) \in \{X, Y, Z\} \times \{+, -\}$, one for each qubit. We can estimate the circuit output by

$$\rho_{\mathcal{S}} = \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} D_s, \tag{4.30}$$

and we can define an (unbiased) estimator for the expectation value of an observable A as [196, 197]

$$\bar{A} = \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \operatorname{tr}(D_s A). \tag{4.31}$$

In practice, we do not perform the measurements independently from one another: instead of randomly drawing a new set of effects from Eq. (4.28) for each run of the circuit, which is impractical for large-scale problems, we run the circuit with the same effects several times. This means that we need to correct the formula for the variance to address this. Calling $\mathcal{M}$ the set of different circuit settings (i.e. of effects), and $\mathcal{S}_m$ the set of outcomes for each setting $m \in \mathcal{M}$, we collect a total number of $|\mathcal{M}||\mathcal{S}|$ shots. Denoting by ξ_m the average outcome for a given setting $m \in \mathcal{M}$, The appropriate way to compute the variance in this case is [58]

$$\operatorname{Var} A = \frac{1}{|\mathcal{M}|^2 |\mathcal{S}_m|^2} \sum_{m \in \mathcal{M}} \sum_{s \in \mathcal{S}_m} (s - \xi_m)^2 + \frac{1}{|\mathcal{M}|^2} \sum_{m \in \mathcal{M}} (\xi_m - \langle A \rangle)^2. \tag{4.32}$$

The DU circuit, for several values of h, was implemented and run with up to 91 superconducting transmon qubits and 4095 two-qubit gates on IBM's 'Eagle' processor `ibm_strasbourg` [198]. The experiment was then extended beyond the analytically tractable points, by tuning b away from $\pi/4$. When decomposed into the native gate set of the quantum processor, each two-qubit block of Eq. (4.12) is transpiled into a sequence that includes two entangling echoed cross-resonance gate (ECR) gates. Thus, each odd (even) layer is built from two entangling layers of parallel ECR gates on odd (even) neighbouring qubit pairs, interleaved with layers of single-qubit gates.

4.3 Noise mitigation

The results are affected by various sources of noise and unwanted interference during the execution of the circuit. For the purposes of this section, we distinguish three types of noise.

- State preparation and measurement (SPAM) errors arise when the qubits fail to be initialised the intended state, for example due to thermal excitations, or because operations on one qubit affect the state of another during preparation or measurement.
- Coherent errors occur in the gate operations, a unitary type of error, which can involve small unexpected rotations acting anywhere on the qubits. A typical situation where this arises is when frequencies of oscillator qubits are misaligned: this results in systematic unitary over- or under-rotations of their states [199].
- Incoherent errors occur because of the entanglement of the qubits with the external environment.

The first technique used to suppress errors is *Pauli twirling*: for every parameter configuration, several instances of circuits that implement the same global unitary but differ in their single-qubit gate layers are run. In other words, independent, random, single-qubit gates are introduced into the logical circuit in such a way that the effective logical circuit remains unchanged, but the noise is effectively converted into a stochastic Pauli channel [200], i.e. a channel E that maps an n-qubit state ρ to

$$E(\rho) = \sum_{P \in P^{\otimes n}} c_P P \rho P^{\dagger}, \tag{4.33}$$

where $P^{\otimes n}$ is the set of the 4^n n-qubit Pauli operators and the coefficients c_P are a probability distribution over $P^{\otimes n}$. Pauli twirling transforms coherent errors into incoherent errors: it does not suppress errors per se, but transforms them into errors that are easier to characterise and mitigate, e.g. with the techniques we will introduce later.

There exist standard readout error mitigation techniques that address SPAM errors such as the one presented in Refs. [201, 202]. In this instance, this kind of error was characterised for all qubits by measuring the single-qubit $\langle 0| Z_i |0\rangle$ expectation values while twirling the readout [203]: deviations from the expected value (of zero) were used to correct the final result.

4.3.1 Tensor-network error mitigation

The techniques introduced in the previous section do not eliminate all errors: the leftover noise is tackled by the tensor-network error mitigation (TEM) method, introduced in [58]. The basic idea behind TEM is to construct an approximate and efficient noise-cancelling map as a TN which maps the noisy state produced by the quantum device to the ideal noiseless state. We assume that the calculation ideally performed by the circuit is described by a sequence $(\mathcal{U}_l)_{l=1}^L$ of unitary layers, and $(\mathcal{N}_l)_{l=1}^L$ are quantum channels describing the effects of noise on each layer, such that the ideal and noisy circuits are described by

$$\mathcal{C} = \mathcal{U}_L \circ \cdots \circ \mathcal{U}_2 \circ \mathcal{U}_1, \qquad \tilde{\mathcal{C}} = \mathcal{U}_L \circ \mathcal{N}_L \circ \cdots \circ \mathcal{U}_2 \circ \mathcal{N}_2 \circ \mathcal{U}_1 \circ \mathcal{N}_1, \tag{4.34}$$

respectively. The TEM algorithm works by constructing an MPO that, applied to the final state of the noisy circuit written as a TN, reconstructs the ideal state, as follows:

$$\begin{aligned}\rho_{\text{ideal}} &= \mathcal{C} \circ \tilde{\mathcal{C}}^{-1}(\rho_{\text{noisy}}) = \\ &= \mathcal{U}_L \circ \cdots \circ \mathcal{U}_2 \circ \mathcal{U}_1 \circ \mathcal{N}_1^{-1} \circ \mathcal{U}_1^{-1} \circ \mathcal{N}_2^{-1} \circ \mathcal{U}_2^{-1} \circ \cdots \circ \mathcal{N}_L^{-1} \circ \mathcal{U}_L^{-1}(\rho_{\text{noisy}}).\end{aligned} \tag{4.35}$$

Constructing the TEM map $\mathcal{M} = \mathcal{C} \circ \tilde{\mathcal{C}}^{-1}$ as in Eq. (4.35) by concatenating the constituent maps layer by layer, in order, would lead to a complexity growing exponentially in the number of layers: it would be equivalent to simulating the circuit exactly. However, the computation is made efficient via the recurrence relation

$$\mathcal{M}_l = \mathcal{U}_l \circ \mathcal{M}_{l-1} \circ \mathcal{N}_l^{-1} \circ \mathcal{U}_l^{-1}, \qquad \mathcal{M}_0 = \mathbb{1}, \tag{4.36}$$

since every unitary layer $\mathcal{U}_l$ and corresponding inverse noisy layer $\mathcal{N}_l^{-1} \circ \mathcal{U}_l^{-1}$ approximately cancel each other. Both $\mathcal{U}_l$ and $\mathcal{N}_l^{-1}$ allow a TN representation in the form of an MPO of bond dimension 4, so the map $\mathcal{M}$ is constructed via recurrent conventional TN contractions of MPOs. In practice, however, we need to compress the MPO at each iteration, therefore in the end we obtain an approximate map capturing the most significant contributions.

In order to apply the TEM map to the output, we first need an MPO representation of the final noisy state; this can be done with the IC measurements described in Section 4.2.2. Since we chose an IC POVM which is a product of single-qubit effects, each operator D_k of the dual frame (which is also a product of single-qubit operators) can be immediately written in MPS form, as the vectorisation of an MPO. In principle, we could then take the average of these MPSs as in Eq. (4.30), apply the TEM map $\mathcal{M}$ to the result and in the end compute the expectation value on the error-mitigated state. Realistically, we want to avoid adding MPSs together, as it can be a computationally heavy operation (see e.g. [62, Sec. 4.3]), so we are better off calculating the expectation value as

$$\langle A \rangle = \sum_{i=1}^{M} \operatorname{tr}\big(A\mathcal{M}(D_{s_i})\big), \tag{4.37}$$

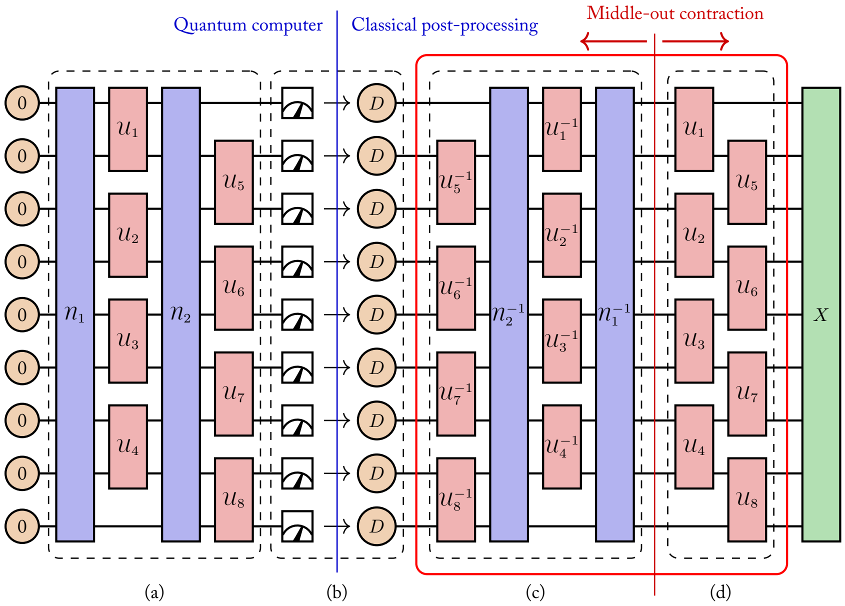

Figure 4.3. Schematic of the TEM algorithm. (a) Circuit on the quantum hardware (which also includes the gates used to create the initial state from $|0\rangle^{\otimes N}$), with noise channels $\mathcal{N}_l$ acting before the unitary layers $\mathcal{U}_l$. (b) Measurement with an IC POVM to whose outcomes we assign a dual operator D. (c) Inverse of the noisy circuit. (d) Ideal circuit. The TEM map, outlined in red, is constructed as a sequence of contractions and compression of the layers, from the middle out. In the end, the observable is applied to the error-mitigated data.

that is applying the TEM map to the individual MPS of the duals, computing the expectation value for each term separately and then taking the average at the very end.

Once we know the TEM map $\mathcal{M}$, we can also use it to understand how to optimise the measurement procedure, i.e. how to adjust our POVM in order to capture as much information as possible not on the final state in general, but on the data needed for estimating the observable of interest. Provided the compression errors are reasonably small, the average value of the TEM-modified observable $\mathcal{M}^\dagger(A)$ in the noisy state is the same as the noiseless estimation of the original observable A. As explained in Ref. [204], the main contribution in $\mathcal{M}^\dagger(A)$ is the rescaled original observable A, i.e. $\mathcal{M}^\dagger(A) \approx cA$ for some $c \geq 1$. In the so-called *typicality scenario*, the dynamics leads to each Pauli-string component of the observable branching into a vast number of strings, that is, if $A = \sum_{p \in \mathrm{P}^{\otimes n}} a_p p$ then $\mathcal{M}^\dagger(p)$ is a linear combination of many Pauli strings (with non-zero coefficient). In this scenario, the Pauli strings in A get damped by the noise by a factor of the same order of magnitude as a randomly chosen Pauli string. If A has a low Pauli weight, then the main contribution to the TEM-modified observable has the same low Pauli weight, meaning it can be efficiently estimated with high accuracy via qubit-wise IC measurements such as the POVM in Eq. (4.28). This means, moreover, that the remaining part of the TEM-modified observable, $\mathcal{M}^\dagger(A) - cA$, can be essentially neglected as it does not contribute significantly to the average value [204]. This is the reason why we choose a POVM which is skewed towards X on the target qubit.

In our specific case, once we obtain $\mathcal{M}^\dagger(X_n)$ through a TN calculation, we look at its decomposition into Pauli strings. We observe that discarding Pauli strings with weight greater than 2 allows us to reproduce the ideal and noisy signals in the DU regime. In non-DU circuits, the signal is simulable in the same way by keeping Pauli strings with a slightly greater Pauli weight, up to about 10. The non-typicality of the observable leads to high-Pauli-weight components in $\mathcal{M}^\dagger(X_n) - cX_n$ that cannot be neglected; however, the estimation of these contributions using Eq. (4.28) leads to a large variance, which is exponential in the Pauli weight [205]. In this case, more advanced and experimentally demanding measurement techniques [197, 205], that adapt more to the modified observable, may be employed.

4.3.2 Noise model

The (approximate) noise-cancelling map $\mathcal{M}$ relies on a representative model of the device noise: to this purpose we assume the noise follows the *sparse Pauli-Lindblad (SPL) model* [183]. In this model, a noise channel $\mathcal{N}$ is determined by a Lindblad generator

$$\mathcal{L} = \sum_{P \in \mathcal{K}} L_P, \qquad L_P(\rho) := \lambda_P (P \rho P^\dagger - \rho) \tag{4.38}$$

where $\lambda_p \geq 0$ and $\mathcal{K}$ is a subset of all n Pauli strings. Each individual L_P in the sum in Eq. (4.38) commutes with the others, so the channel $\mathcal{N} = \exp \mathcal{L}$ can be written as

$$\mathcal{N} = \bigcirc_{P \in \mathcal{K}} \exp L_P, \qquad (\exp L_P)(\rho) = \frac{\mathrm{e}^{-2\lambda_P} + 1}{2} \rho + \frac{\mathrm{e}^{-2\lambda_P} - 1}{2} P \rho P^\dagger. \tag{4.39}$$

The terms in $\mathcal{K}$ are chosen to reflect the noise interactions in the quantum processor, and their number, which determines the model complexity and expressiveness, is usually assumed to scale polynomially in n and therefore allows us to describe the noise of the full device by a *sparse* representation, given by the small set of non-negative coefficients λ_P.

Our quantum circuits consist of two unique entangling ECR layers, and we assume that the noise channel associated with the respective layer is identical whenever the layer appears in the circuit: for this reason we only need to characterise *two* noise layers. In particular, we choose $\mathcal{K}$ so that it contains Pauli strings which are non-trivial on single qubits or on consecutive pairs of qubits only, for a total of $3n + 9(n - 1)$ potentially non-zero parameters λ_P. We refer the reader to the Supplementary Material of Ref. [2] for a full account of the noise-learning protocol that allows us to estimate the λ_P parameters. Essentially, they are obtained by fitting the SPL model to fidelities obtained from cycle benchmarking circuits [206]. In addition, since this procedure does not allow us to learn all parameters independently, we calibrated them on the Clifford circuit representing the kicked Ising model at $h = 0$, $b = J = \frac{\pi}{4}$ as well.

Once the two sets of λ_P parameters, for each of the two ECR layers, are determined, we need to write them as MPOs in order to use them in the TEM algorithm, as well as in the MPS simulation. Since the noise channel consists essentially of Pauli strings, a natural basis for states and operators to write the TNs in is the PTM representation (see Section 1.4), where for one qubit we take as a basis set for the space of states the four Pauli matrices $\{I, X, Y, Z\}$, and for n qubits we simply take the tensor product of the single-qubit basis with itself n times. This basis is particularly convenient also because the noise channels become diagonal matrices, as illustrated in Ref. [58], thus both the TN form of the noise channel $\mathcal{N}$ and its inverse can be trivially written: the result is an MPO of bond dimension 4.

4.3.3 Performance of error mitigation

Applying an error mitigation technique, such as TEM, to the output of a quantum computation results in an associated computational cost, usually referred to as *sampling overhead*, which is a multiplicative factor in the number of circuit executions needed to get the same estimation accuracy as in the noiseless scenario. The sampling overhead is known to scale *exponentially* in both the circuit volume and the error rate [207, 208]. Despite the exponential scaling, error mitigation may still be feasible if the error rate is small enough. Ref. [204] reviews TEM alongside two other established mitigation strategies, zero-noise extrapolation (ZNE) [56] and probabilistic error cancellation (PEC) [183, 185]. These three techniques have in common their reliance on noise characterisation for every essential circuit layer: error mitigation then boils down to inverting the modelled noise. Calling ε the average error rate per qubit per layer and V the circuit volume, i.e. the number of qubits times the number of layers, Ref. [204] shows how, assuming an SPL noise model for all three techniques, the sampling overhead for PEC scales as $\exp(2\varepsilon V)$ while ZNE and TEM scale as $\exp(\varepsilon V)$.[3] This is explained by the fact that while ZNE and TEM are tightly bound to a specific observable which gets measured, PEC is not, therefore it is more general. Moreover, ZNE and TEM differ in a prefactor to $\exp(\varepsilon V)$, which scales polynomially in εV for ZNE and is constant for TEM. Assuming noise is weak, TEM actually achieves the universal lower cost bounds [207] for the estimation of Pauli observables in large-scale circuits.

