Certifying quantum states and their properties

through correlations, communication and tomography

DISSERTATION

zur Erlangung des Grades eines Doktors der Naturwissenschaften

vorgelegt von

Carlos de Gois

eingereicht bei der Naturwissenschaftlich-Technischen
Fakultät der Universität Siegen

Siegen, 2025



Betreuer und erster Gutachter

Prof. Dr. Otfried Gühne
Universität Siegen

Zweite Gutachter

Ass. Prof. Dr. Marco Túlio Quintino
Sorbonne Université

Tag der mündlichen Prüfung

16. Mai 2025

Abstract

Quantum systems offer significant advantages in information processing and communication, but leveraging these benefits requires reliable methods for extracting relevant information. Due to the probabilistic nature of quantum mechanics, the no-cloning theorem, and the existence of incompatible observables, measuring quantum states is an inherently indirect and often costly process. As engineered quantum systems scale beyond the laboratory, the ability to efficiently certify quantum states and their properties becomes indispensable.

This thesis examines quantum state certification from three perspectives: correlations, communication, and tomography. Correlations play a key role in security protocols, while communication scenarios highlight potential applications of quantum systems in information transmission. Beyond practical applications, they can also reveal fundamental capabilities and limitations of quantum information. Tomography, on the other hand, provides techniques for reconstructing unknown quantum states and is essential not only for validating quantum devices, but also for extracting results from simulations and computations. Our main contributions include:

- Establishing a connection between sum-of-squares uncertainty relations and graph-theoretic quantities.
- Constructing a complete set of criteria for certifying entanglement dimension with an untrusted adversary.
- Proving an equivalence between classical and quantum communication with entangled parties.
- Identifying a class of communication games for which having an entanglement dimension larger than the communication dimension can be useful.
- Providing rigorous statistical guarantees for quantum state tomography experiments applicable to any experimental setup.
- Connecting the optimal measurement choices for partial state tomography to graph covering problems.
- Proposing a protocol for characterizing one- and two-body observables in native fermionic simulators.

Zusammenfassung

Quantensysteme bieten bedeutende Vorteile in der Informationsverarbeitung und Kommunikation, doch die Nutzung dieser Vorteile erfordert zuverlässige Methoden zur Extraktion relevanter Informationen. Aufgrund der probabilistischen Natur der Quantenmechanik, des No-Cloning-Theorems und der Existenz inkompatibler Observablen ist die Messung von Quantenzuständen ein grundsätzlich schwieriger Prozess. In dem Maße, in dem technische Quantensysteme über den Laborbereich hinausgehen, wird die Fähigkeit, Quantenzustände und ihre Eigenschaften effizient zu zertifizieren, unverzichtbar.

In dieser Arbeit wird die Zertifizierung von Quantenzuständen aus drei Perspektiven untersucht: Korrelationen, Kommunikation und Tomographie. Korrelationen spielen eine Schlüsselrolle in Sicherheitsprotokollen, während Kommunikationsszenarien mögliche Anwendungen von Quantensystemen bei der Informationsübertragung aufzeigen. Über praktische Anwendungen hinaus können sie auch grundlegende Fähigkeiten und Grenzen der Quanteninformation aufzeigen. Tomographie hingegen liefert Techniken zur Rekonstruktion unbekannter Quantenzustände und ist nicht nur für die Validierung von Quantengeräten essenziell, sondern auch für die Extraktion von Ergebnissen aus Simulationen und Berechnungen. Unsere Hauptbeiträge umfassen:

- Herstellung einer Verbindung zwischen “sum-of-squares” Unschärferelationen und graphentheoretischen Größen.
- Konstruktion eines vollständigen Satzes von Kriterien für die Zertifizierung von Verschränkungsdimensionen mit einem nicht vertrauenswürdigen Gegenspieler.
- Nachweis der Äquivalenz zwischen klassischer und Quantenkommunikation mit verschränkten Teilnehmern.
- Identifizierung einer Klasse von Kommunikationsspielen, für die eine Verschränkungsdimension, die größer ist als die Kommunikationsdimension, nützlich sein kann.
- Strenge statistische Garantien für Quantenzustands-Tomographie-Experimente, die für jeden Versuchsaufbau anwendbar sind.
- Verbindung der optimalen Messverfahren für die partielle Zustandstomographie mit Überdeckungsproblemen in Graphen.
- Vorschlag für ein Protokoll zur Charakterisierung von Ein- und Zweikörper-Observablen in nativen fermionischen Simulatoren.

This thesis is dedicated to all of those who expected to be acknowledged. Your influence is deeply felt but tragically uncredited.

The fact that we so precisely describe it does not at all mean that we know it. It is like an argument from Spinoza, that lens grinder who polishes glass precisely in order that we are able to examine every thing more closely, who creates an incredibly difficult language in order to express his thought because it is said: seeing is knowing.

I want to know, and not give in to logic. What do I care about a proof from the outside, framed as a geometric argument? It provides merely a semblance of logical consequence and of an order pleasing to the mind. There's A, and after A comes B, first definitions, then axioms and numbered theorems, some supplementary conclusions — and you might have the impression that such command is reminiscent of a wonderfully sketched etching in an atlas, where with letters particular sections are marked, where everything seems so clear and transparent. But we still don't know how it all works.

(**Olga Tokarczuk** — Flights)

Preface

By the early 20th century, physics faced a crisis. Black-body radiation, the photoelectric effect, and measurements of atomic spectra not only revealed fundamental limitations in the theories of the time, but also made it clear that a radically different understanding of nature was required. Over the course of three decades, these observations led to the development of quantum mechanics — from scratch, to the formalized and unified theory that now underpins modern physics. Despite its unparalleled success in explaining experimental observations, quantum mechanics carried profound conceptual puzzles. Among those most affected by them was Einstein, who to the end of his life believed the theory was incomplete and immortalized his opinion in the often misinterpreted sentence "Jedenfalls bin ich überzeugt, daß der nicht würfelt.".

During and after World War II, physics research became increasingly focused on practical applications, ranging from nuclear weapon development to solid-state devices. Debates on foundational issues in quantum theory were largely sidelined, and work in this area came to be seen as marginal at best, or unscientific at worst. Tides changed during the 80s and 90s with a series of discoveries on the uses of quantum systems for information transmission and processing. Even though the early debates on quantum foundations were not the primary motivation for the development of quantum information, it soon became clear that the phenomena enabling these advancements were precisely those that had troubled the theory's pioneers in its early years.

Quantum information is the study of how information can be encoded, processed, and decoded using quantum systems. Performing complex operations, such as quantum computation or information compression, is only as useful as our ability to extract the result. But the very same reasons that make quantum theory interesting — its probabilistic nature, the no-cloning theorem, and the existence of incompatible observables — also massively complicate information readout. As the size of controlled quantum systems scale up, the complexity of characterizing and verifying their states also grows, posing significant challenges in theory and practice. This motivates many ongoing research programs focused on characterizing quantum states, including certification and benchmarking, which ensure that a quantum system operates as intended, and quantum state

tomography, which aims to reconstruct the full quantum state from measurement data. Beyond these, other critical areas of investigation include the development of efficient entanglement characterization methods, which provide scalable ways to verify specific quantum properties without full state reconstruction, and research into partial tomography and shadow tomography, promising to enable the estimation of many observables with significantly reduced measurement overhead.

This thesis is concerned with the certification of quantum states and their properties from three complementary perspectives: correlations, communication, and tomography. Like much of modern quantum information research, these problems are addressed from an operational perspective, yet they occasionally offer deeper insights into the nature of quantum theory itself.

This interplay between practical methods and foundational principles is particularly evident in Part I. There, the exploration begins in Chap. 2 with the fundamental limits of measuring incompatible observables, a topic that not only sheds light on the intrinsic uncertainty of quantum systems but also finds practical applications in shadow tomography and entanglement certification. Subsequently, Chap. 3 addresses the certification of entanglement properties in scenarios involving untrusted parties, a critical challenge in quantum communication and cryptography. Finally, in Chap. 4 the investigation turns to the possibilities and limitations of using quantum systems in communication protocols, and the cost of simulating quantum behaviours using classical systems.

Part II shifts focus toward practical methods for learning quantum systems, but remains substantially connected to questions about the fundamental limits of quantum measurements. It starts with a practical and accessible method for constructing reliable estimates of quantum states from experimental data, introduced in Chap. 5. Motivated by the unavoidably large cost of reconstructing entire quantum states from measurement data, Chap. 6 investigates optimal solutions to the problem of determining only the most relevant subsystems of a quantum state. On a similar vein, Chap. 7 considers the same question in the context of fermionic quantum simulators.

A significant portion of the techniques explored in quantum state characterization heavily rely on convex optimization and other computational techniques. Implementing these techniques in practice often requires specialized knowledge or a considerable effort. Closing this thesis, Part III discusses the basic principles and some implementation details behind a programming package for quantum information currently under development.

Contents

Abstract iii

Preface ix

1 Introduction 1
- 1.1 Quantum information . . . 1
  - 1.1.1 States and measurements . . . 2
    - 1.1.1.1 Tomography . . . 6
  - 1.1.2 Correlations . . . 8
    - 1.1.2.1 Entanglement . . . 9
    - 1.1.2.2 Steering . . . 14
    - 1.1.2.3 Communication . . . 16
- 1.2 Graphs . . . 19
- 1.3 Mathematical optimization . . . 22
  - 1.3.1 Semidefinite programming . . . 22
  - 1.3.2 Integer programming . . . 24
- 1.4 Thesis contributions . . . 25

I Correlations 29

2 Uncertainty relations from graph theory 31
- 2.1 Bounding the sum of squares . . . 32
- 2.2 Consequences of graph properties . . . 34
- 2.3 Discussion . . . 36

3 High-dimensional steering certification 39
- 3.1 The steering scenario . . . 41
- 3.2 Certification hierarchy . . . 43
- 3.3 Example application . . . 45

3.4 Symmetrization . . . 48
3.5 Robustness of high-dimensional steering . . . 49
3.6 Proof of completeness . . . 50
3.7 Discussion . . . 53

4 Entanglement-assisted communication 55
4.1 Entanglement-assisted P&M scenarios . . . 57
4.2 Simulating quantum with classical communication . . . 59
4.3 Advantage of higher-dimensional entanglement . . . 60
4.4 Discussion . . . 65

II Tomography 67

5 Confidence regions 69
5.1 Quantum state tomography . . . 72
5.2 Bernstein inequalities . . . 74
5.3 Confidence regions . . . 76
5.4 Performance comparison . . . 79
5.5 Usage example . . . 82
5.6 Discussion . . . 87

6 Partial state tomography 89
6.1 Reduced density matrices . . . 90
6.1.1 Relation to shadow tomography . . . 91
6.2 Minimizing the number of measurement settings . . . 94
6.2.1 Covering algorithms . . . 96
6.3 Application to Pauli settings . . . 100
6.4 Discussion . . . 102

7 Partial tomography of fermionic simulators 105
7.1 Measuring fermionic coherences . . . 106
7.2 Measurement schemes for two-point correlators . . . 108
7.2.1 Example: Single-body reduced density matrix . . . 110
7.3 Measurement schemes for four-point correlators . . . 112
7.3.1 Example: Two-body reduced density matrix and lattices . . . 115
7.3.1.1 Example solutions for four-point correlators . . . 118
7.4 Discussion . . . 120

III Tools **123**

8 `Ket.jl` **125**
8.1 `Julia` 126
8.2 Symmetric extensions hierarchy 128
8.3 Inner symmetric extensions hierarchy 131
8.4 Genuine multipartite entanglement 132
8.5 Other methods 134
8.6 Discussion 135

Conclusion **137**

List of publications **139**

Bibliography **141**

1 Introduction

> *. . . he gathered the men of the village in his little room, and he demonstrated to them, with theories that none of them could understand, the possibility of returning to where one had set out by consistently sailing east. The whole village was convinced that José Arcadio Buendía had lost his reason, when Melquíades returned to set things straight. He gave public praise to the intelligence of a man who from pure astronomical speculation had evolved a theory that had already been proved in practice, although unknown in Macondo until then, and as a proof of his admiration he made him a gift that was to have a profound influence on the future of the village: the laboratory of an alchemist.*
>
> **— Gabriel García Márquez**
> One Hundred Years of Solitude

At first sight, the laboratory of a quantum information theorist is far less exciting than that of a 19th century astronomer. Instead of sextants, heliometers and astronomical tables, one finds linear operators, partial traces and a computer keyboard. Yet, hidden within these abstract tools lie some of the most counterintuitive and profound ideas, ones that have reshaped our understanding of reality and revolutionized the foundations of physics. In this chapter, we introduce the essential mathematical and conceptual tools that will become our laboratory in the chapters that follow.

1.1 Quantum information

Quantum information theory provides a set of tools for describing information processing in quantum systems. Our starting point are quantum states and measurements, which

define how information is represented and extracted from quantum systems. Some of the most striking quantum effects are direct consequences of these definitions. Among them are entanglement, steering, and quantum communication (Sec. 1.1.2) — which largely underpin the emergent quantum technologies. A more extensive and pedagogical treatment of these ideas can be found in [NC10; Bar09].

1.1.1 States and measurements

State vector

Just as in the formalism of analytical mechanics, where the configuration of a system is fully described by generalized coordinates and conjugate momenta, the state of a quantum system is described by a state vector. Using the notation introduced by Dirac, $|\cdot\rangle$ represents a state vector, and to distinguish one *ket* from another we include symbols or indices to specify the state. According to one of the postulates of quantum mechanics, quantum states are unit vectors in a complex Hilbert space, $\mathcal{H}$. [1] The existence of quantum superpositions directly follow from this postulate: Let $|\psi_i\rangle \in \mathcal{H}$ be quantum states indexed by i. Since these belong to a vector space,

Superposition

$$\alpha_1 |\psi_1\rangle + \alpha_2 |\psi_2\rangle + \cdots + \alpha_n |\psi_n\rangle \in \mathcal{H} \tag{1.1}$$

as well. As long as the superposition is also a unit vector, it is a valid quantum state.

Qubit

To be more explicit, consider a basis $\{|i\rangle\}_{i=1}^{d}$ for a d-dimensional Hilbert space $\mathcal{H}^d$. Then any state vector $|\psi\rangle \in \mathcal{H}^d$ can be expressed as $|\psi\rangle = \sum_{i=1}^{d} \alpha_i |i\rangle$, with $\alpha_i \in \mathbb{C}$ satisfying $\sum_i |\alpha_i|^2 = 1$ to ensure that $\langle\psi|\psi\rangle = 1$. Drawing from an analogy to classical bits, two-dimensional quantum states are called qubits (or qu*d*its for their d-dimensional analogues). They are most often expressed in the computational basis $\{\ |0\rangle,\ |1\rangle\ \}$, where $|0\rangle = \begin{pmatrix}1 & 0\end{pmatrix}^\intercal$ and $|1\rangle = \begin{pmatrix}0 & 1\end{pmatrix}^\intercal$. Per the discussion above, any state $|\psi\rangle \in \mathcal{H}^2$ can be written as $|\psi\rangle = \alpha_1 |0\rangle + \alpha_2 |1\rangle = \begin{pmatrix}\alpha_1 & \alpha_2\end{pmatrix}^\intercal$.

Density operator

An equivalent, but often more convenient formulation of quantum states is in terms of density operators (equiv. density matrices). While a state vector represents one single,

[1] As is usual in quantum information — though not exclusive [BV05] — we are going to divert unnecessary mathematical complications by considering only finite-dimensional Hilbert spaces. For our purposes, this is essentially done without loss of generality, since the phenomena we will be interested in, as well as their practical implementations, do not require infinite dimensional spaces.

unique quantum state, a density operator is used to describe an ensemble $\{ p_i, \mid \psi_i\rangle \}$ of states via

$$\rho = \sum_i p_i \, |\psi_i\rangle\langle\psi_i| \, , \tag{1.2}$$

where $p_i \geq 0$ and $\sum_i p_i = 1$. The properties entailed by the definition are that a density operator is normalized ($\mathrm{tr}(\rho) = 1$) [2] and positive semidefinite ($\rho \succcurlyeq 0$). [3] Any operator with these properties represents a valid quantum state, and every postulate of quantum theory can be equivalently expressed in terms of density operators. We will denote the set of density operators on a Hilbert space $\mathcal{H}$ by $\mathcal{D}(\mathcal{H})$. Explicitly,

$$\mathcal{D}(\mathcal{H}) = \{ \, \rho \in \mathcal{L}(\mathcal{H}) \mid \mathrm{tr}(\rho) = 1, \rho \succeq 0 \, \} \, , \tag{1.3}$$

where $\mathcal{L}(\mathcal{H})$ is the set of linear operators acting on $\mathcal{H}$. Composite systems are described by a density operator in $\mathcal{H} = \mathcal{H}_1 \otimes \ldots \otimes \mathcal{H}_n$, where $\otimes$ denotes the tensor product. Thus, an n-qubit state is an operator in $\mathcal{D}\left[\left(\mathcal{H}^2 \right)^{\otimes n} \right]$. Likewise, it can be represented by a unit-trace, $2^n \times 2^n$ Hermitian matrix.

Composite system

Pure states and mixtures

It is important to notice that the sum in Eq. (1.2) is not the same as a superposition [Eq. (1.1)]. For example, suppose a source of quantum states is designed to generate $|\psi_1\rangle$, but due to imperfections, it may instead output another state $|\psi_2\rangle$ with some probability p_2. The average state is then described by the density operator $p_1 \, |\psi_1\rangle\langle\psi_1| + p_2 \, |\psi_2\rangle\langle\psi_2|$, which represents a statistical mixture rather than a superposition. Crucially, ρ contains no coherences (off-diagonal terms) between $|\psi_1\rangle$ and $|\psi_2\rangle$, and thus cannot exhibit interference phenomena involving these two states. When the differences between state vectors and density matrices need to be emphasized, the former is often called a pure state ($\mathrm{rank}\, \rho = 1$), and the latter is said to be mixed ($\mathrm{rank}\, \rho > 1$).

Observable

To start appreciating the convenience brought by the density matrix representation, suppose we want to measure some quantity A on a state. From another postulate of quantum theory, A must be a Hermitian operator (i.e., an observable), and when measured on a state $|\psi\rangle$, its expected value is given by $\langle A \rangle = \langle \psi | A | \psi \rangle$. Measured on the ensemble $\{ \, p_i, \mid \psi_i \rangle \, \}$,

$$\langle A \rangle = \sum_i p_i \, \langle \psi_i | A | \psi_i \rangle \, . \tag{1.4}$$

[2] Restricted to finite dimensions, the trace operation is equivalent to summing the diagonal elements of the matrix representation of ρ, or whatever is the argument to tr. This operation is linear and basis independent.

[3] A Hermitian operator $A : \mathcal{L}(\mathcal{H}) \mapsto \mathcal{L}(\mathcal{H})$ is said to be positive semidefinite, or PSD, if and only if $x^\dagger A x \geq 0$ for every vector x in $\mathcal{H}$. This happens if and only if every eigenvalue of A is real and nonnegative.

Comparing this to Eq. (1.2) leads to a remarkably simple form, $\langle A \rangle = \mathrm{tr}(A\rho)$.

In composite systems, we can also discuss what happens if only one of the subsystems is measured. For the sake of concreteness, suppose $\rho \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$. Let $\{ |\psi_i\rangle\langle\psi_j| \}_{i,j=1}^{d_A}$ be a basis for the linear operators on $\mathcal{H}_A$, and $\{ |\varphi_\alpha\rangle\langle\varphi_\beta| \}_{\alpha,\beta=1}^{d_B}$ one for $\mathcal{H}_B$. Then, any linear operator on $\mathcal{H}_A \otimes \mathcal{H}_B$ — therefore also a quantum state — can be written as

$$\rho = \sum_{ij\alpha\beta} \rho_{ij\alpha\beta} \, |\psi_i\rangle\langle\psi_j| \otimes |\varphi_\alpha\rangle\langle\varphi_\beta| \, . \tag{1.5}$$

Now suppose that an experimenter in the laboratory associated to $\mathcal{H}_A$ wants to measure an observable A on her subsystem. Because nothing is done on $\mathcal{H}_B$, the global observable is $A \otimes \mathbf{1}_B$, where $\mathbf{1}_B$ is the identity operator on $\mathcal{H}_B$. Using the rule described above,

$$\langle A \otimes \mathbf{1}_B \rangle_\rho = \mathrm{tr} \left(\sum_{ij\alpha\beta} \rho_{ij\alpha\beta} \, A \, |\psi_i\rangle\langle\psi_j| \otimes |\varphi_\alpha\rangle\langle\varphi_\beta| \right) \tag{1.6}$$

$$= \mathrm{tr} \left[A \left(\sum_{ij\alpha} \rho_{ij\alpha\alpha} \, |\psi_i\rangle\langle\psi_j| \right) \right] . \tag{1.7}$$

Partial trace

Defining the partial trace (over $\mathcal{H}_B$) as

$$\mathrm{tr}_B(\rho) = \sum_{ij\alpha} \rho_{ij\alpha\alpha} \, |\psi_i\rangle\langle\psi_j| \tag{1.8}$$

Reduced state

and comparing to Eq. (1.7) shows that $\mathrm{tr}\left[(A \otimes \mathbf{1}_B)\rho\right] = \mathrm{tr}\left[A \, \mathrm{tr}_B(\rho)\right]$. This justifies calling $\mathrm{tr}_B(\rho)$ the reduced state (of ρ, on $\mathcal{H}_A$), and it represents the local information contained on the $\mathcal{H}_A$ subsystem alone. [4]

As a trivial example, if $\rho = \rho_A \otimes \rho_B$, then $\mathrm{tr}_B(\rho) = \rho_A$ and similarly for ρ_B. This also implies that $\rho = \mathrm{tr}_B(\rho) \otimes \mathrm{tr}_A(\rho)$. These are called product states, and they are entirely defined by their local components. On the other hand, some quantum states behave in the extreme opposite way. A typical example is the two-qubit state $|\varphi^+\rangle = 1/\sqrt{2}\,(|00\rangle + |11\rangle)$. [5] Direct computation yields $\mathrm{tr}_A(|\varphi^+\rangle\langle\varphi^+|) = \mathrm{tr}_B(|\varphi^+\rangle\langle\varphi^+|) = 1/2$. In this sense, the reduced states of $|\varphi^+\rangle$ contain no information about it at all. This fact is closely related to the difference between separable and entangled states, which will be defined in Sec. 1.1.2.1. Notice that while this formulation of quantum theory might feel more general

[4] This definition generalizes in the natural way for the partial trace over another or multiple subsystems.

[5] $|00\rangle$ is a common shorthand notation for $|0\rangle \otimes |0\rangle$.

than the usual, their equivalence follows from the purification theorem [Sch36; HJW93; Gis89]: Any mixed state can be represented as the partial trace of a pure state in a larger Hilbert space.

When a measurement is performed, an outcome is obtained. The possible outcomes associated to an observable are its eigenvalues. By the spectral theorem, if an observable A is Hermitian on $\mathcal{H}$, its eigenvectors form an orthonormal basis for $\mathcal{H}$ and the associated eigenvalues are real. Thus, we can write $A = \sum_a \lambda_a |a\rangle\langle a|$, and the projectors $\Pi_a = |a\rangle\langle a|$ are such that $\Pi_a \Pi_{a'} = \delta_{a,a'} \Pi_a$. These are called projective measurements (also known as sharp, or von Neumann measurements). The probability of observing outcome a when A is applied to a state ρ is $p(a) = \mathrm{tr}(\Pi_a \rho)$. One of the most notable aspects of quantum theory is that performing a measurement disturbs the state. Lüders rule — popularly referred to as the "wave function collapse" — states that, post measurement, $\rho \mapsto \Pi_a \rho \Pi_a / \mathrm{tr}(\rho \Pi_a)$. As an example, for a pure state and a measurement with rank-1 projectors, the resulting state is one of the eigenvectors of A, therefore repeating the measurement could only again result in a. A far-reaching consequence of this fact, together with the existence of noncommuting observables, is that determining the state of a quantum system requires access to many samples of ρ (see Sec. 1.1.1.1).

Projective measurement

In much the same way as mixed states can be a more convenient representation than pure states, positive-operator-valued measures can be handier than (but in a similar sense, equivalent to) projective measurements. Also called POVMs or generalized measurements, they are described by a set $M = \{E_a\}_a$ of positive semidefinite operators obeying the completeness relation $\sum_a E_a = \mathbf{1}$. These requirements are there to enforce consistency in the probabilistic interpretation of quantum measurements, and they guarantee that $p(a) \geq 0$ and $\sum_a p(a) = 1$, for any possible state ρ. Differently from projective measurements, the measurements effects E_a do not need to be orthogonal projections, meaning that the possible outcomes may smear into one another.

Generalized measurement

A useful intuition is to interpret POVMs as noisy projective measurements. [6] Consider a sharp measurement $M' = \{\, |0\rangle\langle 0| \,, |1\rangle\langle 1| \,\}$. Performed on a state ρ, outcome a happens with probability $p(a) = \mathrm{tr}(|a\rangle\langle a| \, \rho)$, where $a \in \{\, 0, 1 \,\}$. Suppose instead that our experimental apparatus is faulty, giving the opposite result with probability $(1 - \eta)$. Then, $p(a) = \mathrm{tr}(\eta \, |a\rangle\langle a| + (1 - \eta) \, |a \oplus 1\rangle\langle a \oplus 1|)$. A measurement describing this situation is

[6] Notice that this intuition is not always correct, as not every generalized measurement can be interpreted in this way. In particular, noisy projective measurements are not extremal, but there are examples of extremal POVMs which are nevertheless not projective [Osz+17].

$M = \{\, \eta\,|0\rangle\langle 0| + (1-\eta)\,|1\rangle\langle 1|\,, \eta\,|1\rangle\langle 1| + (1-\eta)\,|0\rangle\langle 0|\,\}$. For any $0 < \eta < 1$, the effects in M are not orthogonal projectors, but they do satisfy the conditions to be a POVM.

1.1.1.1 Tomography

Quantum information processing relies on the precise preparation and control of increasingly complex quantum systems. Ultimately, this ability will enable the development of quantum communication systems, quantum metrology and quantum computing. To ensure the correct functioning of a protocol, we must be able to guarantee the quantum states being used have the intended properties. Most thoroughly, one might want to determine a full classical description of the underlying quantum state, or in other words, to perform quantum state tomography (QST).

Because quantum states cannot be copied [WZ82; Die82] and noncommuting observables cannot be simultaneously measured to arbitrary accuracy, performing quantum state tomography requires repeated state preparations. [7] Depending on the platform, this might mean, for example, running a quantum circuit multiple times, or waiting for enough detection events in an optical setup. To further complicate things, a classical description of a high-dimensional state is a very large object, and it is costly to determine both in theory and in practice. Consider, for example, that a state of k qubits is a $2^k \times 2^k$ Hermitian, unit-trace matrix, which has $4^k - 1$ real parameters. Therefore, to fully determine ρ, one must obtain an exponentially large number of matrix elements. Loosely interpreting each of them as a random variable, one might expect that many samples — even exponentially many — are necessary to get enough statistical significance. This is indeed true, and significant effort is thus put on investigating the least demanding procedures for quantum state tomography.

Tomography protocols depend on the available resources: Do we only have access to local operations on each qubit, or can we act globally as we please? Can we use generalized measurements, or are only projective measurements feasible in our setup? Is a quantum memory available to store multiple state preparations and measure them collectively, or must we measure each sample as soon as they arrive? At this point, most platforms can only perform projective measurements on independent samples of the state (though proof-of-principle experiments with generalized measurements do exist, see e.g. [Str+22]). A

[7] That is, we usually assume that a source of independent and identically distributed samples of ρ is available.

convenient way to abstractly describe a tomography procedure starts by letting s denote labels for the different measurement settings that are available. Each setting corresponds to a collection of measurement effects $\{ E_{a|s} \}_a$ each associated with a possible outcome a, which occurs with probability $p_{a|s} = \mathrm{tr}\left(E_{a|s}\rho\right)$ (see Sec. 1.1.1). For convenience, we may collect the probabilities of setting s in the vector $\vec{p}_s = (p_{a|s})_a$. In this way, one can also interpret a measurement as a map $M_s : \mathcal{D}(\mathcal{H}) \mapsto \mathbb{R}^a$. Multiple settings can be combined into a single vector $\vec{p} = \bigoplus_s \vec{p}_s$, which in turn can be associated with the entire measurement map $M : \rho \mapsto \vec{p}$, summarizing all the settings. When a measurement map M^+ is such that $M^+ M\rho = \rho$ for any $\rho \in \mathcal{D}(\mathcal{H})$ exists (that is, M admits a left-inverse), then the measurement is said to be tomographically complete.

Tomographically complete measurement

Now suppose we have measured n_s of the total N state preparations with setting s and counted $c_{a|s}$ occurrences of outcome a. Naturally, the frequency of this outcome is $f_{a|s} = {}^{c_{a|s}}/_{n_s}$, and it converges to the actual value $p_{a|s} = \mathrm{tr}\left(E_{a|s}\rho\right)$ with increasing n_s. Doing the same for each setting, we obtain the collection of frequencies $\vec{f}_s$, which we may also combine into a single vector $\vec{f}$, as we did for $\vec{p}$. If we then apply an inverse measurement map M^+ (namely, an estimator) to the measurement data, this yields a point estimate $\hat{\rho}(\vec{f}) = M^+\vec{f}$ for the true state ρ. A common choice for M^+ is the pseudo inverse, which corresponds to the free least squares estimator, but other choices, including nonlinear estimators, also exist (see [AKG19] for a comparison between several common choices).

State estimate

As an example, let ρ be a single qubit, and assume our measurement settings are the Pauli operators, $s \in \{ X, Y, Z \}$. These are tomographically complete, and it is possible to explicitly parameterize the state as

$$\rho = \frac{1}{2}\left[\mathbf{1} + \mathrm{tr}(X\rho)X + \mathrm{tr}(Y\rho)Y + \mathrm{tr}(Z\rho)Z\right]. \tag{1.9}$$

Each of the expectation values above can be estimated from the experimental frequencies as $\hat{x} = f_{1|X} - f_{-1|X}$, and similarly for Y and Z. As long as they are close enough to the true value, our estimate $\hat{\rho} = \frac{1}{2}\left(\mathbf{1} + \hat{x}X + \hat{y}Y + \hat{z}Z\right)$ should be close to the fiducial state ρ. But how can we guarantee that $\hat{\rho}$ is close to ρ, without knowing ρ?

Accuracy

Ultimately, the goal of quantum state tomography is to provide rigorous guarantees for the accuracy of $\hat{\rho}$. This is usually expressed in terms of some distance $\|\hat{\rho} - \rho\|_\star$ to the true state. Being a probabilistic statement, it usually comes in the form $\mathbb{P}\left[\|\hat{\rho} - \rho\|_\star < \epsilon\right] \geq 1 - \delta$, where $1 - \delta$ is the confidence level and $\star$ is often chosen to be the Euclidean, spectral or trace norm. This means that if the tomography experiment were to be performed

asymptotically many times, in a fraction $1-\delta$ of the experiments, the true state ρ would be ϵ-distant from the estimator we obtained. Quantum state tomography with single copy measurements requires at least $\Omega(d^3/\epsilon^2)$ samples in trace distance (up to $\log$ factors) [Haa+16] (see also [NL25] for a more fine grained analysis in terms of types of measurements). This can be improved using collective measurements, but only down to $\Omega(d^2/\epsilon^2)$, therefore showing that tomography is an inherently costly procedure. Apart from asymptotic bounds, one may also be interested in giving statistical guarantees in practice. This is the topic of Chap. 5, where we study ways of constructing confidence regions that are versatile, performant, and easy to use.

Sample cost

Many approaches to alleviate the requirements of state tomography exist in the literature. They can be grouped in roughly two directions: The first is to impose prior constraints on the state, such as it being pure, permutationally symmetric, a matrix product state etc. (see, for example, [Blu12; Blu10; Sha+13; Gro+10; Mor+12; Cra+10]). Alternatively, we may still be state agnostic, but refrain from having a complete description of ρ and settle for only estimating, for example, the two-body marginals of the state [BBO20; CW20; Gar+20], or the expectation values of a set of local observables [Aar18; HKP20]. This second direction consists of a zoo of related tasks going by the name of partial state tomography, overlapping tomography, shadow tomography, and classical shadows. They will be compared and explained in Sec. 6.1.1, with partial tomography being our central theme in Chaps. 6 and 7.

Partial tomography

1.1.2 Correlations

The most profound departure from classical intuition brought by quantum theory is the existence of uniquely quantum correlations. These correlations emerge directly from the definition of quantum states and measurements, and not only lead to perplexing phenomena with no classical analogue, but also form the backbone of the emerging quantum technologies. This subsection introduces central aspects of quantum correlations, and serves as our starting point for the contributions presented in Part I of this thesis. Entanglement — the quintessential resource underpinning much of quantum information processing — sets the stage for the introduction of quantum steering, a phenomenon so strange that it led key figures in the development of quantum mechanics to question the theory's completeness. Quantum correlations can also enable communication protocols that are impossible of less efficient classically, and the definition of common communication scenarios is our last topic in this section.

1.1.2.1 Entanglement

Named by Schrödinger in a letter to Einstein [Sch35], entanglement has long been regarded as one of the most strikingly nonclassical features of quantum theory. It arises directly from the postulates of quantum theory, and leads to the consequence that composite systems cannot always be described in terms of the states of their individual parts. Instead, they can be entangled, in which case the system can only be accurately described as a whole. Entanglement plays a crucial part in quantum information and has profound implications for our understanding of how information and correlations are carried by physical systems. Unlike classical correlations, which can always be understood as arising from shared randomness, quantum entanglement defies such an interpretation. It enables phenomena such as quantum teleportation and dense coding [Ben+93; BW92] and serves as an essential resource in quantum cryptography [Gis+02]. Motivated by the foundational interest and the possibilities of practical applications, the development of quantum information theory brought insights into the structure of entanglement, and led to a better understanding of its various forms, classifications, and applications.

Its simplest form appears in bipartite systems. An inspiration for its definition can be found in a state preparation procedure [Wer89b]: Suppose two spatially separated laboratories are each equipped with a device that prepares quantum states, and share a source of classical randomness (that is, they have access to local operations and classical communication, or LOCC, protocols). This source is such that a random number i is generated with probability p_i and upon this result the laboratories locally prepare the states ρ_i^A and ρ_i^B. Now if the parties locally measure their systems with the measurements $M_A = \{ A_a \}$ and $M_B = \{ B_b \}$, respectively, the probability of obtaining the outcomes a and b is

$$p(a,b) = \sum_i p_i \operatorname{tr}\Big(A_a \rho_i^A\Big) \operatorname{tr}\Big(B_b \rho_i^B\Big) = \operatorname{tr}(A_a \otimes B_b\, \rho), \tag{1.10}$$

where the last equality holds for any state of the form $\rho = \sum_i p_i \rho_i^A \otimes \rho_i^B$. This leads us to define that a bipartite state $\rho \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ is separable if and only if it can be written as

Bipartite entanglement

$$\rho = \sum_i p_i \rho_i^A \otimes \rho_i^B, \tag{1.11}$$

where $\sum_i p_i = 1$ and, for any i, $p_i \geq 0$, and the operators ρ_i^A (ρ_i^B) are in $\mathcal{D}(\mathcal{H}_A)$ ($\mathcal{D}(\mathcal{H}_B)$). A state that is not separable is entangled.

When the state is pure, separable states reduce to product states, $|\psi\rangle_{AB} = |\psi\rangle_A \otimes |\psi\rangle_B$. When a pure state is not product, it is entangled. This procedure also clarifies that separable states are *not* uncorrelated. Rather, they may exhibit the correlations introduced by the shared classical randomness, and therefore are said to be classically correlated. Entangled states, conversely, produce correlations that are not explainable by separable states, and are thus genuinely quantum.

In practice, we are often interested in determining whether a given state ρ is entangled. Even though a large set of criteria exist [GT09; Hor+09], both analytical and numerical, they often only provide necessary conditions, or require numerical computations that scale unfavourably with the system's dimension, or only hold for special cases such as pure states or particular families of states.

One of the preferred choices of entanglement criteria is the positive partial transpose (PPT) or Peres-Horodecki criterion [Per96; HHH96]. It is a rather simple test based on the partial transposition operation. Taking $\{ |\psi_i\rangle\langle\psi_j| \}$ as a basis for the linear operators in $\mathcal{H}_A$ and $\{ |\varphi_\alpha\rangle\langle\varphi_\beta| \}$ for $\mathcal{H}_B$, the partial transposition on $\mathcal{H}_B$ is given by

PPT criterion

$$\rho^{\mathsf{T}_B} = (\mathbf{1}_A \otimes T)\,\rho = \sum_{ij\alpha\beta} c_{ij\alpha\beta}\, |\psi_i\rangle\langle\psi_j| \otimes (|\varphi_\alpha\rangle\langle\varphi_\beta|)^{\mathsf{T}}, \tag{1.12}$$

where T stands for the transposition map and $c_{ij\alpha\beta}$ define ρ in the bases stated above. Now suppose that ρ is separable, and that we transpose in $\mathcal{H}_B$ (the argument is equivalent for transposing the the other subsystem). Then $\rho^{\mathsf{T}_B} = \sum_i p_i \rho_i^A \otimes \left(\rho_i^B\right)^{\mathsf{T}}$. Because ρ_i^B are quantum states and the transpose preserves its trace and eigenvalues, it follows that $\left(\rho_i^B\right)^{\mathsf{T}}$ is also a valid state and thus $\rho^{\mathsf{T}_B} \succeq 0$ and normalization is also maintained. Consequently, all separable states have positive partial transpose, which also implies that if the partial transpose of some ρ has negative eigenvalues, the state must be entangled.

This argument, posed by Peres [Per96], gives a sufficient condition for entanglement certification. Shortly after, Horodecki et al. proved that when the product of the local dimensions $d_A \cdot d_B \leq 6$, the PPT criterion is also necessary. In larger dimensions, though, one can find states that are entangled but have a positive partial transpose. In these cases, we must resort to more powerful tools.

Symmetric extension

A prominent example is the symmetric extensions, or Dohery-Parrilo-Spedalieri (DPS) hierarchy [DPS04; DPS02]. Their core observation was that if $\rho \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ is separable, then we can take any number $N \geq 1$ of copies of the part in $\mathcal{H}_B$ and write

$$\Xi = \sum_i p_i \rho_i^A \otimes (\rho_i^B)^{\otimes N}. \tag{1.13}$$

So if we choose an N and manage to show there exists no operator $\Xi \in \mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B^{\otimes N})$ such that when we trace out all but one of the subsystems in $\mathcal{H}_B$, we recover ρ,

$$\rho = \mathrm{tr}_{B_i \,:\, i \in [N] \setminus m} (\Xi), \quad \forall m \tag{1.14}$$

then ρ must be entangled. Here the notation $[N]$ is a shorthand for $\{1, \ldots, N\}$ and the operator $\setminus$ is a set difference, thus $B_i : i \in [N] \setminus m = B_1 \ldots B_{m-1} B_{m+1} \ldots B_N$. When the operator Ξ with the properties above exist for a given N, the test is inconclusive, thus this condition is only sufficient for certifying entanglement. However, it turns out that for increasing N this becomes a complete hierarchy of tests. More precisely, ρ is separable if and only it is symmetrically extendible for all natural numbers N [DPS04; Wer89a; CFS02]. Though interesting, this is of limited practical use for separability certification, but in fact the hierarchy can also be modified to approximate the set of separable states from the inside, thus providing sufficient conditions for separability [NOP09a].

Because the conditions in Eq. 1.14 can be expressed as a semidefinite program, they can in principle be efficiently tested using one of a range of conic solvers (see Sec. 1.3.1). However, because of the exponentially increasing dimensions in terms of N, a naive implementation of the hierarchy will not get us very far. A common way to mitigate this issue is to observe that the subsystems in $B_1 \ldots B_N$ can be freely permuted and assumed to be rank-one projectors, thus part of the extension operators can be described in the symmetric subspace $\mathrm{Sym}_N[\mathcal{L}(\mathcal{H}_B^{\otimes N})]$, which is only of dimension $\binom{d_B+N-1}{N}$ [Wat18, Sec. 7.1]. For numerical tests of the DPS hierarchy, using this fact not only enables us to reduce the number of variables, but also to cut off a large amount of constraints, since the partial trace in Eq. (1.14) needs only to be enforced for one choice of m. Practical matters, such as how to obtain an isometry to $\mathrm{Sym}_N[\mathcal{L}(\mathcal{H}_B^{\otimes N})]$, will be further discussed in Sec. 8.2.

Another direction in solving the entanglement problem is to consider particular classes of states. An important example is that of pure bipartite states, for which entanglement

Schmidt decomposition

can be decided by the Schmidt decomposition (see e.g. [NC10]): Any bipartite pure state $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ can be decomposed in the form

$$|\psi\rangle = \sum_{i=1}^{d} \eta_i \, |i_A\rangle \otimes |i_B\rangle \,, \tag{1.15}$$

where $\{|i_A\rangle\}_{i=1}^{d_A}$ and $\{|i_B\rangle\}_{i=1}^{d_B}$ are orthonormal bases for the subsystems in $\mathcal{H}_A$ and $\mathcal{H}_B$. The Schmidt coefficients $\eta_i \geq 0$ are real numbers, $\sum_i \eta_i^2 = 1$, and there are at most $d = \min\{d_A, d_B\}$ terms. The number of nonzero coefficients in the decomposition is its Schmidt rank, denoted $\mathrm{sr}(|\psi\rangle)$. Since a pure state is either product or entangled, $|\psi\rangle$ is entangled if and only if $\mathrm{sr}(|\psi\rangle) > 1$.

