



Capturing dynamics with noisy quantum computers

Dechant, D.S.

Citation

Dechant, D. S. (2026, February 17). *Capturing dynamics with noisy quantum computers*. Retrieved from <https://hdl.handle.net/1887/4290771>

Version: Publisher's Version

[Licence agreement concerning inclusion of doctoral
thesis in the Institutional Repository of the University
of Leiden](#)

License: <https://hdl.handle.net/1887/4290771>

Note: To cite this publication please use the final published version (if applicable).

CHAPTER 4

Mitigating Shot Noise in Local Overlapping Quantum Tomography with Semidefinite Programming

4.1 Introduction

Accurately characterizing quantum states is a crucial subroutine for many tasks in quantum computing, and can be achieved by quantum state tomography methods [9, 160, 161]. It has applications ranging from certifying quantum devices [162] to the design and execution of variational quantum algorithms [21]. However, the full characterization of a quantum state becomes impractical for larger systems. This is due to an exponential growth in the number of required measurement settings, which correspond to distinct Pauli strings. To address this challenge, several alternative strategies have been proposed, such as classical shadow tomography [163], which significantly reduce the number of required measurement settings. However, even with these techniques, each measurement setting must typically be repeated for a number of shots, N_{meas} , to achieve the desired accuracy, with the statistical uncertainty scaling as $O(1/\sqrt{N_{\text{meas}}})$. Since the measurements collapse the quantum state, the state must be re-prepared for each shot.

Many applications, such as ground state optimization of local Hamiltonians [39, 164], rely on the tomography of k -qubit reduced density matrices

The contents of this chapter have been published in Ref. [90].

(RDMs), which describe a subsystem of the full quantum state [165]. The estimation of all k -qubit RDMs of an n -qubit system can be achieved by using $\binom{n}{k} \cdot 3^k$ distinct measurement settings [166–168]. However, even with exact estimates of RDMs, the full characterization of a general quantum state would entail resolving the quantum marginal problem [169], a task known to be QMA-complete [170, 171]. These extensive measurement requirements pose a bottleneck for applying variational quantum algorithms in scenarios requiring high-precision results, such as in fields like computational chemistry, where a precision of 10^{-3} Hartree [172] (chemical accuracy) is often the target.

Recent strategies like overlapping tomography [164, 168, 173] have emerged to minimize the total number of measurement settings while maintaining high precision. These methods achieve this by parallelizing measurements on non-overlapping subsystems and efficiently organizing information from measurements on overlapping subsystems. However, current methods, while promising, often overlook higher-order correlations and compatibility constraints among RDMs, potentially limiting their accuracy.

In this chapter, we propose a novel hierarchy of data-driven semidefinite programs (SDPs) to estimate a set of overlapping reduced density matrices (RDMs) from quantum measurements. Our approach focuses specifically on random Pauli string measurements of n -qubit states with fixed locality. We leverage the inherent higher-order correlations present in quantum state measurement data—information that would be lost when estimating overlapping RDMs independently. This enables us to tighten the uncertainty intervals of the RDM estimates for a given number of measurement shots, with particular advantages in low-shot regimes. Such SDP relaxations are built on re-imposing partial compatibility with the quantum marginal problem [171, 174, 175]. We further constrain each RDM to satisfy the physical requirements of a valid density matrix, namely unit trace and positive semidefiniteness. This comprehensive approach yields two key benefits: it resolves the compatibility issues between overlapping RDMs that arise in linear inversion [176], while simultaneously enhancing the global consistency of the entire set of RDMs.

Our method demonstrates superior performance compared to standard tomography methods across two numerical benchmarks. In the first evaluation, we assess the ability of our method to estimate ground-state RDMs and energies of the 1D XY model with open boundary conditions by using random Pauli string measurements. Under the same measurement budget, our approach achieves more accurate energy estimates compared to conventional tomography methods. We further validate our technique through application to algorithmic cooling [177–179], a practical use case in near-term quantum computing where RDMs inform quantum circuit design. In this application domain, our method again demonstrates measurable advantages over traditional approaches that rely on independent RDM reconstructions.

This chapter is organized as follows: Sec. 4.2 introduces quantum tomography,

semidefinite programming, and related literature. Sec. 4.3 presents the proposed SDP-based reconstruction method and numerical results. In Sec. 4.4, we apply our method to algorithmic cooling. Finally, Sec. 4.5 provides conclusions and future directions.

4.2 Background and preliminaries

4.2.1 Notations

We consider an n -qubit system whose Hilbert space is $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$. For our mathematical representation, we outline the relevant definitions and notations below. Let $\sigma_{0,1,2,3}$ denote, respectively, the identity matrix \mathbb{I} and the Pauli X , Y , Z matrices. For an n -qubit system, the identity matrix is denoted by \mathbb{I}_n . A vector $\mathbf{i} = (i_1, i_2, \dots, i_n) \in \{0, 1, 2, 3\}^n := \mathcal{I}_n$ is used to represent a specific Pauli basis $\sigma_{\mathbf{i}}$, which is a shorthand for the Pauli string $\sigma_{i_1} \otimes \sigma_{i_2} \otimes \dots \otimes \sigma_{i_n}$, in $\{\mathbb{I}, X, Y, Z\}^n$. For simplicity, we use this notation throughout the text. Note that Pauli strings form an orthogonal basis of the \mathbb{R} -vector space of Hermitian matrices. We denote the set $\{1, \dots, m\}$ by $[m]$. For a Hermitian matrix A , the notation $A \succeq 0$ indicates that A is positive semidefinite, i.e., all its eigenvalues are non-negative. We denote by ρ_{AB} the density matrix that describes a quantum state in the Hilbert space $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$, where $A = \{a_1, \dots, a_{|A|}\}$ and $B = \{b_1, \dots, b_{|B|}\}$ label distinct sets of qubits in the respective subsystems. The partial trace over subsystem A is defined as:

$$\text{Tr}_A [\rho_{AB}] := \sum_k (\langle k | \otimes \mathbb{I}_B) \rho_{AB} (|k \rangle \otimes \mathbb{I}_B), \quad (4.1)$$

where $\{|k\rangle\}$ forms an orthonormal basis of the Hilbert space \mathcal{H}_A .

4.2.2 Quantum state tomography

Quantum state tomography [9, 160, 161] is a cornerstone of quantum information science, enabling the reconstruction of quantum states through systematic measurements on an ensemble of identical quantum states. For an n -qubit quantum state, its density matrix, represented by ρ , can be fully characterized by the following relation:

$$\rho = \frac{1}{2^n} \left(\sum_{\mathbf{i} \in \mathcal{I}_n} C_{\mathbf{i}} \sigma_{\mathbf{i}} \right), \quad (4.2)$$

$$C_{\mathbf{i}} = \text{Tr}[\rho \sigma_{\mathbf{i}}]. \quad (4.3)$$

Here, $C_{\mathbf{i}}$ constitutes an element of the corresponding n -qubit Bloch vector, with $C_{\mathbf{0}} = 1$. In practical settings, $C_{\mathbf{i}}$ is not directly accessible, instead an estimate $\hat{C}_{\mathbf{i}}$ is computed by averaging over a finite number of measurements $N_{\text{meas},\mathbf{i}}$ performed in the basis $\sigma_{\mathbf{i}}$:

$$\hat{C}_{\mathbf{i}} = \frac{1}{N_{\text{meas},\mathbf{i}}} \sum_{k=1}^{N_{\text{meas},\mathbf{i}}} m_{\mathbf{i}}^{(k)}, \quad (4.4)$$

where $m_{\mathbf{i}}^{(k)} \in \{1, -1\}$ represents the k -th measurement outcome. In measuring a particular Pauli string $\sigma_{\mathbf{i}}$, one obtains a binary string $s_{\sigma_{\mathbf{i}}} = \{0, 1\}^n$, from which the outcome

$$m_{\mathbf{i}}^{(k)} = (-1)^{|s_{\sigma_{\mathbf{i}}}|}, \quad (4.5)$$

is calculated, with $|s_{\sigma_{\mathbf{i}}}|$ denoting the string Hamming weight, defined as the total count of 1s in $s_{\sigma_{\mathbf{i}}}$. We denote the reconstructed state as $\hat{\rho}$, characterized by the $4^n - 1$ different $\{\hat{C}_{\mathbf{i}}\}_{\mathbf{i} \in \mathcal{I}_n \setminus \{\mathbf{0}\}}$. Due to redundancies [180], it is sufficient to simulate measurements in 3^n different bases, each corresponding to a specific Pauli string $\sigma_{\mathbf{i}}$ with $\mathbf{i} \in \{1, 2, 3\}^n$, to estimate all $\hat{C}_{\mathbf{i}}$. In essence, if a string contains a 0, then the corresponding qubits are not measured, which is equivalent to measuring the RDM of the complementary subsystem. The measurement data collected from the other measurements is sufficient for reconstructing this RDM. The reconstructed state $\hat{\rho}$ takes the form:

$$\hat{\rho} = \frac{1}{2^n} \left(\mathbb{I}_n + \sum_{\mathbf{i} \in \mathcal{I}_n \setminus \{\mathbf{0}\}} \frac{1}{N_{\text{meas},\mathbf{i}}} \sum_{k=1}^{N_{\text{meas},\mathbf{i}}} m_{\mathbf{i}}^{(k)} \sigma_{\mathbf{i}} \right). \quad (4.6)$$

We will refer to this method as the *standard quantum state tomography*. According to the Chernoff-Hoeffding bound, achieving an additive error ϵ in each $\hat{C}_{\mathbf{i}}$ w.r.t. the true values $C_{\mathbf{i}}$ requires on the order of $N_{\text{meas}} \approx 4 \log(2)n/\epsilon^2$ shots [168] for a constant failure probability. If each $\hat{C}_{\mathbf{i}}$ is estimated with N_{meas} shots, then $3^n N_{\text{meas}}$ total measurements are required, because there are 3^n non-identity Pauli strings, rendering this standard quantum state tomography infeasible for larger n . For instance, for a 50-qubit system this would result in $3^{50} \approx 7.18 \times 10^{23}$ measurement settings, and even with the Google Sycamore chip's $4\mu\text{s}$ [157] readout time, performing only one shot per setting would take about 10^{10} years.