4.4 Classical simulation

We benchmark the noisy and the noise-mitigated results against purely classical approximate simulations of the noisy and noiseless circuits, respectively, in both the Schrödinger and Heisenberg picture by using TN techniques. In order to perform a simulation which is as close as possible to what really happens on the quantum computer, we do not use the 'theoretical' gates of Eq. (4.10) but the actual sequence of native gates used by the device to perform the computation: we start from a $|0\rangle^{\otimes N}$ state and apply every gate exactly as the computer does, including the gates that prepare the initial state in Eq. (4.16). This is particularly necessary because the noise model is meant to be applied to the actual gates used by the device.

Each gate of the circuit, being a local single-qubit or two-qubit operator, has a straightforward representation as a tensor with either two or four legs, respectively. The gates can be then naturally grouped into layers according to the structure of the circuit: starting from the first executed gate we collect all gates that commute with it (and that can then be applied in parallel, at least theoretically) and create an MPO by multiplying all the gate tensors together, and we proceed this way until we arrive to the end of the circuit. Due the periodicity of the circuit, we need of course to do this only for the first two Floquet steps (an 'odd' and an 'even' step). We obtain five MPO layers for each Floquet step: some of them are trivial MPOs of bond dimension 1, when they are composed of single-qubit gates only, while the MPOs containing ECR gates have bond dimension 2. In the noisy simulation, we also need to insert the MPOs of the noise channels (as seen in Section 4.3.2) before the corresponding ECR layer. At each step the MPO-MPS multiplication is performed with the density matrix algorithm [74].

---

[3] In our experiment, we have seen how information in the DU case does not spread throughout the whole circuit, but is limited to a smaller region. Even in the non-DU case, the light cone of the observable is smaller than the whole circuit area roughly by half. In such cases where the causality-covered part C of the circuit is significantly smaller than the whole volume, $\varepsilon|C|$ should be used instead of εV.

We simulate the unitary evolution of an n-qubit pure state in the Schrödinger picture by writing the (pure) state in the computational basis $\{|0\rangle, |1\rangle\}$ of eigenstates of Z for each qubit. The initial state $|\psi_0\rangle$ of Eq. (4.16) has bond dimension 2. In this situation, in order to speed up the computation, we do not apply the gates one by one, but we compose them first into MPOs, as showed in Fig. 4.4, by grouping them around the ECR gates. This operation can be performed exactly, since the maximum bond dimension of the result is 2, due to the ECR gates alone: single-qubit gates do not affect the entanglement and can be grouped, arbitrarily, with the ECR preceding or following them, without affecting the bond dimensions of the resulting MPOs. Once we group the gates into MPOs, to advance the circuit by a Floquet step we need only apply two MPOs to the MPS of the current state: this is an efficient compromise between applying many small objects to the MPS (individual gates) and applying a few large ones (MPOs representing larger chunks of the circuit).

We simulate the noisy evolution alternating the layers of the ideal circuit and the noise channels: since in this configuration the states are matrices and circuit layers are channels, the state is vectorised into an MPS in the PTM basis as

$$\rho = \sum_{i_1,\dots,i_N} M^{(1,i_1)} \cdots M^{(N,i_N)} \sigma_{i_1}^{(1)} \otimes \cdots \otimes \sigma_{i_N}^{(N)} \tag{4.40}$$

where $(\sigma_0^{(k)}, \sigma_1^{(k)}, \sigma_2^{(k)}, \sigma_3^{(k)}) = (I_k, X_k, Y_k, Z_k)$ are the Pauli matrices on the k-th qubit; similarly, operators such as the noise channels or the gate operations $\rho \mapsto U\rho U^\dagger$ are transformed into MPOs acting on this MPS.

Here, since we find noise layers between unitary layers, we do not merge the unitary gates as before, but we stop at the noise layers, as in Fig. 4.4, obtaining a few more MPOs per time step. In the PTM representation the unitary layers, containing at most one ECR gate per qubit, have a maximum bond dimension of 4; the noise layers are MPOs of bond dimension 4, too. A priori, the latter may lead to a (potential) additional increase of the bond dimension during the simulation with respect to the ideal case. Throughout the simulations, both ideal and noisy, we set an upper threshold χ_{max} for the bond dimension: if during the MPO-MPS multiplication the bond dimension trespasses this limit, the result is truncated to χ_{max}.

We expect a great amount of entanglement to build up across the qubits, therefore the Schrödinger-picture simulations will likely be very expensive—if at all feasible—from a computational point of view. Because in the mixed-state simulation the local dimension of the MPS is 4 instead of 2 as for pure states, we will not even attempt a full-state simulation, in the Schrödinger picture, of the noisy device. In our case, we start from the observable X_n at the end of the circuit and, step by step, apply the adjoints of the layers until we reach the beginning, then we close the TN contraction with the

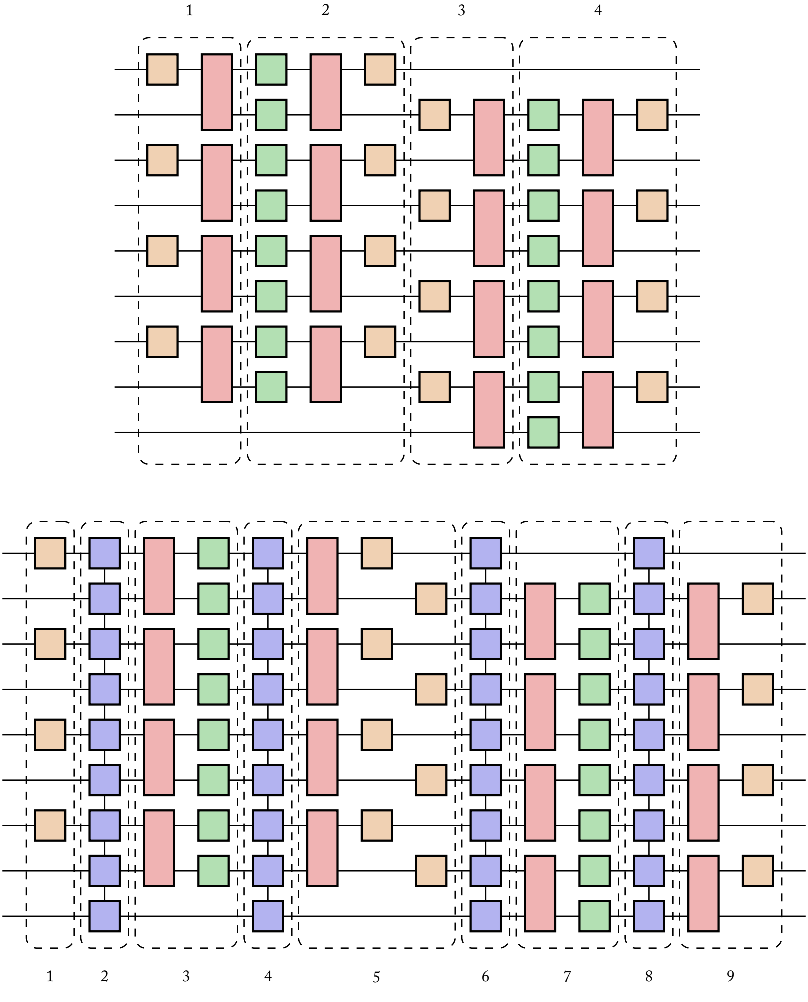

Figure 4.4. Structure of the layers within two consecutive Floquet steps, with different gate-grouping strategies for ideal (*top*) and noisy (*bottom*) circuits. Each block stands for the respective set of native gates obtained after transpiling the circuit.

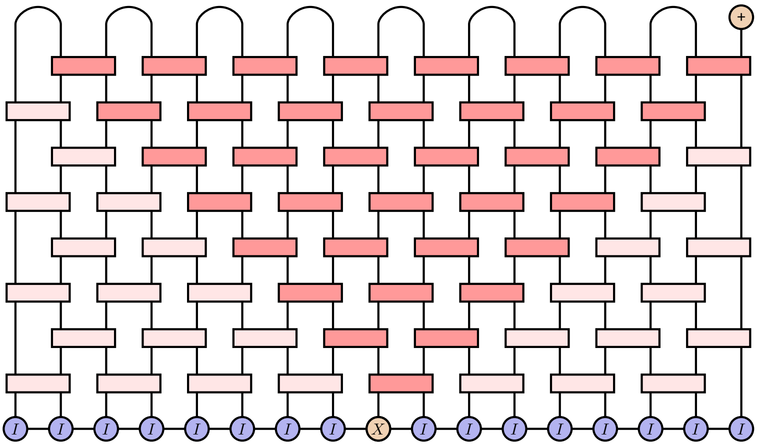

Figure 4.5. The brickwork structure of the unitary circuit and the locality of the observable lead to a significant simplification of the circuit when it is simulated in the Heisenberg picture: of the original circuit, only the gates in the 'light cone' (represented in a darker shade of red) act non-trivially on the MPS of the observable.

MPS of the initial state. By running the simulation in the Heisenberg picture, instead, we expect a significant decrease of the complexity, due to the particular form of the observable and the brickwork pattern of the entangling gates: in this picture, in fact, roughly half of the 'bricks' do nothing on the observable, as illustrated in Fig. 4.5. As identity terms are invariant under unitary gates and noise channels, the entanglement remains confined to the light cone of the evolution, whose extension grows linearly as the circuit layers are applied progressively.

This simplification happens, strictly speaking, only in the unitary case. The MPOs of the noise channels are a composition, as in Eq. (4.39), of two-qubit Pauli operators over many pairs of (adjacent) qubits, therefore in general we may expect a noise layer MPO to create correlations throughout the whole array of qubits; however, if noise is sufficiently small this will be a small perturbation and we can still treat the light-cone simplification of Fig. 4.5 as a first approximation.

4.5 Results

The performance of the quantum computer, together with the TEM correction, was benchmarked at the DU point $b = J = \pi/4$ with $h \in \{0.05, 0.1, 0.15\}$ and different system sizes of 51, 71 and 91 qubits: the results are illustrated in Figs. 4.7 to 4.9 at the end of this chapter.

The correlation function $C_n(t)$ was measured for various numbers of qubits n and time steps t. The observed result was negligible for $t \neq n$, and a finite signal was only measured along the light-cone boundary, i.e. for $t = n$. However, as a consequence of noise, the measured values on the light cone boundary decayed quicker with t than the exact evolution from Eq. (4.24).

The error-mitigated data show good agreement with the theoretical curve, albeit with some deviations. Even when integrability is broken for $h > 0$, the expected behaviour was closely recovered for all considered values of h, although small deviations appear in the mitigated results at larger system sizes. As the system size increases, the circuit depth also increases, and at these larger circuit volumes, errors in the noise model can accumulate, leading to a residual bias in the mitigated results [209]. These are likely a consequence of, for instance, imperfections in noise learning, incorrect model assumptions, or even the increased instability coming from the longer runtimes associated with larger circuit volumes [210]. With the TN simulations of the noisy circuit we can check whether the assumptions on the structure of the noise and its learning process were successful in defining an effective model, by comparing the simulated result with the actual raw data from the quantum computer, i.e. the measurement of the observable directly on the quantum device, evolved with the noisy circuit. The result for the DU circuits is shown in Figs. 4.13 to 4.15 at the end of this chapter. It should be noted that since the $h = 0$ circuits were used to fit the parameters of the noise model, the simulations agree perfectly in those cases with the raw data, and this matching should not be counted towards the ability of the model to describe the noise.

We can observe how the situations in which the error-mitigated data is most distant from the expected curve correspond to the case where the noisy simulation is farther from the raw data. This may happen from one of the reason listed before for the failures of the noise model, or because of an unexpected dependence of the model from the h parameter. Another important factor is the fact the actual noise on the physical device may not be constant in time, whereas in the sparse Pauli-Lindblad model we have no time dependence: as the circuits get executed one after another, the noise can change and the learned parameters might not capture its current effect anymore. In particular, the choice of performing additional noise learning on the $h = 0$ circuits only is inevitably a trade-off between having a model that adheres better to the individual circuits, but has to be re-calibrated each time, and a less detailed model that captures only the most relevant features, but requires less effort to be defined. In the end, even with an unlimited bond dimension, the accuracy of the TEM method is generally limited by the accuracy of the quantum component and its noise model. In [2, Supplemental Information, Sec. III.G] we performed a systematic analysis of the effect that the discrepancies in the noise model had on the error mitigation. We confronted, for the experiments at the DU point, the relative error between the unmitigated outcomes and the simulated results with the relative error between the mitigated outcomes and the exact theoretical prediction. We concluded that the relative error post-mitigation is consistent with the relative error between the noisy experiment and the expected noisy outcomes. In this sense, TEM performed as well as could be expected within the margin of error relative to the accuracy of the noise models used.

For the quantum component of the workflow, the traces of data shown in Figs. 4.7 to 4.12, including the learning of the noise model and the error-mitigation in post-processing, took a wall-clock time of 2 hours and 18 minutes for $N = 51$, 2 hours and 55 minutes for $N = 71$ and 3 hours and 24 minutes for $N = 91$. As the number of gates in the temporal light cone affecting the noisy signal (650 ECR gates for $N = 51$, 1260 for $N = 71$, and 2070 for $N = 91$) increases, the signal-to-noise ratio decreases, and therefore the statistical errors in the mitigated values get larger. This is especially reflected in the size of the error bars in Fig. 4.12, where each data point is estimated from 262 144 individual shots. A complete account of the measurement statistics can be found in the Supplementary Material of [2].

After assessing the accuracy of the mitigated results for analytically solvable DU circuits, we perturb away from the DU points. While working with the same initial state and observable as before, the local expectations values $\langle X_n(t) \rangle$ lose their interpretation as auto-correlation functions when away from the DU points. In Figs. 4.10 to 4.12, at each of the previously considered values of h, we report the change of $\langle X_n(t) \rangle$ for $n = t = (N - 1)/2$ as we perturb the transverse field b away from dual unitarity. Figures 4.16 to 4.18 compare instead the noisy simulations to the raw data; the same considerations made for the DU case hold here, too.

In the absence of exact analytical solutions and at a scale beyond brute-force state-vector simulations, these error-mitigated data can only be compared to approximate classical methods. We compare therefore the results to TN simulations as introduced in Section 4.4. Across the different parameters, the experimental data show strong agreement with the Heisenberg-picture simulations with some deviations arising at larger circuit volumes, but large disagreements with the Schrödinger-picture simulations. Heisenberg-picture simulations display a higher convergence rate than the Schrödinger-picture ones; therefore, in Figs. 4.10 to 4.12 we employ a bond dimension $\chi = 600$ for the former and $\chi = 1500$ for the latter. While dynamics in the Heisenberg picture seem to be converging on classical computers, simulations in the Schrödinger picture with the same convergence rate become unaffordable at the scale of this experiment. In the quantum-classical workflow, instead, in order to produce the error-mitigated data points for $N = 91$ the middle-out contraction for TEM employs a bond dimension of $\chi = 70$. Next, we analyse the convergence of the TN simulations in more depth.

4.5.1 Convergence of tensor-network classical simulations

We assess the reliability of the classical simulations by running each TN simulation with a progressively increasing bond dimension χ of the MPS representing either the pure state or the observable. As χ is gradually increased from run to run, we observe that the difference between the simulated expectation values from subsequent runs decreases, and the convergence is faster in the Heisenberg-picture simulations than in the Schrödinger-picture ones, as shown at the end of this chapter in Fig. 4.19 (a) and (b). In fact, noiseless simulations in the Heisenberg picture show that the absolute difference between the results with bond dimension χ and $\chi + 100$ does not change more than 10^{-2} for $100 \leq \chi \leq 500$ and no more than 10^{-3} for $500 \leq \chi \leq 900$, see Fig. 4.19 (a). We reached a bond dimension of 900 in the 51- and 71-qubit circuits, while in the 91-qubit case we stopped at 600 due to the expensive memory requirements.