Schmidt number

Only pure states have a Schmidt decomposition. An interesting way to generalize it to mixed states $\rho \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ is to consider every possible ensemble of pure states $\{\, p_i, \; |\, \psi_i\rangle \,\}$ that decompose it into $\rho = \sum_i p_i \, |\psi_i\rangle\langle\psi_i|$. The Schmidt number of ρ is then defined as [TH00]

$$\mathrm{sn}(\rho) = \min_{\{|\psi_i\rangle\}_i} \{\max_i \, [\mathrm{sr}(|\psi_i\rangle)]\}, \tag{1.16}$$

In other words, the Schmidt number is the largest Schmidt rank that appears in the "worst case" pure state decomposition of ρ, and in this sense it quantifies the minimum entanglement dimension necessary to prepare ρ. Based on this definition, we can consider the Schmidt classes $\mathcal{S}_k = \{\, \rho \mid \mathrm{sn}(\rho) \leq k \,\}$, which are the subsets of states with maximal Schmidt number $k \geq 1$. It will later become useful to observe that these sets are convex, their extremal points are pure states, $\mathcal{S}_1 = \mathtt{SEP}$ and $\mathcal{S}_{k-1} \subset \mathcal{S}_k$ for any $k \geq 2$ [TH00].

Schmidt classes

Entanglement quantifier

More than deciding entanglement, the Schmidt number can also quantify it. An entanglement measure is a function $E : \mathcal{D}(\mathcal{H}) \to [0, \infty]$ that fulfils certain requirements. Usually we would want that $E(\rho) = 0$ if ρ is separable, and that under any LOCC channel Λ the measure does not increase, that is, $E(\Lambda \circ \rho) \leq E(\rho)$.[8] Convexity is also often desirable, so that $E(\sum_i p_i \rho_i) \leq \sum_i p_i E(\rho_i)$, meaning that entanglement does not increase under statistical mixing. Depending on the intended use of the measure, properties such as additivity or dimension independence can be of interest [Bru02; ES14].

There is in fact an entire zoo of measures that satisfy these properties. As expected for any function that summarizes a complicated property of operators in $\mathcal{D}(\mathcal{H})$ into real number,

[8] Some authors impose the stronger condition that the measure is non-increasing under stochastic LOCC transformations, or even that the measure does not increase on average. A thorough discussion is out of our scope, and we refer the interested reader to [ES14].

entanglement measures do not necessarily agree with each other in the ordering they impose, and the choice of a meaningful measure depends on the application.

It is not difficult to see that the Schmidt number satisfies the required properties, except perhaps that it is non-increasing under LOCC. To see this is indeed the case, consider first the case of two pure bipartite states $|\psi\rangle$ and $|\varphi\rangle$ with Schmidt coefficients $\psi = (\psi_i)_{i=1}^d$ and $\varphi = (\varphi_i)_{i=1}^d$. A seminal result due to M. Nielsen [Nie99] states that ψ can be transformed into φ by means of LOCC if and only if ψ majorizes φ, that is,

$$\sum_{i=1}^{r} \left(\psi_i^{\downarrow}\right)^2 \leq \sum_{i=1}^{r} \left(\varphi_i^{\downarrow}\right)^2, \ \forall r \in \{1, \ldots, d\}, \tag{1.17}$$

where the $\downarrow$ means the coefficients are to be taken in non-increasing order. For this reason, states for which $\eta_i = 1/\sqrt{d}$ for all i are said to be maximally entangled. A consequence of Eq. (1.17) is that a state $|\psi\rangle$ cannot be transformed into any state $|\varphi\rangle$ with $\mathrm{sr}(|\varphi\rangle) > \mathrm{sr}(|\psi\rangle)$, or rather, that the Schmidt rank of a state is non-increasing under LOCC. That the same is true for the Schmidt number follows from convexity and the fact that the extremal points of $\mathcal{S}_k$ are pure states. For further properties see [TH00; SBL01].

Maximally entangled state

The definition of Schmidt number is thus consistent and meaningful, but it does not indicate how to actually compute the Schmidt number of an arbitrary state ρ. Interestingly, it is possible to reduce this problem into a separability certification by embedding ρ in an extended Hilbert space. The key fact comes from Hulpke et al. [Hul+04] and was recently also observed in [Wei+20]. This follows from the fact that any bipartite pure state $|\psi\rangle_{AB} = \sum_{i=1}^k w_i |\alpha_i\rangle_A |\beta_i\rangle_B$ with Schmidt rank at most k can be expressed as

Embedding

$$|\psi\rangle_{AB} = \Pi_k^\dagger \left(|\alpha'\rangle_{AA'} \otimes |\beta'\rangle_{B'B}\right), \tag{1.18}$$

where $\Pi_k = \mathbf{1}_A \otimes \sum_{i=1}^k |ii\rangle_{A'B'} \otimes \mathbf{1}_B$ is an unnormalized, k-dimensional maximally entangled state on the auxiliary spaces, $|\alpha'\rangle_{AA'} = \sum_i w_i |\alpha_i\rangle_A |i\rangle_{A'}$ and $|\beta'\rangle_{BB'} = \sum_i |i\rangle_{B'} |\beta_i\rangle_B$. By convexity, this means that any $\rho \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ with Schmidt number k or less can be embedded in an unnormalized operator $\omega/k \in \mathcal{D}(\mathcal{H}_{AA'} \otimes \mathcal{H}_{B'B})$ which is separable against the $AA' \mid B'B$ bipartition. In practical terms, given a state ρ and a putative Schmidt number k, we can test whether there exists an ω satisfying these conditions and such that $\rho = \Pi_k^\dagger \omega \Pi_k$. If not, then $\mathrm{sn}(\rho) > k$, and we can proceed to test for Schmidt number $k+1$. This effectively transforms the Schmidt number certification problem into a series of entanglement problems, albeit in growing dimensions. Any of many tools for

entanglement certification can be used to solve this problem, and in particular the PPT criterion and the symmetric extensions hierarchy discussed above. For a discussion on how to cast these as a semidefinite program and implementation details, see Secs. 1.3 and 8.2.

As rich and nuanced as it is, bipartite entanglement is but the simplest possible entanglement structure. Matters become significantly more complicated as soon as more particles are introduced. Multiple "kinds" of entanglement are then possible [GT09; Hor+09; SU08; Nav+20; Kra+21], the strongest form of them being called genuine multipartite entanglement (GME). Let $\rho \in \mathcal{D}\left(\bigotimes_{i=1}^{N} \mathcal{H}_i\right)$, and let $A_i \mid B_i$ be bipartitions of the subsystems, that is, $A_i \subset \{1, \ldots, N\}$, and $B_i = \{1, \ldots, N\} \setminus A_i$. Then, ρ is said biseparable if it can be decomposed in the form

Genuine multipartite entanglement

$$\rho = \sum_i p_i \rho_i^{A_i|B_i}, \tag{1.19}$$

where the states $\rho_i^{A_i|B_i}$ are separable with respect to the bipartition $A_i \mid B_i$, and $\{p_i\}$ is a probability distribution. A state that is not biseparable is genuinely multipartite entangled. Many central questions regarding the characterization of GME states are still open, some of which interact with important problems in tensor analysis and network entanglement [WGE16; GT09; Hor+09; Han+22; HG23; Neu+25]. Genuine multipartite entangled states will appear in our discussion on quantum state tomography (Sec. 5.5) and some methods to certify GME will be discussed in Sec. 8.4.

1.1.2.2 Steering

Schrödinger's letter to Einstein was in fact provoked by what came to be known as the EPR argument [EPR35]. Their core idea was that if two parties share an entangled state, a measurement on one particle can steer the other into a state that depends on the choice of the measurement. Einstein derisively referred to this instantaneous influence as a "spooky action at a distance", taking it as a clear sign that quantum mechanics was an incomplete theory.

The idea of local completions of quantum theory in the EPR setting was only formalized much later [WJD07], and now plays a central role in the study of quantum correlations [CS16; Uol+20]. Illustrated in Fig. 1.1, the simplest steering scenario comprises two distant parties. Usually named Alice and Bob, they are allowed the share a bipartite

Steering scenario

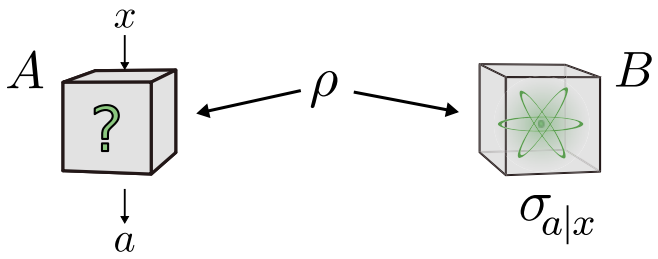


Fig. 1.1.: In a steering scenario, one side of a bipartite quantum state ρ is measured with uncharacterized settings labelled by x. The set of conditional states $\sigma_{a|x}$ produced on Bob's side is known as a state assemblage. A key feature of steering is that this assemblage may exhibit correlations that cannot be explained by a local hidden state model, implying that the original shared state ρ must have been entangled.

quantum state ρ. She wants to convince Bob that she can steer his state by locally measuring her share of ρ. He does not trust her, and wants to check by himself that his observations are not explainable by local states. So they do the following experiment: Suppose she has n_x possible measurements to choose from, and upon a choice $x \in \{1, \ldots, n_x\}$ she reads out a measurement outcome $a \in \{1, \ldots, n_a\}$ Measured on ρ, this would yield a conditional state $\sigma_{a|x} = \mathrm{tr}_B \left[\rho \left(M_{a|x} \otimes \mathbf{1}\right)\right]$, which he is allowed to characterize. [9] The collection $\sigma = \{\sigma_{a|x}\}$ is called a steering assemblage.

Steering assemblage

After many rounds of the experiment, she informs him of her choices of measurements and her outcomes on each round. Bob then tries to explain his observations in the following way: Suppose they only shared classical randomness, denoted by λ. Then, it might be that λ sneakily instructed his apparatus to produce a (local) state σ_λ, and hers to output a with a probability $p(a \mid x, \lambda)$. Could it not be the case that this explains their observations, and that there was no shared ρ after all? To prove so, he must describe the assemblage elements as

Local hidden states model

$$\sigma_{a|x} = \int_\Lambda d\lambda p(\lambda) p(a|x,\lambda) \sigma_\lambda, \tag{1.20}$$

where the integration happens over the possible probability distributions for the hidden variables λ. Assemblages that cannot be explained with such a local hidden states (LHS) model are called steerable, and so is a state from which a steerable assemblage can be obtained.

[9] Recall that $M_{a|x} \succcurlyeq 0$ are measurement effects, and for any measurement x we have $\sum_a M_{a|x} = \mathbf{1}$. Notice also that $\sigma_{a|x}$ is not normalized, and that the reduced state $\sum_a \sigma_{a|x} = \rho_B = \mathrm{tr}_A(\rho)$ is independent of Alice's choice of measurement.

Steering is closely related to other genuinely quantum properties of nature. For example, it follows directly from the definitions of entanglement and of joint measurability that both are necessary for steerability. More thought is necessary to also see that entanglement is not sufficient, though measurement incompatibility is, in the sense that for any incompatible set of measurements there exists a state which can be steered with it [QVB14; UMG14; Uol+15]. On the practical side, much of what can be done with nonlocality [Bru+14; Uol+20] — entanglement verification, key distribution, quantum metrology, randomness generation, etc. — can also be achieved with steering, which has less stringent experimental requirements. The price to pay is that we must trust one of the parties, but trust is often a reasonable assumption, such as in cryptographic protocols where a trusted local lab interacts with an untrusted quantum server.

Although steering can certify entanglement in a semi-device-independent manner, relatively little attention has been given to whether it can also be used to *characterize* entanglement. For instance, can steering be used to determine the exact entanglement class of a state, or to infer specific properties such as its dimensionality? A recent investigation [Des+21] has begun exploring this direction but relies on manually constructed steering inequalities, without providing a systematic method for entanglement characterization. We will return to this question in Chap. 3, where we develop a systematic approach to determining the entanglement dimension of an assemblage.

1.1.2.3 Communication

Communication fundamentally involves the transmission of information that is encoded in physical systems. Prepare and measure (P&M) scenarios are an abstraction of communication processes, distilling a wide range of protocols into a unified description. In recent years, they have become a staple ingredient in quantum information protocols such as semi-device-independent dimension certification [Gal+10; WCD08; BQB14] and the analysis of quantum communication networks [BBP15; Wan+19], while also providing means for self-testing states [Tav+18; MO21], distributing quantum keys [PB11], certifying randomness [Pas+15], and playing a central part in the discussion on the informational foundations of quantum theory [Paw+09; CMG15].

The simplest case consists of a preparation device that encodes information in a physical system, then feeds it into a communication channel, which transmits the state to a device that must measure the system to decode the message. This abstract formulation strips away the complexities of specific implementations of the devices, and lets us focus on

the fundamental limits of communication. By focusing on these basic operations, P&M scenarios can be used to compare different communication resources on an equal footing. Beyond their role in abstracting communication protocols, they also offer a minimal setting to study the core elements of physical theories, namely, state preparation and measurements. An interesting feature in that regard is that, unlike in nonlocality or steering, P&M protocols can behave nonclassically even in the absence of entanglement. This may arise from other quantum features, such as measurement incompatibility [CHT20] and the nonorthogonality of quantum states [BNV13].

Prepare and measure scenario

Prepare-and-measure scenarios are usually categorized in terms of the resources available: Can quantum states be prepared, or only classical states? Are the preparation devices correlated? May they share an entangled state, or only classical randomness? To formalize these instances, suppose Alice can provide a classical input $x \in \mathcal{X}$ to the preparation device. The set of inputs is assumed to be finite, and without loss of generality we can label the possible inputs as $\mathcal{X} = \{1, \ldots, n_X\}$. Her input may influence the probability $p(m \mid x)$ with which a state labelled by m, from a set of possible preparations $\mathcal{S}$, is produced. For example, if the task is to directly communicate a message, x can be loosely interpreted as the message Alice wants to encode and communicate to Bob, and $m \in \mathcal{S}$ as the carrier of that information. So that communication is not trivialized, the information content in the preparations must be limited. Usually, this is done by limiting the dimension of the communicated systems, but other options exist (see for instance [PPT25]). On the receiving end, Bob selects a measurement setting $y \in \mathcal{Y} = \{1, \ldots, n_Y\}$, capturing the idea that the decoding function can be chosen. Lastly, given the input y and the received message m, the measurement device returns $b \in \mathcal{B} = \{1, \ldots, n_B\}$ with some probability $p(b \mid m, y)$. [10] Over many rounds of communication, with a distribution of choices for x and y, the behaviour $\vec{p} = \{p(b \mid x, y)\}$ associated to the experiment can be determined.

Behaviour

This definition leads to six large classes of P&M scenarios, defined by whether $\mathcal{S}$ contains quantum (denoted ρ_x) or classical states (a), and whether the devices are correlated via classical strategies (λ), a quantum state (ρ_{AB}), or uncorrelated (see Fig. 1.2). In addition to this, a scenario also depends on the number of inputs, measurements and outcomes. To denote each case, we will use the following recipe: First comes the kind of communication, either C for classical or Q for quantum. As a superscript, we describe

Sets of behaviours

[10] An implicit assumption is that there are no side channels. This limits the exchange of information to happen only through m and their pre-shared correlations. Otherwise, communication trivializes correlations [Cha+15].

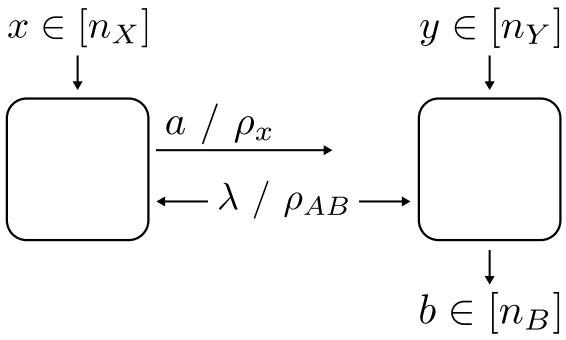


Fig. 1.2.: A preparation device, given an input x, prepares and transmits either a classical (a) or a quantum d-dimensional system (ρ_x) to a measurement device. Upon being given a measurement choice y, it returns an outcome b. The two devices may be uncorrelated, or correlated by a classical (λ) or quantum (ρ_{AB}) resource.

whether the resource is classical shared randomness (λ), quantum states (ρ) or nothing ($\emptyset$). The communication dimension is then the subscript, and lastly, we may add the tuple (n_X, n_Y, n_B), or omit it when it is clear from the context. So, for instance, $Q_2^{\rho}(2,1,2)$ is the set of behaviours achievable with qubit communication, a shared quantum state, two choices of preparation, and one dichotomic measurement. As another example, consider C_d^{λ}. From the discussion above, it follows that

Classical behaviours

$$\vec{p} \in C_d^{\lambda} \quad \Longleftrightarrow \quad p(b \mid x, y) = \sum_{a \in \mathcal{S}} \int_{\Lambda} d\lambda p(\lambda) p(a \mid x, \lambda) p(b \mid a, y, \lambda). \tag{1.21}$$

Any behaviour that cannot be explained by this model (for fixed n_X, n_Y, and n_Y), cannot be reproduced using classical features alone. It also becomes obvious that $C_d^{\lambda}(n_X, n_Y, n_B) \subset C_d^{\rho}(n_X, n_Y, n_B)$, since one can embed classical randomness even in a separable quantum state.

The central goal in P&M scenarios is to understand the relationship between the behaviours achievable with each set of resources. More than just an exercise of foundational interest, this is also of practical significance, as any separation between quantum and classical sets imply there is at least one application for which quantum advantage in communication exists. For a simple concrete case, consider the task of communicating x. This can be represented as a functional on the space of behaviours, $f(\vec{p}) = \sum_b p(b = x \mid x, y)$, that evaluates how well they perform. Dense coding is a P&M protocol that achieves a perfect score using shared entanglement and quantum communication, largely outperforming the best classical behaviours in a similar scenario [BW92]. Plenty of other

Dense coding

concrete examples exist, most prominently random access coding [Amb+99; Amb+08; PŻ10], and communication complexity protocols [Buh+10; Bru+04; TŻB20].

However, investigating the sets of behaviours is significantly more general than considering specific games. In particular, by proving two sets are equal, one indirectly shows the resource in question is not useful in any task. One important result in this direction concerns the comparison between quantum and classical communication in scenarios with a single decoding function ($n_Y = 1$). Surprisingly,

Frenkel-Weiner equivalence

$$Q_d^\lambda(n_X, 1, n_B) = C_d^\lambda(n_X, 1, n_B), \tag{1.22}$$

for any choice of d, n_X, and n_B. This observation can be seen as a generalization of Holevo's bound [Hol73], where more strongly than the channel capacity of a dit and a qudit being equivalent, the entire sets of behaviours are the same. A consequence of this fact is that in any scenario with $n_Y > 1$ measurements which are jointly measurable [DFK19], there cannot be quantum advantage in communication. Measurement incompatibility is therefore a necessary resource [Vie+23], though it is not sufficient [Goi+21; Ege+25]. An extended discussion on the relationships between these sets can be found in [Sou23; Goi21; Pau+22b; Tav+21].

Instances of the C^λ sets [Eq. (1.21)] are the most well studied in the P&M literature. Indeed, they represent the most powerful fully classical behaviours, and any claim of "nonclassicality" must be made in contrast to them. Reasoning about C^λ is also relatively straightforward, as most of the required techniques can be borrowed from the vast literature on nonlocality [Sca19; Bru+14]. On the other hand, the cases involving entanglement instead of classical randomness have only recently began to be systematically investigated [Mor+21; Tav+21; Pau+22b; FW22]. These scenarios are going to be the stars of Chap. 4, and a precise definition of their behaviours is postponed until then.

1.2 Graphs

Graphs are a convenient representation for collections of objects and their relationships [Die25; BM08]. Their combinatorial structure is reflected in graph quantities such as the independence or the chromatic number, which can reveal important properties of the underlying system. For this reason, graphs are used in a wide range of real-world

applications, including social media analysis, scheduling problems, and optimization. Quantum information is no exception, and graphs are used to represent graph states [Hei+06], to encode compatibility relations in quantum contextuality [CSW14], and to model quantum networks [Per+10], among others.

Graph

A graph $G = (V, E)$ consists of a set of vertices V and a set of edges E. In an undirected graph, edges must connect two vertices, thus

$$E \subseteq \{\, \{u, v\} \mid u, v \in V \,\}. \tag{1.23}$$

Hypergraph

Edges with a single vertex are called loops, and graphs where they are forbidden are called simple graphs. A hypergraph generalizes the notion of a graph by allowing edges to connect arbitrary subsets of vertices. Therefore, a hypergraph is a pair $H = (V, E)$ where

$$E \subseteq 2^V \setminus \{\, \emptyset \,\}, \tag{1.24}$$

where 2^V denotes the power set of the vertices. An edge in a hypergraph is thus a subset of the vertices of any size, except zero.

Complement graph

The complement of G, denoted $\overline{G}$, is the graph with the same vertex set V, but with edges only where G does not have them (Fig. 1.3b). More precisely, $\overline{G} = (V, \overline{E})$, where

$$\overline{E} = \{\, \{u, v\} \mid u, v \in V, u \neq v, \{u, v\} \notin E \,\} \tag{1.25}$$

Induced subgraph

The complement graph is useful in studying independence structures and plays a role in problems related to graph colouring and clique detection. A related concept is that of induced subgraphs: Given a subset $S \subseteq V$ of vertices, the subgraph of G induced by S is a graph with vertex set S and all edges in G that have both endpoints in S. Induced subgraphs are useful to analyse local properties of a graph.

Clique number

A clique is a subset of vertices that are pairwise connected. That is, a clique of size k is a subgraph $G' = (V', E')$ where $V' \subseteq V$ has k elements, and $\{\, u, v \,\} \in E'$ for every $u, v \in V'$. The clique number of G, denoted $\omega(G)$, is the size of its largest clique (Fig. 1.3a). Determining $\omega(G)$ is an NP-hard problem [Kar72], meaning that unless $P = NP$, no efficient algorithm exists to compute it for arbitrary graphs. Despite this complexity, cliques can be relevant in quantum information, for example to represent sets of mutually compatible measurements or sets of events that can co-occur in a quantum experiment [CSW14].

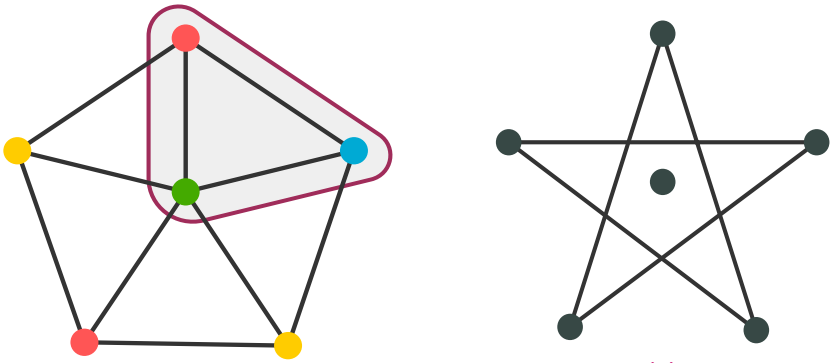


Fig. 1.3.: (a) A graph G with six vertices, where a vertex colouring and a maximum clique are highlighted. The chromatic number $\chi(G)$ is the minimum number of colours needed to colour the vertices such that no connected vertices share the same colour. In this example, $\chi(G) = 4$. The clique number $\omega(G)$ is the size of the largest complete subgraph, which in this case is any of the triangles. (b) The complement of a graph, denoted $\overline{G}$, has the same vertex set as G, but its edges are placed only between nonadjacent vertices of G.

Chromatic number

Vertex colouring is the problem of assigning a number of colours to the vertices of a graph in which no two adjacent vertices share the same colour (Fig. 1.3a). The chromatic number $\chi(G)$ of a graph is the smallest sufficient number of colours for vertex colouring. Determining $\chi(G)$ is also computationally intractable in general. Since a clique of size k requires at least k colours (because any vertex is connected to all others), $\omega(G) \leq \chi(G)$. However, this bound can be arbitrarily bad: There exist constructions, such as the Mycielski graphs [BM08], where $\chi(G)$ can become arbitrarily large while $\omega(G)$ remains constant.

Edge clique cover

Coverings are another category of graph problems that play a key role in optimization. For example, a vertex cover is a set of vertices that touches every edge, and is a special case of the set cover problem. Another covering problem that will prove particularly relevant in this thesis is the edge clique covering problem. Given a graph G, an edge clique cover is a collection of cliques of G whose union contains all edges of G [Orl77]. This problem is studied in combinatorics under the name of covering arrays [CD06], and it has wide-ranging applications in software testing, where test cases correspond to edges that need to be covered by groups of related tests [KKL+10]. In Chaps. 6 and 7, clique covers will also be connected to the parallelization of quantum state tomography protocols. Specifically, cliques can be used to represent sets of jointly measurable observables, and minimizing the clique cover number provides an optimal strategy for parallel measurement execution.

1.3 Mathematical optimization

Optimization is a fundamental concept that appears across a vast range of disciplines, from the physical principles governing natural systems to engineering solutions and the organization of social and economic processes. Nature itself seems to optimize quantities such as energy and entropy, engineers seek to maximize efficiency and reduce costs, logistics systems benefit from the best allocation of resources, and players in the economic market must maximize profit and minimize risks under uncertain data. The study of mathematical optimization dates back to early developments in calculus in the 17th century, gaining momentum with the discovery of linear programming in the 1940s and the creation of computer. Since then, it has grown into a vast research field spanning mathematics, computer science, economics, and operations research.

While optimization in its most general form can be computationally intractable, many strongly structured problems allow for efficient global optimization. A broad and important class is the minimization of convex functions over convex cones, called conic programming. Optimization over the cone of semidefinite matrices, called semidefinite programming (Sec. 1.3.1), has turned into an essential tool in quantum information. Even for nonconvex problems, global optimization is sometimes possible by exploiting discrete structures. A key example is mixed-integer programming, which allows optimization over problems involving both continuous and discrete variables (Sec. 1.3.2). While this has rarely been a tool in quantum information, it can be used to determine certain graph-theoretic quantities that will later appear in the context of quantum state tomography.

1.3.1 Semidefinite programming

Whether we are minimizing the drag force on a bicycle helmet by optimizing its shape, or determining the optimal way to pack kitchen containers inside one another, the underlying optimization problem can be abstractly expressed as

$$p^* = \min \{ f_0(x) \mid f_i(x) \leq b_i, \, \forall i \in \{1, \ldots, m\} \}. \tag{1.26}$$

Feasible set and objective function

Solving it amounts to finding some candidate solution $x^* \in \mathbb{R}^n$ such that the objective function satisfies $f_0(x^*) \leq f_0(x)$, for every x in the feasible set $\{ x \mid \bigcap_{i=1}^m f_i(x) \leq b_i \}$. Semantically, the objective function represents the quantity to be minimized, while the

feasible set encodes the physical constraints (for example, that the shape of the helmet should still fit a head, or that the containers should not be halved).

Finding global optima for nonlinear programs is challenging, and no general-purpose methods are known that are both robust and computationally efficient. To make optimization tractable, one typically restricts attention to specific cases of Eq. (1.26), where additional structure is imposed on the functions. One particularly useful instance of efficiently solvable optimization problems is semidefinite programming (SDP) [VB96; BV04; GM12].

Semidefinite programming is the optimization of a linear function over a spectrahedron — the intersection between the cone of positive semidefinite matrices and a collection of linear inequalities. They can be written as

Semidefinite program

$$\begin{aligned} &\text{given} && c, F_0, \ldots, F_n && \text{(1.27a)}\\ &\max_{x \in \mathbb{R}^n} && c^\intercal x && \text{(1.27b)}\\ &\text{s.t.} && F_0 + \sum_{i=1}^{n} x_i F_i \succeq 0, && \text{(1.27c)} \end{aligned}$$

where $F(x) \equiv F_0 + \sum_{i=1}^{n} x_i F_i \succcurlyeq 0$ is called a linear matrix inequality (LMI), and the positive semidefiniteness constraint ("$\succcurlyeq 0$") forces it to have nonnegative eigenvalues. Furthermore, every matrix F_i in $F(x)$ is symmetric, $x \in \mathbb{R}^n$, and c can specify any linear function on x.

To see that this is indeed an instance of conic programming, one must only notice that positive semidefinite matrices are closed under conic combinations and also convex, since if $F(x) \succcurlyeq 0$ and $F(y) \succcurlyeq 0$, then

$$F[\alpha x + (1-\alpha)y] = \alpha F(x) + (1-\alpha)F(y) \succcurlyeq 0 \tag{1.28}$$

for any $0 \leq \alpha \leq 1$. More than that, it generalizes over linear programming, which can be represented by setting $F_0 = \operatorname{diag}(b)$ and $F_i = \operatorname{diag}(a_i)$ to get the usual $Ax = b$ feasible set defining a polyhedron. In general, $F(x)$ may have a block-diagonal section representing such scalar inequalities, but also positive semidefinite blocks elsewhere, thus it is indeed a spectrahedron. Different from polyhedra, spectrahedra may have some kinds of curved boundaries (for instance, a cylinder [Vin14]). As last interesting comparison with linear programming, notice that, by definition, $F(x) \succcurlyeq 0$ if and only if $v^\intercal F(x) v \geq 0$ for all

Linear programming

real vectors v of appropriate dimension. In this sense, an SDP is a linear program with infinitely many linear constraints on x.

The usefulness of semidefinite programming in quantum information is obvious from the fact that quantum states and measurement effects are Hermitian positive semidefinite matrices. [11] Indeed, semidefinite programs were proven to be useful in entanglement certification (Chap. 8), quantum steering [CS16], characterizing the set of quantum correlations [NPA07; NPA08], and plenty of other problems [SC23; Tav+24].

SDP solvers

Many algorithms and implementations of semidefinite programming solvers exist, both open-source and commercial. As each solver has a different input format, it is convenient to use modelling languages to specify optimization problems in a high-level format. These are either stand-alone programs or packaged as libraries for particular programming languages, and they serve the purpose of internally converting and dispatching optimization models to specific solvers. Semidefinite programs appearing in this thesis were modelled in Julia [Bez+17] using JuMP [Lub+23], and solved using either Mosek [MOS], SCS [ODo+16; ODo21], or Hypatia [CKV21].

1.3.2 Integer programming

Given any convex optimization problem, one can introduce additional constraints requiring certain variables to take integer values. This leads to mixed integer programming (MIP), a widely used and researched generalization of convex optimization [PS98]. Integer programs are generally nonconvex and significantly harder to solve than their continuous counterparts. In fact, even integer linear programming is already an NP-hard problem.

Mixed integer program

Nevertheless, even combinatorially large MIPs can sometimes be solved to global optimality in practice. Consider the program

$$\max_{x\in\mathbb{Z}^n} \; c^\intercal x \tag{1.29a}$$

$$\text{s.t.} \quad Ax \leq b, \tag{1.29b}$$

$$x \geq 0. \tag{1.29c}$$

[11] That they are defined over the complex field while optimization is usually only discussed over the reals is not an issue per se, as one can always map a complex matrix to a real one at expense of additional variables.

Ignoring the integrality constraints ($x \in \mathbb{Z}^n$) leads to a convex relaxation, in this case a linear program. Solving this relaxation leads to an upper bound $p_r^* \geq p^*$ to the optimal value p^*. If, by chance, the optimal x found in the relaxation is already in $\mathbb{Z}^n$, we are done. Otherwise, one can implement a branch-and-bound algorithm on top of the convex solver. Suppose $x_i = \alpha$ is not an integer. Then we create two new linear programs: One with the additional constraint that $x_i \geq \lceil \alpha \rceil$, and another with $x_i \leq \lfloor \alpha \rfloor$. By recursively applying this process to every fractional variables, we construct a search tree. In principle, this would explore the entire search space, but in fact, whenever a branch's optimal solution is poorer than another, it cannot contain an optimal solution, and thus can be pruned.

Relaxation

Branch and bound

In the worst case, the number of branches can still grow combinatorially, so MIP cannot be tractable in the general case. Very often, however, modern solvers can handle very large instances efficiently, leveraging a variety of advanced techniques, such as cutting-plane methods, symmetry detection, and heuristic strategies for variable selection and branching [Lod10]. These solver enhancements significantly reduce the effective search space, enabling the solution of practical problems far beyond what naive combinatorial complexity might suggest. Mixed-integer programs appearing in this thesis were solved with Gurobi [Gur] and SCIP [Bol+24a; Bol+24b], and modelled in Julia [Bez+17] using JuMP [Lub+23].

MIP solvers

1.4 Thesis contributions

Part I: Correlations

Chapter 2: Uncertainty relations from graph theory Given a set $\{ O_i \}$ of observables, how much is $E = \max_\rho \sum_{i=1}^n \operatorname{tr}(O_i \rho)^2$? This may sound arbitrary, but in fact this quantity has appeared in a range of different problems, including bounds on the sample cost of quantum tomography, find ground state properties for particular types of Hamiltonians, characterizing the set of network entangled states, and many others. Our main result is simple but rather surprising: E is bounded by the Lovász number of a graph associated to the anticommutation properties of $\{ O_i \}$.

Chapter 3: High-dimensional steering certification Proving that an assemblage is steerable certifies entanglement even in the situation where one of the parties is untrusted. Can steering also be used to certify stronger properties of the state, such as its entanglement dimension? Our main result in this chapter is a complete

hierarchy of tests that can be formulated as semidefinite programs and be used to determine the Schmidt number underlying any steering assemblage.

Chapter 4: Entanglement-assisted communication In a quantum communication scenario, one qudit of communication assisted by a d-dimensional entangled state can be used to perfectly transmit two dits of information. Does this mean that larger dimensional entanglement is useless in quantum communication protocols? We answer this in the negative, by constructing a class of games in which entanglement of dimension $d + 2$ achieves a better score for any communication dimension d. An auxiliary result of independent interest is that, with unbounded entanglement dimension, any qudit communication behaviour can be achieved by two dits of classical communication.

Part II: Tomography

Chapter 5: Confidence regions Suppose we are given experimental data from a quantum state tomography experiment, and the description of the measurement basis that was used. Can we straightforwardly determine a confidence region $C = \{ \rho \mid \|\hat{\rho} - \rho\| < \epsilon \}$ that contains the true state with high probability, is easy to describe and manipulate, and which is also asymptotically optimal? Based on recent developments in high-dimensional probability theory, we find a procedure that satisfies these criteria, and performs well in simulated tomography experiments.

Chapter 6: Partial state tomography Full state tomography not only is prohibitively expensive in larger dimensions, but is often unnecessary, as many important properties of an n-qubit state can be found in its local neighbourhood or in its few-body reduced states. If only local projective measurements are available — generally the case in practice — what is the minimum number of distinct measurement settings necessary to gather enough data to reconstruct a given list of desired marginals? We show this problem can be mapped to edge clique covering problems, which in particular instances are equivalent to covering arrays. Optimal solutions can be obtained with binary programming, which we compare to a heuristic algorithm and to previous constructions in the literature.

Chapter 7: Partial tomography of fermionic simulators Widely sought for in quantum chemistry, simulations of fermionic systems are classically intractable, as well as in digital quantum computers due to the overhead associated with mapping fermionic statistics to distinguishable qubits. Recent proposals to natively simulate

fermionic systems may provide an alternative direction in the near future. As the result of any computation must be measured, it is crucial to express the readout procedures in terms of the native gates in the simulator, and to understand how scalable they are. Associating this to graph-theoretic quantities and extending the methods of the previous chapter, we show that every single-body fermionic observable, as well as local two-body observables can be efficiently determined.

Part III: Tools

Chapter 8: `Ket.jl` Quantum information theorists often rely on computations, whether for solving optimization problems or gaining numerical insights that guide analytical work. However, implementing even basic operations efficiently can be tedious, and performance is often critical. While Julia excels in numerical computing, it lacked a dedicated quantum information library. `Ket.jl` fills this gap by providing general, fast, arbitrary-precision implementations of key procedures, including subsystem permutations, optimization hierarchies, generation of random objects and much more. In this chapter we introduce Julia and some design principles of `Ket.jl`, before taking a closer look on the implementations of the symmetric extensions hierarchy, and methods for genuine multipartite entanglement certification.

Part I

Correlations

2 Uncertainty relations from graph theory

This chapter is based on *C. de Gois, K. Hansenne, and O. Gühne. "Uncertainty relations from graph theory". In:* Physical Review A *107.6 (2023), p. 062211* [GHG23]. CG and KH equally contributed to these results. These contributions are also discussed in [Han24].

Physical quantities are determined from the results of measurements performed on the state of a system. In quantum theory, observable quantities are described by Hermitian operators, opening up the possibility of having observables that do not commute. A careful analysis of this fact leads to interesting consequences. Consider, for example, a single spin-$1/2$ particle, and let X, Y, Z be the Pauli matrices. Then, for any quantum state ρ, we have $\mathrm{tr}(X\rho)^2 + \mathrm{tr}(Y\rho)^2 + \mathrm{tr}(Z\rho)^2 \leq 1$. Reading between the lines, this is simply telling us that any qubit must be inside the Bloch ball. Consider then the natural extension of this inequality for any set $\mathcal{O} = \{ O_i \}_{i=1}^n$ of observables,

$$\max_{\rho} \sum_{i=1}^{n} \mathrm{tr}\left(O_i \rho\right)^2 \leq \vartheta_{\mathcal{O}}. \tag{2.1}$$

In fact, whenever $O_i O_j + O_j O_i = 0$ for every pair of distinct observables, $\vartheta_{\mathcal{O}} = 1$ [TG05; WW08]. Otherwise, solutions are only available for particular cases or through non-scalable numerical methods [MP14; SŻ19; Abb+16; SDW17; GMR19]. But can we determine $\vartheta_{\mathcal{O}}$ in general, and would that teach us anything useful or interesting?

One useful aspect of $\vartheta_{\mathcal{O}}$ is that it gives a lower bound on the sample cost of shadow tomography (cf. Sec. 6.1.1): Estimating the expectation values of n traceless observables with ± 1 eigenvalues, to error ϵ, and using single-copy measurements requires $\Omega(\frac{n}{\epsilon^2 \vartheta_{\mathcal{O}}})$ samples of ρ [Che+22]. Foundationally, this is connected to the fact that $\vartheta_{\mathcal{O}}$ is also a bound on the minimum variance with which the observables in $\mathcal{O}$ can be simultaneously measured. This holds because for dichotomic observables, $O_i^2 = \mathbf{1}$, so $\min_\rho \sum_i \Delta^2(O_i) = \min_\rho \sum_i \mathrm{tr}(O_i^2 \rho) - \mathrm{tr}(O_i \rho)^2 \leq n - \vartheta_{\mathcal{O}}$. Therefore, in a loose sense, $\vartheta_{\mathcal{O}}$ quantifies how

difficult it is to "group" the observables in $\mathcal{O}$ into a single measurement. Other connections to $\vartheta_{\mathcal{O}}$ appear in ground state problems [GMR19; HO22; AGK24], entanglement in quantum networks [Han+22], monogamy of nonlocality [Kur+11], certification of stabilizer states [BDH24], and the joint numerical range problem [SŻ19].

Motivated by this, in this chapter we will show that Eq. (2.1) can be bounded for any $\mathcal{O}$, either analytically, or by an efficient numerical procedure. Intriguingly, our bound follows from a relation between $\vartheta_{\mathcal{O}}$ and an important quantity in graph theory called the Lovász number [Lov79]. The result is often tight, and when it is not, further connections to graph quantities can be used to control the approximation.

2.1 Bounding the sum of squares

Consider first the case where $\mathcal{O} = \{ O_i \}_{i=1}^n$ contains only dichotomic observables (that is, each O_i has only two distinct eigenvalues). Each O_i is offset and rescaled, if necessary, so that its eigenvalues become ± 1. This can be done without loss of generality, as it can later be reverted. Now let $E_\rho = \sum_{i=1}^n \operatorname{tr}(O_i\rho)O_i$, and consider the quantity $\operatorname{tr}(E_\rho\rho)$ By convexity,

$$\operatorname{tr}\left(E_\rho^2\rho\right) \geq \operatorname{tr}(E_\rho\rho)^2. \tag{2.2}$$

Defining $\vec{o} = (\operatorname{tr}(O_i\rho))_{i=1}^n$, the right-hand-side of the previous inequality becomes $\|\vec{o}\|_2^4$. On the other hand, using that $\operatorname{tr}(O_iO_j\rho)^* = \operatorname{tr}(O_jO_i\rho)$, the left-hand-side becomes

$$\operatorname{tr}\left(E_\rho^2\rho\right) = \sum_{ij} o_io_j \operatorname{tr}(O_iO_j\rho) = \sum_{ij} o_io_j \Re\left[\operatorname{tr}(O_iO_j\rho)\right], \tag{2.3}$$

where $\Re$ stands for the real part of a number. Two features of this sum are noteworthy: when $i = j$, we have $\operatorname{tr}(O_iO_j\rho) = 1$, because the O_i have eigenvalues ± 1. Also, when the observables anticommute ($O_iO_j + O_jO_i = 0$), the term does not contribute to the sum. Interpreting $\operatorname{tr}(O_iO_j\rho)$ as the elements of an $n \times n$ matrix M, we thus have the following properties:

- $M \in \mathbb{R}^{n\times n}$
- $M_{ij} = 0$ for anticommuting O_i and O_j.
- $M_{ii} = 1$.