To ease this exponential scaling in the number of qubits n , various methods have been proposed. Most notably, classical shadow tomography [163] provides a way to extract certain few-body observables with a number of measurements that grows only logarithmically in the number of those observables. Other techniques, such as neural-network-based tomography [181], rely on sampling

protocols and often presume certain structural features of the underlying system. Using global mutually unbiased bases on the full Hilbert space, rather than local Pauli product bases, the number of measurement settings required for quantum state tomography can be reduced to $2^n + 1$ [182, 183]. Furthermore, informationally complete measurements such as symmetric informationally complete positive operator-valued measures (SIC-POVMs) offer an alternative that can reduce the number of measurement settings, potentially down to a single setting, and can be combined with classical-shadow-based tomography schemes [184].

In Sec. 4.2.4, we will present the overlapping tomography method of [168], which can be applied for quantum state tomography of RDMs.

4.2.3 Semidefinite programming

A semidefinite programming (SDP) problem involves optimizing a linear function subject to constraints expressed as linear matrix equalities and inequalities. The feasible regions of an SDP are known as spectrahedra, and efficient algorithms such as interior point methods [185, 186] can be used to solve them. In its standard primal form, an SDP can be written as:

$$\begin{aligned} \min \quad & \langle C, X \rangle \\ \text{s.t.} \quad & \langle A_i, X \rangle = b_i, \quad i = 1, 2, \dots, m, \\ & X \succeq 0. \end{aligned} \tag{4.7}$$

where the cost matrix C , the decision matrix X , and the constraint matrices A_i are Hermitian, b is a real vector, and $\langle \cdot, \cdot \rangle$ denotes the Hilbert-Schmidt inner product, defined as $\langle X, Y \rangle := \text{Tr}[X^\dagger Y]$, with X^\dagger being the Hermitian conjugate of X . SDPs are particularly well-suited for quantum information [174], as density matrices are positive semidefinite. For instance, the quantum marginal problem seeks to determine whether a global state on a full system exists that is consistent with specified marginal states on subsystems, represented as reduced density matrices (RDMs) [171, 174, 175]. This problem can also be formulated as an SDP, as illustrated by the following example.

Assume we have perfect knowledge of the RDMs ρ_{23} , ρ_{13} , and ρ_{12} , acting on subsystems \mathcal{H}_{23} , \mathcal{H}_{13} , and \mathcal{H}_{12} , respectively, and the goal is to find a global

state ρ_{123} on \mathcal{H}_{123} that is consistent with these RDMs:

$$\begin{aligned} & \min_{\rho_{123}} 0 \\ \text{s.t. } & \rho_{123} \succeq 0, \quad \text{Tr}[\rho_{123}] = 1, \\ & \text{Tr}_1[\rho_{123}] = \rho_{23}, \\ & \text{Tr}_2[\rho_{123}] = \rho_{13}, \\ & \text{Tr}_3[\rho_{123}] = \rho_{12}. \end{aligned} \tag{4.8}$$

Here, the minimization of the constant 0 is purely a formalism to match the standard SDP form in Eq. (4.7). If the constraints cannot be satisfied, the SDP is deemed infeasible. Despite the apparent simplicity of compatibility SDPs like Eq. (4.8), the general problem of determining whether a state $\rho_{123\dots n}$ exists is QMA-complete [170, 171], making it computationally intractable even for quantum computers.

To address this complexity, we relax the problem in our method by using standard quantum state tomography estimates and accounting for their error ranges. Additionally, reinforcement learning has been shown to improve the selection of constraints that enforce compatibility with the quantum marginal problem in reconstructed solutions [187].

4.2.4 Overlapping tomography and related literature

The standard quantum state tomography of a full quantum state requires a number of measurement settings that scales as 3^n , where n is the number of qubits [180]. However, if the goal is to estimate only the $\binom{n}{k}$ different k -qubit reduced density matrices (RDMs), each RDM requires only 3^k different measurement settings. The total number of measurements, however, includes an additional overhead due to the number of shots needed to estimate each measurement setting.

Several strategies have been proposed to optimize the selection of measurement settings and minimize the total number of measurements required. In [168], a method called *quantum overlapping tomography* was introduced for estimating k -qubit RDMs. The key idea is to measure each qubit individually in a chosen basis, allowing non-overlapping k -qubit subsystems to be measured in parallel. Overlapping k -qubit subsystems are then reconstructed through classical post-processing. This approach reduces the required number of measurement settings to at most $e^{\mathcal{O}(k)} \log^2(n)$. The measurement selection is determined by a family of hash functions (n, k) , which partition the system into k subsystems, with all qubits in each subsystem measured in the same basis. All k -qubit RDMs can then be reconstructed from this carefully chosen dataset. In a similar way, partitioning-based methods [188] lead to a scaling of the number of measurement settings needed as $\mathcal{O}(3^k \log^{k-1} n)$.

However, in many practical scenarios, such as systems with k -local Hamiltonians and nearest-neighbor interactions, only k -geometrically-local RDMs—those involving neighboring qubits—are required. This further reduces the number of measurement settings needed to a constant scaling with respect to the number of qubits n [164]. Recent advances have shown, both theoretically and experimentally, that leveraging graph theory can optimize this method, allowing even k -qubit RDMs to be estimated with a number of measurement settings that remains constant in the number of qubits n [173].

Other approaches have focused on reducing the statistical uncertainty due to shot noise in various contexts, such as quantum state tomography [189], ground-state estimation [190], and calculations of fidelity and von Neumann entropy [191]. SDP has also been applied to quantum marginal problems for tasks like ground-state estimation [192, 193], incorporating tensor network methods [194], or imposing entropy constraints on RDMs [195].

Our method builds on the principle of measuring the entire quantum system via product measurements on single qubits in the Pauli bases and performing post-processing to estimate k -local RDMs. However, rather than focusing on optimizing the selection of measurement settings, we introduce physicality and consistency constraints on the estimated RDMs. This approach explicitly accounts for the shot noise arising from a limited number of shots per measurement setting, providing a clearer benchmark for practical applications.

4.3 SDP-assisted overlapping tomography

In this section, we introduce the SDP-assisted overlapping tomography framework for estimating local RDMs of n -qubit states described by local Hamiltonians. We begin by presenting the hypergraph representation of local Hamiltonians in Sec. 4.3.1. In Sec. 4.3.2, we detail the core methodology of SDP-assisted overlapping tomography. Finally, to illustrate its effectiveness, we apply the approach to the ground states of the 1D XY model and discuss the numerical results in Sec. 4.3.3.

4.3.1 Local Hamiltonian

The Hamiltonian of an n -qubit k -local system can be represented by a hypergraph $G = (V_G, E_G)$, where each vertex $v_i \in V_G$ corresponds to a qubit, and each hyperedge $e_j \in E_G$ (connecting up to k vertices) represents a local Hamiltonian term. Concretely, the Hamiltonian takes the form

$$H = \sum_{j=1}^m H_j. \quad (4.9)$$

where m is the number of local subsystems, growing polynomially with n , and each H_j acts non-trivially on at most k qubits (i.e., it is k -local). Fig. 4.1 illustrates the hypergraph representation of a k -local Hamiltonian in a many-body system.