As another metric for convergence of the TN simulations, we use the estimated absolute error proposed in Ref. [211], defined by

$$\varDelta_\chi = |\langle X_n \rangle_\chi - \langle X_n \rangle_{\chi \to \infty}|, \tag{4.41}$$

where $\langle X_n \rangle_\chi$ is the average value of the observable in the TN simulation with bond dimension χ, and $\langle X_n \rangle_{\chi \to \infty}$ is obtained by a linear extrapolation of the data as a function of $1/\chi$. It is important to stress that this is a *heuristic* estimate: there is no guarantee that the extrapolated value is actually the value we would get with an unlimited bond dimension, as we do not have any rigorous bounds on the approximation error caused by limiting the bond dimension. However, in conjunction with other types of analyses, it can give us a hint on whether the simulation seems to converge. In practice, in each circuit we identified a specific bond dimension $\bar{\chi}$ such that $\langle X_n \rangle$ increases approximately linearly as a function of χ^{-1} for all $\chi \geq \bar{\chi}$. This happened for all (h, b) points at approximately $\bar{\chi} = 500$, sometimes even at lower bond dimensions. We then fit a straight line to the (at least) three points, i.e. the data with the three greatest bond dimensions. For the circuits under study, we obtain $\varDelta_\chi \sim 10^{-2}$ for $100 \leq \chi \leq 400$ and $\varDelta_\chi \sim 10^{-3}$ for $500 \leq \chi \leq 900$. We illustrate the approach, for some of the circuits under study, in Fig. 4.20: importantly, the results of simulations in the Heisenberg and Schrödinger pictures do not generally agree with each other even within the heuristically determined error bounds, despite the fact that the extrapolations with respect to χ^{-1} seem to be reasonable in both pictures.

For both Schrödinger- and Heisenberg-picture simulations we also monitor how the entanglement entropy of the MPS changes as we increase the bond dimensions, by computing the maximum of Eq. (1.12) over all bipartitions of the MPS. For an MPS with bond dimensions that does not exceed χ, the entanglement entropy is bounded from above by $\log_2 \chi$. Hence, we can safely say that the truncated bond dimension does not significantly affect the results if we see a plateau well below this value. As already pointed out in Section 1.3, this is actually the entanglement of the bipartition of the physical state only when we are considering the ideal simulation in the Schrödinger picture, which is the only case when the MPS represents a pure state; in the other cases, it measures both quantum and classical correlations within the TN, if it represents a mixed state, or the operator-space entanglement entropy, if it represents an operator. In any case, it is a key quantity we want to observe in order to estimate the complexity of the TN calculation.

In Fig. 4.19 (c) we show the maximum of the entanglement entropy among all values of b and h for the 51-qubit circuit; the entropy for the 71- and 91-qubit circuits follows the same trend. We can see that the simulations in the Schrödinger picture achieve the worst possible scaling, hitting the upper bound given by the bond dimension: this is due to the highly-entangling nature of the circuit. The Heisenberg-picture simulations, instead, show an overall lower entanglement entropy. As further confirmation that the Schrödinger-picture simulations are still far from convergence, we show in Fig. 4.21 how after even a few Floquet steps, except in the most trivial case $h = 0$, the simulated state fails to give an expected value $\langle X_t(t) \rangle$ close to the theoretical result.

The particular choice of the observable has also a great role in determining the efficiency of the Heisenberg-picture simulations. In the $b = \frac{\pi}{4}$, $h = 0$ case we have a Clifford circuit, in which case the observable is mapped in the end to a Pauli string which is not trivial on a single qubit only, the $ZIII\cdots$ string, whose entropy is zero. Even when we deviate from these values, by changing b or h or by adding noise, the distribution of the coefficients of each component of the observable remains in all cases peaked around the dominant contribution $ZIII\cdots$; during its evolution the observable picks up other terms that are not trivial on more than one qubit, but their contribution is still not relevant enough to increase the entropy beyond what we can manage with a moderately-sized MPS. Still, the entanglement entropy of both noisy and ideal simulations in the Heisenberg picture grows logarithmically with the bond dimension. This does not contradict the convergence shown in Fig. 4.19 (a) and (b) because we are computing the expectation value of a specific observable in a specific initial state, namely, $|0\rangle\langle 0|^{\otimes N}$. This initial state, in the PTM representation, can be written only with I and Z components on each qubit, so it captures only part of the evolved observable. To take this into account, we recompute the singular values of the MPS of the final observable after projecting it onto these two components (thus eliminating the contributions containing X or Y, which are orthogonal to the initial

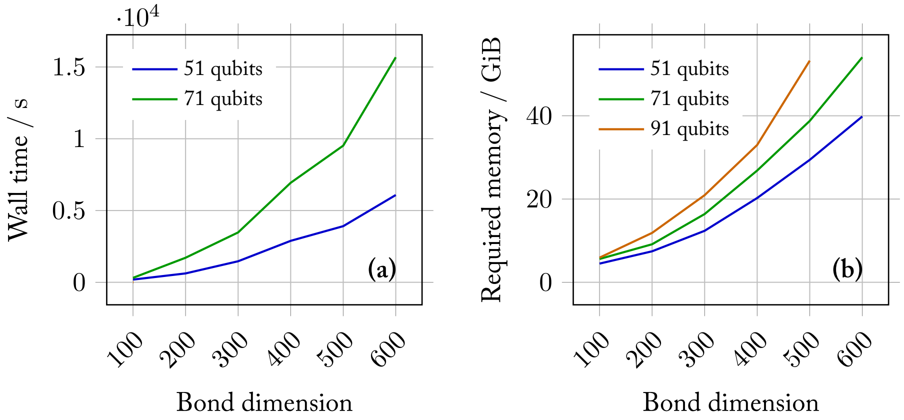

Figure 4.6. Time and memory requirements for Heisenberg-picture TN simulations of the kicked Ising model. (a) Elapsed time in the simulation of the $b = \pi/4 + 0.15$ and $h = 0.15$ ideal circuit, with 16 CPUs and 72 GiB of memory on the HPC cluster Leonardo, with different bond dimensions. (b) Total memory required for the same simulation (without constraints).

state). In this case not only is the entanglement entropy much lower than in the previous cases, but it is also basically constant for all considered bond dimensions in the range from 100 to 900, which substantiates the convergence analysis.

4.5.2 Classical resources

We conclude the analysis with some comments on the resources used for the simulations. The TEM post-processing algorithm was implemented in Python with the Quimb [212] TN library, while the classical TN simulations were implemented in Julia with the ITensor [159] library. Simulations were run on the Karolina, Leonardo high-performance computing (HPC) clusters and on Microsoft Azure cloud virtual machines.

The most expensive part of the simulation is the MPO-MPS multiplication at each layer of the circuit, with time complexity $O(Nd(\chi\chi')^3)$ where N is the number of qubits, d is the local dimension and χ and χ' are the bond dimensions of the MPS and the MPO, respectively. In the Schrödinger-picture simulations, the local dimension is 2 and the MPO encoding the 2-qubit layer has bond dimension 2, whereas in the Heisenberg-picture simulations these numbers become 4 and 4. This means that a

Schrödinger-picture with MPS bond dimension χ_{S} and a Heisenberg-picture simulation with χ_{H} have the same time complexity if $\chi_{\mathrm{S}} = 2^{\frac{4}{3}}\chi_{\mathrm{H}}$. For example, $\chi_{\mathrm{S}} = 1500$ is roughly equivalent to $\chi_{\mathrm{H}} = 600$. Since the circuit of the kicked Ising model creates less entanglement on the observable side than on the state side, it is not surprising that the Schrödinger-picture simulations are still very far from converging even at bond dimension $\chi_{\mathrm{S}} = 1500$, while the Heisenberg-picture simulations with bond dimension $\chi_{\mathrm{H}} = 600$ are already sufficient to obtain reliable results.

Figure 4.6 shows how much wall-clock time the ideal (noiseless) evolution of the observable in the Heisenberg picture—for a specific choice of b and h associated to one of the most expensive simulations—took with fixed number of CPU cores and amount of memory, as well as the overall memory required in an unconstrained setting (no wall-clock time data is available for the 91-qubit circuit). As the plot shows, a 91-qubit simulation with bond dimension 500 already required more than 50 GiB; increasing the bond dimension to 600 would bring the required memory to roughly 90 GiB. We note that, contrary to the statements above, the measured execution time is not proportional to the number of qubits, and its increase does not scale as χ^3, but rather as χ^2. The explanation for both of these facts is that, in the Heisenberg picture, the MPS is almost never at 'full capacity': the information, in fact, spreads linearly along a line from the central qubit, where the observable X_n is supported, with the edges of the circuit being reached only at the very last step in the Heisenberg evolution.

4.5.3 Summary and final comments

The Schrödinger-picture simulations are far from convergence, except in the most trivial cases; our analysis indicates that the MPS representation of the state at the chosen bond dimension, despite the great amount of computational resources required, is still not adequate. The overall low entanglement entropy of the Heisenberg-picture simulations suggests, instead, that those simulations achieve a satisfactory degree of convergence, especially at lower sizes, which is facilitated by the more favourable scaling of the complexity of the calculation with respect to the bond dimension.

For $N = 51$ qubits, in both the non-DU and the DU regimes, the quantum error-mitigated results match the converged Heisenberg-picture simulations closely, highlighting that the TEM method is able to recover unbiased estimators at non-trivial scales. Even at the larger sizes of 71 and 91 qubits we observe that TEM is able to approximately reconstruct the shape of the curves, although with larger discrepancies.

Incorporating the noise model in the classical simulation, moreover, allows us to assess its validity and to understand whether the difference between the expected and the error-mitigated results can be linked to the model failing to describe the actual physical noise. The limitations of the sparse Pauli-Lindblad noise model are known;

still, in our experiment, it provided a good description of the errors in most cases, especially given that it was calibrated only once for all the circuit runs at each different size N. As the engineering of quantum computing platform keeps improving, and the physical noise decreases, we expect that we will need finer, possibly non-Markovian noise models in order to accurately capture the influence of the environment.

These results emphasise the progress of error-mitigated quantum computing in becoming increasingly competitive with widely used classical algorithms, in regimes where we cannot rely on a brute-force exact solution. By considering different choices of observables of the form $P_k P_l$, i.e. product of pairs of Pauli operators, we could study e.g. two-body correlators, magnetisation imbalance (that is, non-uniformity), and the two-body mutual information with the same experimental setup. The auto-correlation functions we studied, in addition, could be used to detect the existence of a ballistic transport regime, where the spin current in the system grows linearly with time [213] and predict the existence of localised phases [214], where thermalisation is extremely slow but not completely arrested, so that the system appears to be localised for a long time.

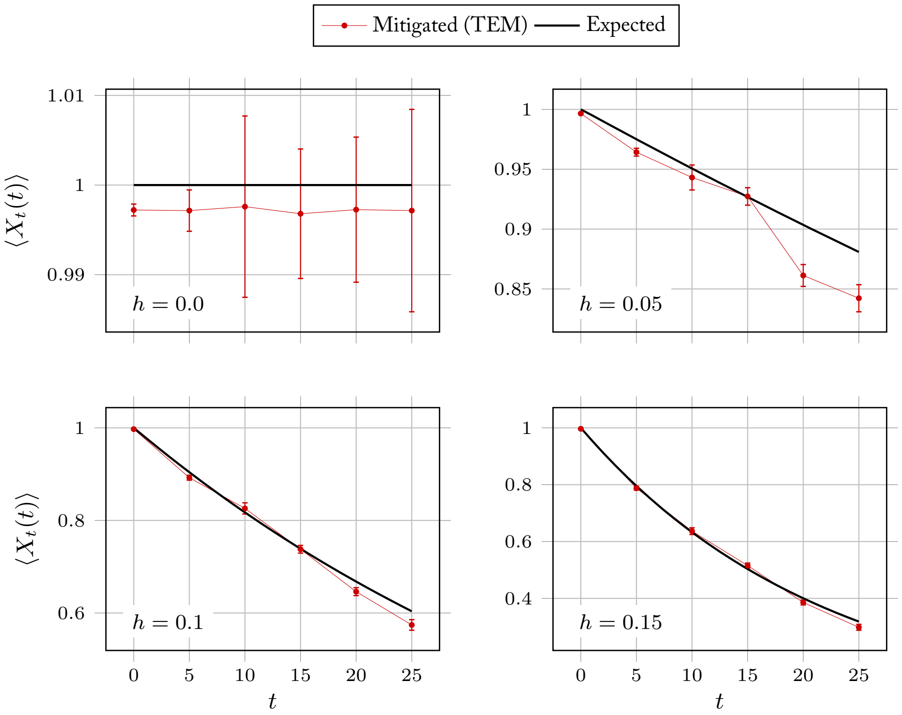

Figure 4.7. Comparison of the error-mitigated data, with error bars indicating one standard error, to the theoretical values at the DU point for $N = 51$ qubits. The measured observable $\langle X_t(t) \rangle$, plotted as a function of the number of circuit steps t, corresponds to the autocorrelation function $C_t(t)$ for increasing values of h and the experimentally probed time steps t, for $b = J = \pi/4$.

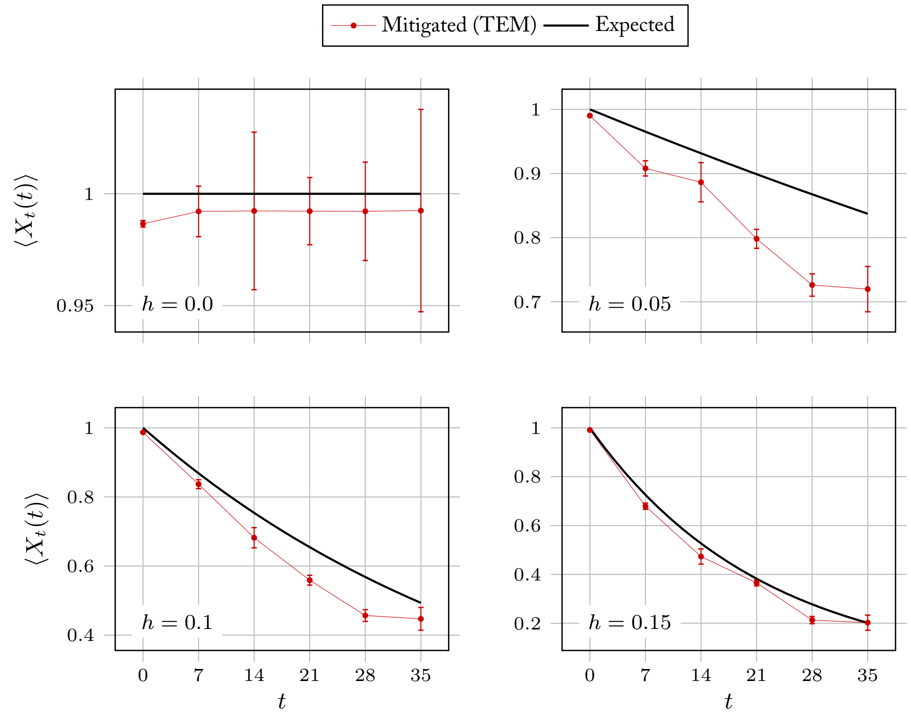

Figure 4.8. Comparison of the error-mitigated data, with error bars indicating one standard error, to the theoretical values at the DU point for $N = 71$ qubits. The measured observable $\langle X_t(t) \rangle$, plotted as a function of the number of circuit steps t, corresponds to the autocorrelation function $C_t(t)$ for increasing values of h and the experimentally probed time steps t, for $b = J = \pi/4$.

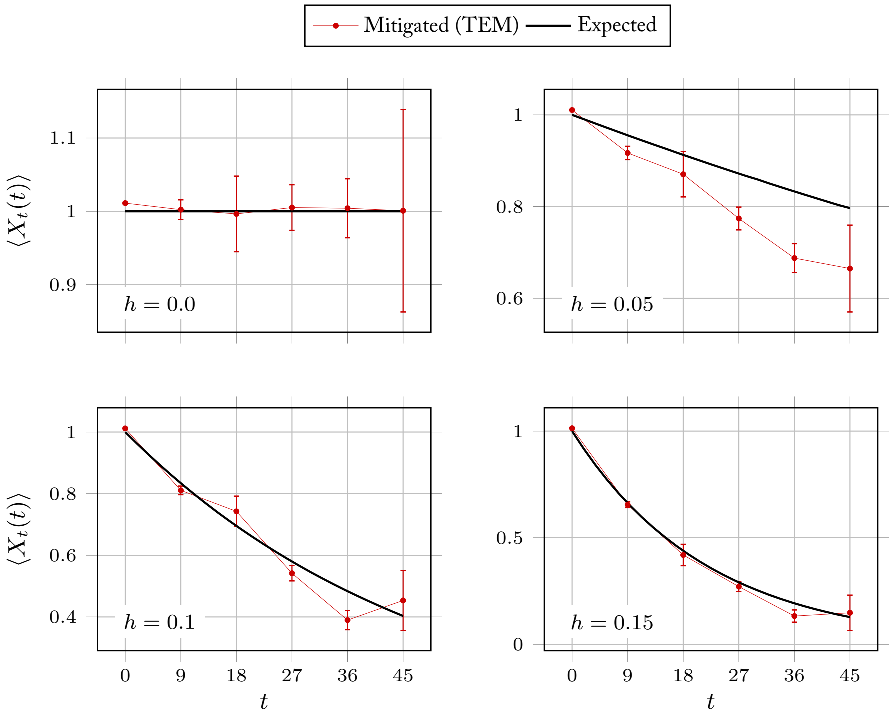

Figure 4.9. Comparison of the error-mitigated data, with error bars indicating one standard error, to the theoretical values at the DU point for $N = 91$ qubits. The measured observable $\langle X_t(t)\rangle$, plotted as a function of the number of circuit steps t, corresponds to the autocorrelation function $C_t(t)$ for increasing values of h and the experimentally probed time steps t, for $b = J = \pi/4$.