- M is positive semidefinite, because the right-hand-side of Eq. (2.3) is equal to $\vec{o}^\intercal M \vec{o}$, which is non-negative for any vector $\vec{o}$ due to Eq. (2.2).

Noticing that $\vec{o}^\intercal M \vec{o} \leq \|\vec{o}\|_2 \lambda_{\max}(M)$, and reassessing Eq. (2.2), we get

$$\sum_{i=1}^{n} \mathrm{tr}(O_i \rho)^2 \leq \|\vec{o}\|_2 \lambda_{\max}(M), \tag{2.4}$$

where

$$\sum_{i=1}^{n} \mathrm{tr}(O_i \rho)^2 \leq \max_{M} \lambda_{\max}(M) \tag{2.5a}$$

$$\text{subject to:} \tag{2.5b}$$

$$M \in \mathbb{R}^{n \times n} \text{ and } M \succeq 0, \tag{2.5c}$$

$$M_{ii} = 1, \forall i \in \{1, \ldots, n\}, \tag{2.5d}$$

$$M_{ij} = 0, \forall \{i, j\} : \{A_i, A_j\} = 0. \tag{2.5e}$$

As interesting as it is, this formulation does not reveal how to actually solve the eigenvalue maximization. In fact, it can often be determined analytically, and otherwise by an efficient numerical procedure, but to understand why we must first detour into a seemingly unrelated quantity.

Channel capacity, in information theory, is the maximum transmission rate achievable in a noisy channel. One way to model a noisy channel is to consider a graph G whose vertices represent words in the communication code, and edges exist between any two words that may get confused when passing through the channel. Independent sets in G can be communicated in parallel without errors, therefore the independence number $\alpha(G)$ is a lower bound on the Shannon capacity of G. This does not rule out the possibility of more sophisticated protocols being able to achieve a better rate by using longer strings of code words. Length two codes can be represented by $G \boxtimes G$, where $\boxtimes$ is the strong product of graphs. In turn, $\sqrt{\alpha(G \boxtimes G)}$ is an upper bound on the Shannon capacity of G, with the actual capacity being by extension $\Theta(G) = \lim_{n \to \infty} \sqrt[k]{\alpha(G^{\boxtimes n})}$ [Sha56].

The computational complexity of determining $\Theta(G)$, and its value even for some surprisingly small graphs, remain unknown. On the other hand, there is an efficient way to bound it from above. The Lovász number $\vartheta(G)$ of a graph, is such that $\alpha(G) \leq \Theta(G) \leq \vartheta(G)$. Perhaps its most recognized definition is in terms of orthogonal representations of graphs

[Lov79], but a plethora of others exist. Among them, a lesser known appears in Sec. 33 in [Knu94]:

$$\vartheta(G) = \max_B \lambda_{\max}(B) \tag{2.6a}$$

$$\text{subject to:} \tag{2.6b}$$

$$B \in \mathbb{R}^{n\times n} \text{ and } B \succeq 0 \tag{2.6c}$$

$$B_{ii} = 1,\ \forall i \in \{1,\dots,n\} \tag{2.6d}$$

$$B_{ij} = 0,\ \forall \{i,j\} \in G. \tag{2.6e}$$

Then let there be a graph $\overline{G}_{\mathcal{O}}$, where $\mathcal{O} = \{ O_i \}_{i=1}^n$ is a set of quantum observables. Its vertex set are the observables themselves, and edges exist between pairs of observables that anticommute. Using $\overline{G}_{\mathcal{O}}$ and comparing to Eq. (2.6) immediately reveals our result:

$$\sum_{i=1}^{n} \operatorname{tr}(O_i\rho)^2 \leq \vartheta(\overline{G}_{\mathcal{O}}). \tag{2.7}$$

As sanity checks, notice that if all observables pairwise anticommute, then $\overline{G}$ is a complete graph. Then B in Eq. (2.6) is the identity matrix, and $\vartheta(\overline{G}) = 1$, which was previously known [TG05; WW08]. Conversely, for an empty graph, the $\vartheta(\overline{G}) = n$, yielding the trivial bound $\sum_i \langle A_i \rangle^2 \leq n$.

A variation of the Lovász number allows us to also bound the expression $\sum_{i=1}^n w_i \operatorname{tr}(O_i\rho)^2$, where $\vec{w}$ is a vector of positive real numbers. Called the weighted Lovász number, $\vartheta(\overline{G}, \vec{w})$ is defined just as in Eq. (2.6), but with $\vec{w}$ in the diagonal of B. Our result for dichotomic observables can then be extended for any set of observables by letting $\vec{w} = (\lambda_{\max}(O_1^2), \dots, \lambda_{\max}(O_n^2))$, so that $\sum_i \operatorname{tr}(O_i\rho)^2 \leq \vartheta(\overline{G}, \vec{w})$. Another possibility is that the observables neither commute nor anticommute, in which case our construction would lead to the trivial bound $\vartheta(\overline{G}) = n$. This can be amended by controlling $\|\{O_i, O_j\}\| < \epsilon_{ij}$ and finding the appropriate changes to Eq. (2.7) [GHG23].

2.2 Consequences of graph properties

Another interesting property of $\vartheta(G)$ is the sandwich theorem [Lov86],

$$\omega(G) \leq \vartheta(\overline{G}) \leq \chi(G), \tag{2.8}$$

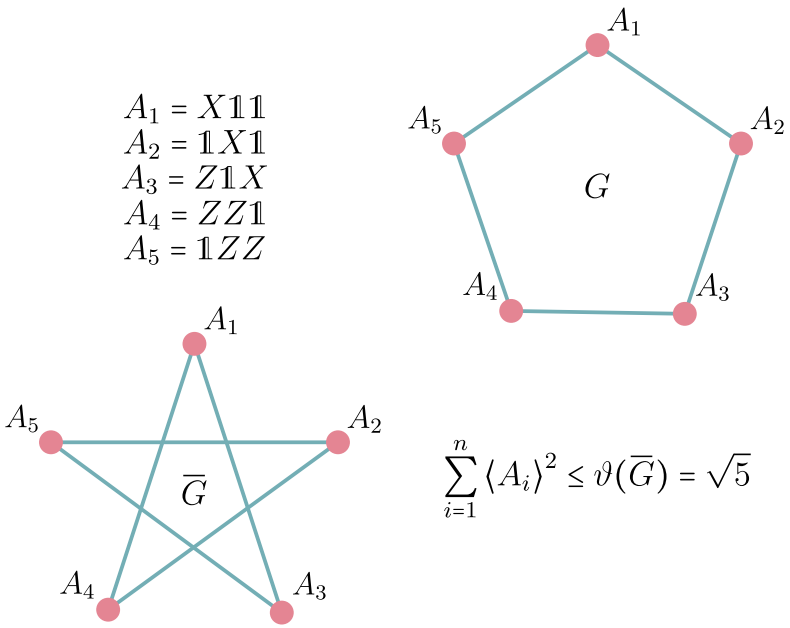


Fig. 2.1.: Given a set of observables $\{O_i\}_{i=1}^n$, we draw a graph $\overline{G}$ with vertices corresponding to each observable and an edge $\{i, j\} \in G$ if and only the observables anticommute, $\{O_i, O_j\} \neq 0$. The Lovász number $\vartheta(\overline{G})$ provides an upper bound to $\sum_i \operatorname{tr}(O_i \rho)^2$. In cases where the observables either commute or anticommute, its complement G is called commutation graph, and $\omega(G)$ is a lower bound on the sum of squares.

where $\chi(G)$ is denotes the chromatic number and $\omega(G)$ is the clique number (see Sec. 1.2). Therefore, the Lovász number provides efficiently computable bounds on two NP-complete quantities.

As discussed above, for pairwise anticommuting observables, $\vartheta = 1$. So if we partition any set of observables into ℓ sets of pairwise anticommuting observables, $\max_\rho \sum \operatorname{tr}(O_i \rho)^2 \leq \ell$ (cf. [Kur+11]). A moment of thought reveals this is simply the chromatic number of the complement of the anticommutation graph (which is the complement of $\overline{G}$). Since there exist graphs for which $\vartheta(\overline{G}) < \chi(G)$, this approach not only is computationally intractable, it can also be rather loose: As a concrete example, Fig. 2.1 illustrates a set of five observables with a star-shaped anticommutation graph. Grouping these observables into pairwise anticommuting subsets yields an upper bound of $\chi(G) = 3$. However, the Lovász number of $\overline{G}$ is $\sqrt{5} \simeq 2.24$. More drastically, it is known that for any $\varepsilon > 0$, there exist graphs satisfying $n^{1-\varepsilon}\vartheta(\overline{G}) < \chi(G)$, where n is the number of vertices in G [Fei95].

This establishes the Lovász bound is a large improvement, but is it tight? A comprehensive characterization can be achieved for the case when the observables are dichotomic and either commute or anticommute: A graph is called perfect if its clique number and its chromatic number are equal for every induced subgraph (and thus equal to the Lovász number). Conjecture in 1961 [Ber61] and finally proven in a whopping 180 pages paper published in 2006 [Chu+06], the strong perfect graph theorem states that a graph is perfect if and only if it contains no cycle with five or more vertices, nor their complements. Interesting perfect graphs include all bipartite graphs, even cycles, and chordal graphs.

Consider $\{ O_i \}$ to be a set of observables that either commute or anticommute — for example, tensors of Pauli matrices — and suppose the associated $\overline{G}$ is perfect. Then $\vartheta(\overline{G}) = \omega(G)$. Since a clique of size $\omega(G)$ in G is a set of pairwise commuting observables, the simultaneous eigenstate of the observables can achieve this value. Therefore, the Lovász bound is tight at least for any such case. However, it can also be loose: For observables $\{O_i\}_{i=1}^5$ with a star anticommutation graph (Fig. 2.1), we can show that $\max_\rho \sum_{i=1}^5 \operatorname{tr}(O_i\, \rho)^2 = \omega(G) = 2$. Indeed, the product observables of the form $\{O_i O_{(i \mod 5)+1}\}_{i=1}^5$ have a fully connected $\overline{G}$, that is, $\sum_{i=1}^5 \operatorname{tr}(O_i O_{i+1})^2 \leq 1$, which can be used to prove the claim.

2.3 Discussion

Plenty of other consequences and applications to deriving entropic steering inequalities, constructing entanglement witnesses, and extensions to observables that do not exactly anticommute are studied in the associated publication [GHG23] and in [Han24]. Nevertheless, many interesting questions are left open.

What is the tight bound in general? Although we had initially believed that it should be $\omega(G)$, this is in fact not the case [XSW24]. Improvements were also found by including further constraints in Eq. (2.6) and formulating the problem as a state polynomial optimization, and graphs with up to seven vertices were extensively analysed [MH24]. On the other hand, the Lovász number is the best that can be done in the sense that for any given anticommutation graph, one can always construct a set of observables satisfying the relations and a state that achieves $\vartheta(\overline{G})$ [XSW24], but for particular sets of observables the bound can be loose. Then is it possible to find the exact bound, and can it be expressed in terms of graph quantities?

Higher order polynomials Another interesting direction for research is to consider higher order polynomials, for example $\max_\rho \sum_i \operatorname{tr}(O_i\rho)^4$ and others. This can be done numerically by extending the ideas in [MH24], and perhaps one could extrapolate from the results an interpretation in terms of other graph quantities.

Further applications Beyond matters of tightness, this result has found interesting applications in shadow tomography [Kin+25], characterizing ground states [AGK24], and certifying stabilizer states [BDH24]. Given its connection to uncertainty relations, which impose fundamental limits on the variances with properties of a state can be simultaneously measured, we expect to find further connections to measurement compatibility, quantum tomography, and entanglement certification. One such case is the application of $\vartheta(\overline{G})$ in constructing witnesses for entanglement dimension [Han24].

Notice also that during the course of publication we became aware of a similar result found in the context of fermionic Hamiltonian optimization [HO22].

3 High-dimensional steering certification

This chapter is based on *C. de Gois, M. Plávala, R. Schwonnek, and O. Gühne. "Complete Hierarchy for High-Dimensional Steering Certification". In:* Physical Review Letters *131 (2023), p. 010201* [Goi+23].

> *"Why are quantum physicists bad drivers?"*
> *"Because they believe in steering without signalling."*
>
> — **Benjamin Yadin**
> Private communication

An ability to produce and control quantum states of increasingly higher dimension can bring us a range of advantages in practice, including an enhanced precision in quantum metrology [GLM06; TA14] and increased communication capacity of quantum channels [BW92; BE95]. More than that, they can also significantly boost the noise robustness of quantum protocols [VPB10; AMP22; Eck+19; Zhu+21; SC15; Mar+15; Qu+22a; Mik+22], which is a crucial feature for overcoming what is one of the main practical challenges in quantum cryptography and the deployment of quantum networks.

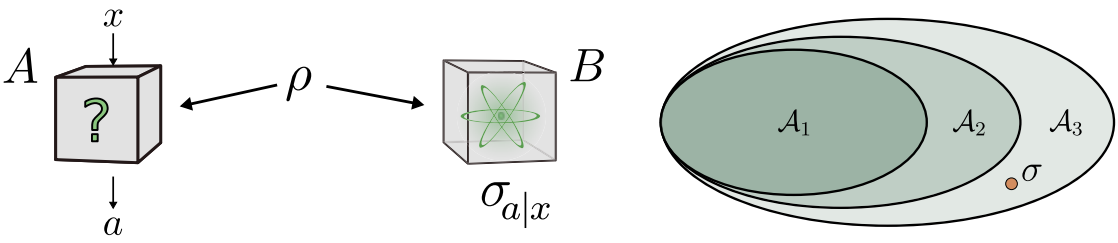
The advantages provided to the applications above often depend on our ability to create and control quantum states that are entangled in many degrees of freedom. An important quantifier of the entanglement dimension is the Schmidt number, introduced in Sec. 1.1.2.1. Operationally, the Schmidt number of a bipartite quantum state ρ corresponds to the minimum local Hilbert space dimension k required for the parties to produce the state. Thus, certifying the Schmidt number has direct implications for verifying the ability of a source of quantum systems to generate high-dimensional entanglement.

Schmidt number certification has been studied in the context of entanglement scenarios for at least the past couple of decades (cf. [TH00; SBL01; Hul+04; Wei+20], to mention a few). However, the possibility of doing the same in quantum steering experiments was only very recently proposed and experimentally verified [Des+21]. Differently to the

entanglement scenario, in a steering experiment only one party, typically Bob, can fully characterize their quantum system, while the other party, Alice, is restricted to black-box measurements, where she can only control the input x and observe the corresponding outcome a (for an introduction to the steering scenario, see Sec. 1.1.2.2). Alice's goal is then to convince Bob that their shared quantum state is of Schmidt number at least k, even though he does not trust her. This setup is particularly relevant for situations where device trust is asymmetric, as it mirrors real-world constraints in a large set of applications, such as randomness certification and quantum key distribution [Uol+20]. High-dimensional steering certification can thus be seen as a fundamental building block for the verification of quantum hardware, where a successful Alice can prove her ability to control and non-trivially manipulate quantum states on a k dimensional Hilbert space.

Designolle et al. [Des+21] studied the construction of steering witnesses that can certify lower bounds for the entanglement dimension in a steering experiment. To do so, they exploited a clever connection between the robustness of steering assemblages and the Schmidt number of the underlying state used in the assemblage preparation. One of the main drawbacks of using steering witnesses is that, like entanglement witnesses [Liu+22], they are always designed to work well for a particular set of resources one wants to certify, but tend to fail otherwise. In their case, the witnesses hold in the case where Alice can only choose between two measurement settings, and were optimized to detect isotropic states measurement with two mutually unbiased bases. This construction was later extended to scenarios with more measurement settings, but the obtained witnesses can be quite weak [Des22; Qu+22b]. Therefore, despite its importance, the certification of high-dimensional steering remains challenging and could so far only be done for particular cases.

In this chapter we will address these limitations by presenting a general and complete hierarchy of semidefinite programs (cf. Sec. 1.3.1) for the certification of high-dimensional steering. Our approach departs from the previous methods [Des+21; Des22; Qu+22b] in two key aspects. First, we base our framework on the characterization of high-dimensional steering assemblages, rather than on constructing steering witnesses. This is a more thorough approach because it accounts for the entire information contained in the assemblage, instead of only the value of a witness, and it also provides insights into the mathematical structure of high-dimensional steering. This can lead to stronger and more general certification criteria, and even though experimentally characterizing an entire assemblage is much more costly than evaluating a witness, the techniques that will be presented can also be used to construct and optimize witnesses for any target



(a) Steering experiment **(b)** k-preparable assemblages

Fig. 3.1.: In a steering experiment, Alice performs local measurements on a bipartite quantum state ρ shared between her and Bob. This results in a set of conditional states on Bob's side [a so-called state assemblage σ, see Eq. (3.1)], which can only be steerable if the shared state ρ were entangled. The central task in high-dimensional steering is to use an assemblage not only to witness entanglement in ρ, but to also quantify it. We say an assemblage is k-preparable assemblage (and write that $\sigma \in \mathcal{A}_k$) if it can be prepared from a state ρ with Schmidt number smaller than or equal to k.

assemblage. Second, our hierarchy is applicable to any number of measurement outcomes and settings, and it is provably complete. This completeness implies that the hierarchy could in principle precisely determine the entanglement dimension certified by a given state assemblage, provided all levels of the hierarchy are considered. Even though it is not possible to do so, we demonstrate that in practice even the first level of the hierarchy can provide significantly tighter bounds on the noise robustness of high-dimensional assemblages than previously known, thereby broadening the scope of high-dimensional steering certification in realistic scenarios.

3.1 The steering scenario

We briefly recall the definition of a steering scenario, as described in Sec. 1.1.2.2 and illustrated in Fig. 3.1a. Steering is a setting where Alice and Bob share parts of a quantum state ρ that they can locally measure. Her goal is to convince Bob that she can "steer" his system by performing local measurements on her share of the state. Similarly to the nonlocality scenario, we consider Alice untrustworthy, therefore her measurements are not characterized. Her only input to the experiment is thus a label $x \in \{1, \ldots, X\}$, indicating a measurement choice. Upon choosing an x, she observes an outcome a from some finite set of classical labels. Bob, on the other hand, has complete control over his laboratory.

Suppose Alice's unknown measurements are denoted by $\mathcal{M}_x = \{M_{a|x}\}$, where $M_{a|x} \succcurlyeq 0$ are the measurement effects and $\sum_a M_{a|x} = \mathbf{1}$. After she chooses x and obtain a, quantum theory describes Bob's resulting, sub-normalized state as

$$\sigma_{a|x} = \mathrm{tr}_A \left[\rho \left(M_{a|x} \otimes \mathbf{1}_B \right) \right]. \tag{3.1}$$

Their collection $\sigma = \{\sigma_{a|x}\}$ is called a steering assemblage. Since Alice cannot communicate with Bob during the experiment, we require the assemblages to be non-signalling: $\sum_a \sigma_{a|x} = \sum_a \sigma_{a|x'}$, for any choice of measurements x and x'.

In Sec. 1.1.2.2 we discussed how whenever an assemblage σ cannot be explained by a classical theory — or more precisely, by the local hidden states model in Eq. (1.20) — then the underlying state ρ must have been entangled. Therefore, any steerable σ certifies entanglement of the shared state. Entanglement, though, comes in many forms. In particular, it can be quantified, but steerability does not provide insights into how entangled ρ is. A topic of recent discussion is whether it is also possible to quantify the entanglement dimension in a steering scenario. The Schmidt number [TH00] is an adequate and useful quantifier for this, as it naturally quantifies the number of entangled degrees of freedom a state has (but notice that other concepts of entanglement dimension also exist [Kra+18]). This means that if a bipartite state of local dimension d has Schmidt number 3, one can say it is only entangled in 3 of its d degrees of freedom and at the same time also certify that $d > 2$.

Let us recall (cf. Sec. 1.1.2.1) that $|\psi\rangle$ is a pure bipartite state with a Schmidt decomposition $|\psi\rangle = \sum_{i=1}^{k} \eta_i |i_A\rangle |i_B\rangle$ with k non-zero coefficients, then it has Schmidt rank k. In the case of mixed states, we say that ρ has Schmidt number $\mathrm{sn}(\rho) = k$ if (i) for any decomposition $\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|$, at least one of the vectors $|\psi_i\rangle$ has Schmidt rank at least k, and (ii) there exists a decomposition of ρ with all vectors $|\psi_i\rangle$ having Schmidt rank at most k. The sets $\mathcal{S}_k = \{\rho \mid \mathrm{sn}(\rho) \leq k\}$, are then convex, their extremal points are pure states, and $\mathcal{S}_{k-1} \subset \mathcal{S}_k$ for all $k \geq 2$, where $\mathcal{S}_1$ is the set of separable states. Extending this definition to the case of steering assemblages [Des+21], we will say that σ is k-preparable if it can be obtained from local measurements on a state $\rho \in \mathcal{S}_k$, and we abuse the notation $\mathrm{sn}(\sigma) \leq k$ to mean so. So for example $\mathrm{sn}(\sigma) \leq k$ if and only if $\sigma_{a|x} = \mathrm{tr}_A \left[\rho \left(M_{a|x} \otimes \mathbf{1}_B \right) \right]$ for some $\rho \in \mathcal{S}_k$. Finally, we also define the sets of assemblages that can be prepared with states in $\mathcal{S}_k$ as $\mathcal{A}_k = \{\sigma \mid \mathrm{sn}(\sigma) \leq k\}$. As illustrated in Fig. 3.1b, the sets $\mathcal{A}_k$ exhibit a nested structure $\mathcal{A}_{k-1} \subset \mathcal{A}_k$, for any $k \geq 2$ [Des+21].

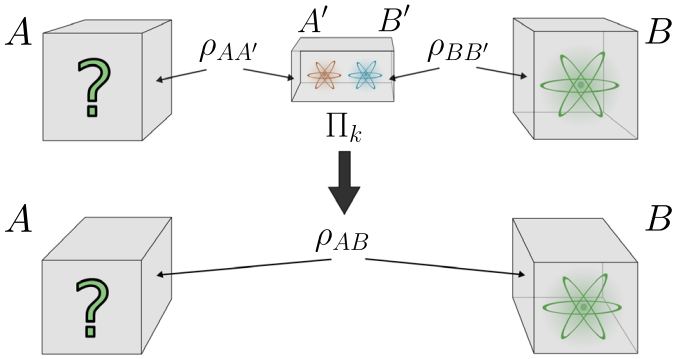


Fig. 3.2.: Any state ρ with Schmidt number k acting on $\mathcal{B}(\mathcal{H}_A \otimes \mathcal{H}_B)$ is equivalent to some separable operator $\sum_i p_i \rho^i_{AA'} \otimes \rho^i_{B'B}$ in an extended Hilbert space, which is then projected with $\Pi_k = \mathbf{1} \otimes \sum_{i=1}^k |i,i\rangle_{A'B'}$ according to Eq. (3.2).

3.2 Certification hierarchy

The central problem in high-dimensional steering is to characterize the sets $\mathcal{A}_k$ of k-preparable assemblages. In practice, this would allow us to determine whether any given assemblage is k-preparable or not. Our main result in this chapter is a hierarchy of tests solving this problem. To construct it, suppose that $\sigma = \{\sigma_{a|x}\}$ is an assemblage in $\mathcal{A}_k$, and so it can be prepared with a state ρ of Schmidt number k. Any state $\rho \in \mathcal{S}_k$ can be obtained from a separable state embedded in an extended Hilbert space [Hul+04]: Following the notation introduced in Sec. 1.1.2.1, we extend Alice's (A) and Bob's (B) subsystems with the k-dimensional subsystems A' and B' and define $\Pi_k = \mathbf{1} \otimes \sum_{i=1}^k |i,i\rangle_{A'B'}$. Then, any pure state $|\psi\rangle_{AB} = \sum_{i=1}^k w_i |\alpha_i\rangle_A |\beta_i\rangle_B$ can be obtained from $|\psi\rangle = \Pi_k^\dagger (|\alpha'_i\rangle_{AA'} \otimes |\beta'_i\rangle_{B'B})$. Consequently, for mixed states,

$$\rho = \Pi_k^\dagger \underbrace{\Big(\sum_i p_i \rho^i_{AA'} \otimes \rho^i_{B'B}\Big)}_{\Omega} \Pi_k. \tag{3.2}$$

This means that any state ρ with Schmidt number k can be associated to an operator Ω which is separable with respect to the $AA' \mid B'B$ partition and normalized to k, at the expense of an entangling projection Π_k over the subsystems in A' and B' (Fig. 3.2).

Substituting Eq. (3.2) into the definition of an assemblage element (Eq. (3.1)), it follows that

$$\begin{aligned}\sigma_{a|x} &= \mathrm{tr}_A\left\{\left[\Pi_k^\dagger\left(\sum_i p_i\rho^i_{AA'}\otimes\rho^i_{B'B}\right)\Pi_k\right]\left[M_{a|x}\otimes\mathbf{1}_B\right]\right\}\\ &= \Pi_k^\dagger\left[\sum_i p_i\mathrm{tr}_A\left(\rho^i_{AA'}\left[M_{a|x}\otimes\mathbf{1}_{A'}\right]\right)\otimes\rho^i_{B'B}\right]\Pi_k\\ &= \Pi_k^\dagger\Big[\underbrace{\sum_i p_i\tau^{i,a,x}_{A'}\otimes\rho^i_{B'B}}_{\Omega_{a|x}}\Big]\Pi_k.\end{aligned} \tag{3.3}$$

Notice that in the second equality we are reusing the notation for Π_k because it anyway acted trivially on the subsystem in A, and that in the third equality we define the operators $\tau^{i,a,x}_{A'}$. These operators are essentially the result of Alice measuring on subsystem A, and therefore their collection is an assemblage. In particular, the non-signalling condition $\sum_{i,a}\tau^{i,a,x}_{A'}=\sum_{i,a}\tau^{i,a,x'}_{A'}$, $\forall x\neq x'$ must hold.

This maps the problem of whether an assemblage $\{\sigma_{a|x}\}$ is k-preparable, to the problem of finding operators $\Omega_{a|x}\in\mathcal{L}\left(\mathcal{H}_k\otimes\mathcal{H}_k\otimes\mathcal{H}_B\right)$ such that $\sigma_{a|x}=\Pi_k^\dagger\Omega_{a|x}\Pi_k$, where $\Omega_{a|x}=\sum_i p_i\tau^{i,a,x}_{A'}\otimes\rho^i_{B'B}$ is a separable operator against $A'\mid B'B$, and its part in the A' space, $\{\tau^{i,a,x}_{A'}\}$, is a non-signalling assemblage.

More than an interesting insight into the structure of k-preparable assemblages, there is a way of determining the existence of the $\Omega_{a|x}$ by means of a complete hierarchy of semidefinite tests. This hierarchy can be seen as a generalization of the symmetric extensions criterion for entanglement certification [DPS04; DPS02] (cf. Secs. 1.1.2.1 and 8.2), but that has the additional non-signalling constraints.

Given $N\geq 1$, we will say $\Omega_{a|x}$ has a *non-signalling symmetric extension* of order N if there exists an operator $\Xi_{a|x}\in\mathcal{B}\left(\mathcal{H}_{A'}\otimes\mathcal{H}^{\otimes N}_{B'B}\right)$ such that

$$\mathrm{tr}_{(B'B)_i\,:\,i\neq m}\left(\Xi_{a|x}\right)=\Omega_{a|x},\ \forall m\in\{1,\ldots,N\}, \tag{3.4a}$$

$$\sum_a\Xi_{a|x}=\sum_a\Xi_{a|x'},\ \forall x,x'. \tag{3.4b}$$

The first condition means the same as in the symmetric extensions hierarchy, therefore ensuring $\Omega_{a|x}$ is separable against $A'\mid B'B$, while the second condition enforces the non-signalling property in the extensions.

With the following theorem, we link the existence of symmetric extensions with that of the $\Omega_{a|x}$ operators described in Eq. (3.3).

Proposition 3.1. *An assemblage $\{\sigma_{a|x}\}$ is k-preparable if and only if there are corresponding operators $\Omega_{a|x} \in \mathcal{B}(\mathcal{H}_{A'} \otimes \mathcal{H}_{B'B})$ such that $\mathcal{H}_{A'}$, $\mathcal{H}_{B'}$ are of dimension k, $\Omega_{a|x}$ has a non-signalling symmetric extension [Eq.* (3.4)*] to any order N, and reproduces the assemblage elements $\sigma_{a|x} = \Pi_k^\dagger \Omega_{a|x} \Pi_k$.*

Proof. The "only if" condition, essential for the applications described ahead, is easily seen to be satisfied by taking $\Xi_{a|x} = \sum_i p_i \tau_{A'}^{i,a,x} \otimes \left(\rho_{B'B}^i\right)^{\otimes N}$, where we, without loss of generality, assumed that $\mathrm{tr}\left(\rho_{B'B}^i\right) = 1$. We postpone the other direction of the proof to Sec. 3.6. □

Theorem 3.1 tells us that, to test if an assemblage $\{\sigma_{a|x}\}$ is not k-preparable, we can check whether its elements have non-signalling symmetric extensions to some order N. That is, given $\{\,\sigma_{a|x}\,\}$, a Schmidt number k and a hierarchy level N, we must search for extensions $\{\,\Xi_{a|x}\,\}$ that satisfy the constraints in Eq. (3.4) together with $\sigma_{a|x} = \Pi_k^\dagger \Omega_{a|x} \Pi_k$, and the normalization $\mathrm{tr}\left(\sum_a \Omega_{a|x}\right) = k$ which comes from the embedding in Eq. (3.2). The normalization comes from the embedding of a Schmidt number k state into a separable operator of trace k in the enlarged space (see Sec. 1.1.2.1).

These tests can be straightforwardly expressed as a semidefinite program (Sec. 1.3.1). If there is no solution to the program, this means the symmetric extension does not exist, therefore $\{\sigma_{a|x}\}$ is not k-preparable. Otherwise, the test is inconclusive, and we must try a higher N. This hierarchy is proven complete in the sense that an assemblage is k-preparable if and only if it admits non-signalling extensions for every natural N. Although of important theoretical significance, in practice one cannot numerically test large values for N. To show the criterion is nevertheless useful even for the smallest possible extensions, we implement and compare it to previous literature in the next section.

3.3 Example application

So that we can compare our technique to previous results [Des+21; Des22; Qu+22b], let us consider the assemblage prepared from the maximally entangled state $|\phi_d^+\rangle =$

$1/\sqrt{d}\sum_{i=1}^{d}|i,i\rangle$ under a set of X mutually unbiased bases (MUB) measurements in dimension d [Dur+10] (see also Sec. 1.1.1),

$$\sigma = \left\{ \left\{ \mathrm{tr}_A \left[\left|\phi_d^+\right\rangle\!\left\langle\phi_d^+\right| \cdot \left(M_{a|x} \otimes \mathbf{1}_B \right) \right] \right\}_{a=1}^{d} \right\}_{x=1}^{X}. \tag{3.5}$$

We are interested in finding at what point on the line between σ and the uniform noise assemblage $\sigma_{\mathbf{1}}$ (the one with elements $1/d^2$), σ becomes k-preparable. Mixing these two assemblages at a visibility η gives us a family of assemblages with elements $\sigma_{a|x}(\eta) = \eta\sigma_{a|x} + (1-\eta)\mathbf{1}/d^2$. Since $\sigma_{\mathbf{1}}$ is always a feasible point in the program described above, one may rewrite it as a maximization on η:

$$\text{given } \{\sigma_{a|x}\},\ k,\ N \tag{3.6a}$$

$$\max_{\{\Xi_{a|x}\}} \eta \tag{3.6b}$$

$$\text{s.t.}$$

$$\Pi_k^\dagger \left[\mathrm{tr}_{(B'B)_i\,:\,i\neq m} \left(\Xi_{a|x} \right) \right] \Pi_k = \sigma_{a|x}(\eta),\ \forall m \in \{1,\ldots,N\} \tag{3.6c}$$

$$\sum_a \Xi_{a|x} = \sum_a \Xi_{a|x'},\ \forall x, x' \tag{3.6d}$$

$$\mathrm{tr}\left(\sum_a \Omega_{a|x} \right) = k. \tag{3.6e}$$

Notice that the variables $\Xi_{a|x}$ implicitly depend on k, which defines their dimension, and that constraint (3.6c) must be satisfied for every a and x. Any solution $\eta^* < 1$ to this maximization problem certifies that the assemblage $\{\sigma_{a|x}\}$ does not have non-signalling symmetric extensions, and therefore that $\mathrm{sn}(\sigma) > k$. In practice one might also be interested in obtaining a series of k-preparability bounds, say L_k, for a particular witness expression $\{F_{a|x}\}$. These bounds can be seen as a generalization of a steering witness, and they must be such that, for any k-preparable assemblage, $\sum_{a,b}\mathrm{tr}\left(F_{a,b}\sigma_{a,b}\right) \leq L_k$. Therefore, any violation of this inequality certifies $\mathrm{sn}(\sigma) > k$. To determine L_k, one can simply substitute the objective function [Eq. (3.6a)] for $\sum_{a,b}\mathrm{tr}\left(F_{a,b}\sigma_{a,b}\right)$, and the right-hand-side of Eq. (3.6c) for variables $\sigma_{a,x}$.

An issue with Eq. (3.6) is that even for small values of the putative Schmidt number k and hierarchy level N, the size of the program increases super-exponentially. This can be partially mitigated by using the same argument as in Sec. 1.1.2.1, allowing us to consider extensions in the symmetric subspace without loss of generality. First, we observe that since the extensions are of the form $\Xi_{a|x} = \sum_i p_i \tau_{A'}^{i,a,x} \otimes \left(\rho_{B'B}^i\right)^{\otimes N}$ then it

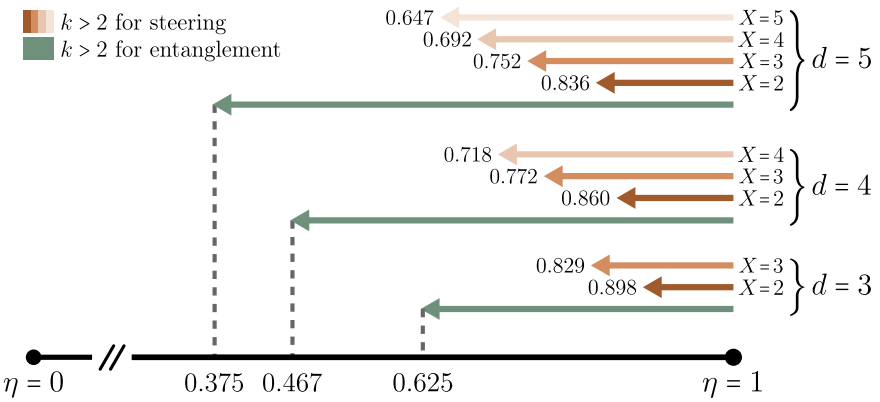


Fig. 3.3.: Schmidt number (k) certification for assemblages prepared with maximally entangled states and X measurements on MUBs. The assemblages were mixed with the white noise assemblage and an upper bound on the minimum visibility such that 2-preparability can be falsified was computed. All results were computed with two copies of the $B'B$ system (i.e., $N = 2$). Increasing N can lead to smaller bounds. For instance, in the $d = 3, X = 3$ case, by setting $N = 3$ we were able to certify that $k > 2$ for all $\eta \geq 0.789$. For comparison, we also show analogous bounds for the entanglement scenario, obtained through a similar procedure: To certify the Schmidt number in an entanglement scenario, the standard symmetric extensions hierarchy can be directly applied to the embedding into a separability problem, as discussed in Secs. 1.1.2.1 and 8.2.

leads to no loss of generality to assume that $\rho^i_{B'B}$ are pure states and thus $\Xi_{a|x}$ acts on $S = \mathcal{H}_{A'} \otimes \mathrm{Sym}_N(B'B)$, where $\mathrm{Sym}_N(B'B)$ is the symmetric subspace of $\mathcal{B}(\mathcal{H}^{\otimes N}_{B'B})$. This space is of dimension $d_S = d_{A'} \times \binom{kd_B+N-1}{kd_B-1}$ (cf. e.g. [Wat18]), considerably reducing the number of variables in our problem. This observation can be put to use by defining the symmetric subspace projector

$$\Pi_S = \mathbf{1}_{A'} \otimes \frac{1}{n!} \sum_{\pi \in S_n} P_\pi, \tag{3.7}$$

where S_N is the N-elements permutation group, and from it constructing a matrix P_S whose d_S rows span the symmetric subspace. Then we can modify the variables in the program above to be the (much smaller) $d_S \times d_S$ matrices $\Xi'_{a|x}$ for each assemblage element, and substitute $\Xi_{a|x}$ for $P_S^\dagger \Xi'_{a|x} P_S$ in all constraints. More than that, due to the permutational symmetry, the constraint in Eq. (3.6c) needs not to be evaluated on every m, but rather for any single one of the $B'B$ subsystems. Thus, this not only reduces the dimension on the optimization variables, but also largely decrease the number of constraints. A method for computing P_S in practice is discussed in Sec. 8.2.

Numerical results obtained with an implementation of these ideas are shown in Fig. 3.3. In addition to the extensions constraints, we have also added the use of positive partial transpose constraints — as is common to accelerate convergence in the symmetric extensions hierarchy [NOP09b]. The results were obtained with the lowest hierarchy level ($N = 2$), and even though Theorem 3.1 only guarantees convergence in the limit $N \to \infty$, they are already of practical significance. Remarkably, this lowest level already provide better bounds than the known values for high-dimensional steering witnesses with $X > 2$ measurements [Des22, Table III], except for the $d = 5$, $X = 3$ case. In comparison to the previous results for $X = 2$ measurements [Des+21], level $N = 2$ of our hierarchy provides slightly worse visibilities, but moving to level $N = 3$ makes up for it. For example, while they find a visibility of 0.886 (0.838) for dimension 3 (4), our method improves it to 0.844 (0.811).

3.4 Symmetrization

Although the computations presented in Fig. 3.3 were computed in a standard computer [except for $(d = 5, X = 5)$], increasing the desired Schmidt number k or the dimension d rapidly leads to instances that are too large to compute without more resources or specialized implementations. An interesting and useful way to reduce the computational cost in particular cases is by exploiting symmetries of the assemblages. This is the case when the measurements in the assemblage preparation are assumed to have some particular symmetry, and may drastically reduce the number of variables and constraints in the problem. Here we show a simple example where we exploit the unitary symmetry of the MUB measurements, but notice that much more thorough investigations on identifying and using measurement symmetries in similar cases exist [Ngu+20].

For simplicity, we follow the example from Sec. 3.3 and assume the assemblage was prepared with MUB measurements. When the dimension d is prime, sets of d MUBs can be obtained through rotations of the Fourier matrix F_d (as defined in [Dur+10, Eq. 5.6]), generated by powers of the Hadamard matrices E_d (ibid., Eq. B.1). In this case, all of our measurements' effects $\{M_{a|x}\}_{a,x}$ can be expressed from a generating set $\{M_{a|1}\}_a$ as

$$M_{a|x} = E_d^{(x-1)} M_{a|1} \left(E_d^{(x-1)}\right)^\dagger, \forall x \in \{1, \ldots, d\}. \tag{3.8}$$

The assemblage's elements then inherit the property $\sigma_{a|x} = (E_d^{(x-1)})^\dagger \sigma_{a|1} E_d^{(x-1)}$, and the relevant constraints in Eq. (3.6) can consequently be rewritten as

$$\Pi_k^\dagger \Big(\mathrm{tr}_{(BB')_i \,:\, i \in [N]\setminus m} \left[{\mathcal{E}_d^{(x-1)}}^\dagger \Xi_{a|1} (\mathcal{E}_d^{(x-1)}) \right] \Big) \Pi_k = \sigma_{a|x}, \tag{3.9}$$

and

$$\sum_a \Xi_{a|1} = \sum_a {\mathcal{E}_d^{(x-1)}}^\dagger \Xi_{a|1} (\mathcal{E}_d^{(x-1)}), \tag{3.10}$$

where $\mathcal{E}_d = \mathbf{1}_{A'} \otimes (\mathbf{1}_{B'} \otimes E_d)^{\otimes N}$, and the equations hold for every x.