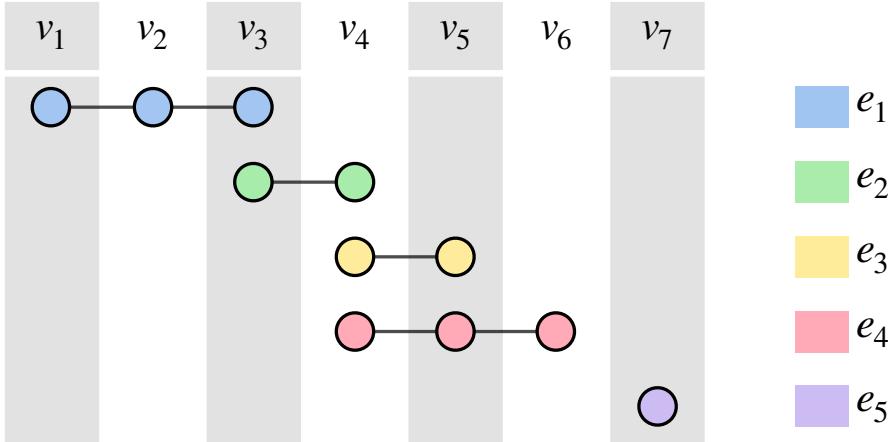


Figure 4.1: Hypergraph representation of the interaction structure of a Hamiltonian $H = \sum_{j=1}^m H_j$. The vertices $V_G = \{v_i\}_{i=1}^n$ correspond to $n = 7$ qubits. Each hyperedge $e_j \in E_G = \{e_j\}_{j=1}^m$ denotes a local Hamiltonian term H_j . In this example, E_G contains the sets $\{v_1, v_2, v_3\}$, $\{v_3, v_4\}$, $\{v_4, v_5\}$, $\{v_4, v_5, v_6\}$, $\{v_7\}$, hence $m = 5$. Hyperedges are shown as horizontal lines linking the relevant vertices, and the corresponding local Hamiltonian term is indicated in the legend on the right.

The hypergraph representation not only provides a visualization of the Hamiltonian structure but also gives a framework for reconstructing the relevant RDMs. Once measurement results are obtained from an ensemble of identical n -qubit states, one can reconstruct a collection of RDM estimates $\hat{\rho}_j$ for each local subsystem. This involves applying Eq. (4.6) to each of the m subsystems:

$$\left\{ \hat{\rho}_j := \frac{1}{2^{|e_j|}} \left(\mathbb{I}_{|e_j|} + \sum_{\mathbf{i} \in \mathcal{I}_{|e_j|} \setminus \{\mathbf{0}\}} \hat{C}_{\mathbf{i}}^j \sigma_{\mathbf{i}} \right) \right\}_{j=1}^m. \quad (4.10)$$

Here, $\hat{C}_{\mathbf{i}}^j$ is the estimator of $C_{\mathbf{i}}^j = \text{Tr} [\text{Tr}_{V_G \setminus e_j} [\rho] \sigma_{\mathbf{i}}] = \text{Tr} [\rho_j \sigma_{\mathbf{i}}]$, where ρ is the true n -qubit global state and ρ_j is the true RDM of the j -th subsystem. Note that, each hyperedge $e_j = \{v_{j_1}, v_{j_2}, \dots\}$ corresponds to the qubits on which H_j acts. These local RDMs are of particular interest because they are sufficient to

estimate the energy expectation values of the local Hamiltonian:

$$\hat{E} = \sum_{j=1}^m \text{Tr}[\hat{\rho}_j H_j]. \quad (4.11)$$

However, the reconstructed local RDMs $\{\hat{\rho}_j\}_{j=1}^m$ can be non-physical in practice; e.g., they may exhibit negative eigenvalues or fail to satisfy mutual compatibility (i.e., overlapping RDMs might not be consistent). While mutual compatibility can only be enforced approximately in real scenarios [187], imposing physicality constraints remains crucial for ensuring reliable energy estimates and other derived properties.

4.3.2 Methodology

In this section, we present the SDP-assisted overlapping tomography framework. Building on standard quantum tomography results, we formulate semidefinite programs (SDPs) that incorporate overlapping-compatibility (OC) and enhanced-compatibility (EC) constraints. These constraints aim to ensure the validity and consistency of the reconstructed states, allowing us to minimize or maximize the energy expectation within the feasible set. In essence, we determine the minimum and maximum energies compatible with the simulated data. By solving these SDP problems, we obtain a set of overlapping local RDMs that are more physically valid and mutually consistent.

Semidefinite relaxations

Relaxations of polynomial optimization problems based on semidefinite constraints play a central role in our method, in a similar spirit as Ref. [192, 196]. In this section, we describe the construction of these constraints.

Consider a local Hamiltonian associated with hypergraph $G = (V_G, E_G)$. For each hyperedge $e_j \in E_G$, we introduce the decision variables for the SDP problem as

$$\left\{ \tilde{\rho}_j := \frac{1}{2^{|e_j|}} \left(\mathbb{I}_{|e_j|} + \sum_{\mathbf{i} \in \mathcal{I}_{|e_j|} \setminus \{\mathbf{0}\}} \tilde{C}_{\mathbf{i}}^j \sigma_{\mathbf{i}} \right) \right\}_{j=1}^m, \quad (4.12)$$

where

$$\tilde{C}_{\mathbf{i}}^j \in \left[\hat{C}_{\mathbf{i}}^j - \epsilon_{\mathbf{i}}^j, \hat{C}_{\mathbf{i}}^j + \epsilon_{\mathbf{i}}^j \right], \quad \forall j \in [m], \forall \mathbf{i} \in \mathcal{I}_{|e_j|} \setminus \{\mathbf{0}\}. \quad (4.13)$$

Here, index j labels the local subsystem associated with the hyperedge e_j , and $\epsilon_{\mathbf{i}}^j$ is a relaxation variable associated to j -th local subsystem. It is proportional

to the variance of $\hat{C}_{\mathbf{i}}^j$, scaled by a coefficient α , i.e., $\epsilon_{\mathbf{i}}^j = \alpha \text{Var}(\hat{C}_{\mathbf{i}}^j)$. The value of α relates to the probability that the unbiased estimator $\hat{C}_{\mathbf{i}}^j$ (of $\text{Tr}[\rho_j \sigma_{\mathbf{i}}]$) lies within the chosen confidence interval, following Chernoff bound arguments. Throughout, tildes ($\tilde{\cdot}$) denote SDP decision variables, while hats ($\hat{\cdot}$) denote results obtained directly from the raw data during standard quantum state tomography.

Moreover, we introduce the semidefinite constraints

$$\tilde{\rho}_j \succeq 0, \quad \forall j \in [m], \quad (4.14)$$

and the overlapping-compatibility (OC) and enhanced-compatibility (EC) constraints

$$\mathcal{R} := \bigcup_{j, j' \in [m]} \mathcal{R}(\tilde{\rho}_j, \tilde{\rho}_{j'}), \quad (4.15)$$

$$\mathcal{G} := \bigcup_{j, j' \in [m]} \mathcal{G}(\tilde{\rho}_j, \tilde{\rho}_{j'}), \quad (4.16)$$

respectively. Specifically, $\mathcal{R}(\tilde{\rho}_j, \tilde{\rho}_{j'})$ is the set of matrix equality constraints ensuring that the local RDMs are consistent on the intersection of their respective supports:

$$\mathcal{R}(\tilde{\rho}_j, \tilde{\rho}_{j'}) := \{\text{Tr}_{e_j \setminus e_{j'}}[\tilde{\rho}_j] - \text{Tr}_{e_{j'} \setminus e_j}[\tilde{\rho}_{j'}] = 0\}, \quad (4.17)$$

while $\mathcal{G}(\tilde{\rho}_j, \tilde{\rho}_{j'})$ is the set of constraints that require the existence of a larger RDM $\tilde{\rho}$ whose marginals agree on the variables indexed by j and j' :

$$\mathcal{G}(\tilde{\rho}_j, \tilde{\rho}_{j'}) := \{\tilde{\rho} \succeq 0, \text{Tr}_{e_j}[\tilde{\rho}] = \tilde{\rho}_{j'}, \text{Tr}_{e_{j'}}[\tilde{\rho}] = \tilde{\rho}_j\}. \quad (4.18)$$

Note that the dimension of $\tilde{\rho}$ can freely be chosen and that the definition of \mathcal{G} in Eq. (4.18) is just one possible choice; in practice, there is a hierarchy of possible relaxations forming a partially ordered set. The optimal choice of \mathcal{G} for a given computational budget is non-trivial, and one can achieve varying degrees of performance by judiciously selecting or refining these constraints [187, 197].