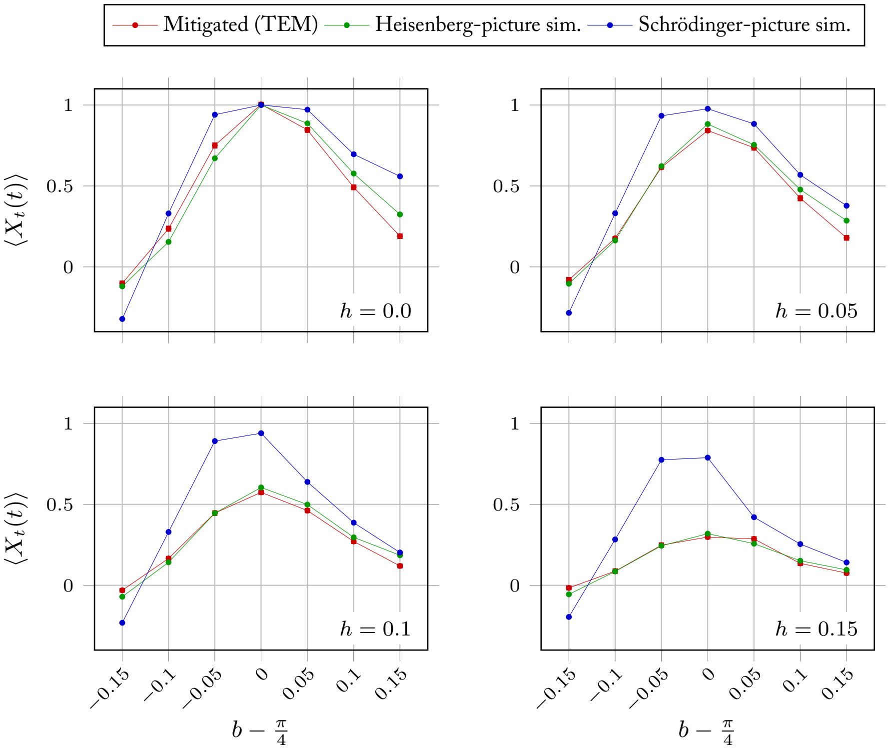

Figure 4.10. Non-dual-unitary circuits beyond exact classical verification at $N = 51$ qubits. Each plot shows the evolution of $\langle X_t(t) \rangle$, with $t = (N-1)/2$, as the transverse field b is tuned away from dual unitarity, for different values of h. Error bars indicate one standard error. The DU points at $b = \pi/4$ correspond to the rightmost points in Fig. 4.7. The results are compared against classical TN simulations in the Schrödinger ($\chi = 1500$) and the Heisenberg ($\chi = 600$) pictures.

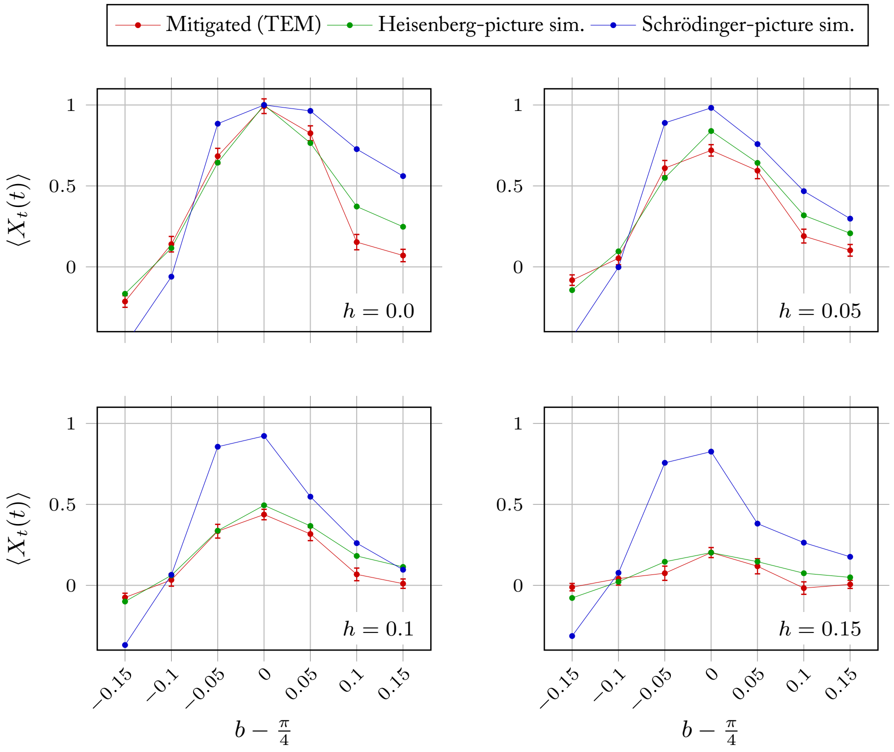

Figure 4.11. Non-dual-unitary circuits beyond exact classical verification at $N = 71$ qubits. Each plot shows the evolution of $\langle X_t(t) \rangle$, with $t = (N-1)/2$, as the transverse field b is tuned away from dual unitarity, for different values of h. Error bars indicate one standard error. The DU points at $b = \pi/4$ correspond to the rightmost points in Fig. 4.8. The results are compared against classical TN simulations in the Schrödinger ($\chi = 1500$) and the Heisenberg ($\chi = 600$) pictures.

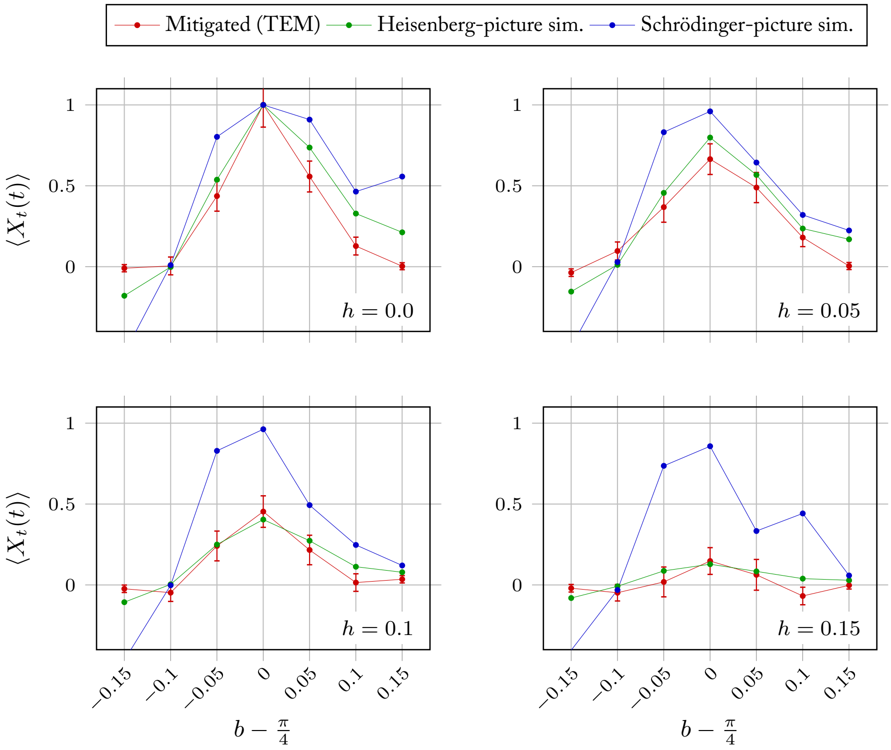

Figure 4.12. Non-dual-unitary circuits beyond exact classical verification at $N = 91$ qubits. Each plot shows the evolution of $\langle X_t(t) \rangle$, with $t = (N-1)/2$, as the transverse field b is tuned away from dual unitarity, for different values of h. Error bars indicate one standard error. The DU points at $b = \pi/4$ correspond to the rightmost points in Fig. 4.9. The results are compared against classical TN simulations in the Schrödinger ($\chi = 1500$) and the Heisenberg ($\chi = 600$) pictures.

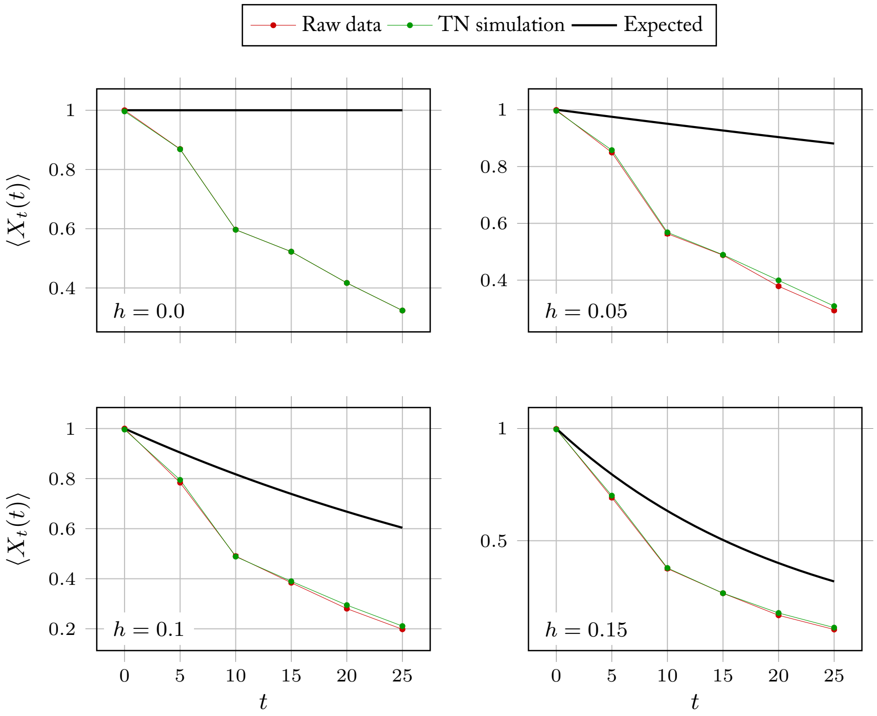

Figure 4.13. TN simulations of the noisy circuit at the DU point for $N = 51$ qubits. The simulated noisy expectation value $\langle X_t(t)\rangle$, plotted as a function of the number of circuit steps t, is compared to the raw data obtained from the quantum device, used for the error-mitigated data of Fig. 4.7. Error bars indicate one standard error (although they cannot be noticed as they are are smaller than the symbol size).

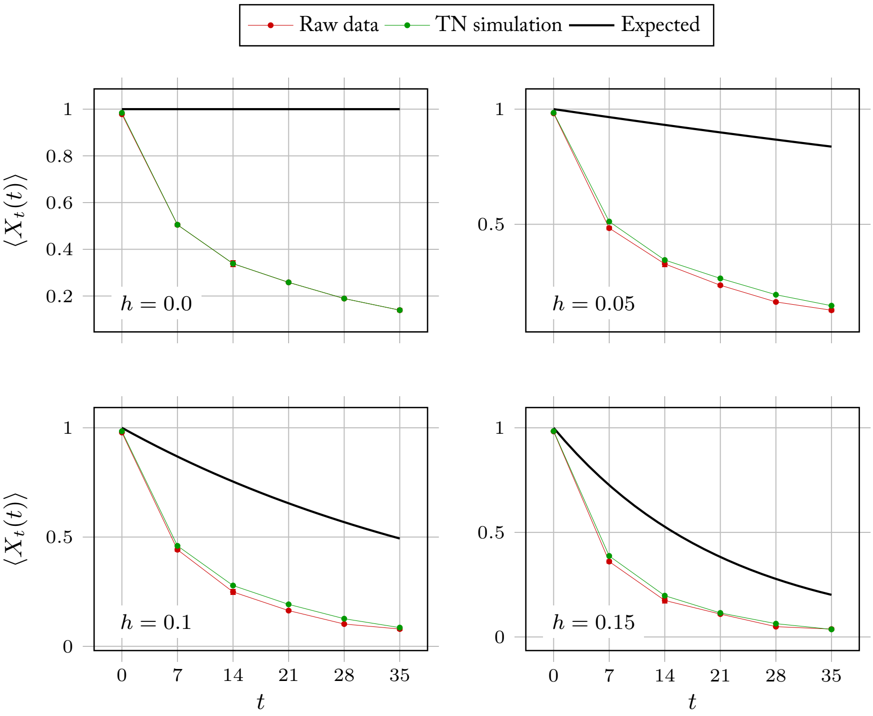

Figure 4.14. TN simulations of the noisy circuit at the DU point for $N = 71$ qubits. The simulated noisy expectation value $\langle X_t(t) \rangle$, plotted as a function of the number of circuit steps t, is compared to the raw data obtained from the quantum device, used for the error-mitigated data of Fig. 4.8. Error bars indicate one standard error (although they cannot be noticed as they are are smaller than the symbol size).

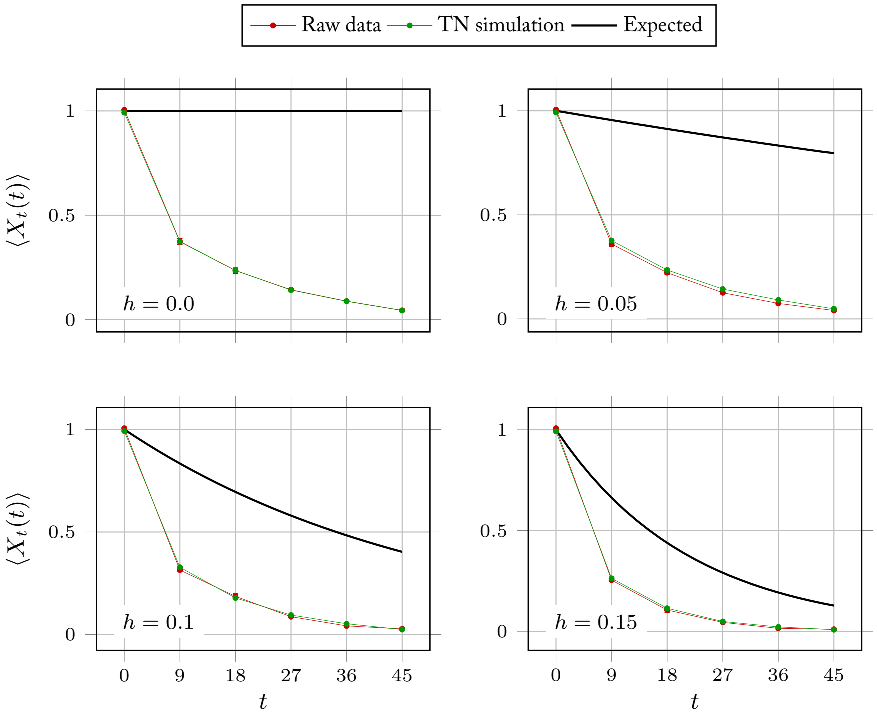

Figure 4.15. TN simulations of the noisy circuit at the DU point for $N = 91$ qubits. The simulated noisy expectation value $\langle X_t(t) \rangle$, plotted as a function of the number of circuit steps t, is compared to the raw data obtained from the quantum device, used for the error-mitigated data of Fig. 4.9. Error bars indicate one standard error (although they cannot be noticed as they are are smaller than the symbol size).