Directly from this simple observation the number of variables that we must account for is reduced by a factor of X — the number of possible measurements. Though not strictly necessary to obtain the results presented in Fig. 3.3, an implementation accounting for these symmetries might lead to a significant speed-up in the running time. As mentioned in the introduction, more sophisticated reductions can be obtained for example by applying the methods of Nguyen et al. [Ngu+20] also to this case.

3.5 Robustness of high-dimensional steering

The optimal visibility η^* described in Sec. 3.3 for an assemblage to be certified can be interpreted as the distance between the assemblage and the set $\mathcal{A}_k$, measured along the line segment connecting σ to the white noise assemblage. A particularly interesting alternative choice of noise model to be considered is that of all k-preparable assemblages. Borrowing ideas from the entanglement and steering robustnesses [VT99; PW15], we define the k-preparability robustness w.r.t. noise model $\mathcal{N}$ as the real number $R_{\mathcal{N}}(\sigma, k) = t^*$ resulting from the following program, in the limit $N \to \infty$.

$$\min t \tag{3.11a}$$

s.t.

$$\frac{\sigma_{a|x} + t\pi_{a|x}}{1+t} = \sigma^k_{a|x} \tag{3.11b}$$

$$\Pi_k^\dagger \Big(\mathrm{tr}_{(BB')_i \,:\, i \neq m} \left[\Xi_{a|x} \right] \Big) \Pi_k = \sigma^k_{a|x}, \tag{3.11c}$$

$$\sum_a \Xi_{a|x} = \sum_a \Xi_{a|x'}, \ \forall x, x'. \tag{3.11d}$$

Since the minimization ranges over $\Xi_{a,x}$, t and the noise operators $\pi_{a|x} \in \mathcal{N}$, this is not a semidefinite program. Nevertheless, when $\mathcal{N}$ is any set defined by a linear matrix inequality (LMI), it can be turned into one. To do so, we rewrite Eq. (3.11b) as

$$\pi_{a|x} = \frac{1}{t}\left[(1+t)\,\sigma^k_{a|x}\right] - \sigma_{a|x} \in \mathcal{N}, \tag{3.12}$$

then redefine $\widetilde{\mathcal{N}} \equiv \left\{t\pi_{a|x} \mid t \geq 0, \pi_{a|x} \in \mathcal{N}\right\}$ and $\widetilde{\Xi}_{a|x} \equiv (1+t)\Xi_{a|x}$. Because $\widetilde{\mathcal{N}}$ is a scalar multiplication of an LMI, it is itself an LMI. And since $\mathrm{tr}\left\{\sum_a \widetilde{\Xi}_{a|x}\right\} = 1 + t$ for any choice of x, an equivalent expression for $R_{\mathcal{N}}(\sigma, k)$ is

$$\min \mathrm{tr}\Big(\sum_a \widetilde{\Xi}_{a|x}\Big) - 1 \tag{3.13}$$

s.t.

$$\Pi_k^\dagger\Big(\mathrm{tr}_{(BB')_i\,:\,i\in[N]\setminus m}\left[\widetilde{\Xi}_{a|x}\right]\Big)\Pi_k - \sigma_{a|x} \in \widetilde{\mathcal{N}} \tag{3.14}$$

$$\sum_a \widetilde{\Xi}_{a|x} = \sum_a \widetilde{\Xi}_{a|x'},\ \forall x, x', \tag{3.15}$$

where any x may be chosen in the objective function.

We have already shown that the set of Schmidt number k assemblages ($\mathcal{A}_k$) can be approximated by means of a similar hierarchy of semidefinite programs. Thus, any choice N for the hierarchy level will provide an upper bound $R^N_{\mathcal{N}}(\sigma, k) \geq R_{\mathcal{N}}(\sigma, k)$. By taking $\mathcal{N} = \mathcal{A}_k$, we can interpret $R_{\mathcal{A}_k}(\sigma, k)$ as the minimal distance between the assemblage σ and the set of Schmidt number k assemblages.

3.6 Proof of completeness

The proof for the "only if" direction of Prop. 3.1 is discussed in the main text, so we are left with proving the other direction. At a first glance, we can identify at least three promising approaches for the proof. One would be to try to adapt the completeness proof for the symmetric extensions hierarchy [DPS04; DPS02], which rely on arguing that a quantum de Finetti theorem [CFS02] can be applied to symmetrically extendable states. A second approach would be to use the recent results on the completeness of bilinear semidefinite programs [Ber+22], which encompass symmetric extensions with additional affine constraints. In even more generality, a similar statement was recently shown to

hold not only for the positive semidefinite cone, but for any cone [AMP22, Thm. 1]. This is the approach we follow.

Before starting, let us briefly introduce some necessary notions to understand the statement in [AMP22, Thm. 1].

Proper cones Let V be a real, finite-dimensional vector space and $C \subset V$ be a convex cone that is also pointed and generating cone. Therefore, C is such that

1. For every $x \in C$ and real $\lambda \geq 0$, we have $\lambda x \in C$,
2. $C \cap -C =$, and
3. $C - C = V$.

A cone with these properties will be called a *proper* cone. An example is the cone of positive semidefinite operators.

Dual cones The dual cone C^* of C is defined as the cone of all linear functions that are positive on elements of C. Therefore, $\psi \in C^*$ if and only if $\psi(x) \geq 0$ for any $x \in C$.

Generalized trace A functional $\varphi \in V^*$ is called a generalized trace on C if $\varphi \in C^*$ and for every $x \in C$, we have $\varphi(x) = 0 \Rightarrow x = 0$. The usual trace of an operator is an example.

Minimal tensor product Let $C_1 \subset V_1$ and $C_2 \subset V_2$ be proper cones over finite-dimensional real vector spaces. The minimal tensor product is defined as

$$C_1 \otimes_{\min} C_2 = \operatorname{conv}(\{x_1 \otimes x_2 : x_1 \in C_1, x_2 \in C_2\}), \tag{3.16}$$

where "conv" is the convex hull. The minimal tensor product can be understood as the set of all separable tensors of C_1 and C_2.

Maximal tensor product The maximal tensor product is defined as

$$C_1 \otimes_{\max} C_2 = (C_1^* \otimes_{\min} C_2^*)^*, \tag{3.17}$$

where C_1^* and C_2^* are the dual cones of C_1 and C_2. It can be understood as the largest possible tensors that are positive on separable combinations of the duals. The maximal tensor product is analogous to the set of entanglement witnesses.

Then let φ_2 be a generalized trace on C_2. Suppose $x \in C_1 \otimes_{\max} C_2$ is such that, for any N, there is an $x_N \in C_1 \otimes_{\max} C_2^{\otimes_{\max} N}$ such that:

1. x_N is invariant with respect to permutations of the N copies of C_2.

2. $(\varphi_2^{N-1} \otimes \mathrm{id}_1 \otimes \mathrm{id}_2)(x_N) = x$, where $\mathrm{id}_1, \mathrm{id}_2$ is the identity map on V_1, V_2 respectively. Here, $(\varphi_2^{N-1} \otimes \mathrm{id}_1 \otimes \mathrm{id}_2)(x_N)$ is a generalized partial trace of x_N where we trace out the first $N-1$ copies of C_2.

Theorem 1 in [AMP22] then states that, under these conditions, we have $x \in C_1 \otimes_{\min} C_2$.

It is now left to show that their result applies in particular to our case.

Proposition 3.1. *An assemblage $\{\sigma_{a|x}\}$ is k-preparable if and only if there are corresponding operators $\Omega_{a|x} \in \mathcal{B}(\mathcal{H}_{A'} \otimes \mathcal{H}_{B'B})$ such that $\mathcal{H}_{A'}$, $\mathcal{H}_{B'}$ are of dimension k, $\Omega_{a|x}$ has a non-signalling symmetric extension [Eq.* (3.4)*] to any order N, and reproduces the assemblage elements $\sigma_{a|x} = \Pi_k^\dagger \Omega_{a|x} \Pi_k$.*

Proof of Thm. 3.1. In our case, C_2 is the usual cone of positive semidefinite operators generated by quantum states, but C_1 is generated by assemblages $\{\sigma_{a|x}\}$ with $a \in \{1, \ldots, A\}$ and $x \in \{1, \ldots, X\}$, that is, C_1 has elements $\{\lambda\sigma_{a|x}\}$ for $\lambda \geq 0$. It can be directly verified that C_1 is a closed, convex, and pointed cone. Letting $V_1 = \mathrm{span}(C_1)$ we get that C_1 is also a generating cone. Furthermore, we choose the generalized trace in C_1 to be defined as

$$\varphi_1(\{\sigma_{a|x}\}) = \sum_a \mathrm{tr}\Big\{\sigma_{a|x}\Big\}. \tag{3.18}$$

Notice that the value of the generalized trace does not depend on x, since every assemblage is non-signalling.

The dual cone C_1^* is generated by functionals of the form

$$\Psi(\{\sigma_{a|x}\}) = \mathrm{tr}\Big\{F\sigma_{a|x}\Big\} \tag{3.19}$$

where $F \geq 0$.

It is easy to see that $\Omega_{a|x} \in C_1 \otimes_{\max} C_2$ because $\mathrm{tr}\Big\{\Omega_{a|x}(E \otimes F)\Big\} \geq 0$ for all $F \geq 0$ and $E \geq 0$. Notice also that non-signalling symmetric extensions of $\Omega_{a|x}$ of order N

coincide with an element $x_N \in C_1 \otimes_{\max} C_2^{\otimes_{\max} N}$ satisfying conditions 1 and 2 above. From [AMP22, Thm. 1], it must then be true that

$$\Omega_{a|x} = \sum_i p_i \tau^i_{a|x} \otimes \rho^i \tag{3.20}$$

where $\tau^i_{a|x}$ is a set of assemblages, ρ^i are quantum states, and $\sum_i p_i = 1$. □

3.7 Discussion

Several implications and potential research directions emerge from the results presented in this chapter. Here, we highlight some of the most notable.

Measurement k-simulability Steering is known to be in one-to-one correspondence with measurement incompatibility [QVB14; Uol+15], a connection which was recently extended to high-dimensional steering and k-simulatability of measurements [Jon+22], as defined in [Ioa+22]. In this sense, the hierarchy presented in this chapter also solves the problem of quantifying measurement incompatibility in terms of a dimension. On a more practical side, it is also possible to use the infeasibility certificates from our hierarchy obtain high-dimensional steering witnesses.

Computational cost Ultimately, the size of the semidefinite program that must be solved in our tests is dictated by the number of measurements and outcomes, the tentative Schmidt number k and the hierarchy level N. Irrespective of the efficiency of semidefinite programming, the Hilbert space dimension grows exponentially with k and N, and the problem becomes intractable. A natural further step is thus to find cheaper relaxations or to adapt the formulation to specific problems by making use of further symmetries, as done for similar problems in [Ngu+20; Agu+18; TRR19; IR21]. These ideas could also be applied to a recently proposed alternative hierarchy for the same problem [DCT25]. Its main advantage is that there is no dependency on k, with the drawback that it is not known to converge. Another direction is to consider extensions on the operators in $\mathcal{H}_{A'}$ instead of $\mathcal{H}_{B'b}$ as we have done. This is interesting because $\mathcal{H}_{A'}$ is of significantly lower dimension, but the separability tests become more complicated because of the (a, x) indices being extended.

Other correlation scenarios Another possible extension comes from the key elements in our proposal, which are the embedding of ρ into a separable state, and the generalization of symmetric extensions to any pair of cones [AMP22]. These same ideas are very likely useful in nonlocality, prepare and measure, and other scenarios.

Inner approximations As an outer approximation to the sets of k-preparable assemblages, our hierarchy can *exclude* an assemblage from $\mathcal{A}_k$ even on a finite hierarchy level N. On the other hand, even though the hierarchy is convergent, it provides no practical way to certify an assemblage *is* k-preparable. In the simpler case of entanglement certification, though, the symmetric extensions hierarchy can be modified to converge to the set of separable states from the inside [NOP09a; NOP09b] (see also Sec. 8.3). Propagated to the high-dimensional steering setting, this could lead to an inner convergent hierarchy.

Binarized measurements Even though the witnesses derived in [Des22] were constructed for multi-outcome measurements, their experiment is in fact implementing binary measurements and post-processing the data. Even though this greatly simplifies experiments with spatial light modulators [Bou+18], it also introduces a mismatch with the theory — an issue that has already been posed some time ago for nonlocality experiments [Cai+16]. In particular, we might ask whether one can harmlessly post-process the results of an experiment with binary outcomes, and use them in a steering test designed for measurements with genuinely multiple outcomes. More precisely, is there an assemblage which is *not* k-preparable, but becomes so in the corresponding binarized scenario? The tests constructed in this chapter, supplemented with the inner hierarchy discussed above, can in principle solve this question.

4 Entanglement-assisted communication

This chapter is based on *C. Vieira, C. de Gois, L. Pollyceno, and R. Rabelo. "Interplays between classical and quantum entanglement-assisted communication scenarios". In:* New Journal of Physics *25.11 (2023), p. 113004* [Vie+23]. These contributions are also discussed in [Sou23].

> *. . . unlike Sokrates, Neville does not resort to puns to account for his mixed-up condition. He simply watches it happen and measures off its three angles: desire moves out from Neville himself, ricochets off Bernard, and bends back to Neville—but not the same Neville. "I become not myself but Neville mixed with somebody." The piece of himself that goes out to Bernard makes Bernard immediately familiar "even before I see who it is." As Sokrates would say, it makes Bernard* oikeios.
>
> **— Anne Carson**
> Eros the Bittersweet: An Essay

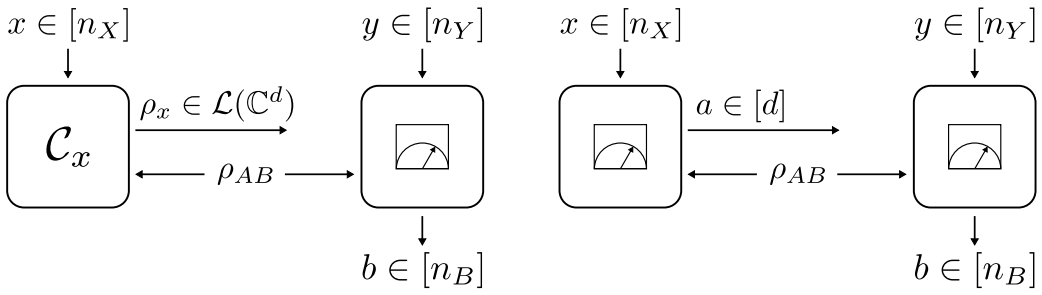
As surprising as it may seem, Anne Carson was not describing entanglement-assisted prepare and measure scenarios, though that is exactly how being half of an entangled pair inside Bob's measurement device must feel like. Preparing and measuring states is one of the building blocks of physics, and describing the expected outcomes is the cornerstone of any physical theory. The prepare and measure (P&M) scenario distils this description to its core principles, allowing classical, quantum and even more exotic theories to be compared on an equal footing (Sec. 1.1.2.3). From a quantum information perspective, they can be used to model and contrast a variety of elementary communication protocols in a unified manner — key examples being dense coding [BW92], random access codes [Wie83; Amb+99], and quantum key distribution [BB14]. This provides a powerful

language to determine which, and *why*, quantum resources may outperform their classical counterparts in practice. Answers to these questions play an important role in the more foundational discussions on informational principles of quantum theory [Paw+09; CMG15].

The most investigated P&M scenario is the one with classical communication and shared randomness (C^λ) [Gal+10; CBB15; Pod+20; Goi+21], and justifiably so: It represents the most powerful, fully classical model of communication, to which any case involving quantum resources should be compared if one wants to discuss what benefits a quantum protocol may provide over its classical counterpart. Moreover, reasoning about C^λ is relatively straightforward, as most of the required techniques can be directly borrowed from nonlocality. Moving beyond this case, a generalization to *entanglement assisted* prepare and measure (EA-P&M) scenarios recently began to be systematically investigated [Mor+21; Tav+21; Pau+22b; FW22]. As in standard P&M scenarios, either quantum or classical communication can be chosen, but this time the devices are correlated through a shared quantum state ρ. Entanglement then enables higher performance in a variety of communication tasks, and allows the reformulation of classic protocols as EA-P&M instances, which brings some level of device independence to them [Mor+21].

Under the hood, any quantum advantage in P&M scenarios mean that, depending on the allowed resources, the sets of possible behaviours can be significantly distinct from C^λ. The central goal in EA-P&M is thus to characterize and relate the quantum (P&M) sets to C^λ and to one another. Understanding how, why, and when these sets differ will shed light on the quantum properties that are practically useful, and provide a clearer distinction on what separates quantum from classical systems.

After introducing the EA-P&M scenarios and some notation in Sec. 4.1, this chapter focuses on two questions concerning the interplay between quantum resources in communication tasks. First, we consider the relationship between classical and quantum communication when entanglement is freely available: What is the cost of simulating entanglement-assisted quantum communication using classical communication, potentially supplemented with additional entanglement? In Sec. 4.2 we demonstrate that any behaviour achievable with d-dimensional quantum communication can be replicated using a d^2-dimensional classical system with entanglement. This goes beyond what is known from dense coding (specifically, that a qudit can match the communication capacity of two dits), as it does not refer to the information capacity, but applies instead to the entire set of correlations. Our second question concerns the case of entanglement-assisted quantum communication, where two parties share an entangled state ρ and are allowed to



(a) Quantum communication **(b)** Classical communication

Fig. 4.1.: In entanglement-assisted prepare and measure scenarios, a shared quantum resource ρ_{AB} can be used to enhance the performance of quantum or classical communication tasks. When communication is quantum, a general behaviour is described by a local, possibly dimension-reducing CPTP transformation ($\mathcal{C}_x$) on one share of the resource, which is then transmitted and measured (see Def. 4.1). When communication is classical, the possible correlations are equivalent to Alice locally measuring ρ_{AB} and communicating her outcome a to Bob, who then uses a and his input y to select a measurement to perform on his share of ρ_{AB} (see Def. 4.2).

communicate using a d-dimensional quantum system. As demonstrated by dense coding, ρ can provide an advantage over shared randomness. Notably, this can occur even when the local entanglement dimension is only as large as the communication dimension d. But what is the precise relationship between these dimensions? Particularly, does having an entangled state ρ with a dimension larger than d always confer an advantage? In Sec. 4.3, we answer this question by showing there are protocols where a $2d$-dimensional entangled resource achieves better performance than any d-dimensional state.

The discussion of EA-P&M scenarios will then continue in the next chapter, where it will be more closely connect to communication complexity problems and to the communication cost of simulating nonlocal behaviours.

4.1 Entanglement-assisted P&M scenarios

We assume the reader is familiar with the standard P&M scenario discussed in Sec. 1.1.2.3, and introduce its entanglement-assisted variants. In place of sharing classical randomness, in an EA-P&M scenario the parties share a quantum state ρ_{AB} that they can exploit to improve performance in communication tasks. Despite being a natural generalization, these have only recently begun to be systematically investigated [Mor+21; Tav+21; Pau+22b; FW22; Pau+22a], and many core questions remain unsolved.

Entanglement-assisted P&M scenarios are split between two cases: quantum or classical communication (Fig. 4.1). In the former, on top of a D-dimensional resource ρ_{AB}, communication is done via a qudit. Protocols in quantum EA-P&M scenarios can be described in the following way [Tav+21; Pau+22b]. Possibly depending on her input x, Alice applies a completely positive and trace preserving (CPTP) map $\mathcal{C}_x : \mathcal{L}(\mathbb{C}^D) \to \mathcal{L}(\mathbb{C}^d)$ to her share of the state ρ_{AB}. This is to say that she may transform her local share of ρ_{AB} in any way she wants, provided that $\rho_A B$ is globally transformed into another quantum state. She then sends her local state to Bob, who reads it out by performing one of the POVMs $\mathcal{M}_y = \{M_{b|y}\}_{b \in n_B}$ on the state $(\mathcal{C}_x \otimes \mathbf{1}_B)(\rho_{AB})$ in his possession. The outcome of this measurement is Bob's output b (Fig. 4.1a). Implicitly, this description depends on the number of inputs for Alice (n_X), and of measurement choices (n_Y) and outcomes (n_B) for Bob. Formally, we thus define the set of behaviours $Q_d^D(n_X, n_Y, n_B)$ as:

Definition 4.1 (Entanglement-assisted quantum behaviors)**.** *A behaviour $\{p(b|x,y)\}_{b,x,y}$ is in $Q_d^D(n_X, n_Y, n_B)$ if there exists*

1. *A state $\rho_{AB} \in \mathcal{L}(\mathbb{C}^D \otimes \mathbb{C}^D)$*
2. *A set of CPTP maps $\{\mathcal{C}_x\}_{x \in [n_X]}$, where $\mathcal{C}_x : \mathcal{L}(\mathbb{C}^D) \to \mathcal{L}(\mathbb{C}^d)$*
3. *A set of POVMs $\{\mathcal{M}_y\}_{y \in [n_Y]}$, where $\mathcal{M}_y = \{M_{b|y}\}_{b \in n_B}$,*

such that, for any $x \in \{0, \dots, n_X - 1\}$, $y \in \{0, \dots, n_Y - 1\}$, and $b \in \{0, \dots, n_B - 1\}$,

$$p(b|x,y) = \mathrm{Tr}\Big[(\mathcal{C}_x \otimes \mathcal{I})(\rho_{AB}) \cdot M_{b|y}\Big]. \tag{4.1}$$

Additionally, by Q_d^E we will mean that ρ_{AB} has any finite dimension, so $Q_d^D \subseteq Q_d^E$, and the cardinalities n_X, n_Y, n_B will often be omitted when they are clear from the context.

As for the classical communication case (Fig. 4.1b), the behaviours can be shown equivalent to a three-step procedure where (i) Alice measures her state with a d-outcome POVM, then (ii) communicates her outcome a to Bob (over a classical channel), who (iii) based on his input y and a, performs a measurement on his share of ρ_{AB}. It might seem odd to restrict Alice to communicate her measurement outcome instead of a general message, but this can actually be done without loss of generality [Tav+21; Pau+22b]. More precisely:

Definition 4.2 (Entanglement-assisted classical behaviors)**.** *A behaviour $\{p(b|x,y)\}_{b,x,y}$ is in $C_d^D(n_X, n_Y, n_B)$ if there exists*

1. *A state* $\rho_{AB} \in \mathcal{L}(\mathbb{C}^D \otimes \mathbb{C}^D)$
2. *A collection of* n_X *POVMs* $\{ M_{a|x} \}_{a=0}^{d-1}$ *on* $\mathcal{L}(\mathbb{C}^D)$
3. *A collection of* $d \cdot n_Y$ *POVMs* $\{ N_{b|a,y} \}_{b=0}^{n_B-1}$ *on* $\mathcal{L}(\mathbb{C}^D)$

such that, for any $x \in \{ 0, \dots, n_X - 1 \}$, $y \in \{ 0, \dots, n_Y - 1 \}$, *and* $b \in \{ 0, \dots, n_B - 1 \}$,

$$p(b|x,y) = \sum_{a=0}^{d-1} \mathrm{Tr}\Big[\rho_{AB}(M_{a|x} \otimes N_{b|a,y})\Big]. \tag{4.2}$$

4.2 Simulating quantum with classical communication

Notation settled, we are ready to work on our first question: What price we must pay to simulate entanglement-assisted quantum communication, if the only available resources are classical communication and pre-shared entanglement? This is directly connected to one of the central objects of study in EA-P&M scenarios, which is to understand the relationships between the sets C_d^D and $Q_{d'}^{D'}$. A subset of this question is answered by dense coding. However, dense coding only applies to a particular EA-P&M task, which is connected to the information capacity but not necessarily to more general correlations. Indeed, even when the channel capacity of the two sets is the same, non-trivial relationships can be found, with the cases involving bits and qubits being by far the most explored [Tav+21; Pau+22b; FW22; Pau+22a]. Past qubits, results are scarce. The next statement fills part of this gap, by showing a non-trivial family (in the sense that both sets have the same channel capacity) of inclusions between dit and qudit entanglement-assisted communication.

Proposition 4.1. *For any choice of* D, d, n_X, n_Y, n_B *we have*

$$C_{d^2}^D(n_X, n_Y, n_B) \subseteq Q_d^{D \cdot d}(n_X, n_Y, n_B), \tag{4.3}$$

and

$$Q_d^D(n_X, n_Y, n_B) \subseteq C_{d^2}^{D \cdot d}(n_X, n_Y, n_B). \tag{4.4}$$

In other words, in *any* EA-P&M scenario, an appropriate amount of additional entanglement can be used to substitute a qubit for two dits of communication and vice versa. The proof of this result, which is detailed in [Vie+23, App. B], relies on the duality between teleportation [Ben+93] and dense coding [BW92; Wer01]. It is done by explicitly

constructing an encoding and decoding of any $C^D_{d^2}$ behaviour into a $Q^{D\cdot d}_d$ behaviour by exploiting dense coding, and of any Q^D_d behaviour into a $C^{D\cdot d}_{d^2}$ behaviour by teleportation and additional encoding and decoding functions.

A corollary of Prop. 4.1 is that, for any choice of d, n_X, n_Y, n_B,

$$C^E_{d^2}(n_X, n_Y, n_B) = Q^E_d(n_X, n_Y, n_B). \tag{4.5}$$

Thus, if entanglement is a free resource between the parties, in the paradigm of P&M scenarios, the communication of one qudit is equivalent to the communication of two dits. From the perspective of the Frenkel-Weiner equivalence of classical and quantum communication for channel ($n_Y = 1$) scenarios (Eq. (1.22)), it is also remarkable that, here, equality holds for any (n_X, n_Y, n_B) setting.Perhaps even more surprisingly, as shown in [Tav+21] and generalized in the next section, dense coding is not the optimal protocol for certain communication tasks in EA-P&M scenarios.

4.3 Advantage of higher-dimensional entanglement

Let us recall the second question posed in the introduction of this chapter: Consider the entanglement-assisted quantum communication scenarios, and suppose the communication dimension is d. Is there, for any d, a protocol where an entanglement resource of dimension larger than d leads to better performance than a d-dimensional entanglement? In other words, is $Q^d_d \subsetneq Q^{d'}_d$ for any d and some $d' > d$?

Similar questions have been posed for many decades. Central to us is the seminal paper by Bennet et al. [BW92], where it is shown that if a maximally entangled pair of qudits is shared between two parties, qudit communication can transmit two dits of information. At first sight, this *dense coding* protocol may seem to contradict Holevo's bound [Hol73], but on a closer inspection it only reveals that the bound cannot be applied when the parties share quantum correlations. The natural next step is to wonder whether more entanglement (albeit with the same communication dimension) would enable more than two dits of communication. Unfortunately, that is not the case. This follows from the fact that, for the game of guessing the preparer's input, the dense coding protocol is optimal and the capacity of a quantum channel can be at most doubled in the presence of entanglement [NS02]. In terms of the sets Q^d_d (which contain the dense coding behaviours) one might therefore be tempted to claim that $Q^d_d = Q^{d'}_d$ for

any $d' \geq d$. Recently, though, in one of the kickstarting studies on EA-P&M scenarios [Tav+21], Q_2^2 was shown to be a strict subset of Q_2^4. Put in words, this means that for qubit communication, a shared state with the same local dimension as the communicated quantum system (Q_2^2) does not provide the entire set of behaviours Q_2^E, and therefore there must exist protocols (for tasks other than simply transmitting the input) where entanglement of higher dimension than the message is useful. This section extends this result from qubits to any communication dimension, answering one of the open problems posed in [Pau+22b].

Consider first the P&M scenario $C_2^{\lambda}(3,1,4)$, namely, the communication of a bit with three possible inputs in the preparation device, and a single, four-outcome measurement in the receiving end. Because shared randomness is present, the set of behaviours for this scenario is a polytope (cf. Sec. 1.1.2.3). Furthermore, the scenario is small enough to be fully characterizable. Listing all of its extremal points and transforming this description into a hyperplane representation [LR15], we obtained the inequalities characterizing the behaviours and found the only non-trivial inequality to be

$$\mathcal{F}_2(\vec{p}) = 2\sum_{i=0}^{2} p(i|i) + \sum_{i=0}^{2} p(3|i) \leq 4, \tag{4.6}$$

the bound holding for any classical behaviour $\vec{p} = \{\, p(b|x) \,\}$. $\mathcal{F}_2$ can be interpreted as the following game. A referee gives an input $x \in \{\, 0,1,2 \,\}$ to Alice. She may send a bit to Bob, and he must output his answer $b \in \{\, 0,1,2,3 \,\}$. Ideally, his goal is to guess Alice's input and output $b = x$. When this happens, they get two points. However, he can also choose a "do not know" answer $b = 3$ which will give them one point, regardless of x. Given a behaviour $\vec{p} = \{\, p(b|x) \,\}$, the value $\mathcal{F}_2(\vec{p})$ is the average number of points earned. [1]

Since only a single bit can be communicated, this game is not trivial, and nonclassical strategies might lead to better performance. Indeed, the introduction of entanglement-assistance enables that to happen. This can be verified numerically by employing the alternated optimization procedure developed in [Tav+21], with which we found a behaviour $\vec{p'} \in C_2^2(3,1,4)$ achieving

$$\mathcal{F}_2(\vec{p'}) = 4.155. \tag{4.7}$$

Code and data are available in [Goi22].

[1] See also [DC21] for similar *ambiguous guessing games*.

Inspired by $\mathcal{F}_2$, let us postulate its natural extension to scenarios having $d+1$ inputs, $d+2$ outcomes but still no choice of measurement,

$$\mathcal{F}_d(\vec{p}) = 2\sum_{i=0}^{d} p(i|i) + \sum_{i=0}^{d} p(d+1|i). \tag{4.8}$$

Proposition 4.2. *The maximum value of $\mathcal{F}_d$ for classical behaviours is $2d$.*

Proof. Let us first show there exists a classical behaviour that reaches $2d$. Any deterministic protocol can be described as an input x that is encoded by a map $\mathcal{E}$ and decoded by $\mathcal{D}$, in time resulting in the output $b = \delta\big(b, (\mathcal{D}\circ\mathcal{E})(x)\big)$. Consider an encoder $\mathcal{E} : \{0,\ldots,d\} \to \{0,\ldots,d-1\}$ such that

$$\mathcal{E}(k) = \begin{cases} k, & \text{if } 0 \leq k \leq d-1 \\ 0, & \text{if } k = d, \end{cases}$$

and the identity decoder $\mathcal{D} : \{0,\ldots,d-1\} \to \{0,\ldots,d+1\}$. Together, they lead to the behaviour $\vec{p}_C$ with elements

$$p_C(b|x) = \delta\big(b, (\mathcal{D}\circ\mathcal{E})(x)\big) = \begin{cases} \delta(b,x), & \text{if } 0 \leq x \leq d-1 \\ \delta(b,0), & \text{if } x = d. \end{cases} \tag{4.9}$$

Then indeed,

$$\mathcal{F}_d(\vec{P}_C) = 2\left(\delta(d,0) + \sum_{i=0}^{d-1}\delta(i,i)\right) + \delta(d+1,0) + \sum_{i=0}^{d-1}\delta(d+1,i) = 2d. \tag{4.10}$$

Proving this is the optimal among all classical behaviours requires a deeper analysis. To start, observe that since $\mathcal{F}_d$ is a linear function and the set $C_d(d+1,1,d+2)$ is a polytope, then the maxima of $\mathcal{F}_d$ must happen on one or more extremal points. Thus, to prove $\vec{p}_C$ is optimal, it suffices to compare it only to other extremal points, which correspond to the deterministic strategies.

Suppose $\vec{p}_D(b|x)$ is an extremal point. Therefore, there are maps $\mathcal{E}$ and $\mathcal{D}$ such that $\vec{p}_D(b|x) = \delta\big(b, (\mathcal{E}\circ\mathcal{D})(x)\big)$. Since the domain of $\mathcal{D}$ has cardinality d, then its image has

cardinality $|\mathrm{Im}(\mathcal{D})| \leq d$. On the other hand, $\mathrm{Im}(\mathcal{D} \circ \mathcal{E}) \subseteq \mathrm{Im}(\mathcal{D})$, from which it follows that $|\mathrm{Im}(\mathcal{D} \circ \mathcal{E})| \leq d$. Our proof now splits in two cases.

Case 1: $d+1 \in \mathrm{Im}(\mathcal{D} \circ \mathcal{E})$

Because $|\mathrm{Im}(\mathcal{D} \circ \mathcal{E})| \leq d$ and $\mathrm{Im}(\mathcal{D} \circ \mathcal{E}) \subset \{0, \ldots, d+1\}$, then there are distinct $k, l \in \{0, \ldots, d\}$ which are not in the image of $\mathcal{D} \circ \mathcal{E}$. Consequently, $p_C(k|k) = p_C(l|l) = 0$, and

$$\mathcal{F}_d(\vec{p}_C) = p_C(d+1|k) + p_C(d+1|l) + \sum_{\substack{i=0 \\ i \neq l,k}}^{d} \Big(2p_C(i|i) + p_C(d+1|i)\Big). \tag{4.11}$$

From normalization, the probabilities appearing in (4.11) are bounded as

$$p_C(d+1|k) \leq 1, \tag{4.12a}$$
$$p_C(d+1|l) \leq 1, \tag{4.12b}$$
$$2p_C(i|i) + p_C(d+1|i) \leq 2\big(p_C(i|i) + p_C(d+1|i)\big) \leq 2. \tag{4.12c}$$

Replacing them in (4.11), we get that $\mathcal{F}_d(\vec{p}_C) \leq 2 + 2(d-1) = 2d$ is the maximum for any deterministic strategy with $d+1$ in its image.

Case 2: $d+1 \notin \mathrm{Im}(\mathcal{D} \circ \mathcal{E})$

In this case, $\mathrm{Im}(\mathcal{D} \circ \mathcal{E}) \subset \{0, \ldots, d\}$, and $p_C(d+1|i) = 0$ for all $i \in \{0, \ldots, d\}$. Thus, $\mathcal{F}_d(\vec{p}_C) = 2\sum_{i=0}^{d} p_C(i|i)$. On the other hand, as $|\mathrm{Im}(\mathcal{D} \circ \mathcal{E})| \leq d$, there must be a $k \in \{0, \ldots, d\}$ such that $k \notin \mathrm{Im}(\mathcal{D} \circ \mathcal{E})$, which implies $p_C(k|k) = 0$. Using that $p_C(i|i) \leq 1$ for all i,

$$\mathcal{F}_d(\vec{p}_C) = 2\sum_{\substack{i=0 \\ i \neq k}}^{d} q(i|i) \leq 2d \tag{4.13}$$

These cases encompass every possibility, and in neither of them $\mathcal{F}_d$ is larger than $2d$, thus concluding the proof. □

Nonetheless, the classical bound for $\mathcal{F}_d$ can be violated with classical communication if a pair of entangled qubits is available. Consider once more the behaviour $\vec{p'} \in C_2^2(3, 1, 4)$

mentioned in Eq. (4.7) and detailed in [Goi22]. This behaviour achieves $\mathcal{F}_2(\vec{p'}) = 4.155$. Now let $\vec{p^*} \in C_d^2(d+1, 1, d+2)$ be a behaviour defined by

$$\vec{p^*}(b|x) = \begin{cases} p'(b|x), & \text{if } 0 \leq x \leq 2 \text{ and } 0 \leq b \leq 3 \\ 0, & \text{if } 0 \leq x \leq 2 \text{ and } b > 3 \\ \delta(b, x), & \text{if } 3 \leq x \leq d. \end{cases} \tag{4.14}$$

It is easy to see that $\vec{p^*}$ indeed is achievable in C_d^2: As Alice can send a dit to Bob, she uses two symbols from this alphabet and the entanglement to simulate the $\vec{p'}(b|x)$ part of the behaviour, while the remaining $d-2$ symbols are used to simulate the $\delta(b, x)$ part of $\vec{p^*}(b|x)$. Playing the $\mathcal{F}_d$ game with this behaviour leads to

$$\mathcal{F}_d\left(\vec{p^*}\right) = \mathcal{F}_2\left(\vec{p'}\right) + 2\sum_{i=3}^{d} p^*(i|i) + \sum_{i=3}^{d} p^*(d+1|i) \tag{4.15}$$

$$= 4.155 + 2\sum_{i=3}^{d} \delta(i, i) + \sum_{i=3}^{d} \delta(d+1, i) \tag{4.16}$$

$$= 2d + 0.155 \tag{4.17}$$

Since, by Prop. 4.2, the maximum of $\mathcal{F}_d$ in C_d is equal to $2d$, it follows that $\vec{p^*} \notin C_d$, thus the following proposition holds.

Proposition 4.3. *For every integer $d \geq 2$, we have $C_d \subsetneq C_d^2$.*

Concatenating Props. 4.3 and 4.1 leads to the main result of this section.

Proposition 4.4. *Consider any entanglement-assisted quantum communication scenario with d-dimensional communication and a shared state of dimension $d+2$. Then, for any $d \geq 2$, there are behaviours in $\bar{Q}_d^{d+2}$ which are not in $\bar{Q}_d^d$.*

Proof. From dense coding we know that $\bar{C}_{d^2} \subseteq \bar{Q}_d^d$, and it is straightforward to see that $\bar{Q}_d^d \subseteq \bar{Q}_{d^2}$. These can be concatenated into $\bar{C}_{d^2} \subseteq \bar{Q}_d^d \subseteq \bar{Q}_{d^2}$.

For scenarios with $n_Y = 1$, a well-known result by Frenkel and Weiner guarantee that $\bar{C}_{d^2} = \bar{Q}_{d^2}$ [FW15] (see also Sec. 1.1.2.3), and with the previous relations we get

$$\bar{Q}_d^d = \bar{C}_{d^2}. \tag{4.18}$$

Finally, combining this last equation with Props. 4.1 and 4.3 results in

$$\bar{Q}_d^d = \bar{C}_{d^2} \subsetneq \bar{C}_{d^2}^2 \subseteq \bar{Q}_d^{2d}. \tag{4.19}$$

□

4.4 Discussion

Following suit to [Pau+22b], this chapter could have been called "Entanglement in prepare-and-measure scenarios: many questions, a few *more* answers". Most nonetheless remain open, but some techniques developed in this chapter may provide a good ground to investigate them. For example, as the proof of Prop. 4.1 is done via an explicit construction, it also provides a blueprint to convert any particular quantum protocol into a classical protocol. The same result can also largely simplify previous questions in EA-P&M scenarios:

Optimizing convex functions in C_d^E and Q_d^E Two noncommutative polynomial optimization models are constructed in [Tav+21, Sec. IV.B]. Notice that these models do not impose a bound on the dimension of the shared resource, as they in fact provide outer approximations with respect to the set of ρ_{AB} states. The difference between C_d^E and Q_d^E is that optimizing over the former requires additional constraints to force the "quantum" communication to be in a diagonal state. However, as a consequence of Prop. 4.1, we know that $C_{d^2}^E = Q_d^E$, so both algorithms solve the same problem and one must simply change the dimension of the communicated system. As the model for Q_d^E involves lower-dimensional variables and fewer constraints when compared to $C_{d^2}^E$, it might be the less computationally demanding solution.

Cap on the entanglement dimension One of the questions posed in [Pau+22b] is whether there is some maximum dimension for ρ_{AB}, after which there are no additional gains. This leads to two questions: Is there a D such that $C_d^D = C_d^{D+k}$, or a D' such that $Q_d^{D'} = Q_d^{D'+k}$, for all k? Proposition 4.1 shows these two questions are equivalent, and gives a precise relationship between D and D'. We deem the C_d^D problem a more promising direction, as it is more directly related to nonlocality scenarios. The fact that an arbitrary amount of entanglement is useful for nonlocality suggests the same might be true for C_d^D [PV10; Slo19; CS20].

Is $Q_d^E \subseteq Q_{d^4}$? Frenkel and Weiner [FW22] showed that, for communication channel scenarios ($n_Y = 1$), it is true that $C_2^D \subseteq C_4$ (with a D^*-dimensional maximally entangled state). They therefore conjectured the more general relationship $C_d^E \subseteq C_{d^2}$, for every d. That being true, a concatenation of Eqs. (4.5) and (1.22) leads to $Q_d^E \subseteq Q_{d^4}$, for every d. This can also be seen as an alternative path to prove their claim.

Likewise, the investigation of "communication channel" scenarios (that is, P&M scenarios with $n_Y = 1$) is still incipient and deserves further investigation. In particular, the fact that there is no choice of measurement greatly limits the space of possible games, and makes the technique used in the proof of the classical bound in Prop. 4.2 be very powerful.

Another interesting direction is to find closer connections between other correlation scenarios, in particular with nonlocality [Ege+25]. This might in turn lead to insights on the long-standing and foundational question about the classical communication cost of reproducing nonlocal correlations, which we will further discuss in the coming chapter.

Part II

Tomography

Confidence regions

5

This chapter is based on *C. de Gois and M. Kleinmann. "User-friendly confidence regions for quantum state tomography". In:* Physical Review A *109 (2023), p. 062417* [GK23].