SDP problem formulation

The relaxations defined in the previous sub-section allow us to specify a feasible set of local RDMs. In this chapter, we focus on finding the minimum and maximum energies consistent with the simulated data used to estimate the tomographic local RDMs $\{\hat{\rho}_j\}$. Specifically, we consider the following optimization problem for determining the optimal state that minimizes (or maximizes)

the energy under these constraints:

$$\begin{aligned}
 & \min(\max)_{\{\tilde{\rho}_j\}} \quad \sum_j \text{Tr} [\tilde{\rho}_j H_j] \quad (4.19) \\
 \text{s.t. } & \tilde{C}_{\mathbf{i}}^j \in \left[\hat{C}_{\mathbf{i}}^j - \epsilon_{\mathbf{i}}^j, \hat{C}_{\mathbf{i}}^j + \epsilon_{\mathbf{i}}^j \right], \quad \forall j \in [m] \text{ and } \forall \mathbf{i} \in \mathcal{I}_{|e_j|} \setminus \{\mathbf{0}\}, \\
 & \tilde{\rho}_j \succeq 0, \quad \forall j \in [m], \\
 & \mathcal{R} \cup \mathcal{G}.
 \end{aligned}$$

Note that Eq. (4.19) does not yet incorporate all the SDP constraints derived from the quantum marginal problem needed to solve the above problem, as it only includes the initial-order EC constraints. Obtaining the exact solution requires higher-order EC constraints, but adding these leads to exponential growth in computational complexity [187, 197]. In large many-body systems, our formulation effectively discards higher-order constraints as n grows but k remains fixed, which loosens the energy bounds. Depending on the specific system, one may selectively include higher-level constraints that are most relevant for the problem at hand, here being the estimation of the minimum and maximum energy compatible with the measurement data [194]. Such strategies (especially the reinforcement-learning-based [187] and the renormalization-based [194] ones) can naturally be incorporated in our framework thereby allowing for even tighter estimates, but this would impair the fairness of the benchmarking of our method. Hence, in the rest of the chapter we proceed with the straightforward approach.

Because shot noise affects the data, setting $\epsilon_{\mathbf{i}}^j = 0$ (i.e., assuming no error) typically makes the SDP infeasible. Allowing a tolerance $\epsilon_{\mathbf{i}}^j$ around the estimated quantities provides a search region for solutions. As more data is collected, the ϵ -values shrink, thereby reducing the volume of the feasible set. Since the true state arises from a valid quantum system (satisfying the quantum marginal problem at all hierarchy levels), the resulting energy bounds become more accurate with increased data.

Studies on ground-state scenario

Semidefinite relaxations have long been used to investigate ground-state properties of many-body systems [192, 196]. We now examine the performance of SDP-assisted overlapping tomography in estimating ground states of local Hamiltonians.

Fig. 4.2 illustrates two feasible sets under the same search region but with or without EC constraints. These feasible sets are defined by all RDMs satisfying both the overlapping- (\mathcal{R}) and enhanced-compatibility (\mathcal{G}) constraints or only the overlapping- (\mathcal{R}) constraints, respectively, while the search region refers to the intersection of positive semidefinite cones restricted by the confidence intervals in Eq. (4.13). Introducing EC constraints narrows the feasible set and

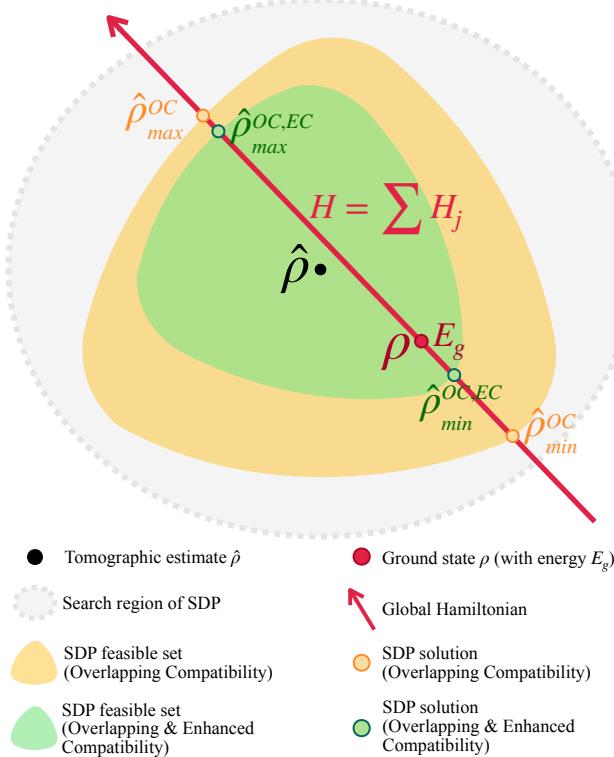


Figure 4.2: Feasible set for the SDP problem. An n -qubit system with RDM ρ and energy E_g (shown in red) undergoes tomography, resulting in an approximate RDM $\hat{\rho}$ with energy \hat{E} (black). The shaded gray region represents the (high-dimensional) search space. The yellow spectrahedron depicts the feasible set with only overlapping-compatibility (OC) constraints, whereas the green spectrahedron also includes enhanced-compatibility (EC).

yields tighter lower and upper bounds on ground-state energy estimate \hat{E}_g :

$$\text{Tr} [H \hat{\rho}_{\min}^{\text{OC,EC}}] \geq \text{Tr} [H \hat{\rho}_{\min}^{\text{OC}}], \quad (4.20)$$

$$\text{Tr} [H \hat{\rho}_{\max}^{\text{OC,EC}}] \leq \text{Tr} [H \hat{\rho}_{\max}^{\text{OC}}]. \quad (4.21)$$

The variable $\epsilon_i^j = \alpha \text{Var}(\hat{C}_i^j)$ in Eq. (4.13) defines the search region explored by the SDP, thereby controlling the confidence interval for the ground-state energy estimate. Different α -values yield different search regions, visualized in Fig. 4.3. The spectrahedra illustrates the feasible set after imposing positive semidefiniteness, OC, and EC constraints. The true state ρ (red) and its

estimated counterpart $\hat{\rho}$ (black) are included for comparison. Crucially, when the search region is large enough within the feasible set to capture the true RDMs ρ , we obtain upper and lower bounds on the ground energy estimate \hat{E}_g :

$$\mathrm{Tr}[H \hat{\rho}_{\min}^{\mathrm{SDP}}(\alpha_0)] \leq \hat{E}_g \leq \mathrm{Tr}[H \hat{\rho}_{\max}^{\mathrm{SDP}}(\alpha_1)], \quad (4.22)$$

where $\alpha_0 > \alpha_1$.

Algorithm 4.1: SDP-assisted tomography on ground states

Input: T identical ground-state preparations, Hamiltonian $H = \sum_{j=1}^m H_j$, hyperparameters $b > \Delta > 0$, where Δ is the tolerance.

Output: Local RDM estimates $\{\hat{\rho}_j^{\mathrm{SDP}}\}_{j=1}^m$ that minimize the energy, and the corresponding minimum energy \hat{E}_g .

Perform standard quantum state tomography to obtain estimated local RDMs $\{\hat{\rho}_j\}_{j=1}^m$ with the precision given by T samples;

repeat

Solve the SDP problem described by Eq.(4.19) with decision variables $\{\tilde{\rho}_j\}_{j=1}^m$ and coefficient $\alpha = b$;

Denote the feasible set as \mathcal{F} ;

$b \leftarrow 2b$;

until $\mathcal{F} \neq \emptyset$;

Initialize $a = 0$;

while $|a - b| \geq \Delta$ **do**

$\alpha_0 \leftarrow a + \frac{|a-b|}{2}$;

Solve the SDP problem described by Eq.(4.19) with decision variables $\{\tilde{\rho}_j\}_{j=1}^m$ and coefficient $\alpha = \alpha_0$;

Denote the feasible set as \mathcal{F} ;

if $\mathcal{F} \neq \emptyset$ **then**

Denote the solution as $\mathcal{S} = \arg \min_{\{\tilde{\rho}_j\}} \sum_j \mathrm{Tr}[\tilde{\rho}_j H_j]$;

$b \leftarrow \alpha_0$;

else

$a \leftarrow \alpha_0$;

$\{\hat{\rho}_j^{\mathrm{SDP}}\}_{j=1}^m \leftarrow \mathcal{S}$;

$\hat{E}_g \leftarrow \sum_{j=1}^m \mathrm{Tr}[\hat{\rho}_j^{\mathrm{SDP}} H_j]$;

We employ a bisection method (Alg. 4.1) to identify the smallest α_0 (within a tolerance Δ_0) for which the feasible set is non-empty, and analogously for α_1 using a maximization variant. This allows us to establish the bounds of Eq. (4.22) on \hat{E}_g .