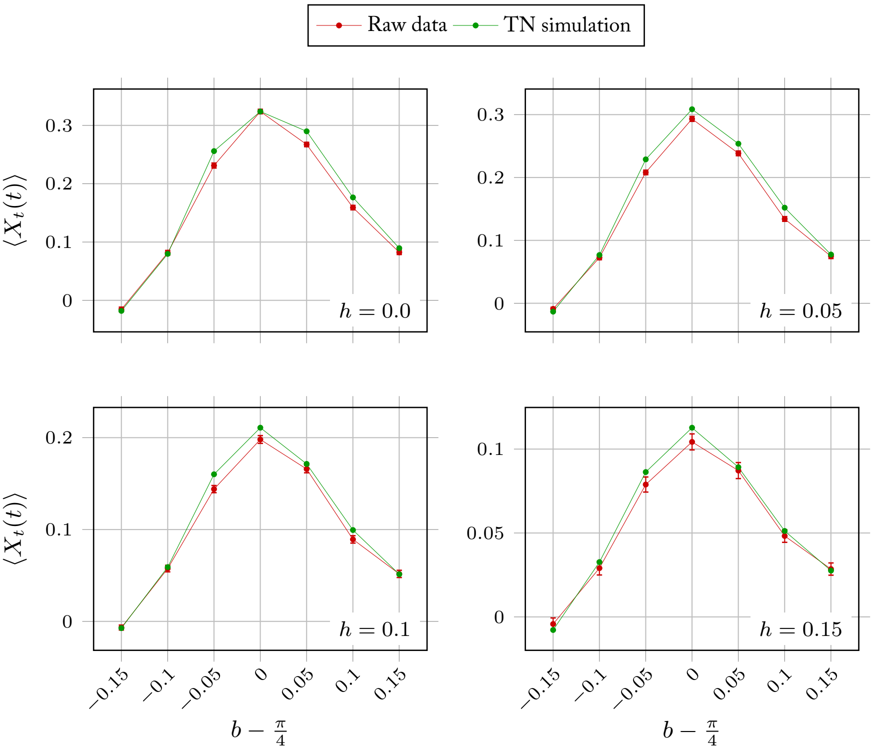

Figure 4.16. TN simulations of the noisy circuit beyond the DU point for $N = 51$ qubits The simulated noisy expectation value $\langle X_t(t) \rangle$, plotted as a function of the parameter b, with $t = (N-1)/2$, is compared to the raw data obtained from the quantum device, used for the error-mitigated data of Fig. 4.10. Error bars indicate one standard error.

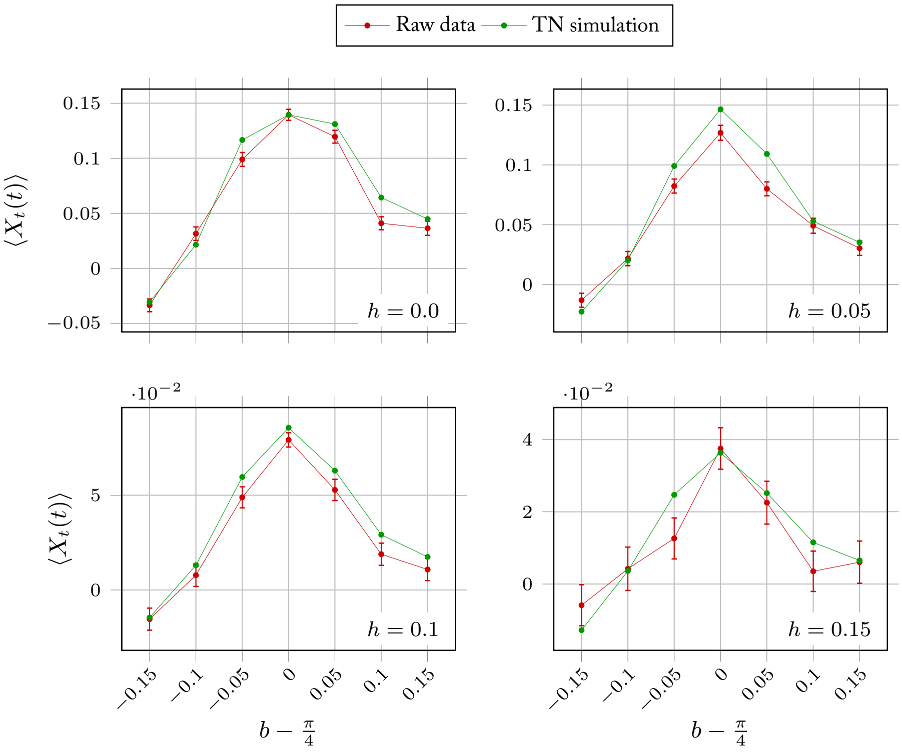

Figure 4.17. TN simulations of the noisy circuit beyond the DU point for $N = 71$ qubits The simulated noisy expectation value $\langle X_t(t) \rangle$, plotted as a function of the parameter b, with $t = (N-1)/2$, is compared to the raw data obtained from the quantum device, used for the error-mitigated data of Fig. 4.11. Error bars indicate one standard error.

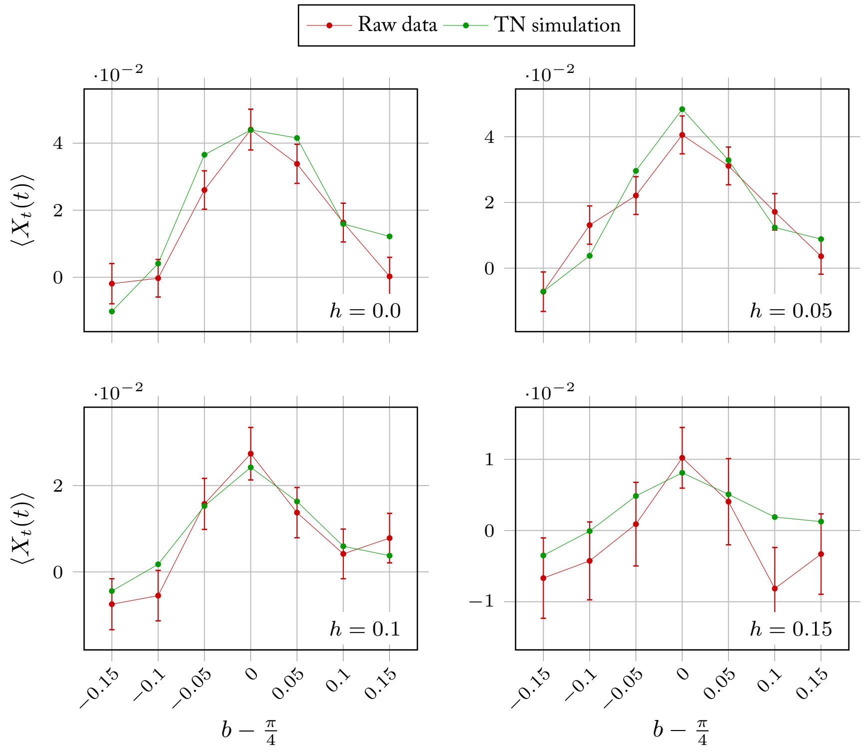

Figure 4.18. TN simulations of the noisy circuit beyond the DU point for $N = 91$ qubits The simulated noisy expectation value $\langle X_t(t) \rangle$, plotted as a function of the parameter b, with $t = (N-1)/2$, is compared to the raw data obtained from the quantum device, used for the error-mitigated data of Fig. 4.12. Error bars indicate one standard error.

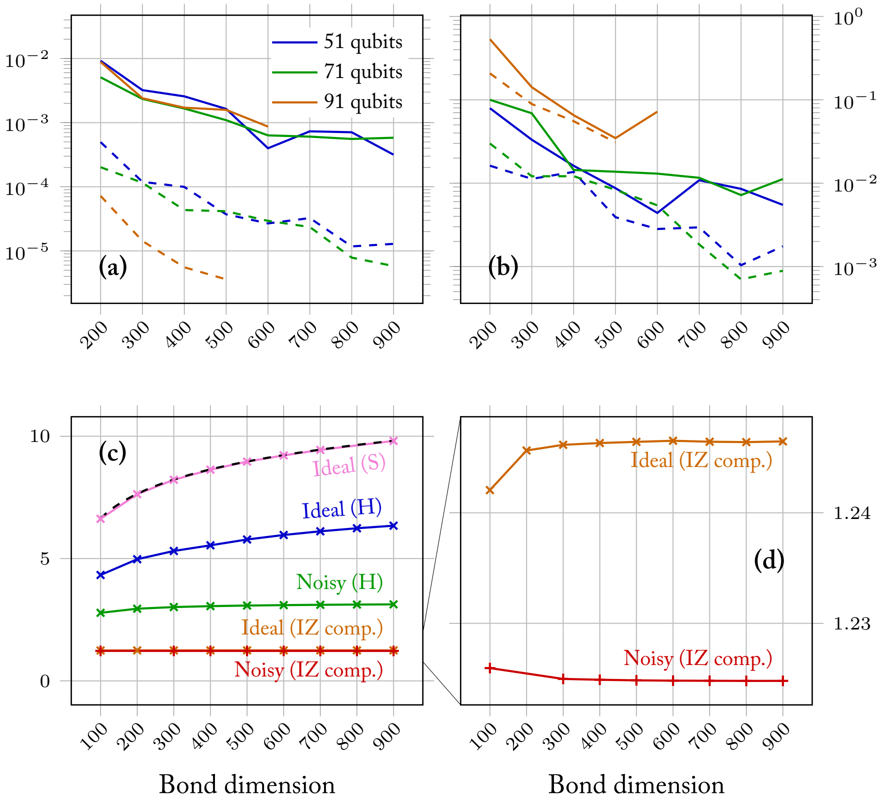

Figure 4.19. Convergence analysis of TN simulations of the kicked Ising model. (a) Maximum of the absolute difference, over all values of b and h, between Heisenberg-picture simulations of $\langle X_n \rangle$ obtained with bond dimension χ and $\chi - 100$. Solid lines represent ideal simulations, while dashed lines represent noisy ones. (b) Same as (a) but relative to the value of $\langle X_n \rangle$. (c) Maximum entanglement entropy, over all values of b and h, of the final MPS in the ideal and noisy simulations of the 51-qubit circuit (H: Heisenberg picture, S: Schrödinger picture). The dashed line indicates the upper bound set by the bond dimension. (d) Focus on the entanglement entropy of the final states after the projection on I and Z components.

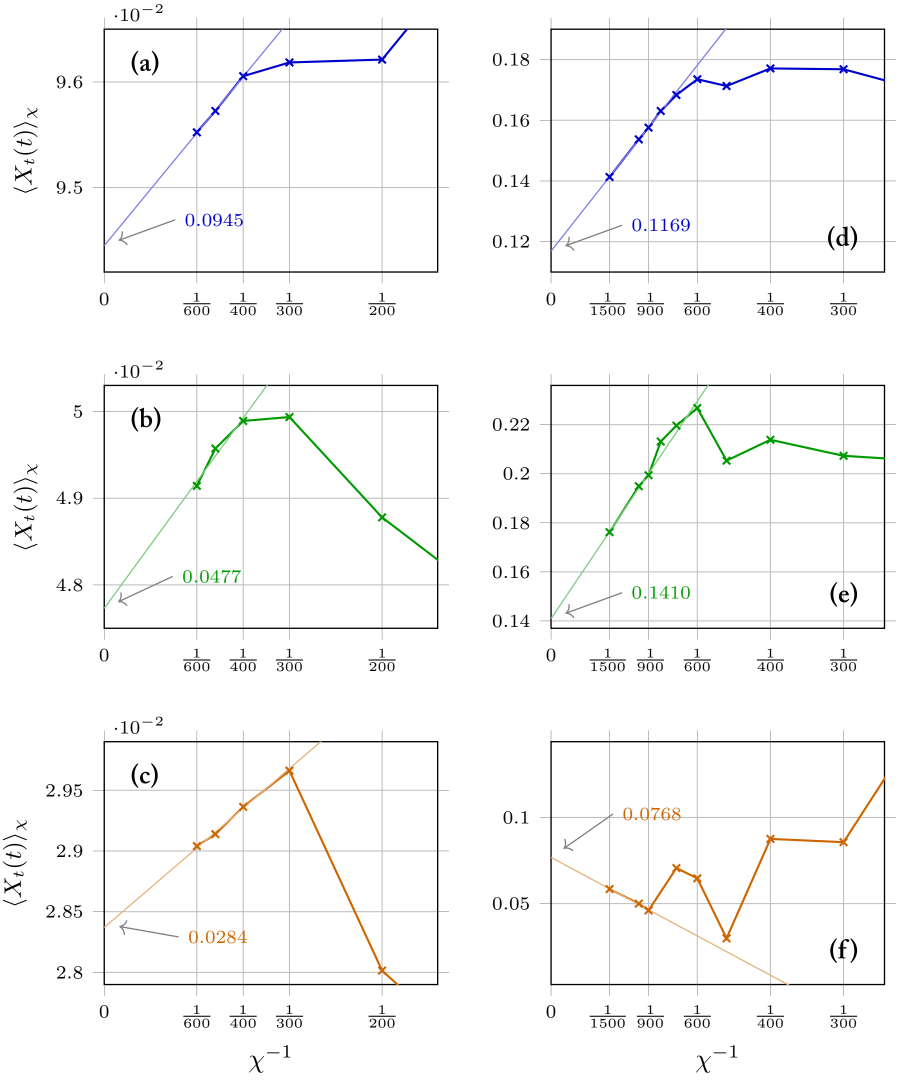

Figure 4.20. Results of TN simulations of the kicked Ising model in the Heisenberg picture (a-c) and Schrödinger picture (d-f) in terms of the inverse of the bond dimension used. Illustrative examples for circuits with $h = 0.15$, $b = \pi/4 + 0.15$ at different circuit sizes: 51 (a, d), 71 (b, e), and 91 (c, e) qubits.

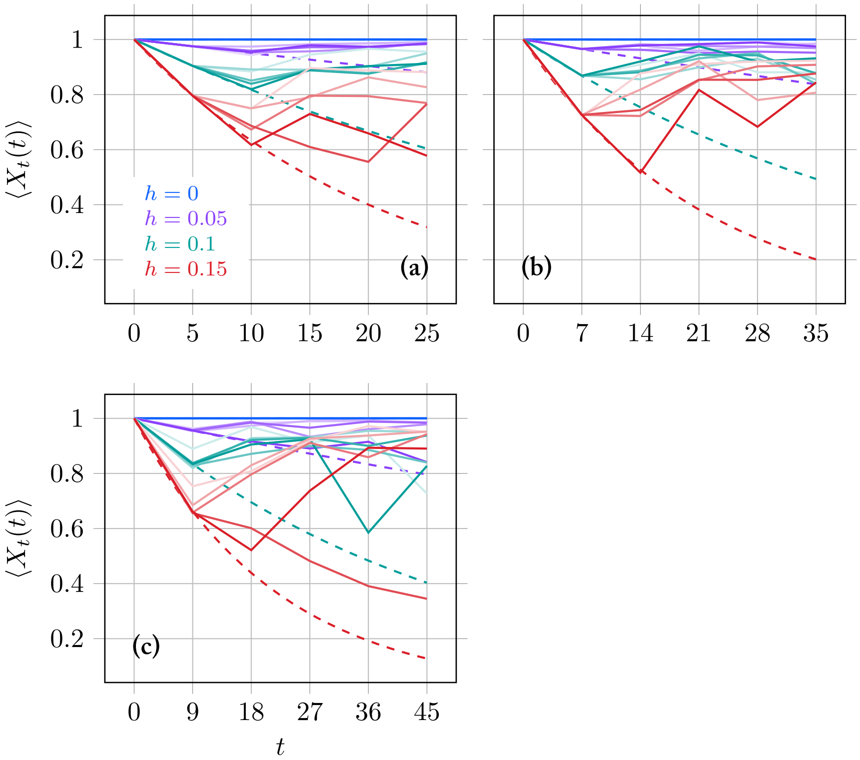

Figure 4.21. (Non-)convergence of Schrödinger-picture simulations of the kicked Ising model at the DU points, shown together with the exact theoretical values (dashed) for 51 (a), 71 (b), and 91 (c) qubits. The opacity of the curves increases together with the bond dimension, from 200 to 1000.

Conclusion

Final remarks

In this PhD thesis we have employed TNs in two challenging realistic settings, where their ability to represent quantum many-body systems has been exploited to simulate the complex behaviour of open systems. We started by reviewing in Chapter 1 the basic theory of MPSs, why and how they are useful for many-body systems. We also described two widespread algorithms for computing the time evolution of MPSs, and, finally, we discussed two ways of using the MPS formalism to represent mixed states. In Chapter 2 we presented the fundamental concepts of the theory of OQSs, before turning to some theoretical and practical techniques for transforming and reshaping the time-evolution problem of OQSs in order to reduce its complexity, with special emphasis on TNs. In our first main example in Chapter 3 we developed the FMC approximation scheme, and implemented it as a TN, in order to numerically simulate the full, non-Markovian dynamics of a fermionic OQS in a generic fermionic environment. In the second main example, presented in Chapter 4, we used TNs to benchmark, instead, a large NISQ device, both as a fully classical simulation tool and as part of a hybrid quantum-classical approach in which TNs formed the basis for an error-mitigation algorithm. These two examples showcase the versatility of TNs and their usefulness in various contexts, from time evolution of quantum systems to data processing in quantum computers. Let us review them in more detail.