> *Abstractions, math, models are fine, but when you're down to it and everybody's hollering for a fix, this is what you do: you go and sit exactly on the target with indifferent shallow trenches for shelter, and you watch it in the silent fire-bloom of its last few seconds, and see what you will see. Chances are astronomically against a perfect hit, of course, that is why one is safest at the center of the target area. Rockets are supposed to be like artillery shells, they disperse about the aiming point in a giant ellipse — the Ellipse of Uncertainty.*
>
> — **Thomas Pynchon**
> Gravity's Rainbow

Although quantum state tomography is widely believed to be less explosive than a V-2, both suffer from the same grim odds of landing on target. This is unrelated to experimental imperfections, but rather arises for the same reason why one cannot prove a die is perfectly unbiased with any finite number of trials. Therefore, as any other experimental result, a state $\hat{\rho}$ estimated from real-world, finite data must be presented with a guarantee of accuracy. The central challenge in quantum state tomography (QST) is that obtaining a good accuracy requires measuring a number of samples that grows exponentially with respect to the number of particles in the system (see Sec. 1.1.1.1). This can be alleviated by assuming prior information on the state — for example that it is pure or permutationally symmetric [Blu12; Blu10; Sha+13; Gro+10; Mor+12; Cra+10] — or by giving up on having a complete description of ρ and only determine, for instance, the two-body marginals of the state [BBO20; CW20; Gar+20; Han+24]

or the expectation values of a collection of local or particularly structured observables [Aar18; HKP20]. For all other cases where one needs a full description of the state and prior assumptions are unjustified, quantum state tomography must be performed, and finding better statistical guarantees may reduce its cost.

The most informative way of expressing the statistical uncertainty in $\hat{\rho}$ is to also provide a region (for example in the state space) that contains the true state with high probability. This can be though of as a higher-dimensional analogue of confidence intervals, and in frequentist statistics they are aptly called confidence regions. In QST, they are understood as a region C that contains the true state ρ with a specified confidence level $1-\delta$, see Fig. 5.1. Then,

$$\mathbb{P}\left[\rho \in C(\text{data})\right] \geq 1-\delta, \tag{5.1}$$

where the data can be, for instance, the outcome frequencies from the experiment. Saying C has confidence level $1-\delta$ means to say that if the entire experiment were to be repeated, the true state ρ would sure be in C in a fraction at least $1-\delta$ of the experiments. Thus, δ is usually chosen to be close to one, and for instance a "3σ" confidence level corresponds to $1-\delta \approx 0.9973$. Generally, confidence levels become higher with an increasing number of samples, or conversely, having a high confidence requires plenty of samples.

Confidence regions can be most conveniently described as a norm region: Say $\hat{\rho}$ is the state estimated from the data (for example via linear inversion) and ρ is the unknown fiducial state. Then, a norm-based region is a statement of the form

$$\mathbb{P}\left[\left\|\rho-\hat{\rho}(\vec{f})\right\|_{\star} \geq \epsilon\right] \leq \delta, \tag{5.2}$$

where $\star$ is a chosen norm, ϵ is the size of the region (essentially, the accuracy), and $1-\delta$ is the confidence level. Given a norm, the next step is to find a precise relation between δ and ϵ, ideally depending explicitly on experimental parameters such as the number of samples and the measurement basis. Then, in practice, one could for example fix a sample budget and a required accuracy, and easily obtain the confidence level of the experiment.

Here we aim to be practically minded: More than investigating fundamental, asymptotic bounds, we are interested in finding a simple way to construct confidence regions that works for different experimental setups. Therefore, such a construction must ideally satisfy three conditions: Since sampling and measuring quantum states is expensive, (1) the regions should be reasonably tight. Because experimental setups may use tai-

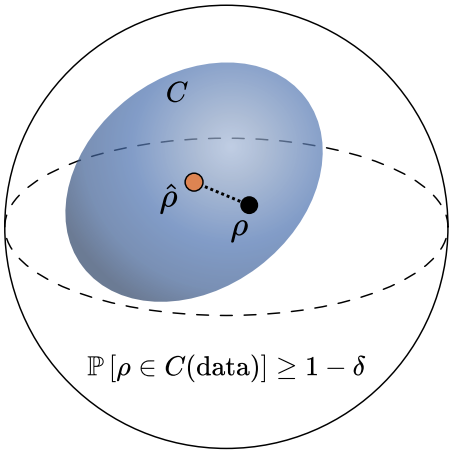


Fig. 5.1.: The objective of quantum state tomography is to accurately reconstruct a quantum state from measurement data. However, in the finite data regime, any estimate $\hat{\rho}$ will inevitably deviate from the true state ρ to some degree. This deviation can be characterized using confidence regions, which provide a probabilistic guarantee on the reliability of the reconstruction. A confidence region (illustrated as the blue ellipsoid) is a subset in some space — such as the state space or the space of Hermitian operators — that contains the true state ρ with high probability. Our main result is the construction of confidence regions in the form of ellipsoids in the space of Hermitian operators. These regions are sample-efficient, broadly applicable, and easy to interpret and report.

lored measurement schemes, it is also important that (2) the confidence regions can be constructed easily, irrespective of which measurement scheme is used. Moreover, (3) it should be possible to report the region associated to the experimental data in a simple and informative way.

A few confidence regions for unconditional, state-agnostic tomography that partially satisfy these desiderata already exist [CR12; STM13; FR16; WSR19; Guţ+20]. Comparing them directly is not always possible, as the constructions can be vastly different. To circumvent this issue, operational tests to benchmark the tightness of a region were recently proposed [AKS23] and two outstanding constructions have been identified — the others being either subsumed by them or performing too poorly. One of them is the polyhedral region of Wang et al. [WSR19], which is simple to adapt for different measurement schemes, satisfying criterion (2), but that does not have an explicit tightness guarantee. As pointed out by the authors, reporting the region is not straightforward, because it requires listing a large number of facets of the polyhedron, and optimizing the most out of it can also become a computationally intensive chore. The second outstanding region is given in terms of the spectral norm $\|\cdot\|_\infty$ [Guţ+20]. It is close to optimal in the asymptotic limit (up to a $\log(d)$ factor) with respect to this norm, so (1) holds, and even

though the region is constructed for a few common classes of measurements, it can be complicated to compute the necessary parameters for arbitrary measurement schemes. Nevertheless, for the measurement schemes that are provided, it is easy to report the region, so criterion (3) is also fulfilled.

Perhaps because of these difficulties, experiments often report a state estimator $\hat{\rho}$ without a confidence region, and then top it up for example with a fidelity to some target state ρ and a confidence interval obtained through bootstrapping. This is a valid approach, but falls short when compared to the more thorough description provided by a confidence region. In fact, a confidence region is the most informative statement that can be made about the experimental data in QST, as from it the confidence intervals for not only the fidelity, but for any function of the state can in principle be computed. Therefore, we ask: Is it possible to construct confidence regions that are simple enough to adapt to any measurement scheme and to describe, but still perform in practice? Our proposal is to try and build a region based on the Frobenius norm $\|\cdot\|_2$, as it is often simpler to manipulate and reason about than other common choices, such as the spectral or the trace norms (Secs. 5.2 and 5.3). Even though one might wonder that they must ultimately compare poorly with the other options (and particularly with the spectral-norm region in [Guţ+20]), in Sec. 5.4 we will show that this is not the case. Finally, in Sec. 5.5 we emphasize its practical usefulness with general guidelines and an example.

5.1 Quantum state tomography

Before starting, let us set up the stage by recalling a few basic concepts of state tomography (see Sec. 1.1.1.1 for further discussion) and fixing some notation. We consider the case where N identically and independently distributed samples of an unknown state ρ are prepared and measured individually. If the measurements provide enough information for a classical description of ρ to be determined, then quantum state tomography is possible, and the measurements are said to be *tomographically complete*. This can either be achieved by using a single, informationally complete measurement, or by collecting results from multiple distinct measurement settings. (The latter is more common in practice, as the implementation of a single tomographically complete measurement is far more challenging than using multiple projective measurements.)

More precisely, consider different measurement settings labelled by s. A setting s is described by a collection of measurement effects $\{ E_{a|s} \}_a$, each associated with a possible

outcome a and such that $E_{a|s} \succcurlyeq 0$, and $\sum_a E_{a|s} = \mathbf{1}$ (for an introductory discussion on quantum measurements refer to Sec. 1.1.1). Measuring s on ρ results in a with probability $p_{a|s} = \mathrm{tr}\Big(E_{a|s}\rho\Big)$. For convenience, we may also describe the outcome probabilities of setting s by the vector $\vec{p}_s = (p_{a|s})_a$, thus one can also think of a measurement as a map $M_s : \mathcal{D}(\mathcal{H}) \mapsto \mathbb{R}^a$. In turn, the $\vec{p}_s$ can be combined into a single vector $\vec{p} = \bigoplus_s \vec{p}_s$, which is obtained from a total measurement map $M : \rho \mapsto \vec{p}$ that summarizes all the settings. When a map M^+ such that $M^+ M\rho = \rho$ for any $\rho \in \mathcal{D}(\mathcal{H})$ exists (that is, M admits a left-inverse), then the measurement is said to be tomographically complete.

Now suppose we have measured n_s of the total N state preparations with setting s and counted $c_{a|s}$ occurrences of outcome a. Naturally, the frequency of this outcome is $f_{a|s} = c_{a|s}/n_s$, and it converges to $p_{a|s}$ with increasing n_s. Doing the same for each setting, we obtain the collection of frequencies $\vec{f}_s$, which we may also combine into a single vector $\vec{f}$, as we did for $\vec{p}$. Then, M^+ yields the point estimate $\hat{\rho}(\vec{f}) = M^+\vec{f}$ for the measurement data. Notice that M^+ can be any left-inverse of M. A common choice is the pseudo-inverse, which corresponds to the free least-squares estimator (for nonlinear estimators, see for instance [AKG19] for a comparison between several common point estimators, [Sch+15] for a cautionary note regarding maximum likelihood estimation, and [Guţ+20] for confidence regions using the projected least squares estimator.)

Central examples of tomographically complete measurement schemes are the Pauli-bases measurement, global Pauli observables, and structured measurements. For Pauli-bases, one considers q qubits, where each is independently measured with one of the settings $s \in \{ X, Y, Z \}$, where X, Y and Z are the Pauli matrices. The outcome on each site is $a \in \{ -1, +1 \}$, and it is handy to write the (projective) measurement effects as $\Pi_{a|s} = (\mathbf{1} + as)/2$. Extending to q parties, a global setting is thus of the form $M_{\vec{s}} = \{ X, Y, Z \}^q$, and cases involving the trivial operation $\mathbf{1}$ on any qubit are disregarded because they can be obtained from the others through marginalization. As the measurement outcomes are registered independently for each qubit, the possible measurement results are the bitstrings $\vec{a} \in \{ -1, +1 \}^q$. Global Pauli observables are similarly constructed, in that one measures tensor products of $\mathbf{1}$, X, Y, and Z. However, they are applied globally, and thus the outcomes are simply $+1$ or -1 for each setting. Lastly, the structured measurements are a large class that encompasses, for example, the symmetric informationally complete measurements [Ren+04], and maximal sets of mutually unbiased bases [Dur+10]. For a precise definition, we refer the interested reader to [Guţ+20; Goi+23] and the references therein.

5.2 Bernstein inequalities

From our previous discussion, it follows that the result of a QST experiment can be interpreted as a sum of random vectors (the outcome frequencies $\vec{f}$) or, alternatively, as a random matrix $\hat{\rho}(\vec{f}) = M^{+}\vec{f}$. Our central goal is to construct an expression similar to Eq. (5.2), so that we may quantify the deviation of the estimate to the true state ρ. Concentration inequalities can provide rigorous guarantees on exactly these types of problem (see Sec. 1.1.1.1). Most well known in the case of scalar variables, these inequalities bound the probability that a random variable deviates from some value, for instance from its mean. They can therefore be though of as non-asymptotic analogues for the law of large numbers. Coming in many forms, there are differences in their assumptions (e.g., from which distribution is the variable sampled), the necessary quantities (the expected value, the variance), their application regime (directly to the random variable, or perhaps to sums or more general functions of multiple variables). Typical examples are Hoeffding's, Bennett's and Bernstein's inequalities, all of which can be applied to sums of (independent) random variables. More recently, similar bounds were found for random matrices or vectors, such as matrix- or vector-Bernstein inequalities [AW02; Gro11; Min17; Tro12; Tro15; Ver18; Wai19]. Their use cases include bounding the deviation of the largest eigenvalue of sums of random Hermitian matrices from its mean, or the deviation of the norm of a vector from its mean. In particular, matrix Bernstein bounds have already been applied to quantum state tomography [Gro+10; Guţ+20], and confidence regions in terms of the spectral norm derived from them.

We wish to instead consider the Frobenius norm as a measure of the distance between the true state and an estimator — this will ultimately result in confidence ellipsoids in the space of Hermitian matrices. Our starting point is the following matrix Bernstein inequality, which will later be cast as a vector inequality:

Proposition 5.1 (Theorem 7.7.1 in [Tro15]). *Let $\{ X_i \}_{i=1}^{N}$ be random Hermitian matrices. Assume that $\mathbb{E}X_i = 0$ and $\lambda_{\max}(X_i) \leq L$, for any i. Consider their sum $\sum_i X_i$ has a bounded variance,*

$$V \succcurlyeq \mathrm{Var}\left(\sum_i X_i\right) = \sum_i \mathbb{E}X_i^2, \tag{5.3}$$

for some Hermitian matrix V.

Define $d = \operatorname{tr}(V)/v$, where $v = \|V\|_\infty$, and $\|\cdot\|_\infty$ is the spectral norm. [1] *Then*

$$\mathbb{P}\left[\lambda_{\max}\left(\sum_i X_i\right) \geq t\right] \leq 4d \exp\left[-\frac{t^2}{2v}\, w\left(\frac{tL}{v}\right)\right], \tag{5.4}$$

with $w(x) = 3/3+x$, holds for all $t > 0$. [2]

Transitions to such kinds of statement into vector bounds can be found at several places in the literature, for example in [Tro15; Min17; Wai19]. The central idea is that for any vector $vecx$, one can consider its Hermitian dilation

$$D_{\text{Herm}} : \vec{x} \mapsto \begin{pmatrix} 0 & \vec{x}^\dagger \\ \vec{x} & 0 \end{pmatrix}. \tag{5.5}$$

Therefore, if $\{ X_i \}_{i=1}^N$ are random vectors, we can apply Prop. 5.1 to bound the concentration of the sum of random Hermitian matrices $D_{\text{Herm}}(X_i)$.

Proposition 5.2. *Let $\{ X_i \}_{i=1}^N$ be random complex vectors. Suppose there are positive real numbers σ and λ such that $\lambda \geq \|X_i\|_2$, and*

$$\sigma^2 \geq \frac{1}{N} \sum_i \mathbb{E}\|X_i\|_2^2, \tag{5.6}$$

for every $i \in \{ 1, \ldots, N \}$. Then, for any $\epsilon > 0$,

$$\mathbb{P}\left[\left\|\frac{1}{N}\sum_i X_i\right\|_2 \geq \epsilon\sigma\right] \leq 8 \exp\left[-\frac{N\epsilon^2}{2}\, w\left(\epsilon\frac{\lambda}{\sigma}\right)\right], \tag{5.7}$$

with $w(x) = 3/3+x$.

Proof. We define the random Hermitian matrices $Y_i = D_{\text{Herm}}(X_i)$. Because $\mathbb{E}X_i = 0$, the Y_i also have zero mean. Moreover, it is also known that for vectors X_i,

$$\lambda_{\max}\left(D_{\text{Herm}}(X_i)\right) \leq \|X_i\|_2, \tag{5.8}$$

[1] The parameter d is also called the intrinsic dimension of a matrix V. It is introduced following the observation that low-rank matrices will concentrate in a small subset of their dimensions, thus concentration bounds depending on the ambient dimension might be wasteful in these cases [Tro15, Chap. 7]. This is particularly important for us in light of Prop. 5.2.

[2] In [Tro15, Theorem 7.7.1] this is proven for $t \geq \sqrt{v} + L/3$. See [GK23, Remark A1] for an argument of why it can be safely extended to any $t > 0$.

where $\|\cdot\|_2$ is the Euclidean norm.

To apply Prop. 5.1 we must still find an upper bound to the variance $V \succcurlyeq \sum_i \mathbb{E} Y_i^2$ and compute its spectra norm. Consider

$$V = \begin{pmatrix} V_1 & 0 \\ 0 & V_2 \end{pmatrix} \succcurlyeq \sum_i \mathbb{E} \begin{pmatrix} X_i^\dagger X_i & 0 \\ 0 & X_i X_i^\dagger \end{pmatrix} = \sum_i \mathbb{E} Y_i^2, \tag{5.9}$$

and rearrange it into

$$\begin{pmatrix} V_1 - \sum_i \mathbb{E}\|X_i\|_2^2 & 0 \\ 0 & V_2 - \sum_i \mathbb{E} X_i X_i^\dagger \end{pmatrix} \succcurlyeq 0. \tag{5.10}$$

Let us then set $V_1 = \sum_i \mathbb{E}\|X_i\|_2^2$ and $V_2 = \sum_i \mathbb{E} X_i X_i^\dagger$. Since V is block-diagonal, $v = \|V\|_\infty = \max\{V_1, \|V_2\|_\infty\}$. But

$$\|V_2\|_\infty = \left\| \sum_i \mathbb{E} X_i X_i^\dagger \right\|_\infty \leq \sum_i \mathbb{E} \left\| X_i X_i^\dagger \right\|_\infty = \sum_i \mathbb{E}\|X_i\|_2^2 \leq V_1, \tag{5.11}$$

so $v = V_1$ and $d = \mathrm{tr}(V)/v = 1 + {}^{\mathrm{tr}(V_2)}/_{V_1} \leq 2$. Substituting into Eq. (5.4), we get

$$\mathbb{P}\left[\lambda_{\max}\left(\sum_i Y_i\right) \geq t\right] \leq 8 \exp\left[-\frac{t^2}{2N\sigma^2}\, w\left(\frac{t\lambda}{N\sigma^2}\right)\right] \tag{5.12}$$

where we used Eq. (5.6) to rewrite $V_1 = N\sigma^2$. Setting $\epsilon = {}^{t}/_{N\sigma}$ and using Eq. (5.8) together with sub-additivity of the norm results in Eq. (5.7). □

5.3 Confidence regions

To achieve our goal, Prop. 5.2 must be tailored into a form that more closely resembles a quantum state tomography experiment. First, we observe that for a fixed number N_s of state preparations, the frequencies vector $\vec{f}_s$ obtained for a given measurement setting s is proportional to a sample from the multinomial distribution $\texttt{Mult}(\vec{p}_s, N_s)$, where $\vec{p}_s = M_s \rho = \left(\mathrm{tr}\left[\rho E_{a|s}\right]\right)_a$. At the end of the day, we want to quantify the Frobenius distance between our data and the fiducial state. This can be written as $\left\|\sum_s M_s^+ \left(\vec{f}_s - \vec{p}_s\right)\right\|_2$. Therefore, we propose to substitute into Eq. (5.7) the random vectors $X_s = G_s(\vec{f}_s - \vec{p}_s)$, where $\vec{f}_s \sim \texttt{Mult}\left(\vec{p}_s, N_s\right)/N_s$ and the G_s are linear maps

determined by the measurement settings used in the experiment. The remaining step is to find a simple expression for σ and λ in terms of the G_s. As shown in the following proposition, this is indeed possible.

Proposition 5.3 (Multinomial estimation)**.** *Let $G_s \in \mathbb{C}^{d\times \ell_s}$ for $s \in \{1,\dots,m\}$, and $N_s \in \mathbb{N}$ be such that $\sum_s N_s = N$. Define $\|G_s\|_{2,\infty}$ as the largest Euclidean norm of the columns of G_s, and set $\sigma^2 = \sum_s {}^{N\|G_s\|_{2,\infty}^2}/_{N_s}$. Assume $\vec{f}_s \sim \texttt{Mult}(\vec{p}_s, N_s)/N_s$ are independent random vectors, where $\vec{p}_s \in \mathbb{R}^{\ell_s}$ is a probability vector. Then*

$$\mathbb{P}\left[\left\|\sum_s G_s(\vec{f}_s - \vec{p}_s)\right\|_2 \geq \epsilon\sigma\right] \leq 8\exp\left[-\frac{3N\epsilon^2}{2(3+\epsilon\eta)}\right], \tag{5.13}$$

where $\eta = 2\max_s(\{\, {}^{N\|G_s\|_{2,\infty}}/_{N_s} \,\})$.

A proof is presented in [GK23, Apps. A.2.a and A.2.b]. In essence, it follows by modelling the multinomial samples as a collection of samples from categorical distributions for which, after a couple of approximations, σ and λ in Prop. 5.2 can be bounded in terms of the G_s.

Equating the right-hand side of Eq. (5.13) to the confidence level $1-\delta$, a confidence region for multinomial samples immediately follows:

$$C_A(\vec{f},\epsilon) = \left\{\vec{p} \;\middle|\; \left\|\sum_s M_s^+(\vec{f}_s - \vec{p}_s)\right\|_2 \leq \epsilon\sigma_A\right\}. \tag{5.14}$$

From Prop. (5.3), $\sigma_A^2 = \sum_s {}^{N\|M_s^+\|_{2,\infty}^2}/_{N_2}$ is essentially a control on the variances, and for the precision direct calculation yields

$$\epsilon = \frac{6}{\eta}\sqrt{u}\left(\sqrt{u} + \sqrt{u+1}\right), \tag{5.15}$$

where $u = {}^{\eta^2\log(8/\delta)}/_{18N}$. These regions are therefore spheres in the space of Hermitian operators. The important parameters are the number of samples obtained per measurement setting N_s, the choice of left-inverses $G_s = M_s^+$ for the measurements M_s used in the experiment, and either the desired confidence level $1-\delta$ or the accuracy ϵ. A common choice of left-inverse is the Moore-Penrose inverse, which corresponds to the least-squares estimator, but other useful options exist [AKG19] and may lead to tighter regions in particular cases [GK23, App. A.3.b]. Fixed the left-inverse, one may proceed to compute σ. In many cases of interest, this is a straightforward calculation that can

	σ_A	σ'_B
Pauli-bases	$d^{1.1610}$	$d^{1.2925}$
Pauli observables	$d^{3/2}$	$d^{3/2}$
Structured measurements	$\sqrt{5/4}\,d$	$\sqrt{3/2}\,d$

Tab. 5.1.: Expressions for the radius $\epsilon\sigma_A$ of confidence region C_A and the largest semi-axis $\epsilon\sigma'_B$ of region C_B depending on the total dimension d of the system. For the Pauli-bases, each qubit is measured with all Pauli matrices X, Y, Z, while for Pauli observables, all combination of Pauli matrices, including $\mathbf{1}$, are measured globally. Structured measurements [Guţ+20] are a general class of qudit measurements which includes, for example, maximal sets of mutually unbiased bases and symmetric informationally complete measurements.

be done analytically [GK23, App. B] (see Tab. 5.1), otherwise symbolic or numerical methods can also be used.

Another possibility is to consider the estimators $G_s = M M_s^+$, which when applied to the data leads to the vector $M \sum_s M_s^+(\vec{p}_s - \vec{f}_s) = M\,(\rho - \hat{\rho})$. We thus define the region

$$C_B(\vec{f}, \epsilon) = \left\{ \vec{p} \;\middle|\; \left\| M \sum_s M_s^+ (\vec{f}_s - \vec{p}_s) \right\|_2 \leq \epsilon \sigma_B \right\}. \tag{5.16}$$

Effectively, this "remeasuring" of the difference yields an ellipsoidal region, with the semi-axes depending on the measurement basis as described in [GK23, Sec. IV]. Notice that the change of estimator leads to different expressions for the σ described in Prop. (5.3), and expressions for typical measurement bases are shown in Tab. 5.1 and computed in [GK23, App. B]. Even though the change of estimator is not expected to change the scaling behaviour, it may give better results in practice, as will be discussed in Sec. 5.4.

Thus far, we constructed confidence regions which are easy to adapt to different experimental configurations, and are easy to report, as only the size $\epsilon\sigma$ and the estimator $\{ M_s^+ \}$ is needed. This almost completes the desiderata posed in the introduction of this chapter. However, it is not a priori clear whether the approximations used for Prop. (5.3) are without too much loss in tightness. In the next section we will show this to be the case, both in asymptotic behaviour and in simulated experiments.

5.4 Performance comparison

The standard way of evaluating the performance of a confidence region is by its asymptotic sample cost for $d \to \infty$, so let us start by investigating the behaviour of the region for some common measurement schemes. Solving Eq. (5.15) for N yields

$$N = \frac{2\sigma(\epsilon\eta + 3\sigma)\log\left({}^{8}/_{\delta}\right)}{3\epsilon^2}. \tag{5.17}$$

In the case of a single measurement settings, $\eta = 2$ and N reduces to $2\log({}^{8}/_{\delta})(2\sigma/\epsilon + 3\sigma^2/\epsilon^2)$. Assuming $\sigma = \Omega(1)$ (i.e., σ's dependency on the quantum state dimension d is larger or equal to a constant), it follows that $\hat{\rho}$ is ϵ-close to ρ in the Frobenius norm and with confidence level $1 - \delta$ for

$$N = O\left[\left(\frac{\sigma}{\epsilon}\right)^2 \log\left(\frac{1}{\delta}\right)\right]. \tag{5.18}$$

Using the expressions for σ shown in Table 5.1, we get that $N = O(d^2/\epsilon^2)$ for a structured measurement and a fixed confidence level, matching the best known sample cost for the Frobenius norm [TO23; KRT17]. In terms of the trace norm, the well known relations $\|X\|_1 = \operatorname{tr}|X|$ and $\|X\|_1 \leq \sqrt{d}\|X\|_{\mathrm{HS}}$, imply that ρ can be ϵ-approximated in trace norm with $N = O(d^3/\epsilon^2)$ samples. This agrees with a lower bound on the sample cost for non-adaptive, single-copy measurements [Haa+16; NL25], therefore our region is asymptotically tight for structured measurements. This extends the results from Ref. [Guţ+20], where optimality was only established for the uniform measurement. For the other measurement schemes considered therein, our results improve over their trace-norm sample costs by a factor or order $\log(d)$.

This argument serves to prove that the approximations used in the derivation of our region are not particularly lossy, as there are measurement schemes for which the region is asymptotically optimal. However, the issue with the asymptotic sample complexity is that it does not predict how well the region performs in practice, since even a method with worse asymptotic complexity can be better in realistic settings, for reasons similar to why galactic algorithms are not useful in practice. To address this concern, Almeida et al. proposed two region-agnostic, non-asymptotic tests to assess the performance of a confidence region via simulated tomography experiments [AKS23]. These tests were then applied to compare the confidence regions available in the literature, and the authors found that the spectral norm region by Guţă et al. [Guţ+20], henceforth C_R,

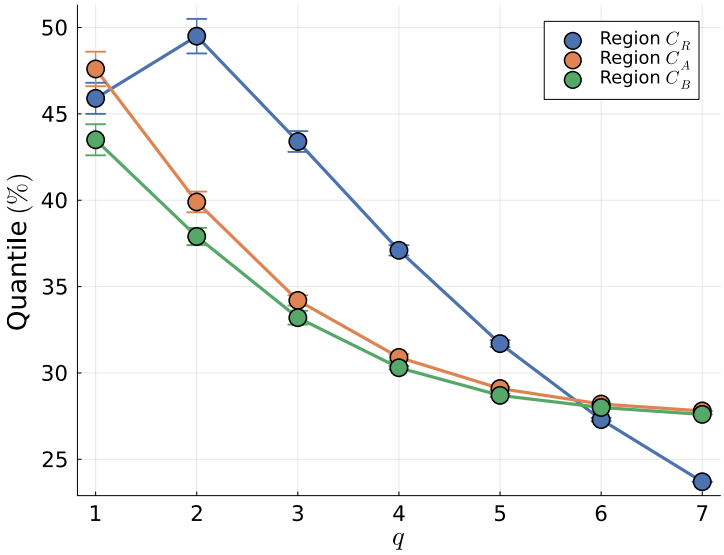


Fig. 5.2.: Tightness of confidence regions C_R, C_A, and C_B in terms of the $1-\delta$ quantiles for the ratio $r = \|\rho - \hat{\rho}\|_\star / (\epsilon\sigma)$. The distance $\|\rho - \hat{\rho}\|_\star$ between the estimated state $\hat{\rho}$ and the fiducial state ρ is compared to the size of the confidence region $\epsilon\sigma$. The $1-\delta$ quantile is a measure of the tightness of the confidence region, and would be 100% for a tight region. Here, the distribution of the ratios r is sampled over $10\,000$ reconstructions of ρ, the confidence level is $1-\delta = 0.99$, the number of simulated state preparations for each reconstruction is $N = 60\,000$, and the Pauli basis measurement scheme is simulated. The procedure was repeated for 50 different initial Haar-random pure states and the largest difference in the ratios is presented as the uncertainty.

has a competitive performance, in particular for the Pauli bases measurement. Thus, we benchmark our region against C_R, which can be written as

$$C_R(\epsilon; \vec{f}) = \{\, \rho \mid \left\| \rho - \hat{\rho}(\vec{f}) \right\|_\infty \leq \epsilon\sigma_R \,\}\,, \tag{5.19}$$

where $\|A\|_\infty$ is the spectral norm, $\epsilon = 6\sqrt{u}(\sqrt{u} + \sqrt{u+1})/\eta$, $u = \eta^2 \log(2d/\delta)/18N$, and d is the dimension of the quantum system (for a discussion region C_R in this particular form, see [GK23, App. C.2]). The constants σ_R and η depend on the chosen measurements and were computed for a few measurement schemes [Guţ+20]. When the measurements are performed in the Pauli bases on q qubits, they are $\sigma_R = 3^{q/2}$ and $\eta = 2\,(4/3)^{q/2}$.

The first test proposed in [AKS23] is the *ratios test*, and it quantifies the tightness of a confidence region. It is best explained by a procedure: First, we postulate a state ρ and simulate a tomography experiment. This can be done by sampling from the multinomial distribution associated with ρ and a chosen set of tomographic measurements. From this sample, an estimate $\hat{\rho}$ is constructed with a suitable estimator. In the form presented

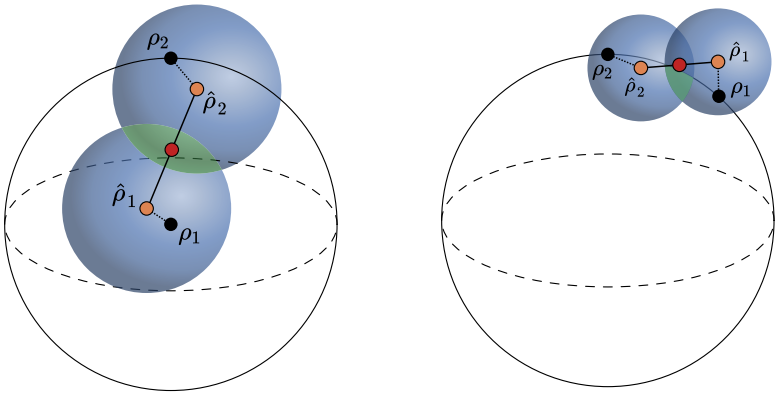


Fig. 5.3.: If the confidence regions intersect in the state space (green regions), typically the average state $\frac{1}{2}(\hat{\rho}_1 + \hat{\rho}_2)$ (red circles) will be in the intersection, as in Fig. 5.3a. This can be used as an ansatz to determine the required number of samples without resorting to optimization. However, in pathological cases, such as in Fig. 5.3b, the average state may fail to be positive. To avoid this possibility, we performed the test using the optimization method.

above, C_R is based on the least squares estimator, and that will be our choice for our regions as well. [3] With the simulated experiment we can compute $r = \|\rho - \hat{\rho}\|_\star/(\epsilon\sigma)$, that is, the ratio of the distance of the estimator to the true state, to the radius of the confidence region. Repeating the procedure with a new sample of the experiment will lead to a slightly different $\hat{\rho}$, which in turn leads to a distribution of ratios. The $1-\delta$ quantile of this distribution is computed and presented as the result of the test. (It goes to one for a tight confidence region.) We implemented this test for C_R, C_A, and C_B using the Pauli bases measurement for up to 7 qubits, and the results are shown in Fig. 5.2. Up to 5 qubits, region C_R performs better, after which it is overtaken by regions C_A and C_B for large systems.

Though the ratios can tell us how tight a confidence region is, it is not particularly informative regarding the number of samples needed to confidently reconstruct a quantum state from measurement data. This can be better addressed by the *intersections test*, which quantified the distinguishing power of the confidence regions [AKS23]. For this test, one chooses two fiducial states ρ_1 and ρ_2 and simulates tomography with N samples for each of the states. The respective confidence regions are then constructed and one verifies whether there exists a common quantum state ρ in both confidence regions. If that is the

[3] Notice that a confidence region in trace norm for the projected least squares estimator is also presented in [Guţ+20].

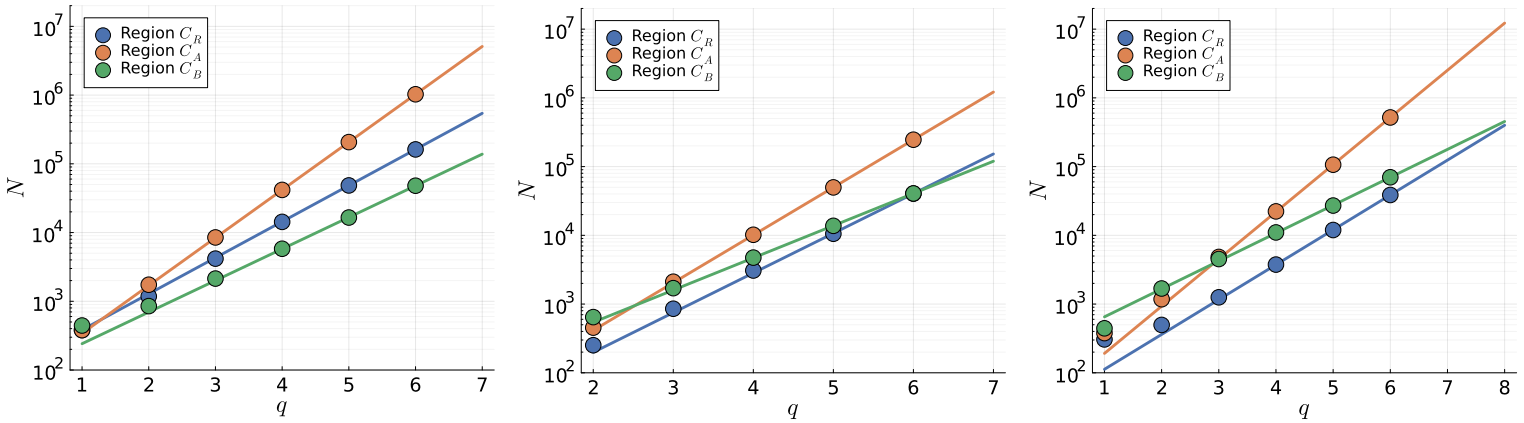
case, then the data is not significant enough to distinguish these two states at the desired confidence level. Clearly, the probability that this happens decreases when increasing the number of samples, as the regions become more concentrated around the corresponding estimate. Expectedly, if there is a non-empty intersection, then $\frac{1}{2}\|\hat{\rho}_1 + \hat{\rho}_2\|_\star$ would be in both confidence regions. However, in pathological cases, $\frac{1}{2}(\hat{\rho}_1 + \hat{\rho}_2)$ may fail to be a quantum state (Fig. 5.3). For this reason, we determine the intersection by means of a convex program: Given the estimates $\hat{\rho}_1$ and $\hat{\rho}_2$, together with a region specified by $\|\cdot\|_\star$ and ϵ,

$$\begin{aligned} \text{find } \rho\colon \quad & \rho \succcurlyeq 0 \\ & \operatorname{tr}(\rho) = 1 \\ & \|\rho - \hat{\rho}_i\|_\star \leq \epsilon, \text{ for } i = 1, 2. \end{aligned} \tag{5.20}$$

This procedure was performed for the three different pairs of states described in Fig. 5.4. For comparison purposes, we have estimated the number of samples N for which the probability of the confidence regions intersecting is approximately $1/2$. While no region outperforms the others in all cases, the results indicate that region C_A is typically weaker than C_R, while C_B is comparable to region C_R and tends to outperform it for larger systems. Lastly, we note that the ellipsoidal shape of region C_B is reminiscent of a Gaussian approximation, and it is possible to show that it outperforms a confidence region based on a Gaussian approximation [AKS23] for a sufficiently large number of qubits [GK23, App. C.1].

5.5 Usage example

In this section we highlight ways of constructing, optimizing, and using the confidence regions. Let us start with the construction. As discussed in Sec. 5.3, reporting these confidence regions in a tomography experiment amounts simply to stating the estimated state $\hat{\rho}$ and the radius of region (or, in the case of region C_B, the semi-axes of the ellipsoid). Consider first the case of a single tomographically complete measurement described by the effects $(E_1, E_2, \ldots)$. Assume N samples were measured and the outcome frequencies $\vec{f} = (f_1, f_2, \ldots)$ obtained. For a chosen confidence level $1 - \delta$, the value of ϵ can be straightforwardly computed from Eq. (5.15). To determine σ, one must write a representation of the measurement map M. That can be done with the help of any vectorization $\vec{v}(A) \in \mathbb{R}^{d^2}$ of the d-dimensional Hermitian operators with the property that



(a) $\mathcal{D}_{\frac{1}{2}}(|0\rangle^{\otimes q})$ and $\mathcal{D}_{\frac{1}{2}}(|1\rangle^{\otimes q})$ **(b)** $|\mathrm{GHZ}\rangle$ and $|\mathrm{W}\rangle$ **(c)** $|\mathrm{GHZ}\rangle$ and $\mathbf{1}/2^q$

Fig. 5.4.: Distinguishing power of the confidence regions C_R, C_A and C_B. The plots show the number of samples N for which the confidence regions of a pair of states, at a confidence level $1-\delta = 0.9$, no longer overlap. We consider different pairs for $q = 1, \ldots, 6$ qubits, namely, (a) $\mathcal{D}_{1/2}(|0\ldots0\rangle)$ and $\mathcal{D}_{1/2}(|1\ldots1\rangle)$ with $\mathcal{D}_t(|\psi\rangle) = t\,|\psi\rangle\langle\psi| + (1-t)\mathbf{1}/2^q$, (b) Greenberger–Horne–Zeilinger (GHZ) state and W state with $|\mathrm{GHZ}\rangle = (|0\ldots0\rangle + |1\ldots1\rangle)/\sqrt{2}$ and $|\mathrm{W}\rangle = (|00\ldots01\rangle + \cdots + |10\ldots00\rangle)/\sqrt{q}$, (c) GHZ state and maximally mixed state. No notable difference was observed when local random rotations were applied simultaneously to both of the fiducial states. We simulate tomography experiments using the Pauli-bases measurement and determine N over 1024 estimation procedures such that in $(50 \pm 2)\%$ of the cases the confidence regions do not overlap. To each dataset we have fitted an exponential model via least-squares minimization.

$\vec{v}(A) \cdot \vec{v}(B) = \mathrm{tr}(AB)$. A common choice is to let $\vec{v}(A)$ simply concatenate the rows of A. Then:

1. Construct a matrix M with rows i given by $\vec{v}(E_i)$.
2. Compute the pseudo-inverse M^+ of M (this can be done numerically, symbolically or analytically).
3. Determine an estimate $\hat{\rho} = \vec{v}^{-1}(M^+ \vec{f})$ for the state.
4. Let σ be the maximal Euclidean norm of the columns of M^+.

This concludes the construction for the spherical region C_A. In the case of the ellipsoidal region C_B, the value of σ must be computed from the norm of MM^+ instead, and the semi-axes of the ellipsoid must be provided (see [GK23, Sec. IV] for an explicit expression.)

For the case of multiple measurement settings, even though our construction allows us to treat each setting independently, one can also represent them as a single generalized measurement and apply the same procedure above. To do so, we assume the measurements settings labelled by $s' = 1, \ldots, m$, are performed at random with probability $q_{s'}$ each. Then, if the operator $E'_{a'|s'}$ describes the measurement outcome a' of measurement setting s', the unified measurement has outcomes $a = (a', s')$ described by the operators $E_a = q_{s'} E'_{a'|s'}$. Correspondingly, the observed frequencies $f'_{a'|s'}$ are replaced by $f_a = q_{s'} f_{a'|s'}$. An apparent issue with this construction is that the number of samples per setting $N_s = q_s N$ ceases to be constant, but is instead modelled as a random variable itself — more precisely, $\vec{q} \sim \mathtt{Mult}(\vec{q}, N)$. However, this can only increase the size of the region [GK23, Rmk. B.1], therefore an experiment in which the number of samples per setting is fixed can still be analysed with the unified measurement construction.

In situations where sampling is particularly expensive, the regions can be optimized by various other means. To start it might be beneficial to go the extra mile and treat the settings separately instead of considering a unified measurement. Perhaps more interestingly, since σ directly depends on the choice of estimator (i.e., the left-inverse of the measurement map M), this can also be optimized [GK23, App. A.3.b]. Complementarily, the distribution of the samples into the different settings can also play a role [GK23, App. A.2.c].