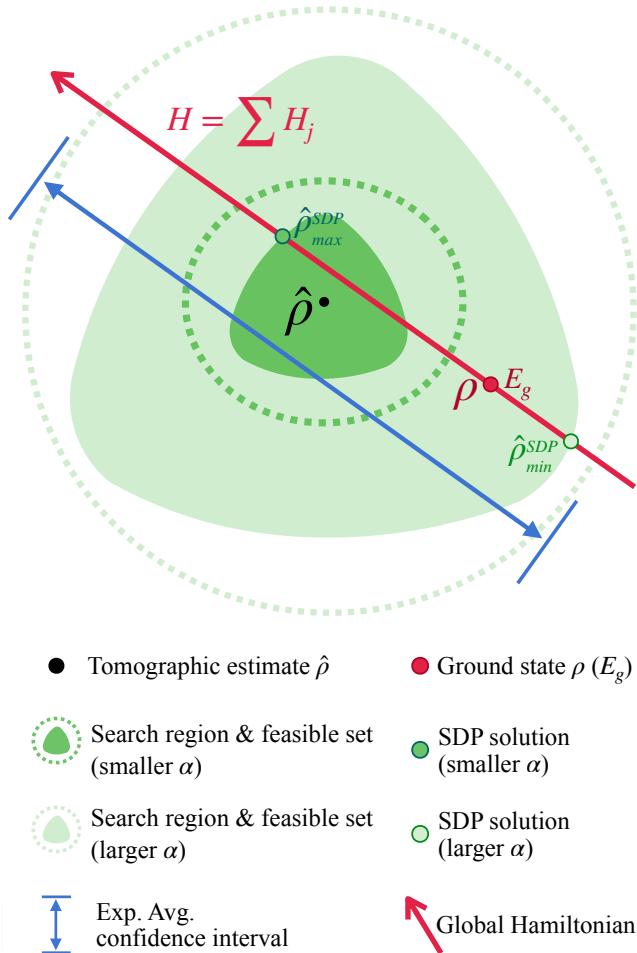


Figure 4.3: An illustrative depiction of the SDP feasibility region. The true RDM ρ with energy E_g of Hamiltonian H (red) and its tomographic estimate $\hat{\rho}$ (black) lie within a search region of adjustable size (dashed circle) determined by the coefficient α . The feasible set under overlapping-compatibility (OC) and enhanced-compatibility (EC) constraints appears as a shaded spectrahedron. The blue interval indicates the confidence interval for the average of the energy obtained by standard quantum state tomography from the simulated data.

4.3.3 Numerical simulations

In this section, we present numerical results for estimating the ground-state energy of the 1D-chain XY model [198]. We also compare the estimation

accuracy achieved by standard tomography to that of the proposed SDP-assisted tomography.

Problem description

The Hamiltonian of the 1D-chain XY model is given by

$$H_C = \sum_{j=1}^{n-1} H_j = J \sum_{j=1}^{n-1} (X_j X_{j+1} + Y_j Y_{j+1}). \quad (4.23)$$

where H_j is the local two-qubit Hamiltonian term capturing the interaction between qubits j and $j+1$. Here, n denotes the total number of qubits, and J is the interaction strength. The subscript 'C' stands for 'chain'. This model is frustrated, meaning that the global ground state of H_C does not simultaneously minimize the energy of each local term H_j .

Although the accuracy of ground-state energy estimation through tomography is fundamentally limited by the number of ground-state samples, introducing additional constraints can significantly tighten the resulting confidence interval. In the context of the XY model H_C , the overlapping- (OC) and enhanced-compatibility (EC) constraints in Eq. (4.19) take the form

$$\mathcal{R}_C := \cup_{j \in [n-1]} \mathcal{R}_C(\tilde{\rho}_j, \tilde{\rho}_{j+1}), \quad (4.24)$$

$$\mathcal{G}_C := \cup_{j \in [n-2]} \mathcal{G}_C(\tilde{\rho}_j, \tilde{\rho}_{j+1}), \quad (4.25)$$

respectively, with

$$\mathcal{R}_C(\tilde{\rho}_j, \tilde{\rho}_{j+1}) := \{\text{Tr}_{\{j\}}[\tilde{\rho}_j] - \text{Tr}_{\{j+2\}}[\tilde{\rho}_{j+1}] = 0\}, \quad (4.26)$$

$$\mathcal{G}_C(\tilde{\rho}_j, \tilde{\rho}_{j+1}) := \{\tilde{\rho} \succeq 0, \text{Tr}_{\{j\}}[\tilde{\rho}] = \tilde{\rho}_{j+1}, \text{Tr}_{\{j+2\}}[\tilde{\rho}] = \tilde{\rho}_j\}. \quad (4.27)$$

Here, $\tilde{\rho}_j$ is the SDP decision variable corresponding to the two-qubit subsystem of qubits $(j, j+1)$, and $\tilde{\rho}$ is a three-qubit SDP variable spanning qubits $j, j+1$, and $j+2$. The subscript in the partial trace indicates which qubit is being traced out.

Note that the energy function can be computed from the 2-qubit RDMs in a natural way, but also equivalently from the 3-qubit RDMs that stem from the ECs, either by taking their partial traces (which are compatible with the 2-qubit RDMs, by construction) or by extending the Hamiltonian terms to a larger Hilbert space (by appropriately tensoring them with identity operators).

Explanation on the results

For the SDP-assisted tomography, we are firstly given a number of identical copies of the XY model ground state, as in Eq. (4.23). We perform a series

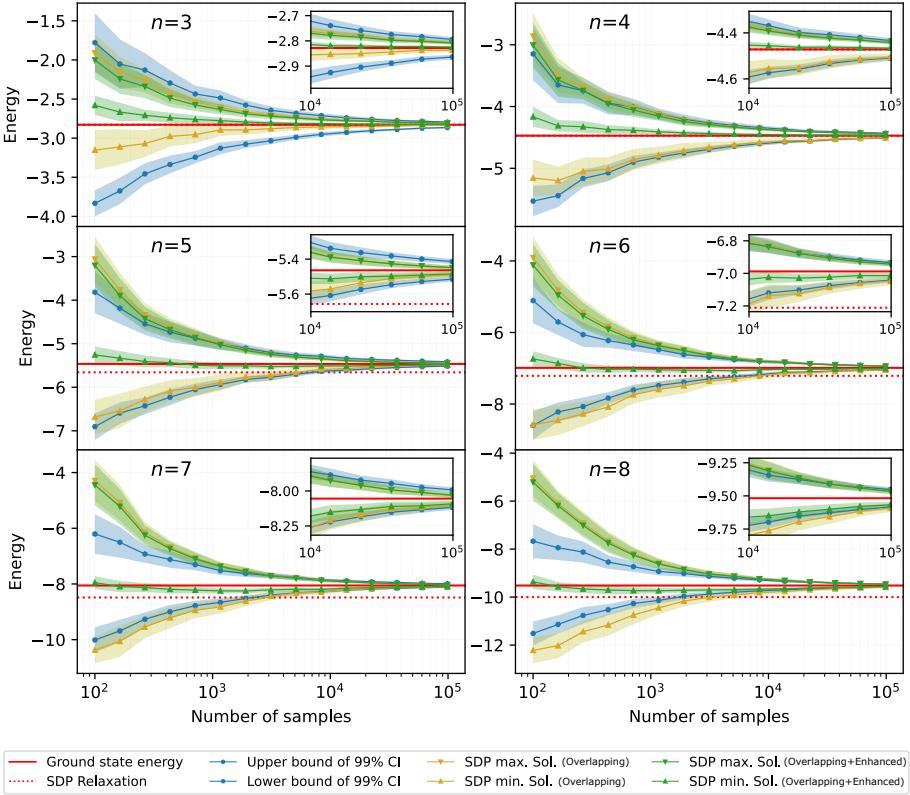


Figure 4.4: Tightening expectation value confidence interval via SDP-assisted quantum tomography. Ground-state energy estimates for the XY model ($J = 1$) with $n = 3, 4, 5, 6, 7, 8$ qubits against the number of state samples. The red solid lines indicate the exact ground-state energy; the red dashed lines are lower bounds from the relaxation method. The blue curves show the 99% confidence interval upper and lower bounds for standard quantum tomography. The yellow (green) curves give the results of SDP minimization with tolerance parameters $\Delta_0 = 0.1$, and maximization with tolerance parameters $\Delta_0 = 0.001$, under OC (OC+EC) constraints, respectively. In each plot, the two green curves can be interpreted as the upper and lower energy bounds using SDP-enhanced quantum tomography.

of measurements in randomly chosen Pauli basis σ_i , with i uniformly sampled from $\{1, 2, 3\}^n$. From these measurement outcomes, we reconstruct the local RDMs via standard tomography and estimate the ground-state energy through Eq. (4.11). We then apply our SDP-based post-processing to refine these ground-state energy estimates.

The plots in Fig. 4.4 show the resulting confidence intervals for the ground-

state energy as a function of the number of samples, for $n = 3, 4, 5, 6, 7, 8$. In each subplot, the red solid line indicates the exact ground-state energy from Eq. (4.23) obtained by exact diagonalization, while the blue region marks the 99% confidence interval from standard tomography. The green curves represent the solutions of Alg. 4.1 applied to SDP minimization and maximization (with OC+EC constraints, Eqs. (4.26) and (4.27)) and two different tolerances $\Delta_0 = 0.1$ and $\Delta_1 = 0.001$. This choice is motivated by the fact that there is a strong asymmetry between upper and lower bounds when tomographing the ground state. In contrast, the yellow curves illustrate the SDP results with only OC constraints (Eq. (4.26)).