With the FMC construct we introduced a way to improve the numerical simulation with TNs of open systems coupled to continuous fermionic environments. First, we explored a method—closely related to the thermofield transformation—to reshape the environments through the TTCF-based equivalence. We were able to absorb the factors depending on chemical potential and temperature into modified spectral densities, which made it possible to merge several environments into one or two effective ones. This simplification leads to a reduction of the number of chains to be determined by means of the TEDOPA chain mapping, that would otherwise require one chain per environment. In principle, this simplification could benefit other simulation schemes as well. Second, after performing the chain mapping on the environment in order to transform it into a one-dimensional system with nearest-neighbour interactions, we developed a technique to approximate the dissipative behaviour of the environments using only a finite number of sites. This construction extends the existing MC construction for bosonic environments to fermionic environments, retaining most of its properties, in particular the *universality* of the closure. The technique, in fact, approximates the behaviour of continuous environments in a way that does not rely on a specific form of the spectral density, but only on its asymptotic behaviour, which is common to all spectral densities in the Szegő class, and does not need ad-hoc discretisation schemes.

In order to test the correctness of the FMC, after rewriting the master equation in the superfermion formalism, we considered Gaussian OQSs, for which we were able to compute the time evolution of expectation values and correlation functions in a numerically exact way. We found that the results aligned very well with the ones obtained from an analogous numerical solution of the environment discretised with TEDOPA.

We then tested the FMC in the TN form as well, analysing several technical steps that must be taken so as to obtain an efficient representation, such as interleaving the two environments and leveraging the number-conservation property of the closure to use sparse tensors in calculations. We compared the TN form of the FMC with a standard TEDOPA simulation (with TNs) and, if available, with the Gaussian solution on several systems and environments, and we found in all situations a satisfactory match between the expectation values of operators measured in the two cases, with relative errors below 10^{-3}.

In the end, we observed that the FMC can in some instances reduce the time complexity of the simulation and proves to be more efficient than standard TEDOPA in simulations involving long chains and moderate levels of entanglement. All these results, starting from the equivalence theorem, show the validity of the FMC method.

In Chapter 4 we discussed an experiment designed to highlight the utility of pre-fault-tolerant quantum processors for studying many-body physical models with non-trivial dynamics. This experiment adds to the growing body of work that leverages classical computation to extend the reach of near-term quantum processors [58, 215–217], advancing the boundaries of quantum simulation.

In particular, we outlined the main features of DU circuits and how they can be used to represent the time evolution of a chaotic physical model, the kicked Ising model. This circuit model induces a significant amount of entanglement in the device, but at the same time it is equipped with certain physical quantities that can be analytically calculated, proving to be an interesting benchmark system for NISQ devices. Next, we illustrated the methodology of the experiment, from the measurement scheme to the noise modelling and, finally, the TN scheme at the heart of the post-processing calculations. This combination allowed us to recover meaningful data from the noisy output of the quantum computer. We then focused in particular on the classical MPS-based simulations that we ran parallel to the hybrid approach, studying different methods of assessing their convergence, their efficiency, and whether they were well suited to give us results that matched the error-mitigated output of the quantum device.

The error-mitigated results, when compared to the expected theoretical curves of auto-correlation functions of the DU kicked Ising model, demonstrated the capability of our quantum-classical hybrid approach to accurately simulate their behaviour, and therefore of how these systems can serve as performance benchmarks for non-Clifford circuits. The fully classical TN simulator, in turn, proved to be a useful method to check the validity of the noise model—by running simulations of the noisy circuit—as well as a way to explore the behaviour of the circuit when perturbed away from the DU points, where no analytical solutions existed for our physical quantities of interest. We observed that, while the simulations of the noisy device converged to a satisfactory degree, a full-state naive simulation of the state of the device is not feasible at the scale of the NISQ device we used, and more clever schemes need to be devised in order to have meaningful results.

Future directions

The set of pseudomodes defined in the FMC could be used in impurity solvers as a direct replacement of the environment, if the original spectral density function already resembles a semicircle. Choosing a suitable initial state for the FMC, however, is not trivial because the pseudomodes directly interact with one another: a state where each pseudomode is independently in a thermal equilibrium state would therefore not be the correct overall thermal equilibrium state. In this regard, the chain mapping (together with the thermofield or T-TEDOPA techniques) also serves the purpose of 'thermalising' the environment, so that the initial state of the FMC modes can be either the empty or the filled state.

There are of course some improvements that can be studied in order to enhance the performance of the FMC method. While the quality of the results obtained by means of the FMC is already satisfying, we can look for a better parameterisation of the closure, e.g. by means of the techniques introduced in Refs. [166, 167], in order to improve the accuracy of the fitting of the asymptotic spectral density, as well as alternative fermion-to-spin representations leading to a decrease of the entanglement in the tensor network. Another area where work can be done is the shape of the TN, as mentioned at the end of Chapter 1. If, for example, we only had one environment, by choosing a 'ladder' structure as in Fig. 1.15 all interactions between sites would be of the nearest-neighbour type and we would have a local dimension equal to that of pure states. However, it is not clear whether the smaller overall bond dimensions would outweigh the greater complexity of the tensor operations, as the individual block of the TN would generally have more indices (legs).

Another open question is how we can extract meaningful physical quantities of the residual environment from the state of the pseudomodes in the MC, if it can be done at all. The truncation of the TEDOPA chain, in fact, breaks the possibility of reversing the chain mapping. A standard TEDOPA chain mapping can always be inverted to recover the original formulation of the continuous environment, the only error being in the finite length of the actual chain used in the simulations (which is however not a significant issue if the chain is long enough to properly contain the dynamics). This makes it possible, for example, to estimate TTCFs involving the original environment mode operators from the TTCFs of the TEDOPA modes. This, however, is not possible anymore once we introduce the MC, since the transformation from the residual environment to the pseudomodes is not invertible, and we cannot recover the mode operators of the former from the latter. This precludes us, currently, from computing useful quantities such as correlations between the open system and the whole environment with the MC scheme.

Concurrently, we continue to search for new physical systems where the FMC can improve the quality of the numerical simulations to analyse physically relevant quantities, such as spectral functions of interacting systems or the Kondo phenomenon. A particular class of systems where this could happen are systems that need to be described with mixed states from the start, for example because of driving terms acting on the open system: in this case a standard TEDOPA simulation, with mixed states, would be very demanding because of the high local dimension and the additional classical correlations, so the FMC would undoubtedly provide a great simplification.

As far as the experiment in Chapter 4 is concerned, finally, the approach we explained could pave the way to a new series of quantum simulations of many-body dynamics on quantum processors that are able to outmatch classical simulators, already before the advent of fault-tolerance as hardware advances towards lower error rates [218], more stable noise [210] and faster speeds [219]. In principle, the hybrid approach to quantum computation with TEM is agnostic to what is actually computed on the device. In its current formulation, TEM it is of course limited to circuits that are either one-dimensional or that can be easily flattened to a one-dimensional structure. It works best when the device noise is low and when, as in the DU circuits we encountered, the information does not significantly spread across all qubits, taking into the account both the initial state and the measured observable. More complex observables would need a more complicated, possibly non-local POVM. In this regard, we can read the TEM-transformed observable in order to find the most significant contributions e.g. in terms of Pauli string components: with this information we can then tailor the POVM to the noisy circuit, obtaining an optimised measurement routine that can to capture the most significant components, simultaneously keeping the variance small.

As more and more developed engineering techniques reduce the noise in more sophisticated quantum devices, we expect that a coarse Markovian noise model, such as the sparse Pauli-Lindblad one that we used, will not accurately capture the influence of the environment anymore. We must therefore look for finer models which include non-Markovian effects and that, at the same time, can be efficiently learned and have a simple TN representation.

The auto-correlation functions we considered, moreover, could be used to extract transport properties [213] and predict the existence of localised phases [214] in further experiments of condensed matter physics. With different choices of observables of the form $P_k P_l$, with P a Pauli operator, we could study for example two-body correlators, magnetisation imbalance, and two-body mutual information with the same experimental setup. Such observables are likely to show a similar qualitative behaviour to our auto-correlator as functions of time at the DU point and for the initial state we used. However, their dependence from h might be different and offer new physical insight.

Bibliography

Own works

[1] D. Ferracin, A. Smirne, S. F. Huelga, M. B. Plenio and D. Tamascelli. 'Spectral density modulation and universal Markovian closure of fermionic environments'. *The Journal of Chemical Physics* 161.17, 174114 (2024).

[2] L. E. Fischer, M. Leahy, A. Eddins, N. Keenan, D. Ferracin, M. A. C. Rossi et al. 'Dynamical simulations of many-body quantum chaos on a quantum computer'. To appear in: *Nature Physics* (2026). Preprint: arXiv: `2411.00765`.

[3] [SW] D. Ferracin and D. Tamascelli, *markovian_closure_fermions*. LIC: GPL-3.0-or-later. URL: `https://github.com/phaerrax/markovian_closure_fermions`.

References

[4] A. Georges, G. Kotliar, W. Krauth and M. J. Rozenberg. 'Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions'. *Reviews of Modern Physics* 68 (1 1996), pp. 13–125.

[5] W. M. C. Foulkes, L. Mitas, R. J. Needs and G. Rajagopal. 'Quantum Monte Carlo simulations of solids'. *Reviews of Modern Physics* 73 (1 2001), pp. 33–83.

[6] G. Carleo and M. Troyer. 'Solving the quantum many-body problem with artificial neural networks'. *Science* 355.6325 (2017), pp. 602–606.

[7] M. J. Hartmann and G. Carleo. 'Neural-Network Approach to Dissipative Quantum Many-Body Dynamics'. *Physical Review Letters* 122, 250502 (25 2019).

[8] R. P. Rundle and M. J. Everitt. 'Overview of the Phase Space Formulation of Quantum Mechanics with Application to Quantum Technologies'. *Advanced Quantum Technologies* 4.6, 2100016 (2021).

[9] J. Richter and R. Steinigeweg. 'Combining dynamical quantum typicality and numerical linked cluster expansions'. *Physical Review B* 99, 094419 (9 2019).

[10] R. Orús. 'Tensor networks for complex quantum systems'. *Nature Reviews Physics* 1.9 (2019), pp. 538–550.

[11] M. Fannes, B. Nachtergaele and R. F. Werner. 'Finitely correlated states on quantum spin chains'. *Communications in Mathematical Physics* 144 (3 1992), pp. 443–490.

[12] I. Affleck, T. Kennedy, E. H. Lieb and H. Tasaki. 'Rigorous results on valence-bond ground states in antiferromagnets'. *Physical Review Letters* 59 (7 1987), pp. 799–802.

[13] A. Klümper, A. Schadschneider and J. Zittartz. 'Matrix Product Ground States for One-Dimensional Spin-1 Quantum Antiferromagnets'. *Europhysics Letters* 24.4 (1993), p. 293.

[14] A. K. Kolezhuk, H.-J. Mikeska and S. Yamamoto. 'Matrix-product-states approach to Heisenberg ferrimagnetic spin chains'. *Physical Review B* 55 (6 1997), R3336–R3339.

[15] S. R. White. 'Density matrix formulation for quantum renormalization groups'. *Physical Review Letters* 69 (19 1992), pp. 2863–2866.

[16] U. Schollwöck. 'The density-matrix renormalization group'. *Reviews of Modern Physics* 77 (1 2005), pp. 259–315.

[17] S. Rommer and S. Östlund. 'Class of ansatz wave functions for one-dimensional spin systems and their relation to the density matrix renormalization group'. *Physical Review B* 55 (4 1997), pp. 2164–2181.

[18] J. Eisert, M. Cramer and M. B. Plenio. 'Colloquium: Area laws for the entanglement entropy'. *Reviews of Modern Physics* 82 (1 2010), pp. 277–306.

[19] G. Vidal, J. I. Latorre, E. Rico and A. Kitaev. 'Entanglement in Quantum Critical Phenomena'. *Physical Review Letters* 90, 227902 (22 2003).

[20] F. Verstraete and J. I. Cirac. 'Matrix product states represent ground states faithfully'. *Physical Review B* 73, 094423 (9 2006).

[21] G. Vidal. 'Efficient Classical Simulation of Slightly Entangled Quantum Computations'. *Physical Review Letters* 91, 147902 (14 2003).

[22] G. Vidal. 'Efficient Simulation of One-Dimensional Quantum Many-Body Systems'. *Physical Review Letters* 93, 040502 (4 2004).

[23] F. Verstraete and J. I. Cirac. *Renormalization algorithms for quantum many-body systems in two and higher dimensions*. 2004. arXiv: `cond-mat/0407066`.

[24] G. Vidal. 'Class of Quantum Many-Body States That Can Be Efficiently Simulated'. *Physical Review Letters* 101, 110501 (11 2008).

[25] B. Swingle. 'Entanglement renormalization and holography'. *Physical Review D* 86, 065007 (6 2012).

[26] G. Evenbly and G. Vidal. 'Tensor Network States and Geometry'. *Journal of Statistical Physics* 145.4 (2011), pp. 891–918.

[27] M. Levin and C. P. Nave. 'Tensor Renormalization Group Approach to Two-Dimensional Classical Lattice Models'. *Physical Review Letters* 99, 120601 (12 2007).

[28] Y. Levine, O. Sharir, N. Cohen and A. Shashua. 'Quantum Entanglement in Deep Learning Architectures'. *Physical Review Letters* 122, 065301 (6 2019).

[29] E. Stoudenmire and D. J. Schwab. 'Supervised Learning with Tensor Networks'. In: *Advances in Neural Information Processing Systems*. 30th Conference on Neural Information Processing Systems (NIPS 2016). Ed. by D. Lee, M. Sugiyama, U. Luxburg, I. Guyon and R. Garnett. Vol. 29. Barcelona, Spain: Curran Associates, Inc., 2016.

[30] M. Lubasch, P. Moinier and D. Jaksch. 'Multigrid renormalization'. *Journal of Computational Physics* 372 (2018), pp. 587–602.

[31] N. Gourianov, P. Givi, D. Jaksch and S. B. Pope. 'Tensor networks enable the calculation of turbulence probability distributions'. *Science Advances* 11.5, eads5990 (2025).

[32] H.-P. Breuer and F. Petruccione. *The Theory of Open Quantum Systems*. Oxford: Oxford University Press, 2002.

[33] A. W. Chin et al. 'The role of non-equilibrium vibrational structures in electronic coherence and recoherence in pigment-protein complexes'. *Nature Physics* 9 (2013), p. 113.

[34] J. Prior, A. W. Chin, S. F. Huelga and M. B. Plenio. 'Efficient Simulation of Strong System-Environment Interactions'. *Physical Review Letters* 105, 050404 (5 2010).

[35] A. Nüßeler et al. 'Fingerprint and Universal Markovian Closure of Structured Bosonic Environments'. *Physical Review Letters* 129, 140604 (14 2022).

[36] J. Friedel. 'On some electrical and magnetic properties of metallic solid solutions'. *Canadian Journal of Physics* 34.12A (1956), pp. 1190–1211.

[37] A. M. Clogston et al. 'Local Magnetic Moment Associated with an Iron Atom Dissolved in Various Transition Metal Alloys'. *Physical Review* 125 (2 1962), pp. 541–552.

[38] L. P. Kouwenhoven et al. ‘Electron Transport in Quantum Dots’. In: *Mesoscopic Electron Transport* (Curaçao, Netherlands Antilles, 25th June–5th July 1996). Ed. by L. L. Sohn, L. P. Kouwenhoven and G. Schön. Springer Netherlands, 1997, pp. 105–214.

[39] A. Dorda, M. Nuss, W. von der Linden and E. Arrigoni. ‘Auxiliary master equation approach to nonequilibrium correlated impurities’. *Physical Review B* 89, 165105 (16 2014).

[40] F. Schwarz, I. Weymann, J. von Delft and A. Weichselbaum. ‘Nonequilibrium Steady-State Transport in Quantum Impurity Models: A Thermofield and Quantum Quench Approach Using Matrix Product States’. *Physical Review Letters* 121, 137702 (13 2018).