A state estimate together with a confidence region is the most informative statement about a tomography experiment. Therefore, at least in principle, they can be used to compute

error bars for any function of the state by determining the minimum and the maximum of $f(\rho)$, over every $\rho \in C$ with $\rho \succcurlyeq 0$. Since the confidence regions for QST are generally convex, any linear function $f(\rho)$ can be maximized and minimized using standard convex programming solvers, immediately returning the error bars. Alternatively, convex (concave) functions f can be minimized (maximized), or for more general functions one must resort to other techniques. Common functions include the expectation value of observables, the fidelity to a pure state, the best-case fidelity to a mixed state, the negativity of entanglement, etc.

Similarly, the same idea can be used to perform a hypothesis test for ρ to be outside a convex set $\mathcal{S}$, by verifying that the confidence region does not contain any state in $\mathcal{S}$. As a concrete example, consider the problem of certifying that a state is genuinely multipartite entangled (GME) [SU08; GT09]. This cannot be done simple to verifying the estimate $\hat{\rho}(\vec{f})$ is GME. Rather, we must show that its entire confidence region only contains GME states. [4]

Let us recall, from Sec. 1.1.2.1, that genuine multipartite entanglement is a property of a state that cannot be written as a convex combination of biseparable states. More precisely, let ρ_{ABC} be a three-partite state with the subsystems labelled by A, B and C. Then ρ is said to be genuine multipartite entangled if and only if [GT09; Hor+09; SU08]

$$\rho_{ABC} \neq p_1\rho_{A|BC} + p_2\rho_{AB|C} + p_3\rho_{AC|B}, \tag{5.21}$$

where (p_1, p_2, p_3) forms a probability distribution and $\rho_{X|Y}$ stands for a state which is separable across the bipartition $X|Y$, that is, $\rho = \sum_i p_i \rho_X^i \otimes \rho_Y^i$. This definition extends naturally for q-partite systems, in which case all $\binom{q}{2}$ bipartitions must be considered. The use of biseparability in this definition can be relaxed to that of a bipartite state with positive partial-transpose (PPT), leading to the PPT mixtures [JMG11]. Although the set of PPT mixtures is a superset of the biseparable states, it is frequently a good approximation to the former [JMG11]. Most importantly, it has a finite semidefinite representation given by

$$\left\{ \rho = \sum_i \omega_i \;\middle|\; \omega_i \succcurlyeq 0 \;\forall i, \quad \omega_i^{\mathsf{T}_i} \succcurlyeq 0 \;\forall i, \quad \sum_i \mathrm{tr}(\omega_i) = 1 \right\} \tag{5.22}$$

[4] Another common approach to entanglement certification in practice is to use entanglement witnesses [Ter00], thus avoiding the need for tomography altogether, and drastically reducing the sample cost. These tests heavily rely on prior information about the state, such that a suitable witness can be constructed. However, if the state is randomly sampled, effectively detecting its entanglement requires $e^{\Omega(d)}$ witnesses to be tested [Žni+07; Liu+22], and that is impractical enough to justify tomography.

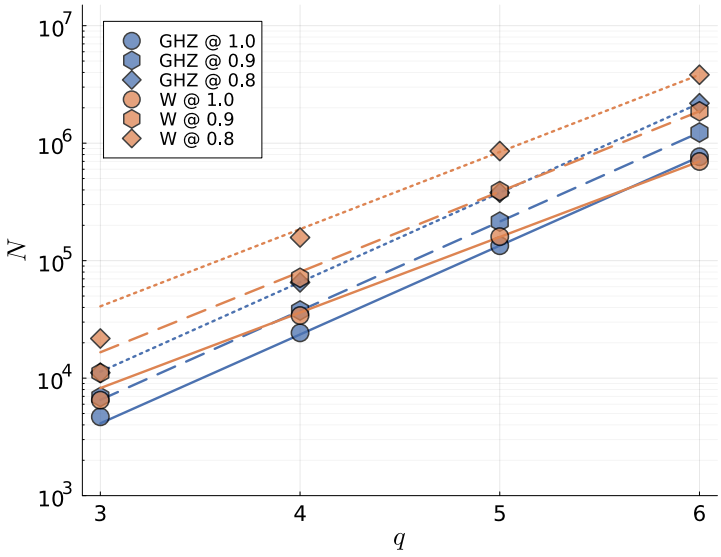


Fig. 5.5.: Number of samples N required to certify genuine multipartite entanglement by means of the confidence region C_B for confidence level $1-\delta = 0.9$. We simulate tomography experiments using local symmetric informationally-complete measurements and determine N over 1024 estimation procedures, such that in $(50 \pm 2)\%$ of the cases the confidence region contains no state which is a PPT mixture. Results are shown for $\mathcal{D}_t(|\mathrm{GHZ}\rangle)$ and $\mathcal{D}_t(|\mathrm{W}\rangle)$ with $t = 1.0, 0.9, 0.8$ and $\mathcal{D}_t(|\psi\rangle) = t\,|\psi\rangle\langle\psi| + (1-t)\mathbf{1}/2^q$. The lines correspond to least-squares fits assuming an exponential model. For reference, the largest values for t such that a GHZ (W) state is a PPT mixture are 0.429, 0.467, 0.484, 0.493 (0.479, 0.474, 0.422, 0.353) for $q = 3, 4, 5, 6$ qubits, respectively.

Here, i runs through all possible bipartitions of q qubits and by $\rho_i^{\mathsf{T}_i}$ we mean a partial transposition with respect to the cut imposed by the ith bipartition. The constraints enforce that $\rho = \sum_i \omega_i$ is a PPT mixture, as one can interpret each ω_i as a subnormalized PPT state [JMG11; Tav+23].

Our goal is to certify that a confidence region contains only genuinely entangled states, and we use region C_B as an example. Since it can be represented by the conic constraint, it can be directly added to Eq. (5.22). Then, a feasible solution would certify there exists a state ρ which is not genuine multipartite entangled inside the confidence region. In this case, we cannot guarantee with the chosen confidence that the true state ρ is genuine multipartite entangled.

As a demonstration, we consider $q = 3, \ldots, 6$ qubits, each measured with local symmetric informationally complete measurements [Ren+04]. This scheme is experimentally feasible [Str+22] and has the same performance as the Pauli bases measurement with respect to C_A and C_B [GK23, App. B.2]. Notice that the constants necessary

to evaluate the reference confidence region C_R are not available for this measurement scheme. With these measurements, we simulated a tomography experiment using Greenberger–Horne–Zeilinger states $|\mathrm{GHZ}\rangle = \frac{1}{\sqrt{2}}(|0\rangle^{\otimes q} + |1\rangle^{\otimes q})$ and W states $|\mathrm{W}\rangle = \frac{1}{\sqrt{d}}(|00\cdots001\rangle + |00\cdots010\rangle + \cdots + |10\cdots000\rangle)$ as the fiducial states. The number of samples N was increased until the associated confidence region no longer intersected the set of PPT mixtures, and the results are shown Fig. 5.5. As expected, we observe an exponential increase of the sample cost N with the number of qubits, and generally more samples are required for noisier states. Overall, the number of samples is moderate for small numbers of parties, with millions of state preparations required for six qubits.

5.6 Discussion

We have constructed confidence regions for quantum state tomography which can be used for any measurement set, are easy to describe and to report, are asymptotically tight and can perform better than the best regions to date in simulated experiments for high-dimensional systems. At their core is the concentration bound presented in Prop. 5.3, which can be applied to any linear function of multinomial data. This opens up a range of possible extensions of our results.

Extension to other "tomographies" When full state tomography is unnecessary or infeasible, one must fall back to partial (Chap. 6) or shadow tomography (Sec. 6.1.1). Even further away, it is also an interesting and important problem to characterize other structures such as quantum processes and higher-order operations [Ant+24]. For any of these cases, it is highly likely that the ideas used in this chapter can be adapted and extended.

Improvements on the constants Even though the constants appearing in the respective Bernstein bounds are not important for asymptotic statements, they are crucial in practice: If too large, the regions would have the same behaviour of a galactic algorithm, rendering them useless. Any improvements on the constants of the matrix Bernstein inequality used in Prop. 5.2, for example, would directly lead to tighter regions. Given the nature of the problem, only marginal gains should be expected, but even a linear improvement on the sample cost can turn an infeasible experiment feasible.

Poissonian distribution of samples In optical setups, the number of samples is in itself a random variable. This happens because the photon generation processes are stochastic, and the variable under the control of an experimenter is the time instead of the number of shots. Would there be significant gains in reconsidering Prop. 5.3 in terms of the Poissonian instead of the multinomial distribution?

Nonlinear estimators A similar transition to the projected least squares estimator as used in [Guţ+20] could be applied to our regions, but do our analysis hold also for larger classes of nonlinear estimators?

Partial state tomography

6

This chapter is based on *K. Hansenne, R. Qu, L. T. Weinbrenner, C. de Gois, et al. "Optimal overlapping tomography". In:* arXiv:2408.05730 *(2024)* [Han+24]. Section 6.1.1 is not discussed in the associated publication, as well as the binary program for clique covering without predefined cliques [Eq. (6.6)], and the heuristic algorithm presented in Sec. 6.2.

Plenty of interesting quantum states do not require full-body information to be determined. This is true for almost all pure states [LPW02; JL05], and states prepared by shallow quantum circuits [YW23]. Moreover, physical Hamiltonians are often dominated by short-range interactions. As a result, some ground state properties may, in certain cases, be largely determined by local, few-body quantities [Col63; Maz12]. Prompted by these observations, recent efforts were turned to the problem of obtaining only a subset of few-body marginals of a larger quantum state. The crucial question in *partial state tomography* is then whether — and to what extent — it is more scalable than full state tomography. Naturally, the figure of merit for scalability may vary depending on what is the most costly resource in an experiment. Here we will focus on analysing the number of distinct measurement settings required to solve a partial tomography task, which can also be used as proxy to optimize the sample cost (cf. for example [CW20, App. A] or [VK24]).

For a concrete example, suppose we are working with a platform where only local Pauli measurements on single copies of an n-qubit state ρ are available. This means we may only measure one of the settings from $\{X, Y, Z\}^{\otimes n}$ on each sample of ρ. Additionally, let us set ourselves the goal of finding a classical description for every k-body marginal of ρ. Since each k-body marginal is defined by the expectation values of $\{X, Y, Z\}^{\otimes k}$, and there are $\binom{n}{k}$ marginal states to be reconstructed, a naive count suggests that $3^k\binom{n}{k}$ measurement settings are necessary. However, this measurement scheme could be highly redundant, as it might be that the 3^k settings required for each marginal might appear repeatedly over the global settings. This is indeed the case, and as recently observed

in [CW20; BBO20; Gar+20], a scaling of $e^{O(k)} \log n$ is possible. For the case of two-body marginals, explicit constructions are provided and have a good scaling, but are suboptimal. More recently, the case of nearest-neighbour marginals on some common lattices was also analysed [Ara+22].

Our main concern in this chapter is to find the optimal solution to quantum overlapping tomography in general settings. In doing so, we will also extend the solution to cases where only a subset of (as opposed to all the) marginals are required, or measurements other than local Pauli measurements are available. After formalizing the problem (Sec. 6.1) and discussing its close ties to shadow tomography (Sec. 6.1.1), we derive the main result of this chapter: An equivalence between the minimal measurement schemes and the graph clique covering problem (Fig. 6.1 and Sec. 6.2). This equivalence can not only be used to improve over the previous constructions and generalize them to any k, but also to solve cases where only specific marginals of the state are wanted, as opposed to every k-body marginal. Many methods to compute these solutions are discussed in Sec. 6.2.1 and applied to a series of cases, which we compare to the previous literature in Sec. 6.3.

6.1 Reduced density matrices

During this chapter we will use the case of local Paulis as a running example. Generalizations to other measurement bases should be mostly obvious, and comments will be provided when appropriate.

Let σ_s, for $s \in \{ 0, X, Y, Z \}$ be the Pauli matrices. An n-qubit state can then be decomposed as

$$\rho = \frac{1}{2^n} \sum_{\vec{s} \in \{ 0,X,Y,Z \}^{\times n}} \underbrace{\operatorname{tr} \left[\rho \left(\sigma_{s_1} \otimes \ldots \otimes \sigma_{s_n} \right) \right]}_{c_{\vec{s}}} \sigma_{s_1} \otimes \ldots \otimes \sigma_{s_n}. \tag{6.1}$$

The k-body marginal [1] of qubits $K \subset \{ 1, \ldots, n \}$, where $|K| = k$, is then simply $\rho^K = \operatorname{tr}_{\{ 1,\ldots,n \} \setminus K} (\rho)$. Its matrix elements are the expectation values

$$\operatorname{tr} \left(\rho^K \bigotimes_{i=1}^{k} \sigma_{s_i} \right), \quad \forall \vec{s} \in \{ 0, X, Y, Z \}^{\otimes k}. \tag{6.2}$$

[1] Synonymously, the k-body reduced density matrix, or k-RDM.

In practice, they are obtained by multiplying the local outcomes obtained in the global measurements. Settings involving σ_0 thus do not need to be physically measured, as we can compute their expectation values from the remaining data, so in effect we must only obtain 3^k expectation values for each k-body marginal.

Naively, for a given marginal K, we could simply take these 3^k settings and pad the remaining $n-k$ qubits with identities. Assume however that we want to obtain ℓ distinct k-body marginal of ρ. With identity padding, we would need $\ell 3^k$ measurement settings to gather complete data. However, this feels rather wasteful, and we could try to somehow parallelize the data acquisition. For example, in the "embarrassingly parallel" cases — where only the marginals $\{ K_i \}_{i=1}^{\ell}$ with $K_i \cap K_j = \emptyset$ are wanted — it is easy to construct a set of 3^k settings, which could provide a reduction up to the order of k/n settings.

More interesting and relevant cases arise when the marginals overlap, as for example in the case of reconstructing every k-body marginal, or every nearest neighbour states on a lattice. Appropriately, this became known as the *overlapping tomography* problem. Schemes that can solve the case of every k-body marginal with a number of settings scaling as $e^{O(k)} \log n$ are known to exist [CW20; BBO20], and for $k = 2$ an explicit construction with $6\lceil \log_2 n \rceil + 3$ settings is provided. A better construction requires only $6\lceil \log_3 n \rceil + 3$ settings [Gar+20], but as we will discuss, is also not optimal (Sec. 6.3).

Before continuing our discussion, we shall briefly detour into a comparison of partial tomography to one of its trendy siblings, shadow tomography. Uninterested readers may move on to Sec. 6.2 without loss of continuity.

6.1.1 Relation to shadow tomography

Shadow tomography was first proposed [Aar18] with the following motivation: Since full state tomography is astronomically infeasible for large systems, requiring an exponentially increasing number of samples with respect to the number of qubits n, then what could we learn instead with only $n^{O(1)}$ copies of ρ? This question was formalized as the shadow tomography problem:

Definition 6.1 (Shadow tomography)**.** *Given a set $\mathcal{O} = \{ O_i \}_{i=1}^{L}$ of observables and copies of an unknown d-dimensional quantum state ρ, estimate $\hat{o}_i \in \mathbb{R}$ such that*

$$\mathbb{P}\left[|\hat{o}_i - \operatorname{tr}(O_i \rho)| \leq \epsilon \right] \geq 1 - \delta, \quad \forall i \in \{ 1, \ldots, L \} . \tag{6.3}$$

Because of a seemingly common confusion, we emphasize that shadow tomography is best understood not as a solution or a technique, but simply as a problem, much in the same way as full state tomography also is.

A straightforward and experimentally feasible procedure to solve shadow tomography is to simply measure the observables independently, thus requiring $O(L \log L/\epsilon^2)$ samples and no quantum memory. Aaronson's main contribution in [Aar18] was to construct a protocol that can do the same, for *any given* set of observables, with sample complexity $\tilde{O}(\epsilon^{-5} \log^4 L \log d \log 1/\delta)$. [2] Thus, sub-linearly on the number of observables L, and yet with only an $O(n)$ dependence. This protocol is however far from feasible, requiring entangled measurements on multiple copies of ρ, not to mention an exponentially long classical processing. Unfortunately, there is a sense in which entangled measurements are unavoidable, as for any $L > 1$, there exists a set of L observables on n qubits requiring $\tilde{\Theta}\left(\min(\{ L, 2^n \}/\epsilon^2)\right)$ state samples for any protocol that is limited to single-copy, possibly adaptive measurements [Che+22].

An alternative direction is to instead consider a variant of the shadow tomography problem, henceforth called *restricted shadow tomography*, in which the quantities to be estimated must be taken from a restricted family of observables (see e.g. [HKP20; Kin+25; CGY24]). Though it may seem limiting at first, this specialization still bears significant practical relevance depending on which family is chosen. A prominent example is the case of low-weight Pauli observables. Denoted $\mathcal{P}_k^{(n)}$, these are the sets of n-qubit Pauli strings acting non-trivially on k qubits,

$$\mathcal{P}_k^{(n)} = \{ P \in \{ \mathbf{1}, X, Y, Z \}^{\otimes n} \ : \ |P| = k \}, \tag{6.4}$$

where $|P|$ counts the number of characters in P that are not equal to $\mathbf{1}$. So, for example, $XX\mathbf{1} \in \mathcal{P}_2^{(3)}$, but not XXX and $X\mathbf{11}$. [3] Knowing $\mathcal{P}_k^{(n)}$ up to a small error gives us enough information to determine every k-body marginal of the underlying state. Therefore, the Pauli shadow- and Pauli partial tomography problems are essentially equivalent — and of course, any solution to partial tomography also solves a related shadow tomography problem. Their motivations are however different: While shadow analyses mostly focus on asymptotic sample costs and the question of what could in principle be efficiently measured, partial tomography is more firmly grounded on practical concerns. Consequently, to *perform* partial state tomography, more than just a sample

[2] Subsequent improvements on the exponents were found for instance in [AR19; BO21].

[3] The shorthand XXX is to be interpreted as $X \otimes X \otimes X$ and accordingly for any other string.

cost scaling, one must find feasible measurement schemes, meaningful quantum state estimators and efficient, non-asymptotic confidence regions.

At the risk of taking another tangent, it is worthwhile to briefly comment on a second related topic. *Classical shadows* [HKP20] are probably the most discussed solution to restricted shadow tomography problems, in particular for the estimation of low-weight Pauli observables. They are, however, best understood as solutions to a slightly different problem (cf. for example [CGY24; GPS24]):

Definition 6.2 (Oblivious shadow tomography)**.** *Given N copies of a state ρ, measure them and store a classical description $\mathcal{C}$ of the results. Afterwards, given a set $\mathcal{O} = \{ O_i \}_{i=1}^{L}$ of observables, estimate every $\mathrm{tr}(O_i\rho)$, with accuracy ϵ, and solely from $\mathcal{C}$.*

Crucial to this definition is that $\mathcal{O}$ is not known beforehand, therefore the measurements (and consequently, the *classical shadow* $\mathcal{C}$) cannot depend on the observables. Notice that oblivious shadow tomography also solves related shadow tomography tasks, but not necessarily vice-versa. Expectedly, a solution with manageable sample cost for any $\mathcal{O}$ should not possible, as it would also solve full state tomography. Instead, one might hope to find protocols that give good estimates for restricted families of observables. From a theoretical standpoint, it is interesting to ask whether an efficient solution is possible, or if the lack of knowledge about $\mathcal{O}$ necessarily results in different complexities compared to shadow tomography. On the other hand, potential use cases for oblivious shadow tomography protocols are arguably exotic, as we usually know what we want to measure, and can carefully devise an efficient protocol for that.

To get a better feeling for how it relates to shadow tomography and why we do not simply call it a day on partial tomography, let us consider the protocol presented in [HKP20]. It starts with the choice of a set of measurements $\mathcal{M}$ that is tomographically complete. For each copy of ρ, a random $M \in \mathcal{M}$ is then measured, and the outcome is registered as a bit string $|a\rangle$. Since $\mathcal{M}$ is tomographically complete, we can consider the expected value of the measurement outcomes as a measurement channel, and compute its inverse $\mathcal{M}^{-1}$. This can in turn be applied to each measurement outcome, the result $\mathcal{C} = \{ \mathcal{M}^{-1} \circ |a\rangle\langle a| \}_{i=1}^{N}$ being the classical shadow of ρ. Ultimately, which observables can be efficiently estimated from $\mathcal{C}$ must depend on what measurement ensemble $\mathcal{M}$ was used. Huang et al. [HKP20] show that when $\mathcal{M}$ is the set of single-qubit Pauli measurements, then low-weight Paulis can be estimated with a good scaling, and when

$\mathcal{M}$ are global Clifford measurements, then observables with a small Hilbert-Schmidt norm are a good match.

One technical but important difference between this and the earlier shadow tomography results is that [HKP20] provide explicit, non-asymptotic sample costs, bringing their proposal closer to practical implementation (though the median of means estimators they have used is not expected to be the best choice in general). Thus, together with [CW20; Gar+20; BBO20], it can be considered as one of the first protocols to perform scalable partial tomography of every k-body marginal using independent Pauli measurements. However, in its original proposal, classical shadows' need for randomization and its indirectness may be unnecessary complications when compared to the others. Notice though that randomization can be circumvented with a small classical processing price [HKP21], and the description of classical shadows can be greatly simplified by reformulating them in terms of generalized measurements [ASS21; Ngu+22].

Many other interesting aspects of shadow tomography and classical shadows are unfortunately out of our scope, and we direct interested readers to [Kin+25] for a discussion of the classical processing efficiency, [Inn+23] for an analysis of statistical estimators other than the median of means, and [Str+22] for experimental comparisons of full state tomography with diverse aspects of shadow tomography.

6.2 Minimizing the number of measurement settings

Going back to our running example, which we later generalize, assume we can only perform local Pauli measurements $\{ X, Y, Z \}$ on single copies of an n-qubit state ρ, and our goal is to measure a classical description of every two-body marginal of ρ. Each marginal must thus be obtained by measuring $\{ X, Y, Z \}^{\otimes 2}$ on the respective parties, and the possible measurement settings are elements from $\{ X, Y, Z \}^{\otimes n}$. Locally, X, Y and Z are mutually exclusive, but the local settings on two distinct qubits are always compatible. These relations can be represented in a graph (Sec. 1.2) with $3n$ vertices — which we label with X_i, Y_i, Z_i for $i \in \{ 1, \ldots, n \}$ — and edges between any two vertices that have distinct subscripts (see Fig. 6.1). We will call this the "possible measurements graph", for obvious reasons, and call it K. In graph-theoretic lingo, they are particular instances of complete multipartite graphs, namely, the ones with n parties and three vertices per party, $K_{n,3}$.

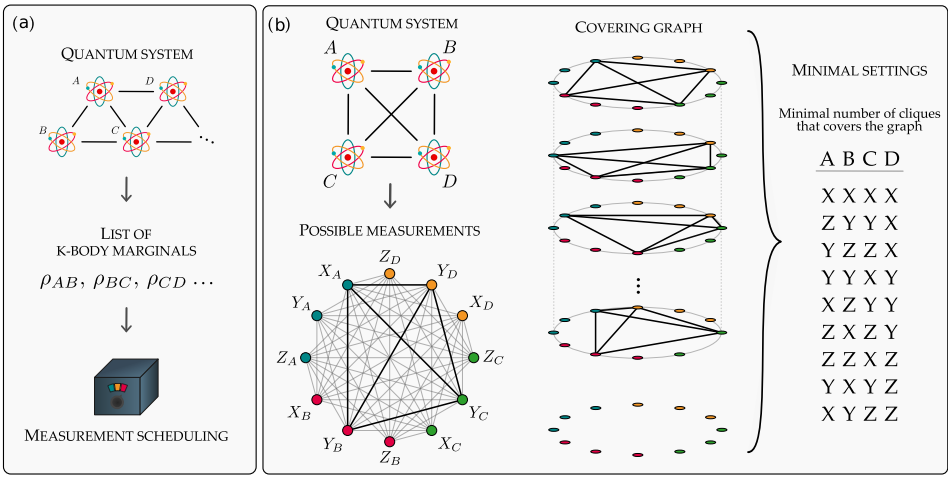


Fig. 6.1.: (a) Partial tomography is the task of reconstructing a subset of the marginals of a state. This can only be accomplished if a set of measurements that is tomographically complete on each required marginal. (b) Minimizing the number of settings can be mapped to a graph covering problem. Using local Pauli measurements as an example, we start from a connectivity graph dictating the marginals to be reconstructed — in this example, every two-body marginal. A "possible measurements" graph is used to represent the possible measurement settings and their compatibility relations. Since only compatible settings are connected, any complete subgraph (a clique) is a valid measurement. Finding a minimal amount of cliques that cover every edge of the graph (Sec. 6.2) yields an optimal set of measurements which are tomographically complete on the required marginals.

With the measurements graph expressing which measurements can be done, now we also need a graph representation of which measurements *must* be done to obtain complete data. In our running example, the "required measurements graph" G is simply the same graph: Every necessary setting (the ones in $\{ X, Y, Z \}^{\otimes 2}$) corresponds to an edge in $K_{n,3}$, thus if we obtain the expectation values associated to each edge, we have complete data to describe the marginals. (This is largely a coincidence, though, as whenever $k \neq 2$ or less than every reduced state is demanded, we will have $G \neq K$.)

Constructions in hand, the solution to the minimal measurement scheduling problem becomes obviously equivalent to finding the minimal number of cliques of K whose edges cover every edge in G. When $K = G$, this reduces to an edge clique covering problem, which has a long history in graph theory [Orl77], though the case of $K \neq G$ seems to be less studied. In turn, some instances of edge clique covering are equivalent to covering arrays [Hou+06; Mal09; RC09], a kind of combinatorial design [Col10; TI13] which is widely used in software testing [KKL+10]. This is true, for instance, in the Pauli tomography example, thus several constructions can be immediately borrowed

therefrom [Col; Tor]. Several connections between partitioning of quantum observables and graph problems or combinatorial designs have recently been reported in the literature on variational quantum eigensolvers, and error mitigation in quantum circuits. Of special note are the similar, independent results presented in [VYI20; BW24; VK24] (see Sec. 6.4).

Before discussing solutions to the edge clique cover problem, it is appropriate to generalize our representation to cases other than Pauli measurements, $k = 2$, and every marginal. Changing the last condition, by for instance asking only for the nearest-neighbours marginals on a lattice, amount to removing the edges between the unwanted marginals in G, and K remains unchanged. To represent $k > 2$, one may promote both G and K to hypergraphs (Sec. 1.2) so their edges can now represent sets with more than two elements, and even mixed numbers of elements. Lastly, the choice of Pauli measurements influences our representation because it determines the compatibility relations (expressed by the edges in K), and the set of k-body tomographically complete measurements, which is related to the edges of G in the obvious way. Other choices of measurement bases can thus be encoded by modifying these features appropriately.

6.2.1 Covering algorithms

Having mapped the measurement scheduling problem to covering problems, we now turn to the question of finding the smallest set of measurements allowing for the reconstruction of the desired marginals. This is simpler to discuss in terms of the graph quantities (see Sec. 1.2 for nomenclature, and the previous section for how it related to physical quantities). To avoid overcomplicating the notation, we will once more go back to the case of two-body marginals, from which the generalization to larger k (via hypergraphs) follows straightforwardly.

We let K and G be given graphs, and C_K be the set of maximal cliques of K. [4] As is common in graph theory, $E(G)$ and $V(G)$ denote the edge and vertex sets of G, respectively. The edge clique covering (ECC) problem is to find the smallest subset of C_K that covers all edges of G. Is it more common to consider the case $G = K$, but for us the more general case $G \neq K$ is also of interest, otherwise we will only be able to express the problem of obtaining *every* k-body marginal.

[4] Whether these cliques must be given or not as an input depends on the algorithm, as we will discuss in due time.

This problem is only well posed when $E(G)$ is a subset of $E(K)$, otherwise no covering exists. When $G = K$, the solution of this problem is equivalent to a (minimal) covering array [Col04; TI13; CD06]. Covering arrays also have a long history in the combinatorial designs and the software testing literature [Col04]. Nevertheless, few optimal solutions are known, and in practice one resorts to heuristic methods. Let us look in more detail at three exact and one heuristic methods, which will be used in this and the next chapters.

The first exact formulation is done in terms of a binary program (Sec. 1.3). Given C_K (the set of maximal cliques of a graph K, representing the possible global measurement settings) and G, we let z_c be binary optimization variables, where c is a label for each clique in C_K. Semantically, $z_c = 1$ means clique c is part of the solution, and $z_c = 0$ that it is not. Any clique that is part of the solution will be considered to cover all edges in G that are contained in it. Naturally, every edge of G must be covered by at least one clique, and we wish to minimize the number of variables z_c that are set to one. Altogether, we arrive at the following binary program:

$$\min_{\{z_c\}} \sum_c z_c \tag{6.5a}$$

s.t.

$$z_c \in \{0,1\}, \quad \forall c \in C_K \tag{6.5b}$$

$$\sum_c z_c e_c \geq 1, \quad \forall e \in E(G) \tag{6.5c}$$

$$\text{where } e_c = \begin{cases} 1 \text{ if } e \in c \\ 0 \text{ if } e \notin c \end{cases}. \tag{6.5d}$$

Notice that the values of e_c are constants defined by the graphs K and G and can be determined before setting up the model. Together, constraints 6.5c and 6.5d ensure that every edge of G is covered by at least one clique in the solution, by encoding the fact that if e is not in c, then z_c does not affect the coverage of e, otherwise it does.

Binary programs can be solved to optimality by any linear programming solver supplemented with a branch-and-bound algorithm. A plethora of highly optimized commercial and open-source solvers are available out-of-the-box, as well as user-friendly interfaces for specifying the optimization model (see Sec. 1.3). Even so, in this model the number of optimization variables is equal to the number of maximal cliques in K, which can in some instances be exponential on the number of parties. As in some cases the number of cliques in an optimal solution is expected to grow much more slowly than the total

number of cliques (for example, when only local marginals are required), it is worth considering an alternative formulation for this situation.

In the second formulation, which is also a binary program, we start with a guess N_c for an upper bound on the number of cliques in the solution. Letting $c \in \{ 1, \ldots, N_c \}$ label the cliques in the solution, we instantiate three groups of binary variables with different meanings:

1. $x_{v,c}$ for $v \in V(K) \rightarrow$ Is vertex v of K in clique c?

2. $y_{e,c}$ for $e \in E(K) \rightarrow$ Is edge e of K in clique c?

3. $z_c \rightarrow$ Is clique c of K in the solution?

These variables must be coupled to ensure that the solution is expressed in terms of the cliques of K while simultaneously covering the edges of G.

Let us first consider the "clique finding" constraints, and deal with the covering later. Suppose first that $z_c = 0$. Then, because $x_{v,c}$ is a vertex of c and c is not in the solution, vertex $x_{v,c}$ cannot be in the solution. This holds for any v, and can be represented by $x_{v,c} \leq z_c$. Now assume $e = \{ u, v \} \in E(K)$. Given a clique c that does not contain vertex v, it cannot contain the edge e (otherwise it would not be a maximal clique). Thus, we enforce that $y_{e,c} \leq x_{u,c}$ and $y_{e,c} \leq x_{v,c}$. On the other hand, if and both $x_{u,c} = x_{v,c} = 1$, then edge $y_{e,c}$ must be in the solution, and to express that we can say $x_{u,c} + x_{v,c} - 1 \leq y_{e,c}$. Lastly, when $e = \{ u, v \} \notin E(K)$, then $x_{u,c} + x_{v,c} \leq 1$.

The model still misses the edge covering constraints, which fortunately take a rather simple form: For each edge e_G of the graph G that we wish to cover, it must be true that

$\sum_c y_{e_G,c} \geq 1$. In other words, any valid solution must contain at least one clique c that contains e_G. Putting it all together, we get:

$$\min_{\{z_c, x_{v,c}, y_{e,c}\}} \sum_c z_c \tag{6.6a}$$

s.t.

$$x_{v,c} \leq z_c \tag{6.6b}$$

$$\left.\begin{aligned} y_{e,c} &\leq x_{u,c} \\ y_{e,c} &\leq x_{v,c} \\ x_{u,c} + x_{v,c} - 1 &\leq y_{e,c} \end{aligned}\right\} \quad \text{if } e = \{u, v\} \in K \tag{6.6c}$$

$$x_{u,c} + x_{v,c} \leq 1 \quad \text{if } e = \{u, v\} \notin K \tag{6.6d}$$

$$\sum_c z_c y_{e,c} \geq 1 \quad \text{if } e \in G \tag{6.6e}$$

This is again a binary optimization problem, but this time the cliques of K are not required in the input, and the number of variables is dependent on our initial guess N_c, instead of the total number of cliques in K. When no good guess N_c is known, and the user happens to use something smaller than the optimal solution, the program will be infeasible. [5]

Integer programming is occasionally highly sensitive to the model, in the sense that whether a solution will be found or not can strongly depend on how the program is posed, even though they are computing the same quantity. Since both models above lead to optimal solutions, we have used them interchangeably depending on which of them performed better for each instance. Because of how the branch-and-bound algorithm is

[5] This formulation can also be expressed as a constraint satisfaction program. Constraint programming (CP) is different from mixed integer programming (MIP) in the sense that it focus on finding a solution that satisfies logical constraints, rather than optimizing a quantity. Though there is some superficial similarities in their formulations, the algorithms used CP are largely different from the ones use in MIP. While integer programming solvers rely on linear programming relaxations and branch-and-bound algorithms, CP solvers often use backtracking algorithms and constraint propagation. Because of this, CP is usually better at handling logical constraints on discrete variables, while MIP is more suited to handling linear constraints and optimization. Equation (6.6) can be easily translated to a CP model, by reinterpreting the semantics in terms of implications. A CP solver would then try to find an attribution of variables that satisfy these logical relationships, or prove that none exists. By running through possible guesses for N_c, this could be used to find an optimal solution. However, many CP solvers are efficient for finding feasible solutions, but may struggle in proving unsatisfiability. In the course of this project, we have also implemented variations of problem as a CP model using the CP-SAT solver [PD24]. We have found that the CP model is typically much better than the MIP model in finding a good solution, provided that we leave some slack on N_c. Though this might be relevant in practice, for our purposes the MIP model and the heuristic approach discussed below were enough to obtain all the results we will discuss.

structured, interrupting the process before it finds an optimal covering will return a valid, but potentially suboptimal solution, as well as a lower-bound on the minimum number of necessary cliques. This can be useful for problems that have a prohibitive running time, and we will see examples of this in Chap. 7.

In Chap. 7 we will also discuss solutions obtained with a fast and scalable, though suboptimal, heuristic algorithm. It is a trivial generalization of the procedure proposed in [CGM16] to the case $K \neq G$: Starting with an empty set C, we want to populate it with cliques from K until a covering of G is found. We select a random edge $\{ u, v \}$ from G to be covered (currently, they are all uncovered). By assumption, $\{ u, v \} \in E(K)$ as well, otherwise the covering is infeasible. Growing a clique on K containing u and v can be done by observing that if there exists a $w \in V(K)$ in the neighbourhoods of *both* u and v, then it is safe to add w to the growing clique. For the next step, we consider the neighbourhood $N(u) \cap N(v) \cap N(w)$ of the vertices in the growing clique. A new vertex is selected from it, and the procedure continues until the intersection becomes empty, as it means no further vertices can be added to the clique. However, we are still missing the most crucial step: *Which* of the vertices in the shared neighbourhood should we select? As our goal is to cover G, a sensible choice is to choose a vertex such that the new edges will cover the maximum number of yet uncovered edges of G. This rule preserves the efficiency of the procedure, but at the cost of optimality. Clique completed, we add it to C, select another random uncovered edge of G, and repeat until G is entirely covered.

The quality of the solution will crucially depend on the two heuristics implicit above: random selection of the next edge of G to be covered, and the choice of the next vertex going into the growing cliques. An empirical analysis suggests that, at least for the case where $G = K$, they often lead to satisfactory solutions [CGM16]. Because of the randomness introduced by the first heuristic, in practice we have run the algorithm many times and selected the best result.

6.3 Application to Pauli settings

To facilitate a comparison of these observations with previous literature, let us focus on the problem of constructing the measurement scheme for every two-body marginal of an n-qubit state. Using the binary programming model outlined in Equation 6.5 and leveraging known results from covering arrays, we show the minimal solution sizes for

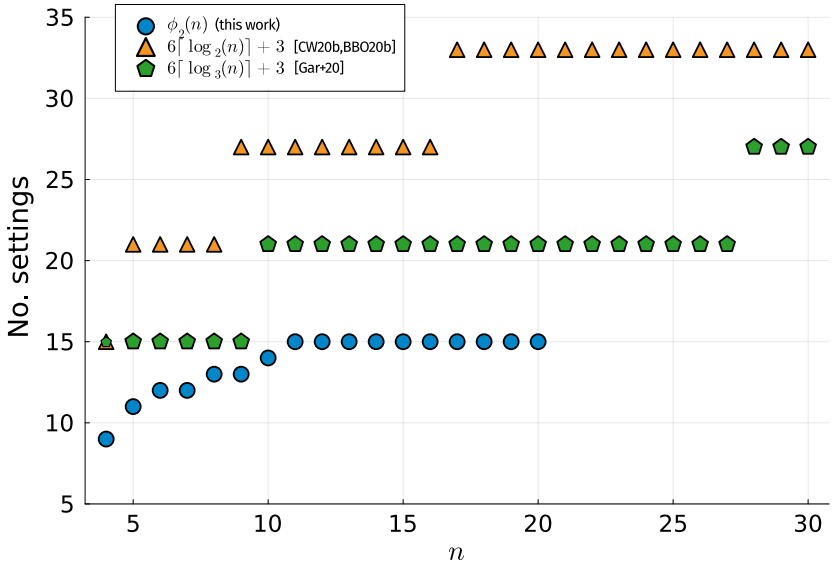


Fig. 6.2.: Number of measurement settings for two-body partial tomography of n qubits according to different constructions. Circles denote the minimal construction possible, as obtained from the equivalence to edge clique covering discussed in Sec. 6.2, and also using known optimal solutions to covering arrays [Kok+20]. Triangles and pentagons stand for the constructions in [CW20; BBO20] and [Gar+20], respectively.

systems with up to 20 qubits in Fig. 6.2. See also Fig. 6.1 for the minimal solution of the $n = 4$ case, and the associated publication for other examples [Han+24].

Our approach demonstrates a two- to three-fold improvement over previous constructions. However, the binary program encounters convergence issues for larger systems. In such cases, the heuristic algorithm or the constraint programming versions should be employed, and an example comparing the heuristic method to the exact method will be presented for a similar problem in the subsequent chapter.

We emphasize that our method is also applicable for higher k-body marginals and any specified set of marginals. For example, when $k = 3$, the number of measurement settings required for 4, 5 and 6 qubit are 27, 33, and 33, respectively. Furthermore, through additional graph-theoretic arguments, it can be shown that all nearest-neighbour marginals in planar lattices can be obtained with a small, constant number of measurement settings, regardless of the system size [Han+24, App. A]. The actual settings for that can be easily constructed from any valid covering array for the k-body marginals, where k is the maximum degree of the graph representing the lattice.

While this theoretically resolves the problem of finding minimal measurement sets for any partial tomography task, practical limitations prevent the attainment of the most optimal solutions for very large systems. Techniques used in software testing could lead to further refinements of the heuristic and constraint programming methods, and bridge the gap between theory and practical applications.

6.4 Discussion

Many interesting aspects of partial tomography discussed in the associated publication were not covered in this chapter, some highlights being:

Recursive constructions Obtaining the optimal measurement scheduling is computationally costly, but known solutions can be recursively extended to accommodate more parties. When only the marginals described by edges on a planar graph are required, one can show that the number of settings is constant irrespective of the number of particles in the system.

Non-Pauli measurements Relaxing the restriction to Pauli measurements and considering instead any local project measurements, we demonstrate that only 3^k measurement settings are sufficient, regardless of the number of qubits. This possibility was more thoroughly analysed in the $k = 2$ case, where we optimize the measurement vectors so that the variances induced on the measurement results are minimized.

Experimental verification An experimental verification with 6-photon Dicke states was conducted, successfully obtaining every two-body marginal. This was done for both the optimal Pauli solution, which has 12 settings, and for the case of nine projective measurements mentioned above. The experimental results seem to corroborate the trade-off between fewer settings and a larger variance.

During the course of this work, we became aware of similar results in the literature on variational quantum eigensolver and error mitigation. Many of their conclusions are applicable to our work and vice versa. Of particular significance are studies where the problem of partitioning observables was also mapped to graph-theoretic quantities [VYI20] and in the $k = 2$ case to covering arrays [BW24]. A noteworthy approach to the problem by using the sample cost as the central figure of merit is discussed in [VK24], where it is shown that there are cases where it is advantageous to have a larger set of

measurement settings than the minimal one. Their constructions are also expressed in terms of graph-theoretic quantities, in this case a vertex multicolouring problem. This discussion underscores the interdisciplinary nature of our findings and their relevance to other research areas.

7 Partial tomography of fermionic simulators

This chapter is based on *A. Ghoshal, C. de Gois, K. Hansenne, O. Gühne, et al. "Partial tomography of fermionic simulators". In:* (In preparation) *(2025)* [Gho+25]. AG and CG equally contributed to these results.