Comparing the yellow and green curves marked with triangles shows that adding EC constraints substantially tightens the lower bound obtained by SDP minimization. When the sample size is small, the tomographic RDM $\hat{\rho}$ can significantly diverge from the true RDM ρ , creating bias in the SDP energy bound in the presence of large statistical noise. Together with the fact that the probability of failure in estimating the true ground state is higher in extremely-low-shot regimes, this fact makes it much more likely for the green curves marked with triangles to exceed the exact ground-state energy. However, as the number of samples grows, the SDP-assisted approach narrows the confidence interval more effectively than standard tomography. This improvement is especially clear in the zoomed-in subplots, where the green curve, representing SDP-assisted (OC+EC) bounds, spans a narrower range for $\langle H \rangle$ than the blue curve from standard quantum state tomography. Notably, SDP-assisted tomography can reduce the required number of samples by a factor of 10^1 to 10^2 compared to standard quantum state tomography, while achieving the same level of precision in lower-bounding the energy.

The red dashed lines represent lower bounds obtained via an existing relaxation method [196]. Such approaches optimize the ground-state energy under only a subset of the constraints that we use, thus guaranteeing a strict, albeit potentially conservative, bound. While such relaxations have successfully established lower bounds for the ground-state energies of local Hamiltonians in many-body settings [187, 194, 196, 199–202], the accuracy of these bounds tends to degrade for larger problem sizes. This shortfall arises from omitting higher-level EC constraints to keep the computational effort manageable.

All numerical simulations were performed using Qiskit [203] for quantum simulations and CVXPY [204, 205] to solve SDP formulations using the SCS optimizer.

4.4 Application

In this section, we demonstrate how to integrate our SDP-assisted tomography into a variational procedure for preparing and characterizing low-energy

states of local Hamiltonians. Specifically, we embed SDP-assisted tomography within the *algorithmic cooling* (AC) method [177–179] to heuristically minimize the energy of a target quantum Hamiltonian. Sec. 4.4.1 introduces AC, and Sec. 4.4.2 explains how SDP-assisted tomography is incorporated into the AC workflow. Numerical results for the 1D-chain XY model are presented in Sec. 4.4.3, comparing the performance of AC both with and without semidefinite programming.

4.4.1 Algorithmic cooling

Algorithmic cooling (AC) aims to prepare a low-energy state $|\psi\rangle$ of an n -qubit quantum system governed by a local Hamiltonian H , using near-term quantum devices. The AC method is adapted to the practical capabilities of the experimental setup, which can differ considerably across platforms. Thus, we do not assume access to arbitrary unitaries; instead, we denote by \mathbf{h} the set of Hermitian operators h for which e^{-iht} can be implemented natively on the device. The unitary e^{-iht} is implemented by turning on an interaction given by h for an amount of time t . We also assume, without loss of generality, that each $h \in \mathbf{h}$ does not commute with H (Hamiltonian of XY model). Fig. 4.5 outlines the AC workflow.

4

Optimization step

We now describe the optimization routine in AC, which incrementally builds a shallow circuit to prepare a low-energy state for local Hamiltonians. We begin by introducing the cooling principle.

Consider a Hamiltonian H of the form in Eq. (4.9) and an operator $h \in \mathbf{h}$ such that $[H, h] \neq 0$ and, for simplicity, $h^2 = \mathbb{I}$. The last assumption can be relaxed [151]. Given an initial quantum state $|\psi_0\rangle$,

$$\begin{aligned} |\psi_0^h(t)\rangle &= e^{-ith} |\psi_0\rangle \\ &= (\cos(t)\mathbb{I} - i \sin(t)h) |\psi_0\rangle, \end{aligned} \tag{4.28}$$

we define the energy function:

$$\begin{aligned} E_h(t) &= \langle \psi_0^h(t) | H | \psi_0^h(t) \rangle \\ &= \cos^2(t) \langle H \rangle + i \cos(t) \sin(t) \langle [h, H] \rangle + \sin^2(t) \langle h H h \rangle. \end{aligned} \tag{4.29}$$

By using trigonometric identities, one finds

$$E_h(t) = \langle H \rangle + A \sin^2(t) + \frac{B}{2} \sin(2t), \tag{4.30}$$

where $\langle \cdot \rangle$ indicates the expectation value in $|\psi_0\rangle$, $A := \langle h H h - H \rangle$, and $B :=$

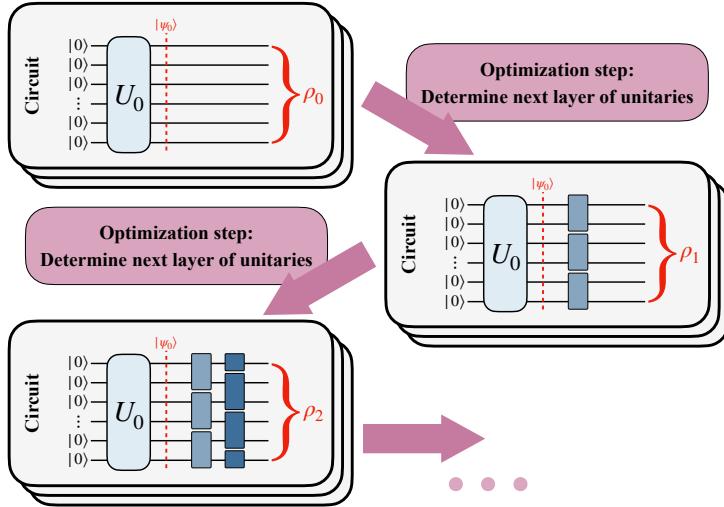


Figure 4.5: Conceptual diagram of algorithmic cooling (AC) Algorithmic cooling iteratively constructs a quantum circuit that prepares a low-energy state of a local Hamiltonian. Initially, the circuit is the identity, and the trial state is $|\psi_0\rangle = U_0 |0\rangle^{\otimes n}$. Each iteration generates an ensemble of identical states for tomography, yielding local RDMs that guide the choice of a new layer of gates (see Sec. 4.4.1), which is appended to form the next circuit.

$i\langle [h, H] \rangle$. The minimum of $E_h(t)$ occurs at

$$t^* = \frac{1}{2} \arctan \frac{-B}{A}, \quad (4.31)$$

yielding

$$E_h(t^*) = \langle H \rangle + \frac{1}{2} \left(A - \sqrt{A^2 + B^2} \right). \quad (4.32)$$

Note that there is a slight ambiguity in the definition of t^* given in Eq. (4.31): the minimum is achieved at any t^* satisfying $\cos(2t^*) = A/\sqrt{A^2 + B^2}$ and $\sin(2t^*) = -B/\sqrt{A^2 + B^2}$. Choosing the opposite sign leads to a maximum at $t^* + \pi/2$,

$$E_h(t^* + \pi/2) = \langle H \rangle + \left(A + \sqrt{A^2 + B^2} \right) / 2. \quad (4.33)$$

Therefore, if A and B can be accurately estimated, applying the interaction given by h (or $-h$) for a period of time $|t^*|$ will decrease the energy of $|\psi_0\rangle$ by an amount of $\frac{1}{2} (\sqrt{A^2 + B^2} - A)$. Since H is local, the value of A and B can

be extracted from measurements on only a constant number of qubits:

$$A = \sum_{j: \text{supp}(H_j \cap h) \neq \emptyset} \langle h H_j h - H_j \rangle, \quad (4.34)$$

$$B = \sum_{j: \text{supp}(H_j \cap h) \neq \emptyset} i \langle [h, H_j] \rangle. \quad (4.35)$$

Here, $\text{supp}(H_j \cap h) := \text{supp}(H_j) \cap \text{supp}(h)$, and $\text{supp}(\cdot) \subseteq [n]$ denotes the qubits on which an operator acts non-trivially.

Because each update step is chosen to locally decrease energy, the energy function forms a monotonically decreasing sequence (assuming we have access to exact RDMs $\{\rho_j\}_{j=1}^m$) that is bounded by below, guaranteeing convergence (albeit not necessarily to the true ground state). The algorithmic cooling approach is thus a heuristic strategy that can fail for several reasons: \mathbf{h} may be too small to reach the global ground state from the initial state; choosing unitaries in certain orders could lead to local minima; or noisy RDM estimates might produce suboptimal energy updates.

Nonetheless, AC provides a flexible, greedy technique for variationally reducing energy. Integrating SDP-assisted tomography can enhance the accuracy of local RDM estimates, thereby improving the reliability and efficiency of each iteration.

Algorithmic cooling circuit compilation

This subsection explains how we construct the layout of the variational circuit by integrating the optimization step from Sec. 4.4.1 into the algorithmic cooling (AC) method. Note that there are many ways to construct the heuristic, each with a different performance depending on the problem. Here, we outline a specific one for illustrative purposes.

We begin with an empty circuit, $\mathcal{U}_C = \mathbb{I}$, and an initial trial state $|\psi_0\rangle$. At each step, we consider all available operators $h \in \mathbf{h}$, selecting the one that yields the greatest energy reduction. We then append the corresponding unitary $U_k = e^{-i h_k t_k^*}$ to the circuit, ensuring a strictly decreasing energy at each iteration. Although this approach guarantees monotonic improvement, it may not fully exploit device capabilities.