[41] A. J. Daley et al. ‘Practical quantum advantage in quantum simulation’. *Nature* 607.7920 (2022), pp. 667–676.

[42] D. A. Abanin et al. ‘Observation of constructive interference at the edge of quantum ergodicity’. *Nature* 646.8086 (2025), pp. 825–830.

[43] H. Gharibyan et al. *Heuristic Quantum Advantage with Peaked Circuits*. 2025. arXiv: `2510.25838`.

[44] J. Ruane et al. *Quantum Index Report 2025*. MIT Initiative on the Digital Economy, Massachusetts Institute of Technology, 2025.

[45] E. Knill, R. Laflamme and W. H. Zurek. ‘Resilient Quantum Computation’. *Science* 279.5349 (1998), pp. 342–345.

[46] A. Y. Kitaev. ‘Fault-tolerant quantum computation by anyons’. *Annals of Physics* 303.1 (2003), pp. 2–30.

[47] D. Aharonov and M. Ben-Or. ‘Fault-Tolerant Quantum Computation with Constant Error Rate’. *SIAM Journal on Computing* 38.4 (2008), pp. 1207–1282.

[48] F. Arute et al. ‘Quantum supremacy using a programmable superconducting processor’. *Nature* 574.7779 (2019), pp. 505–510.

[49] J. Clarke and F. K. Wilhelm. ‘Superconducting quantum bits’. *Nature* 453.7198 (2008), pp. 1031–1042.

[50] S. Debnath et al. ‘Demonstration of a small programmable quantum computer with atomic qubits’. *Nature* 536.7614 (2016), pp. 63–66.

[51] N. Friis et al. ‘Observation of Entangled States of a Fully Controlled 20-Qubit System’. *Physical Review X* 8.2, 021012 (2018).

[52] A. Noiri et al. 'Fast universal quantum gate above the fault-tolerance threshold in silicon'. *Nature* 601.7893 (2022), pp. 338–342.

[53] M. T. Mądzik et al. 'Precision tomography of a three-qubit donor quantum processor in silicon'. *Nature* 601.7893 (2022), pp. 348–353.

[54] J. Preskill. 'Quantum Computing in the NISQ era and beyond'. *Quantum* 2, 79 (2018).

[55] X. Mi et al. 'Information scrambling in quantum circuits'. *Science* 374.6574 (2021), pp. 1479–1483.

[56] Y. Kim et al. 'Evidence for the utility of quantum computing before fault tolerance'. *Nature* 618.7965 (2023), pp. 500–505.

[57] S. Chen, J. Cotler, H.-Y. Huang and J. Li. 'The complexity of NISQ'. *Nature Communications* 14.1 (2023), pp. 6001–6006.

[58] S. Filippov, M. Leahy, M. A. C. Rossi and G. García-Pérez. *Scalable tensor-network error mitigation for near-term quantum computing*. 2023. arXiv: `2307.11740`.

[59] B. Bertini, P. Kos and T. Prosen. 'Exact Correlation Functions for Dual-Unitary Lattice Models in $1+1$ Dimensions'. *Physical Review Letters* 123, 210601 (21 2019).

[60] B. Bertini, P. Kos and T. Prosen. 'Exact Spectral Form Factor in a Minimal Model of Many-Body Quantum Chaos'. *Physical Review Letters* 121, 264101 (26 2018).

[61] I. L. Markov and Y. Shi. 'Simulating Quantum Computation by Contracting Tensor Networks'. *SIAM Journal on Computing* 38.3 (2008), pp. 963–981.

[62] U. Schollwöck. 'The density-matrix renormalization group in the age of matrix product states'. *Annals of Physics* 326.1 (2011), pp. 96–192.

[63] M. B. Hastings. 'An area law for one-dimensional quantum systems'. *Journal of Statistical Mechanics: Theory and Experiment* 2007.08 (2007), P08024.

[64] P. Calabrese and J. Cardy. 'Evolution of entanglement entropy in one-dimensional systems'. *Journal of Statistical Mechanics: Theory and Experiment* 2005, P04010 (2005).

[65] M. Fagotti and P. Calabrese. 'Evolution of entanglement entropy following a quantum quench: Analytic results for the XY chain in a transverse magnetic field'. *Physical Review A* 78.1 (2008).

[66] S. Paeckel et al. 'Time-evolution methods for matrix-product states'. *Annals of Physics* 411, 167998 (2019).

[67] A. J. Ferris and G. Vidal. 'Perfect sampling with unitary tensor networks'. *Physical Review B* 85, 165146 (16 2012).

[68] E. M. Stoudenmire and S. R. White. 'Real-space parallel density matrix renormalization group'. *Physical Review B* 87.15, 155137 (2013).

[69] P. Jordan and E. Wigner. 'Über das Paulische Äquivalenzverbot'. *Zeitschrift für Physik* 47 (9 1928), pp. 631–651.

[70] S. B. Bravyi and A. Y. Kitaev. 'Fermionic Quantum Computation'. *Annals of Physics* 298.1 (2002), pp. 210–226.

[71] G. C. Wick, A. S. Wightman and E. P. Wigner. 'The Intrinsic Parity of Elementary Particles'. *Physical Review* 88.1 (1952), pp. 101–105.

[72] G. C. Hegerfeldt, K. Kraus and E. P. Wigner. 'Proof of the Fermion Superselection Rule without the Assumption of Time-Reversal Invariance'. *Journal of Mathematical Physics* 9.12 (1968), pp. 2029–2031.

[73] N. Friis. 'Reasonable fermionic quantum information theories require relativity'. *New Journal of Physics* 18.3, 033014 (2016).

[74] L. Ma, M. Fishman, E. M. Stoudenmire and E. Solomonik. 'Approximate Contraction of Arbitrary Tensor Networks with a Flexible and Efficient Density Matrix Algorithm'. *Quantum* 8, 1580 (2024).

[75] M. Stoian, R. M. Milbradt and C. B. Mendl. 'On the Optimal Linear Contraction Order of Tree Tensor Networks, and Beyond'. *SIAM Journal on Scientific Computing* 46.5 (2024), B647–B668.

[76] H. Zassenhaus. 'Über Lie'sche Ringe mit Primzahlcharakteristik'. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg* 13 (1 1939), pp. 1–100.

[77] M. Suzuki. 'Decomposition formulas of exponential operators and Lie exponentials with some applications to quantum mechanics and statistical physics'. *Journal of Mathematical Physics* 26.4 (1985), pp. 601–612.

[78] M. Suzuki. 'Fractal decomposition of exponential operators with applications to many-body theories and Monte Carlo simulations'. *Physics Letters A* 146.6 (1990), pp. 319–323.

[79] E. M. Stoudenmire and S. R. White. 'Minimally entangled typical thermal state algorithms'. *New Journal of Physics* 12.5, 055026 (2010).

[80] C. Lubich, I. V. Oseledets and B. Vandereycken. 'Time Integration of Tensor Trains'. *SIAM Journal on Numerical Analysis* 53.2 (2015), pp. 917–941.

[81] J. Haegeman, M. Mariën, T. J. Osborne and F. Verstraete. 'Geometry of matrix product states: Metric, parallel transport, and curvature'. *Journal of Mathematical Physics* 55.2, 021902 (2014).

[82] J. Haegeman, C. Lubich, I. Oseledets, B. Vandereycken and F. Verstraete. 'Unifying time evolution and optimization with matrix product states'. *Physical Review B* 94, 165116 (16 2016).

[83] M. Yang and S. R. White. 'Time-dependent variational principle with ancillary Krylov subspace'. *Physical Review B* 102.9, 094315 (2020).

[84] A. J. Dunnett and A. W. Chin. 'Efficient bond-adaptive approach for finite-temperature open quantum dynamics using the one-site time-dependent variational principle for matrix product states'. *Physical Review B* 104, 214302 (21 2021).

[85] P. Secular et al. 'Parallel time-dependent variational principle algorithm for matrix product states'. *Physical Review B* 101, 235123 (23 2020).

[86] G. Vidal and R. F. Werner. 'Computable measure of entanglement'. *Physical Review A* 65, 032314 (3 2002).

[87] T. Prosen and I. Pižorn. 'Operator space entanglement entropy in a transverse Ising chain'. *Physical Review A* 76, 032316 (3 2007).

[88] A. H. Werner et al. 'Positive Tensor Network Approach for Simulating Open Quantum Many-Body Systems'. *Physical Review Letters* 116, 237201 (23 2016).

[89] M. Kliesch, D. Gross and J. Eisert. 'Matrix-Product Operators and States: NP-Hardness and Undecidability'. *Physical Review Letters* 113.16, 160503 (2014).

[90] M. Schmutz. 'Real-Time Green's Functions in Many Body Problems'. *Zeitschrift für Physik B Condensed Matter* 30.1 (1978), pp. 97–106.

[91] M. Brenes et al. 'Tensor-Network Method to Simulate Strongly Interacting Quantum Thermal Machines'. *Physical Review X* 10, 031040 (3 2020).

[92] D. J. Strachan, A. Purkayastha and S. R. Clark. 'Extracting dynamical maps of non-Markovian open quantum systems'. *The Journal of Chemical Physics* 161.15, 154105 (2024).

[93] J. Dereziński. 'Introduction to Representations of the Canonical Commutation and Anticommutation Relations'. In: *Large Coulomb Systems: Lecture Notes on Mathematical Aspects of QED*. Ed. by J. Dereziński and H. Siedentop. Springer Berlin Heidelberg, 2006, pp. 63–143.

[94] N. T. Vidal, M. L. Bera, A. Riera, M. Lewenstein and M. N. Bera. 'Quantum operations in an information theory for fermions'. *Physical Review A* 104, 032411 (3 2021).

[95] S. F. Huelga and M. B. Plenio. 'Vibrations, quanta and biology'. *Contemporary Physics* 54.4 (2013), pp. 181–207.

[96] P. M. Harrington, E. J. Mueller and K. W. Murch. 'Engineered dissipation for quantum information science'. *Nature Reviews Physics* 4.10 (2022), pp. 660–671.

[97] A. O. Caldeira and A. J. Leggett. 'Influence of Dissipation on Quantum Tunneling in Macroscopic Systems'. *Physical Review Letters* 46 (4 1981), pp. 211–214.

[98] A. J. Leggett et al. 'Dynamics of the dissipative two-state system'. *Reviews of Modern Physics* 59 (1 1987), pp. 1–85.

[99] M. Thorwart, E. Paladino and M. Grifoni. 'Dynamics of the spin-boson model with a structured environment'. *Chemical Physics* 296.2 (2004), pp. 333–344.

[100] J. Kondo. 'Resistance Minimum in Dilute Magnetic Alloys'. *Progress of Theoretical Physics* 32.1 (1964), pp. 37–49.

[101] P. W. Anderson. 'Localized Magnetic States in Metals'. *Physical Review* 124 (1 1961), pp. 41–53.

[102] M. T. Mitchison and M. B. Plenio. 'Non-additive dissipation in open quantum networks out of equilibrium'. *New Journal of Physics* 20.3, 033005 (2018).

[103] G. Wójtowicz, J. E. Elenewski, M. M. Rams and M. Zwolak. 'Open-system tensor networks and Kramers' crossover for quantum transport'. *Physical Review A* 101.5, 050301 (2020).

[104] M. Zwolak. 'Analytic expressions for the steady-state current with finite extended reservoirs'. *The Journal of Chemical Physics* 153.22, 224107 (2020).

[105] N. Friis, A. R. Lee and D. E. Bruschi. 'Fermionic-mode entanglement in quantum information'. *Physical Review A* 87, 022338 (2 2013).

[106] S. Banerjee. *Open Quantum Systems: Dynamics of Nonclassical Evolution*. Texts and Readings in Physical Sciences. Springer Singapore, 2018.

[107] F. B. Anders. 'Steady-State Currents through Nanodevices: A Scattering-States Numerical Renormalization-Group Approach to Open Quantum Systems'. *Physical Review Letters* 101, 066804 (6 2008).

[108] S. Schmitt and F. B. Anders. 'Comparison between scattering-states numerical renormalization group and the Kadanoff-Baym-Keldysh approach to quantum transport: Crossover from weak to strong correlations'. *Physical Review B* 81, 165106 (16 2010).

[109] F. Wegner. 'Flow-equations for Hamiltonians'. *Annalen der Physik* 506 (1994), pp. 77–91.

[110] S. Kehrein. *The Flow Equation Approach to Many-Particle Systems*. Springer Tracts in Modern Physics. Springer Berlin, 2006.

[111] H. Wang and M. Thoss. 'Numerically exact quantum dynamics for indistinguishable particles: The multilayer multiconfiguration time-dependent Hartree theory in second quantization representation'. *The Journal of Chemical Physics* 131, 024114 (2009).

[112] G. Cohen, E. Gull, D. R. Reichman and A. J. Millis. 'Taming the Dynamical Sign Problem in Real-Time Evolution of Quantum Many-Body Problems'. *Physical Review Letters* 115, 266802 (2015).

[113] Q. Dong et al. 'Quantum Monte Carlo solution of the dynamical mean field equations in real time'. *Physical Review B* 96, 155126 (2017).

[114] M. Ridley, V. N. Singh, E. Gull and G. Cohen. 'Numerically exact full counting statistics of the nonequilibrium Anderson impurity model'. *Physical Review B* 97, 115109 (2018).

[115] Y. Tanimura. 'Numerically "exact" approach to open quantum dynamics: The hierarchical equations of motion (HEOM)'. *The Journal of Chemical Physics* 153.2, 020901 (2020).

[116] A. Strathearn, P. Kirton, D. Kilda, J. Keeling and B. W. Lovett. 'Efficient non-Markovian quantum dynamics using time-evolving matrix product operators'. *Nature Communications* 9.1, 3322 (2018).

[117] A. D. Somoza, O. Marty, J. Lim, S. F. Huelga and M. B. Plenio. 'Dissipation-Assisted Matrix Product Factorization'. *Physical Review Letters* 123, 100502 (10 2019).

[118] A. D. Somoza, N. Lorenzoni, J. Lim, S. F. Huelga and M. B. Plenio. 'Driving force and nonequilibrium vibronic dynamics in charge separation of strongly bound electron-hole pairs'. *Communications Physics* 6.1 (2023), p. 65.

[119] M. Cygorek et al. 'Simulation of open quantum systems by automated compression of arbitrary environments'. *Nature Physics* 18 (2022), pp. 662–668.

[120] V. Gorini, A. Kossakowski and E. C. G. Sudarshan. 'Completely positive dynamical semigroups of N-level systems'. *Journal of Mathematical Physics* 17.5 (1976), pp. 821–825.

[121] G. Lindblad. 'Completely positive maps and entropy inequalities'. *Communications in Mathematical Physics* 40.2 (1975), pp. 147–151.

[122] W. G. van der Wiel et al. 'The Kondo Effect in the Unitary Limit'. *Science* 289.5487 (2000), pp. 2105–2108.

[123] S. Sasaki et al. 'Kondo effect in an integer-spin quantum dot'. *Nature* 405 (2000), pp. 764–767.

[124] P. Etchegoin. 'Molecular transistors as open fermionic quantum systems'. *Physica A: Statistical Mechanics and its Applications* 313.3 (2002), pp. 488–502.

[125] L. Fritz and M. Vojta. 'Kondo screening in unconventional superconductors: The role of anomalous propagators'. *Physical Review B* 72, 212510 (21 2005).

[126] C. Romeike, M. R. Wegewijs, W. Hofstetter and H. Schoeller. 'Kondo-Transport Spectroscopy of Single Molecule Magnets'. *Physical Review Letters* 97, 206601 (2006).

[127] P. Strasberg, G. Schaller, N. Lambert and T. Brandes. 'Nonequilibrium thermodynamics in the strong coupling and non-Markovian regime based on a reaction coordinate mapping'. *New Journal of Physics* 18.7, 073007 (2016).

[128] K. Ptaszyński. 'Non-Markovian thermal operations boosting the performance of quantum heat engines'. *Physical Review E* 106, 014114 (1 2022).

[129] A. W. Chin, S. F. Huelga and M. B. Plenio. 'Quantum Metrology in Non-Markovian Environments'. *Physical Review Letters* 109, 233601 (23 2012).

[130] A. Smirne, J. Kołodyński, S. F. Huelga and R. Demkowicz-Dobrzański. 'Ultimate Precision Limits for Noisy Frequency Estimation'. *Physical Review Letters* 116, 120801 (12 2016).