Quantum simulation of fermionic systems has long been a central challenge physics, chemistry and materials science. While digital quantum computers often rely on fermion-to-qubit mappings [Tac+20; Llo96] (commonly the Jordan-Wigner or the Bravyi-Kitaev transformations [JW28; BK02]), these encodings incur significant overhead due to the nonlocal structure of fermionic operators. Simulating fermionic systems natively, in contrast, would enable a direct representation of fermionic statistics without the mapping overhead. This can lead to more efficient simulations of fermionic many-body systems, such as strongly correlated electronic systems and molecular structures, with applications in quantum chemistry and materials science [Arg+19; GAN14].

A recent proposal in this direction is the programmable neutral-atom fermionic quantum processor introduced in [Gon+23]. Therein, fermionic modes are encoded in the internal states of neutral atoms trapped in optical tweezers. This architecture naturally supports a universal gate set that guarantee Fermi statistics on hardware level, thus sidestepping the overhead associated with mapping fermionic statistics onto qubits. The gate set proposed in the paper contains two types of operations. The interaction gates $\mathcal{U}_{i,j}^{(\mathrm{int})}(\theta) = e^{-i\theta n_i n_j}$ impart a common phase conditioned on the occupation of both modes, and the tunnelling gates

$$\mathcal{U}_{i,j}^{(t)}(\theta_1, \theta_2, \theta_3) = e^{-i\left[\frac{\theta_1}{2}\left(e^{-i\theta_2} c_i^\dagger c_j + \mathrm{H.c.}\right) + \frac{\theta_3}{2}(n_i - n_j)\right]} \tag{7.1}$$

enable mode hopping and the creation of relative phases. Together, they form a universal set for fermionic quantum computation [Gon+23]. Implementing them in a programmable neutral atom array is challenging, but might offer a direct route to simulating fermionic dynamics without the resource overhead associated to qubit mappings, potentially enabling scalable simulation of large fermionic systems.

Reading out the results of a simulation ultimately requires measuring the system. Many-body systems are prohibitively difficult to fully characterize, and in practice that might also be wasteful, as commonly only a few special properties of the system are of interest. This is particularly true in fermionic states, where much of the relevant physical information is contained in few-body marginals of the state, including ground state properties, excitation energies and correlation functions, among others [Col63; Löw55]. By concentrating only on reduced density matrices (RDMs), one could thus hope to find scalable measurement procedures while still capturing the essential physics. This is indeed the case, and partial state tomography — the task of reconstructing particular marginals of a quantum state — has drawn significant recent attention both in the case of bosonic and fermionic systems [BBO20; Ara+22; Til+21; Jia+20; ZRM21; Kin+25; Han+24; VYI20; CW20; Gar+20].

Here we reassess this problem from the perspective of native fermionic simulators, and derive protocols described directly in terms of occupation measurements and two simple tunnelling gates. Building on top of a straightforward procedure to infer the expectation values of fermionic coherences $c_i^\dagger c_j$ (Sec. 7.1), we investigate ways of parallelizing the measurements of any set of correlators. As an example, we give a procedure to learn every single-body observable on an n-mode system, showing that $\Omega(n)$ different measurement settings are necessary and sufficient (Secs. 7.2). Key to our solution is an analogy to edge colourings of an associated graph, which was also used to solve an analogous problem framed in terms of Majorana operators [Ara+22]. This discussion is then extended to the case of four-point correlators $c_i^\dagger c_j c_k^\dagger c_l$, or two-body observables (Sec. 7.3). Efficient parallelization is connected to a particular type of graph covering (cf. Sec. 6.2.1 and [Han+24; VYI20]), and a procedure to create a sufficient set of measurements for determining the entire two-body reduced density matrix is described and applied. For particular cases of local two-body observables, a simple argument is used to conclude that the number of measurement settings is constant, independently of the number of modes in the system.

7.1 Measuring fermionic coherences

Consider first the task of characterizing modes i and j in a fermionic simulator. This requires measuring the occupation numbers $n_i = c_i^\dagger c_i$ and $n_j = c_j^\dagger c_j$, as well as the

coherence term $c_i^\dagger c_j$. [1] While occupations n_i are directly accessible by projective measurements of mode occupancy, the coherence $c_i^\dagger c_j$ poses a fundamental challenge: $c_i^\dagger c_j$ is non-Hermitian, and thus does not correspond to a physical observable. This contrasts sharply with qubit systems, where σ_X and σ_Y measurements can be implemented via unitary rotations followed by σ_Z readout — a strategy that relies on the target observables being Hermitian.

This can be resolved by decomposing the coherence $c_i^\dagger c_j$ into Hermitian components, which could then be measured individually by finding a transformation from the number basis into these quantities, in terms of the gate set of the simulator. To that end, we define the spin-like operators

$$S_\alpha^{ij} = \begin{pmatrix} c_i^\dagger & c_j^\dagger \end{pmatrix} \sigma_\alpha \begin{pmatrix} c_i \\ c_j \end{pmatrix}, \tag{7.3}$$

where σ_α with $\alpha \in \{X, Y, Z\}$ are the Pauli matrices. Crucially, $S_Z^{ij} = c_i^\dagger c_i - c_j^\dagger c_j$ is diagonal in the occupation basis and is thus directly measurable. However, $S_X^{ij} = c_i^\dagger c_j + c_j^\dagger c_i$ and $S_Y^{ij} = i(c_j^\dagger c_i - c_i^\dagger c_j)$ require indirect measurement strategies. Crucially, they encode the real and imaginary parts of $\langle c_i^\dagger c_j \rangle$ via

$$\langle c_i^\dagger c_j \rangle = \frac{\langle S_X^{ij} \rangle + i \langle S_Y^{ij} \rangle}{2}, \tag{7.4}$$

thus their measurement results can be used to determine $\left\langle c_i^\dagger c_j \right\rangle$.

The next step is to map S_X^{ij} and S_Y^{ij} into occupation measurements by using the universal gate set. That can be done by noticing that $S_X^{ij} = U_X^\dagger S_Z^{ij} U_X$, where $U_X = e^{-i\theta S_Y^{ij}}$ and $\theta = (4m+3)\pi/4$ for integer m. Therefore, U_X rotates S_X^{ij} into the occupation number basis, enabling $\langle S_X^{ij} \rangle$ to be inferred from $\langle n_i \rangle - \langle n_j \rangle$. Similarly, S_Y^{ij} is unitary equivalent to S_Z^{ij} for $U_Y = e^{-i\phi S_X^{ij}}$ and $\phi = (4m+1)\pi/4$. Together, these observations tell us that a two-mode fermionic system can be fully characterized by three measurement settings:

1. Measure the occupation numbers n_i and n_j.

2. Evolve the system with U_X, and then measure the occupation numbers.

[1] Here, $c_i^\dagger$ and c_i are the creation and annihilation operators, which satisfy the canonical fermionic anticommutation relations

$$\{c_i, c_j\} = 0, \quad \{c_i^\dagger, c_j\} = \delta_{ij}, \text{ and } (c_i)^2 = (c_i^\dagger)^2 = 0, \tag{7.2}$$

for all modes i, j. These relations encode the Pauli exclusion principle, determining the allowed fermionic states [BF04].

3. Evolve the system with U_Y, and then measure the occupation numbers.

Up to a statistical significance determined by the number of samples for each setting (see Chap. 5), this yields $\left\langle c_i^\dagger c_i \right\rangle$, $\left\langle c_j^\dagger c_j \right\rangle$ and $\langle c_i \dagger c_j \rangle$ (the other possibilities, for instance $\left\langle c_j^\dagger c_i \right\rangle$, being just complex conjugates).

Importantly, the changes of basis associated to U_X and U_Y can be implemented with the native fermionic gates introduced proposed in [Gon+23]. Explicitly, this can be done with the tunnelling gates in Eq. (7.1) by setting $\mathcal{U}_{i,j}^{(t)}(2\theta, \pi/2, 0)$ and $\mathcal{U}_{i,j}^{(t)}(2\phi, 0, 0)$, respectively.

Building on top of the measurement procedure for coherences between two modes, in the next sections we will extend it the measurement of any set of two- and four-point correlators for an n-mode fermionic simulator.

7.2 Measurement schemes for two-point correlators

Suppose we now have an n-mode fermionic system, and that we want to learn the values of some set $\mathcal{O}_1 = \{ O_i^{(1)} \}_i$ of single-body observables. As any single-body fermionic observable can be decomposed in terms of two-point correlators,

$$O^{(1)} = \sum_{i,j} \alpha_{ij} c_i^\dagger c_j, \tag{7.5}$$

learning $\mathcal{O}_1$ amounts to measuring the expectation values for each $\alpha_{ij} \neq 0$ in the target observables. In particular, having every $\left\langle c_i^\dagger c_j \right\rangle$ is sufficient to determine the 1-RDM $\rho_{ij}^{(1)} = \langle c_i^\dagger c_j \rangle$, or equivalently, the expectation value of any single-body fermionic observable.

We want to find a small set of measurement settings that is sufficient to learn $\mathcal{O}_1$, under the constraint that any measurement must be based directly on the native fermionic gates. This will be done by leveraging the basic procedure to measure coherences, and parallelizing wherever possible. To represent the compatibility relations, we associate to $\mathcal{O}_1$ a graph G in which the vertices $i \in \{ 1, \ldots, n \}$ correspond to the modes of the fermionic system, and an edge $\{ i, j \}$ represents a term $c_i^\dagger c_j$ appearing in the target

observables. [2] For example, the 1-RDM case is associated to the complete graph $G = K_n$, which has n vertices and edges between any two vertices. In other situations, the properties of interest might already be revealed by measuring neighbouring modes in a particular lattice, and G should be constructed accordingly.

It follows from the discussion in Sec. 7.1 that any edge $\{ i, j \}$ in G (that is, the expectation value of $c_i^\dagger c_j$ when $i \neq j$) can be determined from two measurement settings, S_X^{ij} and S_Y^{ij}. Notice that occupation number terms might also be present in $\mathcal{O}_1$, and in that case they can be obtained from the single global setting $n_1 \dots n_n$. On the other hand, the tunnelling gates associated to S_α^{ij} and S_β^{kl} (for $\alpha, \beta \in \{ X, Y \}$) can only be performed in parallel if they are done on non-overlapping modes, that is, $i \neq j \neq k \neq l$. In terms of the graph representation, this means the edges $\{ i, j \}$ and $\{ j, l \}$ cannot share vertices. A straightforward solution for parallelization is thus to colour the edges of G in such a way that no two edges incident to the same vertex share a colour. Each colour then denotes a set of independent pairs of modes, over which the measurements can be parallelized. For each colour, we must measure both S_X^{ij} and S_Y^{ij}, thereby requiring two distinct settings per colour. Letting $\chi'(G)$ denote the minimum number of required colours (known as the *chromatic index*), it follows that at most $2\chi'(G) + 1$ measurement settings are necessary to obtain all the two-point correlators in the set of observables $\mathcal{O}_1$ associated to the graph, where the additional setting accounts for the occupation number measurement.

Edge colouring is well understood in graph theory, and the Vizing-Gupta theorem states that, for any graph, $\Delta \leq \chi' \leq \Delta + 1$, where Δ is the degree of the graph (the maximum number of edges incident to a vertex) [Viz65; Gup78]. In general, the problem of determining whether χ' is Δ or $\Delta + 1$ is NP-complete [Hol81; GJ83]. Settling for a $\Delta + 1$ colouring simplifies the problem, as one can always be constructed for any graph in time complexity $O(|V||E|)$ [MG92], which can be improved for particular classes of graphs (cf. [CK08, Table 1]). Moreover, as the optimality gap is at most one, in the worst case only two extra measurement settings are wasted in doing so. For many physically meaningful cases, χ' is independent of the number of modes in the system: on a square lattice, for example, $\Delta = 4$, therefore only 9 measurement settings are sufficient to fully characterize its nearest-neighbour observables. While this construction gives a sufficient solution, in [Gho+25] we also show that at least

$$1 + \frac{2|E(G)|}{n} \tag{7.6}$$

[2] When using the settings S_X^{ij} and S_Y^{ij}, the measurement data needed to obtain $\langle c_i^\dagger c_j \rangle$ also allows for determining $\langle c_j^\dagger c_i \rangle$, thus these two correlators correspond to the same edge in G.

settings are necessary, where $|E(G)|$ is the number of edges in the graph.

7.2.1 Example: Single-body reduced density matrix

As a concrete example, let us analyse the case of obtaining the entire 1-RDM. Then, $\mathcal{O}_1$ may contain any single-body observable, and G would be the complete graph K_n. It is well-known that $\chi'(K_{2n}) = \Delta$ and $\chi'(K_{2n+1}) = \Delta+1$. To construct the optimal colouring, let us first consider the case of an even number of modes, K_{2n}, and a set of colours labelled by $C = \{ c_k \}_{k=1}^{2n-1}$. The construction can be visualized geometrically in the complex plane (see Fig. 7.1b for an example with six vertices). We start by distributing the vertices $v \in \{ 1, \ldots, 2n-1 \}$ around the unit circle, at the positions $e^{v\frac{2\pi i}{2n-1}}$. The remaining vertex, $v = 2n$, goes at the origin. Edges connecting the central vertex to any other are of the form $\{ v, 2n \}$, for v in $\{ 1, \ldots, 2n-1 \}$, and they must be given different colours since they meet at $2n$. To satisfy that, we attribute c_v for each edge of this form. A total of $2n-1$ edges are now coloured. We proceed by fixing any $v \in \{ 1, \ldots, 2n-1 \}$ and consider the edges of the form $\{ v + \ell \mod (2n-1), f(v,\ell) \}$ where

$$f(v,l) = \begin{cases} (v-l) \mod (2n-1) & \text{if } (v-l) \neq 0 \\ 2n-1 & \text{if } (v-l) = 0 \end{cases}, \tag{7.7}$$

and $\ell \in \{ 1, \ldots, n-1 \}$. From their angles, it is easy to see these edges are perpendicular to $\{ v, 2n \}$, and they all being parallel allows us to colour them with the same colour, c_v.

Proceeding analogously for the other choices of v and noticing that for any $v' \neq v$, the edges $\{ v + \ell \mod (2n-1), f(v,\ell) \}$ and $\{ v' + \ell \mod (2n-1), f(v',\ell) \}$ are distinct, there are now $(n-1)(2n-1)$ newly coloured edges. Summing up with the initial ones, we get $n(2n-1)$ distinct edges. This is equal to the total number of edges in K_{2n}, and the colouring is valid by construction. For the case of an odd number of vertices ($G = K_{2n-1}$), we can simply colour the graph K_{2n} and then remove the extra vertex and its associated edges, ending with $2n-1$ colours.

Similar constructions appear in many forms across the combinatorial designs' literature, where they are also presented as the 1-factorisability of complete graphs, as the "round-robin tournament" problem or as symmetric Latin squares; see for example [CD06, Chap. 5.5] for a more elegant construction using group theoretic methods. As can be expected,

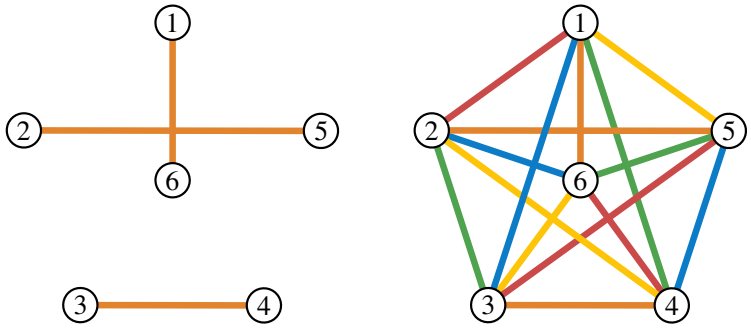


Fig. 7.1.: Colouring of the K_6 graph using 5 colours. In a we pick a colour for the radial edge $\{1,6\}$ then draw all edges that are perpendicular to it. The second step b is to rotate this construction using a different colour. This solution is optimal, and the procedure can be naturally extended to any complete graph with an even number of vertices. Each colour corresponds to two settings; for a, these settings are $S_X^{1,6}S_X^{2,5}S_X^{3,4}$ and $S_Y^{1,6}S_Y^{2,5}S_Y^{3,4}$.

this is not the only possible edge colouring K_n. In fact, the number of non-isomorphic colourings for $n = 4, 6, 8$ and 10 are known to be $1, 6, 6240$ and 1225566720 [OEI].

Most importantly, it is now straightforward to construct a sufficient and economical measurement procedure for learning the 1-RDM of any fermionic system with n modes:

Algorithm 1: Measurements for the 1-RDM (even n)

measure $n_1 \dots n_n$;
for *each* $v \in \{1, \dots, n-1\}$ **do**
  select modes $m_v = \{\{v, n\}\} \cup \{\{v + \ell \mod (n-1), f(v,\ell)\}\}_{\ell=1}^{n/2-1}$;
  apply U_X to each pair in m_v and measure $n_1 \dots n_n$;
  apply U_Y to each pair in m_v and measure $n_1 \dots n_n$;
end

For odd numbers of modes, this can be straightforwardly modified as explained previously. The data collected by these measurements can then be used in Eq. (7.4) to determine the values of $\langle c_i^\dagger c_j \rangle$ for any i and j. In total, this procedure requires $2n-1$ measurements which, together with the lower bound from Eq. (7.6), leads us to conclude that the number of settings for learning all single-body observables is $\Theta(n)$.

As mentioned in the introduction to this chapter, proposals for measuring the fermionic 1-RDM in the Majorana basis exist in the literature. This representation has the advantage of corresponding to Pauli strings, whose compatibility relations can then be used for

grouping the operators in commuting sets, but it has the drawback that they cannot be directly measured in the natural basis of the fermionic simulator we consider. Of special notice are the constructions in [BBO20], where the authors demonstrate that $2n-1$ measurement circuits are sufficient for the reconstruction of the 1-RDM, and [Ara+22], where the problem of nearest neighbours partial tomography on typical lattices is considered, and an edge colouring construction is used to find that $2\chi'(G)+1$ measurement settings are sufficient.

7.3 Measurement schemes for four-point correlators

While single-body observables can capture key properties of fermionic systems, such as the kinetic energy and the occupation numbers of the natural orbitals, they cannot capture information on correlations. Extending measurement protocols to include two-body observables enables the characterization of chemically and physically relevant quantities such as excitation energies, ground-state properties, and spatial correlation functions [Löw55; Col63].

In second quantization, two-body observables can be written as

$$O^{(2)} = \sum_{i,j,k,l} \alpha_{ijkl} c_i^\dagger c_j c_k^\dagger c_l. \tag{7.8}$$

Consequently, measuring a set $\mathcal{O}_2 = \{ O_i^{(2)} \}$ of target two-body observables reduces to estimating their four-point correlators $c_i^\dagger c_j c_k^\dagger c_l$. Determining them for all possible i, j, k and l allows us to compute the expectation value of all two-body observables, or equivalently, describe the 2-RDM of the system, $\rho^{(2)}_{ijkl} = \langle c_i^\dagger c_j c_k^\dagger c_l \rangle$.

Let us consider the four-point correlators case by case. Due to the Pauli exclusion principle, $c_i^\dagger c_j c_k^\dagger c_l$ vanish if $i = k$ or $j = l$. There are three remaining cases in terms of indices in (i, j, k, l):

(i) Two indices are repeated.

(ii) One index is repeated.

(iii) All indices are distinct.

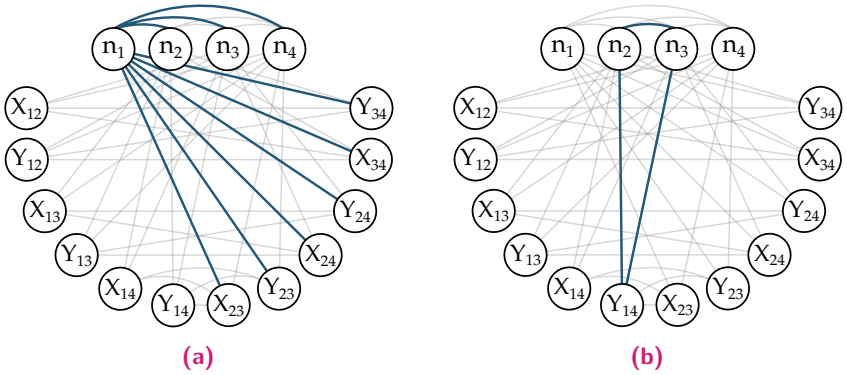


Fig. 7.2.: The measurement graph G_M for a four-mode system, using the shorthand notation $S_\alpha^{ij} = \alpha_{ij}$. Each node represents a one- or two-mode operator. Edges connect nodes if the corresponding operators can be measured in parallel, which happens if they do not share indices. (a) highlights the edges incident to n_1, and (b) shows a clique of G_M, corresponding to the measurement $n_2 n_3 S_Y^{14}$.

The associated four-point correlators are of the form

$$\left\langle c_i^\dagger c_i c_j^\dagger c_j \right\rangle = \langle n_i n_j \rangle \,, \tag{7.9}$$

$$\left\langle c_i^\dagger c_i c_j^\dagger c_k \right\rangle = \left\langle n_i c_j^\dagger c_k \right\rangle, \tag{7.10}$$

$$\left\langle c_i^\dagger c_j c_k^\dagger c_l \right\rangle, \tag{7.11}$$

for distinct $i, j, k, l \in [n]$, where $[n]$ is a shorthand for $\{\, 1, \ldots, n \,\}$.

Correlators of the form in Eq. (7.9) can be determined by joint occupation measurements on modes i and j. For a single repeated index [Eq. (7.10)], measuring $n_i S_X^{jk}$ and $n_i S_Y^{jk}$ suffices. This follows from the basic measurement procedure for two modes discussed in Sec. 7.1. Finally, $\left\langle c_i^\dagger c_j c_k^\dagger c_l \right\rangle$ can be obtained from the measurements of $S_X^{ij} S_X^{kl}$, $S_X^{ij} S_Y^{kl}$, $S_Y^{ij} S_X^{kl}$, and $S_Y^{ij} S_Y^{kl}$ through

$$4 c_i^\dagger c_j c_k^\dagger c_l = S_X^{ij} S_X^{kl} - S_Y^{ij} S_Y^{kl} + i \left(S_X^{ij} S_Y^{kl} + S_Y^{ij} S_X^{kl} \right). \tag{7.12}$$

Our goal is to find a small set of measurements to obtain the expectation values of any given set of four-point correlators. This will be done by considering the measurements' compatibility relations to map the solution to an edge clique cover problem (cf. 6.2.1) Let G_M be a graph representing the feasible measurements in the fermionic simulator. Each

number operator n_i and each operator S_X^{ij} and S_Y^{ij} for $i, j \in [n]$ with $i < j$ corresponds to a vertex in G_M. The reason for taking $i < j$ is that

$$S_X^{ji} = S_X^{ij} \quad \text{and} \quad S_Y^{ji} = -S_Y^{ij}, \tag{7.13}$$

ultimately meaning their expectation values can be post-processed from one another. Vertices in G_M are connected if the corresponding measurements can be performed in parallel. Because of the tunnelling gates acting on i and j simultaneously for any measurement of S_α^{ij} (where $\alpha \in \{ X, Y \}$), vertices S_α^{ij} and S_α^{kl} are only connected when the four indices are distinct [Fig. 7.2a]. This applies similarly for n_i and S_α^{jk}. On the other hand, n_i and n_j are connected for any $i \neq j$. Consequently, by construction, cliques of G_M correspond to measurements that can be performed in parallel [Fig. 7.2b].

Next, we construct the target graph G_T representing which measurements are required to determine the four-point correlators in $\mathcal{O}_2$. The set of vertices of G_T is the same of G_M, and we connect the vertices associated to measurements that must be performed to obtain $\mathcal{O}_2$. For example, $\langle n_i n_j \rangle$ would require us to connect the vertices n_i and n_j in G_T. For $\left\langle n_i c_j^\dagger c_k \right\rangle$ [Eq. (7.10)], two edges are necessary, $\{ n_i, S_X^{jk} \}$ and $\{ n_i, S_Y^{jk} \}$. (Because of the way we label the vertices, if $j > k$, the corresponding vertices would be S_X^{kj} and S_Y^{kj}, and that both $\left\langle n_i c_j^\dagger c_k \right\rangle$ and $\left\langle n_i c_k^\dagger c_j \right\rangle$ are represented by the same two edges is justified by the fact that both correlators can be obtained from the same measurement settings, being complex conjugate of each other.) Lastly, to represent four-point correlators with distinct indices, we draw four edges, corresponding to the four measurements needed in order to reconstruct the correlator, $S_\alpha^{ij} S_\beta^{kl}$ for $\alpha, \beta \in \{ X, Y \}$.

With G_M and G_T in hand, we can construct a small set of sufficient measurement settings to obtain $\mathcal{O}_2$ by finding a small set of cliques of G_M (the possible measurements) that cover every edge of G_T (the required measurements). This is a variant of the edge clique covering problem in graph theory, previously discussed in Sec. 6.2. For small instances, the minimum solution can be found exactly, for example using backtracking algorithms or binary programming, otherwise one can fall back to a variety of heuristic procedures (cf. Sec. 6.2.1). Notice that this and other graph quantities were recently proven useful for optimizing measurement settings in a variety of similar problems, for example [Zha+20; Gok+19; VYI20; BW24; VK24; Han+24].

7.3.1 Example: Two-body reduced density matrix and lattices

As a first example, we consider the problem of learning every two-body observable, which is equivalent to determining all four-point correlators for a system with n modes. To start, notice that G_M depends only on the number of modes in the system and is therefore unaffected by the set of observables we want to measure. On the other hand, G_T could be just the same as G_M, but one can do better: Due to the canonical anticommutation relations, many four-point correlators are equivalent, thus it is unnecessary to estimate all of them. Namely, when two indices are repeated [Eq. (7.9)] we have $\langle n_i n_j \rangle = \langle n_j n_i \rangle$, hence with no loss of generality we let $i < j$. Additionally, when no index is repeated [Eq. (7.11)],

$$\left\langle c_i^\dagger c_j c_k^\dagger c_l \right\rangle = -\left\langle c_k^\dagger c_j c_i^\dagger c_l \right\rangle = -\left\langle c_i^\dagger c_l c_k^\dagger c_j \right\rangle = \langle c_k^\dagger c_l c_i^\dagger c_j \rangle, \tag{7.14}$$

for $i < k$ and $j < l$. So we only need to obtain $\left\langle c_i^\dagger c_j c_k^\dagger c_l \right\rangle$, for example. To determine every four-point correlator it is thus enough to measure

(i) $n_i n_j$ with $i < j$,

(ii) $n_i c_j^\dagger c_k$ with $j < k$,

(iii) $c_i^\dagger c_j c_k^\dagger c_l$ with $i < k < l$ and $i < j < l$,

where $i, j, k, l \in [n]$ are all distinct. Following the discussion in the previous section, these can be translated as a target graph G_T.

A sufficient measurement scheduling can now be obtained using the methods described in Sec. 6.2.1, an edge clique cover of G_M onto G_T can be obtained, providing a sufficient measurement scheduling. As llustrated in Fig. 7.3, while the exact minimum can only be computed for a modest number of modes, the heuristic algorithm can provide good solutions for much larger system sizes. Explicit solutions to a few cases are listed in Sec. 7.3.1.1.

Although the previous example yields the entire 2-RDM, many physical systems are naturally arranged or engineered in periodic geometries, such as two-dimensional square lattices, where interactions are often localized. Often in such cases, the key goal is to characterize the local properties of the system by performing two-body tomography on specific regions of the lattice. For a square lattice for instance, these regions correspond to the square tiles formed by four modes. By focusing on the two-body correlations

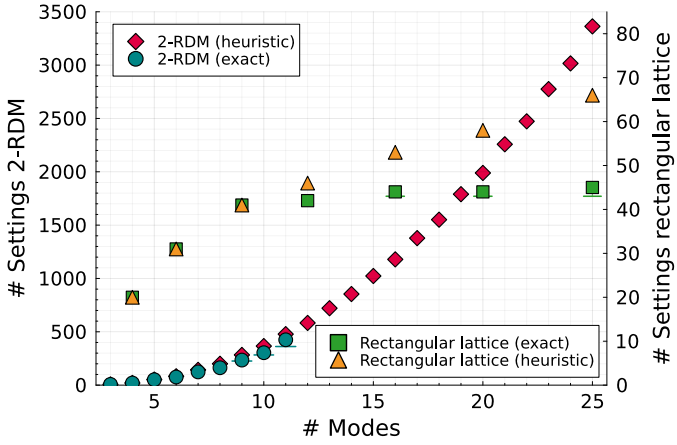


Fig. 7.3.: Number of settings for learning every two-body observable (left axis) and nearest-neighbour observables on rectangular lattices of size (x, x) and $(x, x+1)$ (right axis), as a function of the number of modes n. The exact method builds a minimal solution via binary programming, while the heuristic provides suboptimal solutions but scales efficiently. Error bars indicate cases where convergence was not guaranteed, and a smaller solution down to the bar may exist. An explicit construction for square lattices using 80 settings independently of the system size is also discussed in Sec. 7.3.1. However, the exact numerical solutions suggest a slightly better plateau at around 45 settings.

within each tile, this approach captures essential local features of the system, and can largely simplify the measurement procedure. Most notably, there is a range of cases where characterizing the correlators within individual tiles can be done with a constant number of settings, independent of the total number of modes in the lattice.

This is rather easy to see, and best explained by an example. Consider the 5×6 square lattice depicted in Fig. 7.4. From our previous discussion, the four-point correlators involving modes 1, 2, 7, 8 (region S_1 in Fig. 7.4) can be learned from the 20 measurement settings for the four modes 2-RDM listed in Sec. 7.3.1.1. Since the modes contained in the other regions shaded with the same colour do not overlap with the ones in S_1, they can be measured simultaneously (it suffices to apply the same solution to the corresponding modes in parallel). This is however not true for the regions depicted with different colours in Fig. 7.4b. Though they can themselves be measured in parallel, they overlap with one another thus cannot be parallelized. It is easy to see that any larger rectangular lattice can be coloured with four colours in the same way. With four colours and 20 setting per colour, 80 settings are sufficient, independently of the system size. Although constant, this turns out to not be minimal, as can be seen from the clique covering solutions to this problem, presented in Fig. 7.3, and which suggests a plateau around 40 settings.

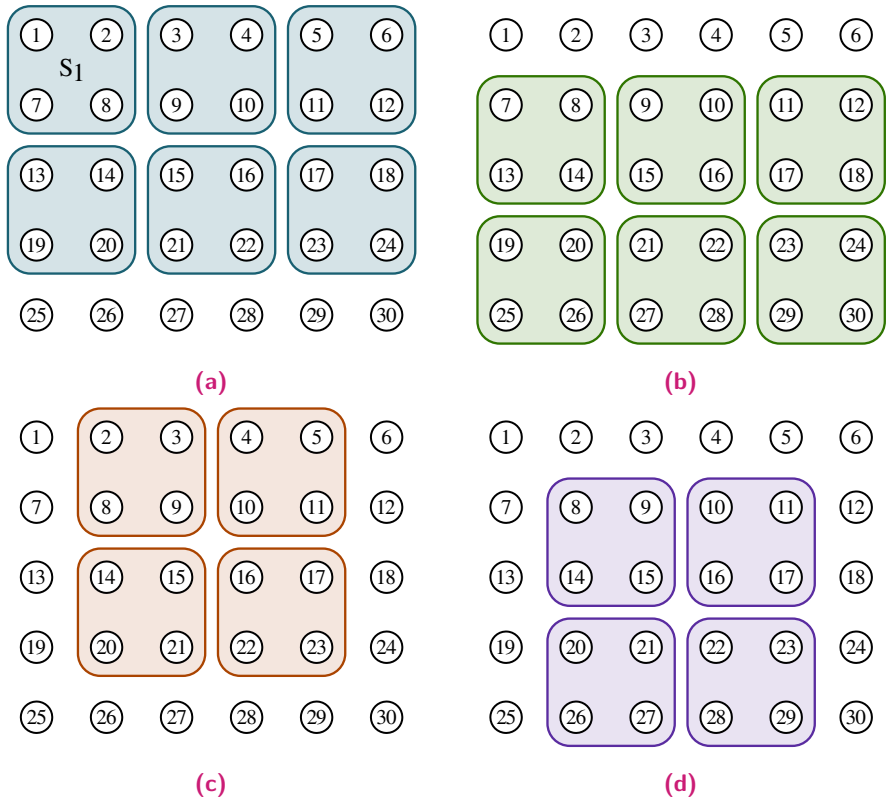


Fig. 7.4.: Non-overlapping tiles of a rectangular lattice. The four-point correlators involving modes in each shaded region can be determined with the 20 measurement settings for the four-mode 2-RDM (Sec. 7.3.1.1). Modes in tiles of the same colour can be measured in parallel, using the same settings on the respective modes. Since any rectangular lattice can be covered with four different colours, all nearest-neighbour observables can be determined with 80 settings.

Lattice	No. settings
Square	80
Triangular	42
Hexagonal	228
Kagome	166

Tab. 7.1.: Constant number of settings for obtaining all four-point correlators of the modes involved in each tile of the lattice. Square and triangular lattices are discussed in the text. The hexagonal lattice is made of six-mode tiles that can be grouped in three sets of non-adjacent tiles. As obtaining all four-point correlators of six modes requires 76 settings (Sec. 7.3.1.1), 228 global settings are sufficient. The Kagome lattice is made of triangular and hexagonal tiles, and obtaining all four-point correlators of each three- and six-modes tiles can be done with 166 measurement settings.

Lastly, notice that the same reasoning extends to other common lattices. For example, the basic measurement for a triangle is equivalent to measuring K_3, which can be done with 7 settings. As six colours are needed to cover the triangular lattice with non-overlapping triangles, 42 measurement settings are sufficient. Further examples are listed in Table 7.1.

7.3.1.1 Example solutions for four-point correlators

Here we list the solutions found with the exact formulation for the 2-RDM problem with 3, 4 and 6 modes. These are of particular interest because, as we will discuss, they can be used as the basis for constructing the coverings of the local partial tomography problem on triangular, rectangular and hexagonal tiles, respectively. Each column shown below corresponds to a measurement setting, with the number on the first row labelling the settings.

2-RDM for 3 modes

1.	2.	3.	4.	5.	6.	7.
n_1	n_1	n_1	n_2	n_2	n_3	n_3
n_2	$X_{2,3}$	$Y_{2,3}$	$X_{1,3}$	$Y_{1,3}$	$Y_{1,2}$	$X_{1,2}$
n_3						

2-RDM for 4 modes

1.	2.	3.	4.	5.	6.	7.	8.	9.	10.
n_1	n_1	n_1	n_1	n_1	n_1	n_2	n_2	n_2	n_2
n_2	n_2	n_3	n_3	n_4	n_4	n_3	n_3	n_4	n_4
$X_{3,4}$	$Y_{3,4}$	$X_{2,4}$	$Y_{2,4}$	$X_{2,3}$	$Y_{2,3}$	$X_{1,4}$	$Y_{1,4}$	$X_{1,3}$	$Y_{1,3}$

11.	12.	13.	14.	15.	16.	17.	18.	19.	20.
n_3	n_3	$X_{1,2}$	$X_{1,2}$	$Y_{1,2}$	$Y_{1,2}$	$X_{1,3}$	$X_{1,3}$	$Y_{1,3}$	$Y_{1,3}$
n_4	n_4	$X_{3,4}$	$Y_{3,4}$	$X_{3,4}$	$Y_{3,4}$	$X_{2,4}$	$Y_{2,4}$	$X_{2,4}$	$Y_{2,4}$
$X_{1,2}$	$Y_{1,2}$								

2-RDM for 6 modes

1.	2.	3.	4.	5.	6.	7.	8.	9.	10.
n_1	n_1	n_1	n_1	n_1	n_1	n_1	n_1	n_1	n_1
n_2	n_2	n_3	n_3	n_4	n_4	n_5	n_5	n_6	n_6
$Y_{3,4}$	$X_{3,6}$	$Y_{2,5}$	$Y_{2,6}$	$X_{2,5}$	$Y_{2,3}$	$X_{2,6}$	$X_{2,3}$	$X_{2,4}$	$X_{2,4}$
$X_{5,6}$	$Y_{4,5}$	$X_{4,6}$	$X_{4,5}$	$Y_{3,6}$	$Y_{5,6}$	$X_{3,4}$	$Y_{4,6}$	$X_{3,5}$	$Y_{3,5}$

11.	12.	13.	14.	15.	16.	17.	18.	19.	20.
n_1	n_1	n_1	n_1	n_1	n_1	n_2	n_2	n_2	n_2
n_6	n_6	n_6	n_6	n_6	n_6	n_3	n_3	n_4	n_4
$Y_{2,4}$	$Y_{2,4}$	$X_{2,3}$	$X_{2,3}$	$Y_{2,3}$	$Y_{2,3}$	$Y_{1,4}$	$Y_{1,5}$	$X_{1,6}$	$Y_{1,6}$
$X_{3,5}$	$Y_{3,5}$	$X_{4,5}$	$Y_{4,5}$	$X_{4,5}$	$Y_{4,5}$	$Y_{5,6}$	$Y_{4,6}$	$X_{3,5}$	$Y_{3,5}$

21.	22.	23.	24.	25.	26.	27.	28.	29.	30.
n_2	n_2	n_2	n_2	n_3	n_3	n_3	n_3	n_3	n_3
n_5	n_5	n_6	n_6	n_5	n_5	n_4	n_4	n_4	n_4
$Y_{1,3}$	$X_{1,4}$	$X_{1,3}$	$X_{1,5}$	$X_{1,6}$	$Y_{1,6}$	$X_{1,2}$	$X_{1,5}$	$X_{1,5}$	$Y_{1,5}$
$X_{4,6}$	$Y_{3,6}$	$X_{4,5}$	$X_{3,4}$	$X_{2,4}$	$Y_{2,4}$	$X_{5,6}$	$X_{2,6}$	$Y_{2,6}$	$X_{2,6}$

31.	32.	33.	34.	35.	36.	37.	38.	39.	40.
n_3	n_3	n_3	n_4	n_4	n_4	n_4	n_5	n_5	$X_{1,2}$
n_4	n_6	n_6	n_5	n_5	n_6	n_6	n_6	n_6	$X_{3,5}$
$Y_{1,5}$	$Y_{1,2}$	$X_{1,4}$	$Y_{1,2}$	$X_{1,3}$	$y_{1,3}$	$Y_{1,5}$	$X_{1,2}$	$Y_{1,4}$	$Y_{4,6}$
$Y_{2,6}$	$Y_{4,5}$	$X_{2,5}$	$X_{3,6}$	$Y_{2,6}$	$y_{2,5}$	$X_{2,3}$	$Y_{3,4}$	$Y_{2,3}$	

41.	42.	43.	44.	45.	46.	47.	48.	49.	50.
$X_{1,2}$	$Y_{1,2}$	$Y_{1,2}$	$X_{1,2}$	$X_{1,2}$	$Y_{1,2}$	$X_{1,2}$	$Y_{1,2}$	$Y_{1,2}$	$X_{1,3}$
$Y_{3,5}$	$X_{3,5}$	$Y_{3,5}$	$X_{4,5}$	$Y_{4,5}$	$X_{4,5}$	$Y_{5,6}$	$X_{5,6}$	$Y_{5,6}$	$X_{2,5}$
$X_{4,6}$	$X_{4,6}$	$Y_{4,6}$	$Y_{3,6}$	$X_{3,6}$	$Y_{3,6}$	$X_{3,4}$	$X_{3,4}$	$Y_{3,4}$	$X_{4,6}$

51.	52.	53.	54.	55.	56.	57.	58.	59.	60.
$X_{1,3}$	$Y_{1,3}$	$X_{1,3}$	$Y_{1,3}$	$Y_{1,3}$	$X_{1,3}$	$X_{1,3}$	$Y_{1,3}$	$Y_{1,3}$	$X_{1,4}$
$Y_{2,5}$	$X_{2,5}$	$X_{2,6}$	$X_{2,6}$	$Y_{2,6}$	$X_{2,4}$	$Y_{2,4}$	$X_{2,4}$	$Y_{2,4}$	$X_{2,3}$
$Y_{4,6}$	$Y_{4,6}$	$Y_{4,5}$	$X_{4,5}$	$Y_{4,5}$	$Y_{5,6}$	$X_{5,6}$	$X_{5,6}$	$Y_{5,6}$	$X_{5,6}$

61.	62.	63.	64.	65.	66.	67.	68.	69.	70.
$X_{1,4}$	$Y_{1,4}$	$X_{1,4}$	$Y_{1,4}$	$Y_{1,4}$	$X_{1,4}$	$X_{1,4}$	$Y_{1,4}$	$Y_{1,4}$	$X_{1,5}$
$X_{2,3}$	$Y_{2,3}$	$Y_{2,5}$	$X_{2,5}$	$Y_{2,5}$	$X_{2,6}$	$Y_{2,6}$	$X_{2,6}$	$Y_{2,6}$	$X_{2,4}$
$Y_{5,6}$	$X_{5,6}$	$X_{3,6}$	$X_{3,6}$	$Y_{3,6}$	$X_{3,5}$	$Y_{3,5}$	$X_{3,5}$	$Y_{3,5}$	$X_{3,6}$

71.	72.	73.	74.	75.	76.
$X_{1,5}$	$Y_{1,5}$	$Y_{1,5}$	$X_{1,5}$	$X_{1,5}$	$Y_{1,5}$
$X_{2,4}$	$Y_{2,4}$	$Y_{2,4}$	$Y_{2,3}$	$Y_{2,3}$	$X_{2,3}$
$Y_{3,6}$	$X_{3,6}$	$Y_{3,6}$	$Y_{4,6}$	$X_{4,6}$	$X_{4,6}$

7.4 Discussion

We have introduced a tomographically complete measurement scheme that can be directly implemented in fermionic simulators using only two-mode interactions. As examples, we analyse the complexity of performing partial tomography in fermionic systems and provide minimal measurement protocols to obtain the single- and two-body reduced

states for any given lattice. Differently to previous works, which have focused on digital quantum simulation of fermionic systems, we emphasize that our methods rely solely on the native operations of a fermionic simulator, thus avoiding the overheads of fermion to qubit encodings and measurement circuits.