To make better use of circuit depth, we group gates into layers V_1, \dots, V_L . Each layer $V_l = U_l^1 \cdots U_l^{k_l} \cdots U_l^{K_l}$ consists of K_l parallel unitaries, with disjoint supports,

$$\text{supp}(U_l^k) \cap \text{supp}(U_l^{k'}) = \emptyset, \quad \forall 1 \leq k, k' \leq K_l. \quad (4.36)$$

The overall circuit is then $\mathcal{U}_C = V_L \cdots V_1$, allowing for a more efficient arrangement while still achieving a steady decrease in energy.

To optimize a layer V_l , one could iteratively:

1. Select an operator h , compute t^* , and update $V_l \leftarrow V_l e^{-iht^*}$.
2. Repeat for the next operator h' , finding t'^* , and so on.

However, each new operator choice requires a full state preparation and measurement cycle. Since H is local, it can become more resource-efficient to collect a suitable set of RDMs for the current layout of V_l and then perform all unitary updates for that layer classically, rather than measuring after every single addition of e^{-iht^*} .

In Alg. 4.2, we present a more refined approach that avoids predefining the circuit layout. Instead, at each iteration, the algorithm greedily selects the operator h^* and corresponding t^* that produce the largest energy decrease. The selection is restricted to operators whose supports are disjoint from any non-trivial gates already in the layer. This ensures the circuit grows incrementally while maintaining a clear and consistent reduction in energy at each step.

Algorithm 4.2: Algorithmic cooling

Input: Initial state $|\psi_0\rangle$, Hamiltonian $H = \sum_{j=1}^m H_j$, operator set \mathbf{h}_0 , maximum number of iterations L .

Output: Quantum circuit \mathcal{U}_C .

Initialize the system to $|\psi_0\rangle$ and set $\mathcal{U}_C = \mathbb{I}$;

$l \leftarrow 0$;

repeat

Prepare an ensemble of identical states $\mathcal{U}_C |\psi_0\rangle$ and perform standard quantum state tomography to obtain local RDMs $\{\hat{\rho}_j\}$;

$\mathbf{h} \leftarrow \mathbf{h}_0$, $V_l \leftarrow \mathbb{I}$;

while $\mathbf{h} \neq \emptyset$ **do**

Estimate the parameters A (Eq. (4.34)) and B (Eq. (4.35)) for every $h \in \mathbf{h}$ with local RDMs $\{\hat{\rho}_j\}$, respectively;

Select the h^* and t^* that maximize the energy decrease

$$\frac{1}{2} (\sqrt{A^2 + B^2} - A), \text{ c.f. Eq. (4.32);}$$

$V_l \leftarrow e^{-i h^* t^*} V_l$;

$$\hat{\rho}_j \leftarrow e^{-i h^* t^*} \hat{\rho}_j e^{i (h^*)^\dagger t^*}, \forall \hat{\rho}_j \in \{\hat{\rho}_j\};$$

$\mathbf{h} \leftarrow \mathbf{h} \setminus \{h \in \mathbf{h} \mid \text{supp}(h) \cap \text{supp}(h^*) \neq \emptyset\}$;

$\mathcal{U}_C \leftarrow V_l \mathcal{U}_C$;

$l \leftarrow l + 1$;

until $l = L$;

4.4.2 Embedding SDP-assisted tomography

In practice, the set of possible operators \mathbf{h} can be extremely large. Rather than testing each $h \in \mathbf{h}$ individually, one can instead perform tomography on the reduced density matrices (RDMs) of relevant qubit subsets. Although tomography itself is expensive, if both the Hamiltonian H and the unitaries to be optimized are local, then the quantities $\langle hH_jh - H_j \rangle$ and $\langle [h, H_j] \rangle$ needed to identify the optimal unitary e^{-iht^*} depend only on the RDMs of $\text{supp}(H_j) \cup \text{supp}(h)$, which is constant for every j . These expectation values from Eqs. (4.34)–(4.35) can be computed via

$$\langle hH_jh - H_j \rangle = \text{Tr}[\rho_r(hH_jh - H_j)], \quad (4.37)$$

$$i\langle [h, H_j] \rangle = i\text{Tr}[\rho_r[h, H_j]], \quad (4.38)$$

where ρ_r is a collection of relevant RDMs of the current global state $|\psi\rangle$,

$$\rho_r = \text{Tr}_{[n] \setminus (\text{supp}(H_j) \cup \text{supp}(h))}(|\psi\rangle\langle\psi|). \quad (4.39)$$

To obtain the relevant RDMs, one first performs state tomography. However, due to shot noise, the raw tomographic local RDMs may become non-physical. To address this, one can formulate an SDP (Eq. (4.19)) to recover locally consistent RDMs that minimize the energy while preserving physicality. These SDP-refined RDMs are then provided to the algorithmic cooling (AC) method (Alg. 4.2), guiding the optimization parameters for each cooling step.

4.4.3 Numerical simulations

We now present numerical simulations of the AC procedure for approximating the ground state of the 1D-chain XY model (a frustrated Hamiltonian), specified by Eq. (4.23). We consider two initial states: (a) the uniform superposition $|+\rangle^{\otimes n}$ and (b) a Hartree–Fock (mean-field) product state, which corresponds to the product state of minimal energy, which can be obtained efficiently in one dimension [206], thereby providing an in principle better initialization. We employ Alg. 4.2 that iteratively prepares the ground state and we restrict the set of available operators \mathbf{h} to geometrically local Pauli operators.

Figs. 4.5a and 4.5b illustrates the results for different pairs (n, n_s) , where n (number of qubits) ranges from 3 to 8 and n_s (number of samples per iteration) takes values $10^1, 10^2, 10^3, 10^4$. Each iteration uses n_s samples per iteration to estimate the necessary parameters for cooling either using standard tomography or SDP-assisted tomography. Within each subplot, the red line marks the exact ground-state energy. Two variants of AC method are simulated: (i) the blue curve corresponds to AC with standard tomography, using tomographically reconstructed local RDMs to compute energy, and (ii) the orange curve depicts

AC with SDP-refined RDMs, still using their reconstructed expectation value for energy. The green curve shows the SDP lower-bound on the energy at each iteration, representing the minimal energy compatible with the measured data under the chosen constraints. Each method is repeated 25 times, and the shaded areas around the curves indicate one-sided standard deviation.

Figs. 4.5a and 4.5b also showcase three noteworthy regimes: For a low number of measurements ($\sim 10^2$), the monotonicity property, especially of the AC with tomography, is broken. This is due to the shot noise being too dominant and confusing the heuristic of the AC. For the $|+\rangle^{\otimes n}$ as an initial state, and for $n_s = 10^2$, we observe a sudden dip in energy for the AC with tomography, followed by a steady increase in energy, as iterations progress. This is due to the AC adding gates that are not sufficiently close to optimal, because of the limitations in precision. In contrast, the behavior of AC with SDP is qualitatively closer to monotonicity, especially in the first few iterations, although ultimately converging to a similar value as the AC with tomography. For the $|HF\rangle$ as initial state, both AC with tomography and AC with SDP have the same phenomena, but the energy scales are much lower as we begin from the lowest-energy product state.

For a higher number of measurements ($\sim 10^3$), the AC with SDP performs better compared to the AC with tomography. This showcases a sweet-spot regime where our approach proves more advantageous, as this behavior is consistent across different system sizes and both initializations considered. However, for an even larger number of measurements ($\sim 10^4$), the performance of the two methods becomes similar. This is likely because, with more measurement data, the estimates provided by standard tomography are increasingly physical, leaving less room for the SDP to correct non-physicality. In the absence of shot noise, the performance of both methods is identical by construction.

It is worth noting that the number of samples used in our numerical simulations is significantly lower than in typical physical experiments, which highlights the efficiency of the approach, though it introduces some limitations. It should be noted that there is no guaranteed advantage of the orange curve (AC with SDP-optimized local RDMs) over the blue curve (AC with local RDMs obtained via tomography) in any single simulation. This is because AC is a heuristic method, and its performance depends on various factors, such as the initial state, measurement settings, and other parameters. However, in some settings, the blue and orange curve have non-overlapping shaded areas. Additionally, by construction, the green curve (AC with SDP-optimized local RDMs and minimized energy) provides a lower bound on the orange curve.