[131] M. C. Gélvez-Rueda et al. 'Inducing Charge Separation in Solid-State Two-Dimensional Hybrid Perovskites through the Incorporation of Organic Charge-Transfer Complexes'. *The Journal of Physical Chemistry Letters* 11 (3 2020), pp. 824–830.

[132] A. M. Alvertis et al. 'Impact of exciton delocalization on exciton-vibration interactions in organic semiconductors'. *Physical Review B* 102, 081122 (8 2020).

[133] J. Lim et al. 'Vibronic origin of long-lived coherence in an artificial molecular light harvester'. *Nature Communications* 6.1, 7755 (2015).

[134] E. Romero et al. 'Quantum coherence in photosynthesis for efficient solar-energy conversion'. *Nature Physics* 10.9 (2014), pp. 676–682.

[135] A. Serafini. *Quantum Continuous Variables. A Primer of Theoretical Methods*. 2nd ed. CRC Press, 2023.

[136] J. Surace and L. Tagliacozzo. 'Fermionic Gaussian states: an introduction to numerical approaches'. *SciPost Physics Lecture Notes* (2022), p. 54.

[137] G. C. Wick. 'The Evaluation of the Collision Matrix'. *Physical Review* 80 (2 1950), pp. 268–272.

[138] D. F. V. James. 'Quantum dynamics of cold trapped ions with application to quantum computation'. *Applied Physics B* 66.2 (1998), pp. 181–190.

[139] S. Barzanjeh et al. 'Optomechanics for quantum technologies'. *Nature Physics* 18.1 (2022), pp. 15–24.

[140] E. Fradkin. *Field Theories of Condensed Matter Physics*. 2nd ed. Cambridge University Press, 2013.

[141] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin and X. Yuan. 'Quantum computational chemistry'. *Reviews of Modern Physics* 92, 015003 (1 2020).

[142] B. M. Garraway. 'Nonperturbative decay of an atomic system in a cavity'. *Physical Review A* 55 (1997), pp. 2290–2303.

[143] G. Pleasance, B. M. Garraway and F. Petruccione. 'Generalized theory of pseudomodes for exact descriptions of non-Markovian quantum processes'. *Physical Review Research* 2, 043058 (4 2020).

[144] D. Tamascelli, A. Smirne, S. F. Huelga and M. B. Plenio. 'Nonperturbative Treatment of non-Markovian Dynamics of Open Quantum Systems'. *Physical Review Letters* 120, 030402 (6 2018).

[145] A. Dorda, M. Nuss, W. von der Linden and E. Arrigoni. 'Auxiliary master equation approach to nonequilibrium correlated impurities'. *Physical Review B* 89, 165105 (16 2014).

[146] A. Dorda, M. Sorantin, W. von der Linden and E. Arrigoni. 'Optimized auxiliary representation of non-Markovian impurity problems by a Lindblad equation'. *New Journal of Physics* 19.6 (2017).

[147] F. Chen, E. Arrigoni and M. Galperin. 'Markovian treatment of non-Markovian dynamics of open Fermionic systems'. *New Journal of Physics* 21.12, 123035 (2019).

[148] M. Cirio et al. ‘Pseudofermion method for the exact description of fermionic environments: From single-molecule electronics to the Kondo resonance’. *Physical Review Research* 5 (3 2023).

[149] F. Mascherpa et al. ‘Optimized auxiliary oscillators for the simulation of general open quantum systems’. *Physical Review A* 101, 052108 (5 2020).

[150] F. Mascherpa, A. Smirne, S. F. Huelga and M. B. Plenio. ‘Open Systems with Error Bounds: Spin-Boson Model with Spectral Density Variations’. *Physical Review Letters* 118, 100401 (10 2017).

[151] A. W. Chin, Á. Rivas, S. F. Huelga and M. B. Plenio. ‘Exact mapping between system-reservoir quantum models and semi-infinite discrete chains using orthogonal polynomials’. *Journal of Mathematical Physics* 51.9, 092109 (2010).

[152] A. Nüßeler, I. Dhand, S. F. Huelga and M. B. Plenio. ‘Efficient simulation of open quantum systems coupled to a fermionic bath’. *Physical Review B* 101, 155134 (15 2020).

[153] D. Tamascelli, A. Smirne, J. Lim, S. F. Huelga and M. B. Plenio. ‘Efficient Simulation of Finite-Temperature Open Quantum Systems’. *Physical Review Letters* 123.9, 090402 (2019).

[154] I. de Vega and M.-C. Bañuls. ‘Thermofield-based chain-mapping approach for open quantum systems’. *Physical Review A* 92, 052116 (5 2015).

[155] I. de Vega, U. Schollwöck and F. A. Wolf. ‘How to discretize a quantum bath for real-time evolution’. *Physical Review B* 92, 155126 (15 2015).

[156] M. P. Woods and M. B. Plenio. ‘Dynamical error bounds for continuum discretisation via Gauss quadrature rules—A Lieb-Robinson bound approach’. *Journal of Mathematical Physics* 57.2 (2016).

[157] L. Kohn and G. E. Santoro. ‘Efficient mapping for Anderson impurity problems with matrix product states’. *Physical Review B* 104, 014303 (1 2021).

[158] L. Kohn and G. E. Santoro. ‘Quench dynamics of the Anderson impurity model at finite temperature using matrix product states: entanglement and bath dynamics’. *Journal of Statistical Mechanics: Theory and Experiment* 2022.6, 063102 (2022-06).

[159] M. Fishman, S. R. White and E. M. Stoudenmire. ‘The ITensor Software Library for Tensor Network Calculations’. *SciPost Physics Codebases* 4 (2022).

[160] J.-H. Jung and J. D. Noh. ‘Guide to Exact Diagonalization Study of Quantum Thermalization’. *Journal of the Korean Physical Society* 76.8 (2020), pp. 670–683.

[161] G. Szegő. *Orthogonal Polynomials*. Revised. Colloquium publications. New York: American Mathematical Society, 1959.

[162] D. Tamascelli. 'Excitation Dynamics in Chain-Mapped Environments'. *Entropy* 22.11, 1320 (2020).

[163] R. Martinazzo, B. Vacchini, K. H. Hughes and I. Burghardt. 'Communication: Universal Markovian reduction of Brownian particle dynamics'. *The Journal of Chemical Physics* 134.1, 011101 (2011).

[164] G. Beylkin and L. Monzón. 'On approximation of functions by exponential sums'. *Applied and Computational Harmonic Analysis* 19.1 (2005), pp. 17–48.

[165] M. Cirio, P.-C. Kuo, Y.-N. Chen, F. Nori and N. Lambert. 'Canonical derivation of the fermionic influence superoperator'. *Physical Review B* 105, 035121 (3 2022).

[166] I. Medina, F. J. García-Vidal, A. I. Fernández-Domínguez and J. Feist. 'Few-Mode Field Quantization of Arbitrary Electromagnetic Spectral Densities'. *Physical Review Letters* 126, 093601 (9 2021).

[167] M. Sánchez-Barquilla and J. Feist. 'Accurate Truncations of Chain Mapping Models for Open Quantum Systems'. *Nanomaterials* 11.8, 2104 (2021).

[168] K. H. Hughes, C. D. Christ and I. Burghardt. 'Effective-mode representation of non-Markovian dynamics: A hierarchical approximation of the spectral density. I. Application to single surface dynamics'. *The Journal of Chemical Physics* 131, 024109 (2009).

[169] N. Anto-Sztrikacs and D. Segal. 'Capturing non-Markovian dynamics with the reaction coordinate method'. *Physical Review A* 104, 052617 (5 2021).

[170] D. Bauernfeind and M. Aichhorn. 'Time dependent variational principle for tree Tensor Networks'. *SciPost Physics* 8, 024 (2020).

[171] T. Lacroix, B. W. Lovett and A. W. Chin. *Connectivity matters: Impact of bath modes ordering and geometry in open quantum system simulation with Tensor Network States*. 2025. arXiv: `2409.04145`.

[172] M. Ganahl et al. 'Efficient DMFT impurity solver using real-time dynamics with matrix product states'. *Physical Review B* 92, 155132 (15 2015).

[173] D. Tamascelli, R. Rosenbach and M. B. Plenio. 'Improved scaling of time-evolving block-decimation algorithm through reduced-rank randomized singular value decomposition'. *Physical Review E* 91, 063306 (6 2015).

[174] D. J. Strachan, A. Purkayastha and S. R. Clark. *Accelerated calculation of impurity Green's functions exploiting the extreme Mpemba effect*. 2025. arXiv: `2510.26651`.

[175] R. Rosenbach, J. Cerrillo, S. F. Huelga, J. Cao and M. B. Plenio. 'Efficient simulation of non-Markovian system-environment interaction'. *New Journal of Physics* 18.2, 023035 (2016).

[176] R. P. Feynman. 'Simulating physics with computers'. *International Journal of Theoretical Physics* 21 (6 1982), pp. 467–488.

[177] B. Skinner, J. Ruhman and A. Nahum. 'Measurement-Induced Phase Transitions in the Dynamics of Entanglement'. *Physical Review X* 9, 031009 (3 2019).

[178] A. Chan, A. De Luca and J. T. Chalker. 'Solution of a Minimal Model for Many-Body Quantum Chaos'. *Physical Review X* 8, 041019 (4 2018).

[179] Y. Li, X. Chen and M. P. A. Fisher. 'Quantum Zeno effect and the many-body entanglement transition'. *Physical Review B* 98, 205136 (20 2018).

[180] B. Bertini, P. Kos and T. Prosen. 'Entanglement Spreading in a Minimal Model of Maximal Many-Body Quantum Chaos'. *Physical Review X* 9, 021033 (2 2019).

[181] B. Bertini, P. Kos and T. Prosen. 'Operator entanglement in local quantum circuits I: Chaotic dual-unitary circuits'. *SciPost Physics* 8.4, 067 (2020).

[182] E. Chertkov et al. 'Holographic dynamics simulations with a trapped-ion quantum computer'. *Nature Physics* 18 (2022), pp. 1074–1079.

[183] E. van den Berg, Z. K. Minev, A. Kandala and K. Temme. 'Probabilistic error cancellation with sparse Pauli-Lindblad models on noisy quantum processors'. *Nature Physics* 19.8 (2023), pp. 1116–1121.

[184] A. Kandala et al. 'Error mitigation extends the computational reach of a noisy quantum processor'. *Nature* 567 (2019), pp. 491–495.

[185] K. Temme, S. Bravyi and J. M. Gambetta. 'Error Mitigation for Short-Depth Quantum Circuits'. *Physical Review Letters* 119, 180509 (18 2017).

[186] A. Foligno and B. Bertini. 'Growth of entanglement of generic states under dual-unitary dynamics'. *Physical Review B* 107, 174311 (17 2023).

[187] G. Casati, B. V. Chirikov, F. M. Izraelev and J. Ford. 'Stochastic behavior of a quantum pendulum under a periodic perturbation'. In: *Stochastic Behavior in Classical and Quantum Hamiltonian Systems*. Ed. by G. Casati and J. Ford. Springer Berlin Heidelberg, 1979, pp. 334–352.

[188] M. Žnidarič. 'Two-step relaxation in local many-body Floquet systems'. *Journal of Physics A: Mathematical and Theoretical* 56.43, 434001 (2023).

[189] M. Bukov, L. D'Alessio and A. Polkovnikov. 'Universal high-frequency behavior of periodically driven systems: from dynamical stabilization to Floquet engineering'. *Advances in Physics* 64.2 (2015), pp. 139–226.

[190] T. Prosen. 'General relation between quantum ergodicity and fidelity of quantum dynamics'. *Physical Review E* 65, 036208 (3 2002).

[191] T. Prosen. 'Chaos and complexity of quantum motion'. *Journal of Physics A: Mathematical and Theoretical* 40.28, 7881 (2007).

[192] L. M. Sieberer et al. 'Digital quantum simulation, Trotter errors, and quantum chaos of the kicked top'. *npj Quantum Information* 5 (1 2019), p. 78.

[193] M. Akila, D. Waltner, B. Gutkin and T. Guhr. 'Particle-time duality in the kicked Ising spin chain'. *Journal of Physics A: Mathematical and Theoretical* 49.37, 375101 (2016).

[194] D. Gottesman. 'The Heisenberg Representation of Quantum Computers'. In: *Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics* (Hobart, Australia, 13th–17th July 1998). Ed. by S. P. Corney, R. Delbourgo and P. D. Jarvis. International Press, 1999, pp. 32–43.

[195] J. Řeháček and M. Paris. *Quantum State Estimation*. Vol. 649. Lecture Notes in Physics. Springer Berlin Heidelberg, 2004.

[196] H.-Y. Huang, R. Kueng and J. Preskill. 'Predicting many properties of a quantum system from very few measurements'. *Nature Physics* 16 (2020), pp. 1050–1057.

[197] A. Glos et al. *Adaptive POVM implementations and measurement error mitigation strategies for near-term quantum devices*. 2022. arXiv: `2208.07817`.

[198] J. Chow, O. Dial and J. Gambetta. *IBM Quantum breaks the 100-qubit processor barrier*. 2021. URL: `https://www.ibm.com/quantum/blog/127-qubit-quantum-processor-eagle`.

[199] S. Bravyi, M. Englbrecht, R. König and N. Peard. 'Correcting coherent errors with surface codes'. *npj Quantum Information* 4, 55 (1 2018).

[200] J. J. Wallman and J. Emerson. 'Noise tailoring for scalable quantum computation via randomized compiling'. *Physical Review A* 94.5, 052325 (2016).

[201] F. B. Maciejewski, Z. Zimborás and M. Oszmaniec. 'Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography'. *Quantum* 4, 257 (2020).

[202] M. R. Geller. ‘Rigorous measurement error correction’. *Quantum Science and Technology* 5.3, 03LT01 (2020).

[203] E. van den Berg, Z. K. Minev and K. Temme. ‘Model-free readout-error mitigation for quantum expectation values’. *Physical Review A* 105.3, 032620 (2022).

[204] S. N. Filippov, S. Maniscalco and G. García-Pérez. *Scalability of quantum error mitigation techniques: from utility to advantage*. 2024. arXiv: `2403.13542`.

[205] G. García-Pérez et al. ‘Learning to Measure: Adaptive Informationally Complete Generalized Measurements for Quantum Algorithms’. *PRX Quantum* 2.4, 040342 (2021).

[206] A. Erhard et al. ‘Characterizing large-scale quantum computers via cycle benchmarking’. *Nature Communications* 10, 5347 (1 2019).

[207] K. Tsubouchi, T. Sagawa and N. Yoshioka. ‘Universal Cost Bound of Quantum Error Mitigation Based on Quantum Estimation Theory’. *Physical Review Letters* 131, 210601 (21 2023).

[208] R. Takagi, S. Endo, S. Minagawa and M. Gu. ‘Fundamental limits of quantum error mitigation’. *npj Quantum Information* 8.114 (2022).

[209] L. C. G. Govia et al. ‘Bounding the Systematic Error in Quantum Error Mitigation due to Model Violation’. *PRX Quantum* 6, 010354 (2025).

[210] Y. Kim et al. ‘Error mitigation with stabilized noise in superconducting quantum processors’. *Nature Communications* 16, 8439 (2025).

[211] T. Begušić, J. Gray and G. K.-L. Chan. ‘Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance’. *Science Advances* 10.3 (2024).

[212] J. Gray. ‘quimb: A python package for quantum information and many-body calculations’. *Journal of Open Source Software* 3.29 (2018).

[213] M. Ljubotina, L. Zadnik and T. Prosen. ‘Ballistic Spin Transport in a Periodically Driven Integrable Quantum System’. *Physical Review Letters* 122, 150605 (15 2019).

[214] D. M. Long, P. J. D. Crowley, V. Khemani and A. Chandran. ‘Phenomenology of the Prethermal Many-Body Localized Regime’. *Physical Review Letters* 131, 106301 (10 2023).

[215] J. Robledo-Moreno et al. ‘Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer’. *Science Advances* 11.25, eadu9991 (2025).

[216] A. Eddins, M. C. Tran and P. Rall. *Lightcone shading for classically accelerated quantum error mitigation*. 2024. arXiv: `2409.04401`.

[217] N. F. Robertson et al. 'Tensor Network Enhanced Dynamic Multiproduct Formulas'. *PRX Quantum* 6, 020360 (2 2025).

[218] J. Stehlik et al. 'Tunable Coupling Architecture for Fixed-Frequency Transmon Superconducting Qubits'. *Physical Review Letters* 127, 080505 (8 2021).

[219] A. D. Rajagopala et al. *Hardware-Assisted Parameterized Circuit Execution*. 2025. arXiv: `2409.03725`.