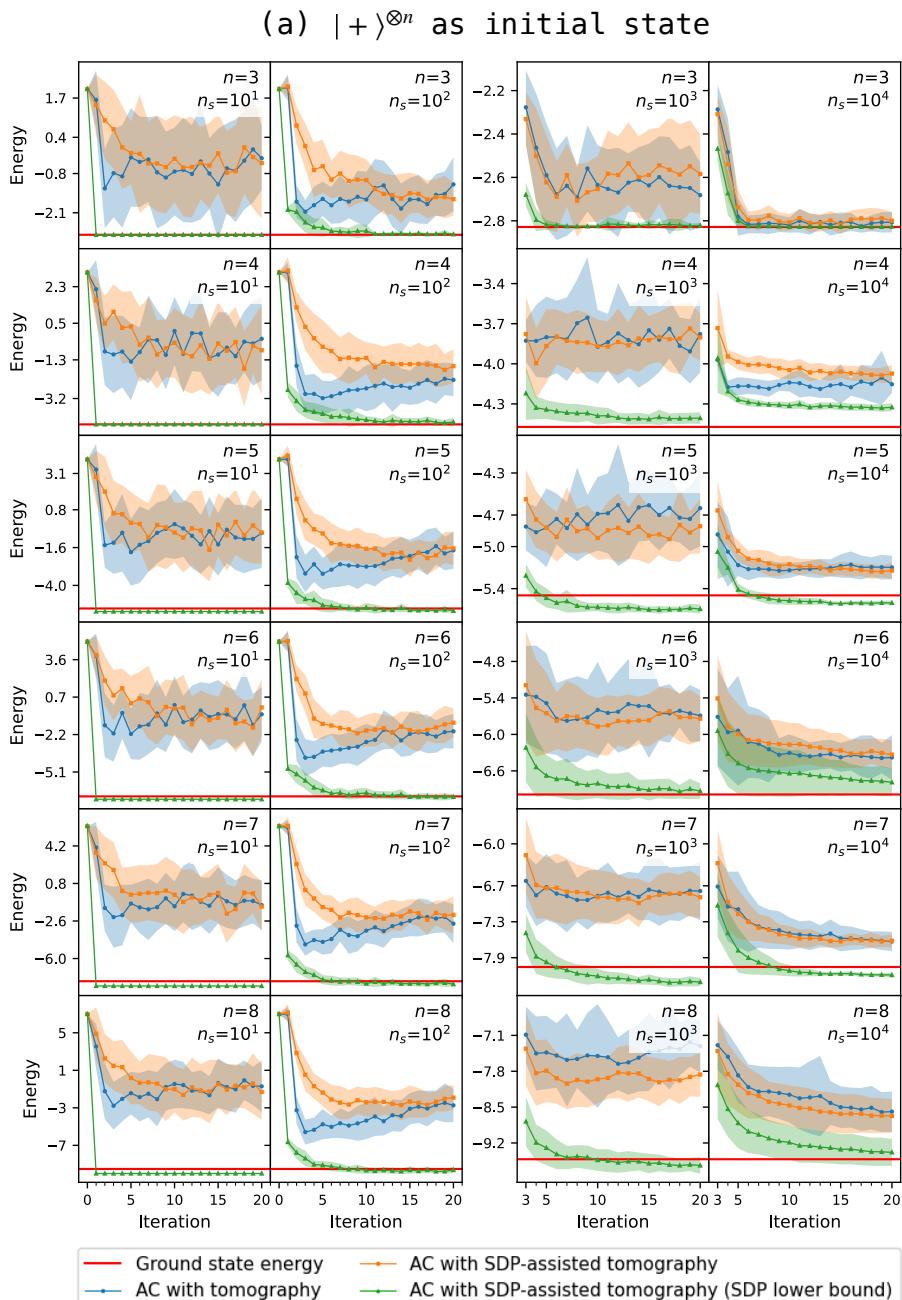


Figure 4.5a: Algorithmic cooling results for a 1D-chain XY model (frustrated Hamiltonian), initialized in $|+\rangle^{\otimes n}$. Energy expectation values are shown as a function of cooling iterations using standard tomography and SDP methods.

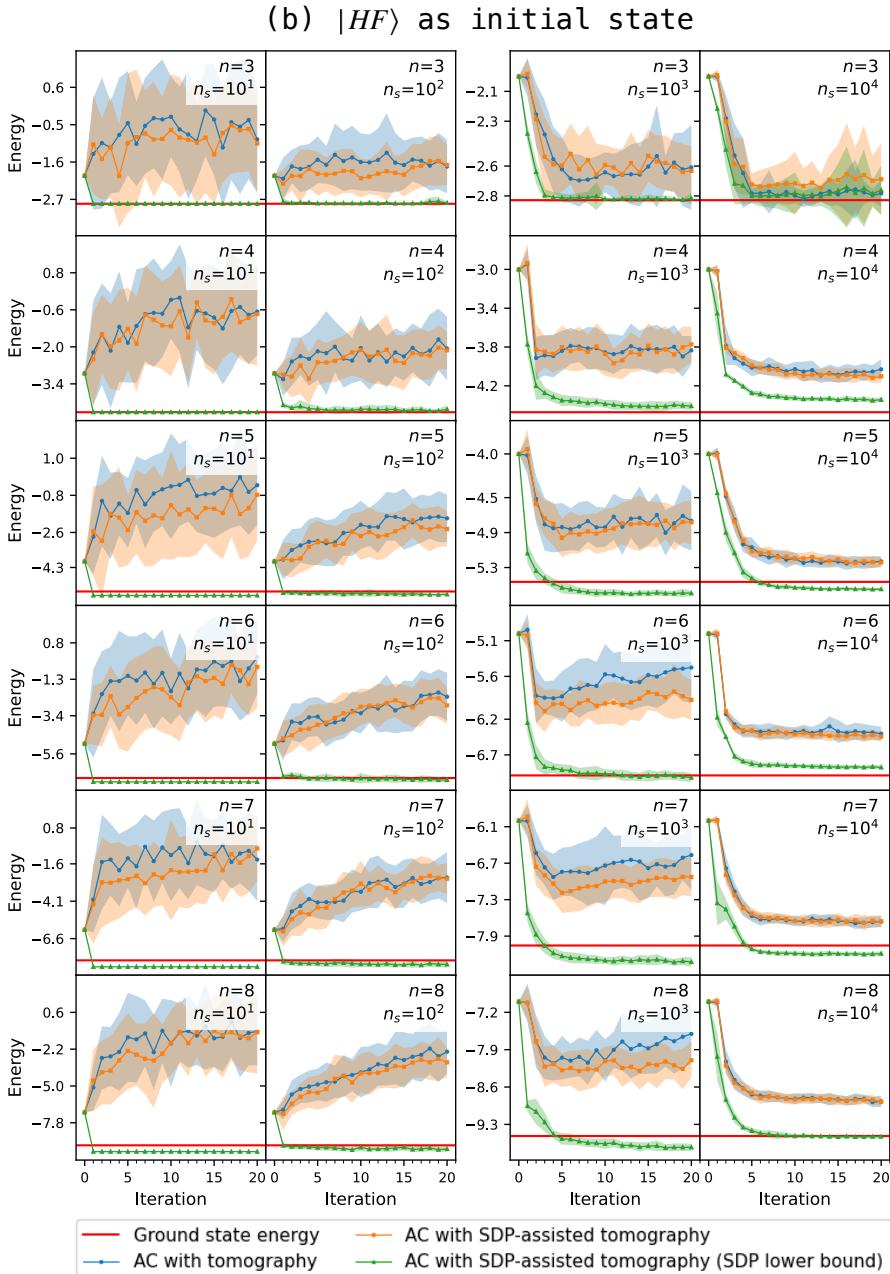


Figure 4.5b: Algorithmic cooling results for a 1D-chain XY model (frustrated Hamiltonian), initialized in $|HF\rangle$. Energy expectation values are shown as a function of cooling iterations using standard tomography and SDP methods.

4.5 Conclusion and outlook

In this chapter, we introduced a semidefinite programming (SDP)-assisted technique for reconstructing overlapping reduced density matrices (RDMs) from experimental measurement data, effectively addressing the challenges posed by shot noise in near-term quantum computing. By enforcing overlapping-compatibility (OC) and enhanced-compatibility (EC) constraints on local RDMs through a polynomial-sized SDP problem, our method alleviates the impact of limited measurements and ensures the consistency of the reconstructed local RDMs. As a result, our approach not only enhances the accuracy of local observable estimates but also provides tighter confidence intervals compared to standard tomographic procedures.

To illustrate its effectiveness, we applied SDP-assisted tomography within a variational quantum algorithm, algorithmic cooling (AC), aimed at heuristically preparing low-energy states of local Hamiltonians. Numerical simulations of a 1D-chain XY model demonstrated the advantages of our approach in terms of both accuracy and resource efficiency. These findings indicate that SDP-assisted tomography can serve as a valuable asset for boosting the performance of variational quantum algorithms and other quantum information processing tasks in the near term.

Looking ahead, a promising direction for future research is the incorporation of additional constraints, such as entropy constraints, into the SDP formulation to further refine the reconstruction process [195]. It would also be interesting to extend the SDP-assisted tomography framework to more complex many-body systems and investigate its behavior on frustration-free Hamiltonians. Applying it to specific quantum chemistry calculations might be promising, as done in a similar way in Refs. [207, 208]. In addition, exploiting the framework to optimize other local observables, such as correlation functions, is another potential area of exploration. Another open question is the effect of systematic errors on the application of our method, and how strategies such as those presented in Ref. [209] can help to reduce them.

Finally, examining how SDP-based methods can be combined with noise mitigation strategies on quantum hardware could offer valuable insights for practical enhancements in near-term quantum computing.