From a technical standpoint, we note that edge graph colouring was also recently used as a technique to tame experimental requirements in the simulation of fermions using digital quantum computers: When the two-point correlators are encoded in the Majorana basis, edge colouring can be used to identify the underlying commutation structure [Ara+22]. Similarly, variants of edge clique covering, as well as other suboptimal solutions to similar problems have recently appeared in several other contexts, which is unsurprising when seen as part of a larger trend of using graph-theoretic quantities to investigate the compatibility of quantum structures [VYI20; GHG23; BW24; VK24; Han+24].

Apart from the obvious questions regarding the optimal scaling of the number of measurement settings, many other problems are left open and worthy of further investigation, including:

Learning n-body observables Our analysis was focused on single- and two-body observables, but it can in principle be extended to the case of n-body observables. In the formulation we have presented, this would lead to more complicated compatibility constraints between the required correlators.

Many-body gates The measurement basis we have proposed is constructed in analogy to a spin-1/2 particle [Eq. (7.3)]. This could in principle be extended to operations involving more than two modes. Would it bring any advantage in practice, or lead to interesting insights?

Sample cost analysis We have used the number of measurement settings as figure of merit. Together with scalar concentration bounds, a union bound can be applied to obtain a sample cost that scales with the number of settings. However, it is often the case that a more detailed investigation of the statistics involved leads to significant improvements in practice, one recent example being [VK24]. Could a similar improvement be found here?

Part III

Tools

Ket.jl

8

This chapter is based on *M. Araújo, P. Brown, S. Designolle, C. de Gois, et al.* Ket.jl: a Julia toolbox for quantum information, nonlocality, and entanglement. *https://zenodo.org/records/14674643. 2025* [Ara+25]. The software package is in continuous development.

> *In Calabozo they met an old man who had never left his village. In spite of this he possessed a laboratory: glass vessels and bottles, metal equipment for measuring earthquakes, humidity, and magnetism. Plus a primitive machine with pointers that moved if anyone in the vicinity told a lie or said something stupid.*
>
> — **Daniel Kehlmann**
> Measuring the World

`Ket.jl` is an attempt to teach quantum information theory to a primitive machine and hope that its pointers do not move. Implemented in Julia — a high-level but high-performance programming language specifically designed for numerical and scientific computing [Bez+17] — Ket is a toolbox for quantum information, nonlocality, and entanglement. Despite its strengths, Julia still lacked dedicated libraries for quantum information. Early attempts either had a different scope or are unmaintained (see [Ara+25] for a list of related libraries). Users were then required to implement their own versions of even basic operations such as subsystem permutation and partial trace. Compared to established libraries in other languages — the most notable examples probably being QETLAB [Joh16] for MATLAB — this gap justified the development of `Ket.jl`. Leveraging Julia's powerful typing system, `Ket.jl`'s functions are designed for generic types, allowing the user to choose for example Int64 or Float64 for efficiency, or larger precision types when needed. One of the strongest points in Julia's ecosystem is its advanced support for formulating mathematical optimization models, due to JuMP [Lub+23]. Wherever possible, `Ket.jl`'s functions can be called on optimization variables, thus be

directly used inside JuMP models. We follow modern software engineering standards, with continuous integration and extensive automated testing covering a large portion of the codebase.

After a brief description of the Julia language (Sec. 8.1), this chapter describes some particularly important algorithms implemented by `Ket.jl` and not easily found elsewhere. Section 8.2 discusses the symmetric extensions hierarchy, a powerful test for entanglement that can straightforwardly be added as a constraint to any user-defined optimization model. The inner symmetric extensions (Sec. 8.3) and genuine multipartite entanglement certification (Sec. 8.4) build on these entanglement tests. Other highlights of `Ket.jl` are more briefly described in Sec. 8.5, and we conclude with a discussion of our prospects in Sec. 8.6. Example applications are provided in the documentation [Ara+25].

8.1 Julia

Julia is a general-purpose programming language specifically designed for numerical and scientific computing. Despite being a relatively recent additional to the computing landscape, it has already gained significant traction both in academia, and in industry. Notable applications of Julia include data analysis at CERN, cosmic microwave background analysis at NASA, satellite design in the Brazilian space research institute, not to mention uses in the pharmaceuticals industry and banking (see recordings of previous JuliaCon conferences, for example [Jul24]).

The two main languages filling a similar space are MATLAB and Python. MATLAB was specifically created for scientific and numerical computing. Its performance is a strong aspect, but comes paired with several drawbacks: it is expensive, clunky, offers limited interoperability with other tools, and it lacks conveniences found in more modern programming languages. Nevertheless, it is a popular choice in the quantum information community, mainly due to the user-created optimization modelling languages YALMIP [Lof04] and CVX [GB14; GB08], and to QETLAB [Joh16].

Python, unlike Julia and MATLAB, was not designed for numerical computation, a domain where it is extremely inefficient. Its rise in scientific programming was largely due to its extensive ecosystem, spearheaded by the development of NumPy [Har+20] and SciPy [Vir+20]. Even though NumPy brought efficient numerical computation to Python, the

user is required to adopt specific practices, such as relying on vectorized operations and avoiding native control structures such as for loops. Other motivations for its use in quantum information are the optimization modelling languages such as CVXPY [DB16; Agr+18] and PICOS [SS22], and some libraries for quantum information, including toqito [Rus20] and Inflation [BWP23]. From a design perspective, Python's object-oriented programming (OOP) paradigm is the idiomatic way of defining user types. This is not always ideal for scientific computing, where an unobtrusive type system combined with polymorphism tends to be more effective and extensible, while requiring far less boilerplate code.

Julia combines the good and ditch the worst aspects from Python and MATLAB:

Convenience Multiple workflows are supported, including an interactive REPL, Jupyter notebooks, and script execution. High-level features such as automatic garbage collection and type inference enable rapid coding. Its native support for calling C and Fortran functions without requiring any glue code enables easy interoperability with a large amount of existing code. The built-in package manager trivializes software distribution for both developers and users.

Speed Julia achieves execution speeds comparable to C and Fortran while maintaining a high-level syntax. It does so even for non-vectorize code. This is due to JIT compilation, type stability and native support for lightweight multithreading.

Expressiveness A powerful type system, first-class functions, multiple dispatch, and metaprogramming enable concise yet expressive code. These capabilities make it particularly well-suited for scientific computing, where rapid prototyping and refactoring are often necessary, without taking a toll in performance.

JuMP Mathematical optimization is one of the strongest points in Julia's ecosystem, largely due to JuMP. Going far beyond CVX and PICOS, JuMP provides a high-level interface for formulating and solving not only convex, but also mixed-integer, and nonlinear optimization problems. Its support for a much broader range of problem types and superior performance made it one of the leading optimization frameworks.

These are core reasons behind the rapid growth of Julia's user base, which we hope to extend to the quantum information community with the conveniences provided by `Ket.jl`.

8.2 Symmetric extensions hierarchy

Symmetric extensions are a common pattern in quantum information, appearing frequently in optimization models. Here we discuss implementation details, and show how the function `_dps_constraints!` can be integrated in any user-defined optimization model. In fact, `Ket.jl` uses this as a building block for computing the entanglement entropy and robustness, determining the Schmidt number of mixed states, and certifying genuinely multipartite entanglement (Sec. 8.4), among others.

We recall from Sec. 1.1.2.1 that if $\rho \in \mathcal{D}(\mathcal{H}_A \otimes \mathcal{H}_B)$ is separable (that is, $\rho = \sum_i p_i \rho_i^A \otimes \rho_i^B$), then for any natural N we can define a Hermitian, positive-semidefinite operator

$$\Xi = \sum_i p_i \rho_i^A \otimes (\rho_i^B)^{\otimes N} \tag{8.1}$$

from which ρ is recovered by tracing out any set of $N-1$ subsystems in $\mathcal{H}_B^{\otimes N}$. Therefore, if for a given state ρ, there is no $\Xi \in \mathcal{L}(\mathcal{H}_A \otimes \mathcal{H}_B^{\otimes N})$ such that

$$\rho = \mathrm{tr}_{B_i \,:\, i \in [N]\setminus m}(\Xi), \quad \forall m, \tag{8.2}$$

then ρ must be entangled. This is a powerful entanglement criterion called symmetric extensions hierarchy, or DPS hierarchy [DPS04; DPS02].

Any conic solver that supports the PSD cone (Sec. 1.3.1) can in principle handle this test. In practice, the dimension of the optimization variable Ξ grows exponentially (it is a Hermitian matrix of size $d_A d_B^N$), as well as the number of constraints. This makes many interesting cases out of reach, in particular larger state dimensions where PPT entanglement appears. However, it is possible to reduce this growth to a polynomial in N by observing part of Ξ is supported on the symmetric subspace.

The symmetric subspace of $\mathcal{H}^{\otimes N}$ is the set of vectors which are invariant under any permutation of the tensor factors:

$$\mathrm{Sym}(\mathcal{H}_B^{\otimes N}) = \{\, v \in \mathcal{H}^{\otimes N} \mid v = P_s v \text{ for every } s \in S_N \,\}, \tag{8.3}$$

where S_N is the set of all permutations of N indexes and P_s the permutation operator associated to the permutation s. For example, the vectors $|00\rangle$, $|11\rangle$, and $|01\rangle + |10\rangle$ are

in $\mathrm{Sym}(\mathcal{H}_2^{\otimes 2})$, but $|01\rangle - |10\rangle$ is not, since it gets a negative sign under a swap operation. A projection onto $\mathrm{Sym}(\mathcal{H}_B^{\otimes N})$ can be obtained by averaging over the P_s,

$$P = \frac{1}{N!} \sum_{s \in S_N} P_s, \tag{8.4}$$

and combinatorial arguments can be used to show the dimension of its image is $\binom{d+N-1}{N}$ [Wat18, Sec. 7.1], so the example above is in fact a basis for $\mathrm{Sym}(\mathcal{H}_2^{\otimes 2})$. Then, by writing Ξ [Eq. (8.1)] in terms of rank-1 projectors and verifying permutational symmetry in the copies of $\mathcal{H}_B$, we can conclude that the $\mathcal{H}_B^{\otimes N}$ subsystems are supported in the symmetric subspace [DPS04], [1] so $\mathcal{L}\left(\mathcal{H}_A \otimes \mathrm{Sym}(\mathcal{H}_B^{\otimes N})\right)$ is of dimension

$$d_{\mathrm{DPS}} = d_A \cdot \binom{d_B + N - 1}{N}. \tag{8.5}$$

However, if in an optimization model we instantiate a matrix variable Ξ' in the symmetric subspace, then the reduction constraint [Eq. (8.2)] must be rewritten in terms of Ξ' (or one must find a representation of the partial trace directly on the symmetric subspace). For that, we need an isometry that embeds $\mathrm{Sym}(\mathcal{H}_B^{\otimes N})$ into the full, d^N-dimensional vector space. The projection P in Eq. (8.4) can be used to obtain an orthonormal basis for $\mathrm{Sym}(\mathcal{H}_B^{\otimes N})$, for instance through a singular value decomposition: Writing $P = U\Sigma V$, one can verify from the number of non-zero values in Σ that the rank of P is d_Ξ, and extract the associated column vectors from U. These vectors span the image of P, namely, the symmetric subspace. Letting them be the columns of a matrix S, we may then substitute any occurrence of Ξ for $S^\dagger \Xi' S$. Altogether, this leads to the symmetrized DPS hierarchy:

$$\text{given } \rho,\ N \tag{8.6a}$$
$$\text{find } \Xi' \in \mathrm{Pos}\left(d_{\mathrm{DPS}}\right) \tag{8.6b}$$
$$\text{s.t.}$$
$$\mathrm{tr}_{B_i\,:\,i \neq 1}\left(S^\dagger\, \Xi' S\right) = \rho, \tag{8.6c}$$

Unfeasibility for any hierarchy level N certifies ρ is entangled. This is often cast as an optimization problem by mixing ρ with some noise $\mathcal{N}$ and asking for the largest $0 \leq \eta \leq 1$

[1] Notice that not every permutationally symmetric density operator is supported in the symmetric subspace. An example is $\rho = {}^1\!/_2\left(|01\rangle\langle 01| + |10\rangle\langle 10|\right)$, which is evidently invariant under permutation, but is not entirely supported in the symmetric subspace because $|01\rangle, |10\rangle \notin \mathrm{Sym}(\mathcal{H}_2^{\otimes 2})$.

such that $\eta\rho + (1-\eta)\mathcal{N}(\rho)$ has a symmetric extension. Any state with η above the obtained value is then entangled. Positive partial transpose constraints are commonly added to the extension and considerably speed up convergence [NOP09b].

`Ket.jl` implements the DPS hierarchy as a side-effecting function that adds the constraints to any JuMP model passed as the input. It has the signature

```
function _dps_constraints!(
    model::JuMP.GenericModel{T},
    ρ::AbstractMatrix,
    dims::AbstractVector{<:Integer},
    n::Integer;
    ppt::Bool = true,
    is_complex::Bool = true,
    isometry::AbstractMatrix = I(size(ρ, 1))
) where {T}
```

Listing 8.1: Signature of DPS constraints function

Calling `_dps_constraints!` will then modify `model` in-place, such that ρ becomes constrained as in Eq. (8.6). Through `ppt`, the user can choose whether to include PPT constraints (over every possible bipartition) or not, and `isometry` takes an isometry is applied to the extensions after the partial trace — this is used, for example, in the function `schmidt_number` to represent the embedding defined in Sec. 1.1.2.1. Lastly, `is_complex` chooses whether Ξ' is a complex Hermitian or a real symmetric variable, which can be useful for the reasons discussed next.

When $\rho \in \mathbb{R}^{d_A d_B \times d_A d_B}$, another reduction of the program size can be found. This is because if ρ has a symmetric extension [Eq. (8.1)] to order N, where in general $\Xi \in \mathbb{C}^{d_A d_B^N \times d_A d_B^N}$, then one can use that $\rho + \rho^* = 2\rho$ to write

$$2\rho = \mathrm{tr}_{B_i\,:\,i\neq 1}\left(\sum_i p_i \rho_i^A \otimes (\rho_i^B)^{\otimes N}\right) + \mathrm{tr}_{B_i\,:\,i\neq 1}\left(\sum_i p_i \rho_i^A \otimes (\rho_i^B)^{\otimes N}\right)^* \tag{8.7a}$$

$$= \mathrm{tr}_{B_i\,:\,i\neq 1}\underbrace{\left(\sum_i p_i \left[\rho_i^A + (\rho_i^A)^*\right] \otimes \left[(\rho_i^B)^{\otimes N} + \left((\rho_i^B)^*\right)^{\otimes N}\right]\right)}_{\Xi'}. \tag{8.7b}$$

Therefore, there is also a $\Xi' \in \mathbb{R}^{d_A d_B^N \times d_A d_B^N}$ that suits ρ. In such cases, instantiating complex Hermitian variables for the extensions is wasteful: most convex programming solvers are implemented on real variables, and the modelling language converts complex variables into a larger set of real variables and constraints. Our implementation in `_dps_constraints!` lets the user choose whether to instantiate complex Hermitian or symmetric real variables, depending on the input ρ. Some functions that rely on it, such as `entanglement_robustness`, do this choice autonomously, which can largely reduce the solution time.

8.3 Inner symmetric extensions hierarchy

Symmetric extensions are outer approximations to SEP which in general only converge for $N \to \infty$, thus in practice they can only *exclude* a state from SEP. Instead, for certifying separability, an inner approximation is needed. Intuition says it should be possible to get one, because if a state ρ has symmetric extensions of order N (i.e. $\rho \in \mathrm{DPS}_N$), then there must be some perturbation that takes it inside the set of separable states (SEP). Moreover, we expect this perturbation to become smaller with increasing N, since DPS_N progressively close in on SEP. Fortunately, Navascués et al. [NOP09a; NOP09b] have already found a closed-form expression for such a perturbation: Assuming $\rho \in \mathrm{DPS}_N$, then

$$\frac{N}{N+d_B}\rho + \frac{d_B}{N+d_B}\,\mathrm{tr}_B(\rho) \otimes \frac{\mathbf{1}_B}{d_B} \in \mathrm{SEP}, \tag{8.8}$$

with d_B being the dimension of $\mathcal{H}_B$. Notice that DPS_N is the set of symmetric extensions without PPT constraints. When enforcing a PPT on the bipartition $AB^{\lceil N/2 \rceil} \mid B^{\lfloor N/2 \rfloor}$ of the extension operator, one can use [2]

$$(1-\epsilon_N)\,\rho + \epsilon_N\,\mathrm{tr}_B(\rho) \otimes \frac{\mathbf{1}_B}{d_B} \in \mathrm{SEP}, \tag{8.9}$$

[2] The superscripts in the bipartition are denoting the number of subsystems in each partition, not the indexing of the subsystems.

where to compute ϵ_N we must consider the zeros of the Jacobi polynomials $P_n^{(\alpha,\beta)}$ and choose [3]

$$\epsilon_N = \frac{d_B}{2(d_B-1)} \min\left\{1-x \mid P^{(d_B-2,\, N \bmod 2)}_{\lfloor N/2 \rfloor+1}(x) = 0\right\}. \tag{8.10}$$

Either case leads to a hierarchy in N which converges to the set of separable states from the inside, but differently from the DPS hierarchy, the approximation levels are not necessarily nested into one another [NOP09a]. A simple rewriting of the DPS criteria then leads to a sufficient criterion to certify separability. For instance, with additional PPT constraints we would write

$$\begin{aligned}
&\text{given } \rho,\ N && \text{(8.11a)}\\
&\text{find } \Xi \in \mathrm{Pos}\left(d_{\mathrm{DPS}}\right) && \text{(8.11b)}\\
&\text{s.t.}\\
&\quad \Xi = S^\dagger\, \Xi' S, && \text{(8.11c)}\\
&\quad \mathrm{tr}_{B_i\,:\,i\neq 1}\left(\Xi\right) = \sigma, && \text{(8.11d)}\\
&\quad \epsilon_N \sigma + (1-\epsilon_N)\,\mathrm{tr}_B(\sigma) \otimes \frac{\mathbf{1}_B}{d} = \rho && \text{(8.11e)}\\
&\quad \left(\mathbf{1}_{AB_1\ldots B_{\lceil N/2 \rceil}} \otimes T\right) \Xi \succcurlyeq 0. && \text{(8.11f)}
\end{aligned}$$

where T is the transposition map. Notice that if Ξ is PPT across every bipartition (as it would be enforced by `_dps_constraints!` with `ppt=true`), then Eq. (8.11f) is naturally also satisfied. The inner DPS hierarchy is implemented in `Ket.jl` as the set of constraints `_inner_dps_constraints!`, which has a similar signature to `_dps_constraints!`, which enables for example the computation of upper bounds on the entanglement robustness to be computed with `entanglement_robustness`.

8.4 Genuine multipartite entanglement

Entanglement is even more nuanced in multipartite systems, as we must consider the separability structure between the possible partitions of the system (Sec. 1.1.2.1). A central concept is that of genuine multipartite entangled states, which are those that

[3] There are many ways to compute zeros of polynomials, for example with root finding algorithms. The Jacobi polynomials, however, have a three-term recurrence relation, and therefore their zeros can be computed by finding the eigenvalues of a particular symmetric tridiagonal matrix [GST07, Chap. 5.3], and that is the method we have used.

cannot be described by a mixture of biseparable states. Notice that if ρ is an n-partite state with $A_1A_2\dots A_n$, which is biseparable with respect to $A_1 \mid A_2\dots A_n$, it might still be highly entangled in the $A_2\dots A_n$ subsystems (in particular, it can even be GME therein). That is why GME is though of as the strongest form of multipartite entanglement.

Recall from Sec. 1.1.2.1 that for an n-partite state $\rho \in \mathcal{D}\left(\bigotimes_{i=1}^{N} \mathcal{H}_i\right)$, we can define $2^{n-1}-1$ bipartitions $A_i|B_i$ of its subsystems, where $A_i \in \{1,\dots,N\}$, and $B_i = \{1,\dots,N\} \setminus A_i$. A biseparable is then of the form $\rho = \sum_i p_i \rho_i^{A_i|B_i}$, where the states $\rho_i^{A_i|B_i}$ are separable with respect to $A_i \mid B_i$, and $\{p_i\}$ is a probability distribution. Genuine multipartite entangled states are those that cannot be written in this form.

Absorbing the probabilities p_i into the states $\rho_i^{A_i|B_i}$ (which will make them sub-normalized) lets us represent set of biseparable states in terms or representations of separability:

$$\text{given } \rho, n \tag{8.12a}$$

$$\text{find } \{\sigma_i \in \mathrm{Pos}(d_\rho)\}_{i=1}^{2^{n-1}-1} \tag{8.12b}$$

$$\text{s.t.}$$

$$\sigma_i \in \mathtt{SEP}(P_i \mid \bar{P}_i), \quad \forall i \tag{8.12c}$$

$$\sum_i \sigma_i = \rho. \tag{8.12d}$$

Notice Eq. (8.12d) guarantees that $\mathrm{tr}\left(\sum_i \sigma_i\right) = 1$, and for any feasible solution one can write the states $\tilde{\sigma}_i = \sigma_i / \mathrm{tr}(\sigma_i)$ and probabilities $p_i = \mathrm{tr}(\sigma_i)$ such that $\rho = \sum_i p_i \tilde{\sigma}_i$.

Turning this into a proper convex program requires substituting `SEP` for some finite-dimensional approximation of the separable states. The simplest choice is introducing PPT conditions, which will lead to a superset of the separable states whose elements are sometimes called "PPT mixtures" [JMG11]. Differently from the bipartite case, where PPT is equal to `SEP` for 2×2 and 2×3 systems, the PPT mixtures are in principle only a superset of the biseparable states, even in the simplest case of $2\times2\times2$ systems, as this would already require testing for separability of 2×4 systems in the bipartitions. Nevertheless, the PPT mixtures have the advantage of being computationally cheap, and a standard computer can test states with six or seven qubits without major efforts. A `ppt_mixture` function is available in `Ket.jl`. Another option is to provide a state and a list of observables, and ask whether there is some witness based only on them that detects GME of the given state. Technically, the function is based on `_decomposable_witness_constraints!`, which can be used to add the constraint of being a PPT mixture to any given optimization variable in a user-defined model (see [JMG11] for details on decomposable witnesses).

A tighter approximation can be obtained by using DPS constraints instead of PPT, that is,

$$\text{given } \rho, n, N \tag{8.13a}$$

$$\text{find } \{\,\Xi_i\,\}_{i=1}^{2^{n-1}-1} \tag{8.13b}$$

$$\text{s.t.}$$

$$\sum_i \mathrm{tr}_{(\bar{P}_i)_\ell\,:\,\ell\neq 1}\left(S_i^\dagger\,\Xi_i' S_i\right) = \rho, \tag{8.13c}$$

$$\tag{8.13d}$$

where the extension operators Ξ_i are of dimension $d_{P_i}\cdot\binom{d_{\bar{P}_i}+N-1}{N}$ if the extensions are done on the $\bar{P}_i$ side of the bipartition, and S_i are the corresponding symmetric subspace isometries (cf. Sec. 8.2). Additional PPT constraints can be included on the extensions to tighten the approximations, cf. Eq. (8.11f). It is also interesting to implement it by extending on the smallest side of the bipartitions, so that the dimension of the extensions is minimized. Nevertheless, this is in general only expected to be feasible for small N and n. An implementation is provided in the function `gme`.

8.5 Other methods

In addition to the functions described in detail earlier, `Ket.jl` provides a variety of powerful tools, often implemented with specialized algorithms for efficiency. Some highlights are:

Nonlocality, steering and entanglement

- Compute the local and quantum bounds (NPA) of multipartite Bell inequalities.
- Determine the nonlocality robustness of a behaviour for different kinds of noise such as random, local and general noise.
- Efficiently convert between the Collins-Gisin, full probability, and correlator representations.
- Compute entanglement measures such as the relative entropy of entanglement, robustness to random, separable or general noise, and the Schmidt number of a state.

- Determine the incompatibility robustness of quantum measurements for different measures of incompatibility such as depolarizing, random, probabilistic, generalized or jointly measurable. This can be used to certify steering as well.

States and operators

- Generate mutually unbiased bases (MUBs) and symmetric informationally complete POVMs (SIC-POVMs) with arbitrary precision.
- Sample random objects such as states, unitaries, isometries, probabilities and POVMs.
- Construct well-known families of quantum states, such as Bell, GHZ, W, Dicke, Horodecki and the super-singlet states.

Multilinear Algebra

- Perform common tensor operations in quantum information such as partial trace, partial transpose and subsystem permutations, on arbitrary matrices, including optimization variables.

These features, combined with the optimized implementation of `Ket.jl`, make it a versatile library for quantum information processing, entanglement theory, and quantum foundations. The package is in continuous development, with many new features and improvements planned for the future.

8.6 Discussion

While `Ket.jl` already provides a powerful and efficient framework for quantum information processing, there is always room for improvement:

Code optimization Optimizations can always be made to further enhance performance. Thorough profiling of the code and benchmarking against other implementations is routinely done, but there is always room for improvement.

Increasing the scope The potential for extended functionality is vast, as features for quantum metrology, computation of quantum key distribution rates, and simulation and analysis of tomography experiments — to mention a few — are still missing and not available elsewhere. This would require specialists in each topic, therefore contributions from the community are being incentivized.

Symmetrization and handling of special cases Scalability remains a fundamental challenge in most techniques to analyse quantum systems, and this is also true for optimization problems. Exploiting symmetries and sparsity to reduce problem size is an important research direction. One example is the well-known use of the symmetric subspace for the implementation of the symmetric extensions hierarchy, as previously described. Similar things can often be done to other problems (e.g. [TRR19; Ngu+20]). Further exploration and integration of these ideas into `Ket.jl` could provide significant computational benefits.

Documentation and examples There is an ongoing need to expand and refine the documentation to make the library more accessible to new users. We also expect that with more adopters, further use cases will be contributed to the documentation in the form of new examples.

Conclusion

Quantum systems promise substantial advantages in processing and transmitting information, but these advantages depend on reliable ways to read out the results. Unlike classical information systems, quantum theory introduces unique challenges in this regard, which stem from the probabilistic nature of measurements, the no-cloning theorem, and the existence of incompatible observables. By investigating quantum state certification through the lenses of correlations, communication and tomography, this thesis contributes to a growing body of literature that confronts the challenge of developing methods to efficiently and reliably learn properties of quantum systems.

As the field of quantum information evolves, scalability stands as the defining obstacle. Opportunities for progress come from a growing range of disciplines. Techniques from data analysis and machine learning are being applied to reduce the computational cost of quantum state reconstruction [Tor+18]. High-dimensional probability tools such as compressed sensing and matrix concentration bounds continue to be used to provide rigorous statistical guarantees for experimental results [KR21]. At the frontier with theoretical computer science, certification protocols are reshaping concepts like verifiability and interactive proofs [Ji+21], and ideas borrowed from distributed computing are pushing quantum information into networked scenarios [Bal+24]. Meanwhile, rapid progress is also coming from synergies with convex algebraic geometry and non-commutative polynomial optimization, which not only have been used to improve quantum information protocols [BFF24], but also to uncover fundamental aspects of quantum theory [Ren+21].

Ultimately, the future of quantum information will be shaped by its ability to account for two parallel threads: the industrialization of quantum technologies, demanding robust, standardized verification for commercial devices, and the theoretical deepening of quantum information as a lens on physics, computation, and complexity. Bridging them requires tools that are both practical and rigorous, and that are becoming increasingly multidisciplinary. This convergence will continue to teach us not only what can be built, but also what we can fundamentally understand about the nature of quantum systems, their limits, and their role in physics and computation.

List of publications

Publications discussed or partially discussed in this thesis:

I C. de Gois, K. Hansenne, and O. Gühne. "Uncertainty relations from graph theory". In: *Physical Review A* 107.6 (2023), p. 062211

II C. de Gois, M. Plávala, R. Schwonnek, and O. Gühne. "Complete Hierarchy for High-Dimensional Steering Certification". In: *Physical Review Letters* 131 (2023), p. 010201

III C. Vieira, C. de Gois, L. Pollyceno, and R. Rabelo. "Interplays between classical and quantum entanglement-assisted communication scenarios". In: *New Journal of Physics* 25.11 (2023), p. 113004

IV C. de Gois and M. Kleinmann. "User-friendly confidence regions for quantum state tomography". In: *Physical Review A* 109 (2023), p. 062417

V K. Hansenne, R. Qu, L. T. Weinbrenner, C. de Gois, H. Wang, Y. Ming, Z. Yang, P. Horodecki, W. Gao, and O. Gühne. "Optimal overlapping tomography". In: *arXiv:2408.05730* (2024)

VI A. Ghoshal, C. de Gois, K. Hansenne, O. Gühne, and H.-C. Nguyen. "Partial tomography of fermionic simulators". In: *(In preparation)* (2025)

Software repositories:

VII M. Araújo, P. Brown, S. Designolle, C. de Gois, Y.-C. Liu, L. Porto, and M. T. Quintino. *Ket.jl: a Julia toolbox for quantum information, nonlocality, and entanglement*. `https://zenodo.org/records/14674643`. 2025

Other publications by the author not discussed in this thesis:

VIII C. Vieira, C. de Gois, S. Designolle, P. Lauand, and M. T. Quintino. "Outcome communication cannot explain quantum correlations". In: *(In preparation)* (2025)

IX C. de Gois and O. Gühne. “A translation of G. Rumer’s *“Zur Theorie der Spinvalenz”*”. In: *(In preparation)* (2025)

X C. de Gois, G. Moreno, R. Nery, S. Brito, R. Chaves, and R. Rabelo. “General Method for Classicality Certification in the Prepare and Measure Scenario”. In: *PRX Quantum* 2 (3 2021), p. 030311

XI G. Moreno, R. Nery, C. de Gois, R. Rabelo, and R. Chaves. “Semi-device-independent certification of entanglement in superdense coding”. In: *Physical Review A* 103 (2 2021), p. 022426

Bibliography

[Aar18] S. Aaronson. “Shadow Tomography of Quantum States”. In: *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2018. Los Angeles, CA, USA: Association for Computing Machinery, 2018, pp. 325–338 (cit. on pp. 8, 70, 91, 92).

[AR19] S. Aaronson and G. N. Rothblum. “Gentle measurement of quantum states and differential privacy”. In: *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2019. Phoenix, AZ, USA: Association for Computing Machinery, 2019, pp. 322–333 (cit. on p. 92).

[Abb+16] A. A. Abbott, P.-L. Alzieu, M. J. Hall, and C. Branciard. “Tight state-independent uncertainty relations for qubits”. In: *Mathematics* 4.1 (2016), p. 8 (cit. on p. 31).

[AKG19] A. Acharya, T. Kypraios, and M. Guţă. “A comparative study of estimation methods in quantum tomography”. In: *Journal of Physics A: Mathematical and Theoretical* 52.23 (2019), p. 234001 (cit. on pp. 7, 73, 77).

[ASS21] A. Acharya, S. Saha, and A. M. Sengupta. “Shadow tomography based on informationally complete positive operator-valued measure”. In: *Physical Review A* 104 (2021), p. 052418 (cit. on p. 94).

[Agr+18] A. Agrawal, R. Verschueren, S. Diamond, and S. Boyd. “A rewriting system for convex optimization problems”. In: *Journal of Control and Decision* 5.1 (2018), pp. 42–60 (cit. on p. 127).

[Agu+18] E. A. Aguilar, J. J. Borkała, P. Mironowicz, and M. Pawłowski. “Connections between Mutually Unbiased Bases and Quantum Random Access Codes”. In: *Physical Review Letters* 121 (2018), p. 050501 (cit. on p. 53).

[AW02] R. Ahlswede and A. Winter. “Strong converse for identification via quantum channels”. In: *IEEE Transactions in Information Theory* 48.3 (2002), pp. 569–579 (cit. on p. 74).

[AKS23] J. O. de Almeida, M. Kleinmann, and G. Sentís. “Comparison of confidence regions for quantum state tomography”. In: *New Journal of Physics* 25.11 (2023), p. 113018 (cit. on pp. 71, 79–82).

[Amb+08] A. Ambainis, D. Leung, L. Mancinska, and M. Ozols. “Quantum Random Access Codes with Shared Randomness”. In: (2008). Preprint at arXiv:0810.2937. arXiv: 0810.2937 (cit. on p. 19).

[Amb+99] A. Ambainis, A. Nayak, A. Ta-Shma, and U. Vazirani. “Dense quantum coding and a lower bound for 1-way quantum automata”. In: *Proceedings of the thirty-first annual ACM symposium on Theory of computing - STOC ’99*. STOC ’99. Atlanta, Georgia, USA: ACM Press, 1999, pp. 376–383 (cit. on pp. 19, 55).

[AGK24] E. R. Anschuetz, D. Gamarnik, and B. T. Kiani. “Bounds on the ground state energy of quantum p-spin Hamiltonians”. In: *arXiv:2404.07231* (2024) (cit. on pp. 32, 37).

[Ant+24] M. Antesberger, M. T. Quintino, P. Walther, and L. A. Rozema. “Higher-Order Process Matrix Tomography of a Passively-Stable Quantum Switch”. In: *PRX Quantum* 5 (2024), p. 010325 (cit. on p. 87).

[Ara+22] B. G. M. Araújo, M. M. Taddei, D. Cavalcanti, and A. Acín. “Local quantum overlapping tomography”. In: *Physical Review A* 106 (2022), p. 062441 (cit. on pp. 90, 106, 112, 121).

[Ara+25] M. Araújo, P. Brown, S. Designolle, C. de Gois, Y.-C. Liu, et al. *Ket.jl: a Julia toolbox for quantum information, nonlocality, and entanglement*. `https://zenodo.org/records/14674643`. 2025 (cit. on pp. 125, 126, 139).

[Arg+19] J. Argüello-Luengo, A. González-Tudela, T. Shi, P. Zoller, and J. I. Cirac. “Analogue quantum chemistry simulation”. In: *Nature* 574.7777 (2019), pp. 215–218 (cit. on p. 105).

[AMP22] G. Aubrun, A. Müller-Hermes, and M. Plávala. “Monogamy of entanglement between cones”. In: *arXiv:2206.11805* (2022) (cit. on pp. 39, 51–54).

[BO21] C. Bădescu and R. O’Donnell. “Improved Quantum data analysis”. In: *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*. STOC 2021. Virtual, Italy: Association for Computing Machinery, 2021, pp. 1398–1411 (cit. on p. 92).

[Bal+24] A. Balliu, S. Brandt, X. Coiteux-Roy, F. d’Amore, M. Equi, et al. “Distributed Quantum Advantage for Local Problems”. In: *arXiv:2411.03240* (2024) (cit. on p. 137).

[BDH24] Z. Bao, P. van Dordrecht, and J. Helsen. “Tolerant testing of stabilizer states with a polynomial gap via a generalized uncertainty relation”. In: *arXiv:2410.21811* (2024) (cit. on pp. 32, 37).

[BE95] A. Barenco and A. K. Ekert. “Dense Coding Based on Quantum Entanglement”. In: *Journal of Modern Optics* 42.6 (1995), pp. 1253–1259 (cit. on p. 39).

[Bar09] S. Barnett. *Quantum Information*. Oxford University Press, 2009 (cit. on p. 2).

[BB14] C. H. Bennett and G. Brassard. “Quantum cryptography: Public key distribution and coin tossing”. In: *Theoretical Computer Science* 560 (2014). Theoretical Aspects of Quantum Cryptography – celebrating 30 years of BB84, pp. 7–11 (cit. on p. 55).

[Ben+93] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, et al. “Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels”. In: *Physical Review Letters* 70 (1993), pp. 1895–1899 (cit. on pp. 9, 59).

[BW92] C. H. Bennett and S. J. Wiesner. “Communication via one- and two-particle operators on Einstein-Podolsky-Rosen states”. In: *Physical Review Letters* 69 (1992), pp. 2881–2884 (cit. on pp. 9, 18, 39, 55, 59, 60).

[BW24] E. van den Berg and P. Wocjan. “Techniques for learning sparse Pauli-Lindblad noise models”. In: *Quantum* 8 (2024), p. 1556 (cit. on pp. 96, 102, 114, 121).

[Ber61] C. Berge. “Farbung von Graphen, deren samtliche bzw. deren ungerade Kreise starr sind”. In: *Wissenschaftliche Zeitschrift* (1961) (cit. on p. 36).

[Ber+22] M. Berta, F. Borderi, O. Fawzi, and V. B. Scholz. “Semidefinite programming hierarchies for constrained bilinear optimization”. en. In: *Mathematical Programming* 194.1-2 (2022), pp. 781–829 (cit. on p. 50).

[Bez+17] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. “Julia: A fresh approach to numerical computing”. In: *SIAM Review* 59.1 (2017), pp. 65–98 (cit. on pp. 24, 25, 125).

[Blu10] R. Blume-Kohout. “Optimal, reliable estimation of quantum states”. In: *New Journal of Physics* 12.4 (2010), p. 043034 (cit. on pp. 8, 69).

[Blu12] R. Blume-Kohout. “Robust error bars for quantum tomography”. In: *arXiv:1202.5270* (2012) (cit. on pp. 8, 69).

[BWP23] E.-C. Boghiu, E. Wolfe, and A. Pozas-Kerstjens. “Inflation: a Python library for classical and quantum causal compatibility”. In: *Quantum* 7 (2023), p. 996. arXiv: 2211.04483 (cit. on p. 127).

[Bol+24a] S. Bolusani, M. Besançon, K. Bestuzheva, A. Chmiela, J. Dionísio, et al. *The SCIP Optimization Suite 9.0*. Technical Report. Optimization Online, 2024 (cit. on p. 25).

[Bol+24b] S. Bolusani, M. Besançon, K. Bestuzheva, A. Chmiela, J. Dionísio, et al. *The SCIP Optimization Suite 9.0*. ZIB-Report 24-02-29. Zuse Institute Berlin, 2024 (cit. on p. 25).

[BM08] J. A. Bondy and U. S. R. Murty. *Graph theory*. Springer Publishing Company, Incorporated, 2008 (cit. on pp. 19, 21).

[BBO20] X. Bonet-Monroig, R. Babbush, and T. E. O’Brien. “Nearly Optimal Measurement Scheduling for Partial Tomography of Quantum States”. In: *Physical Review X* 10 (2020), p. 031064 (cit. on pp. 8, 69, 90, 91, 94, 101, 106, 112).

[Bou+18] F. Bouchard, N. H. Valencia, F. Brandt, R. Fickler, M. Huber, et al. “Measuring azimuthal and radial modes of photons”. In: *Optics Express* 26.24 (2018), pp. 31925–31941 (cit. on p. 54).

[BBP15] J. Bowles, N. Brunner, and M. Pawłowski. “Testing dimension and nonclassicality in communication networks”. en. In: *Physical Review A* 92.2 (2015), p. 022351 (cit. on p. 16).

[BQB14] J. Bowles, M. T. Quintino, and N. Brunner. “Certifying the Dimension of Classical and Quantum Systems in a Prepare-and-Measure Scenario with Independent Devices”. en. In: *Physical Review Letters* 112.14 (2014), p. 140407 (cit. on p. 16).

[BV04] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004 (cit. on p. 23).

[BV05] S. L. Braunstein and P. Van Loock. “Quantum information with continuous variables”. In: *Reviews of modern physics* 77.2 (2005), pp. 513–577 (cit. on p. 2).

[BK02] S. B. Bravyi and A. Y. Kitaev. “Fermionic quantum computation”. In: *Annals of Physics* 298.1 (2002), pp. 210–226 (cit. on p. 105).

[BFF24] P. Brown, H. Fawzi, and O. Fawzi. “Device-independent lower bounds on the conditional von Neumann entropy”. In: *Quantum* 8 (2024), p. 1445 (cit. on p. 137).

[Bru+04] Č. Brukner, M. Żukowski, J.-W. Pan, and A. Zeilinger. “Bell’s inequalities and quantum communication complexity”. In: *Physical review letters* 92.12 (2004), p. 127901 (cit. on p. 19).

[Bru+14] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner. “Bell nonlocality”. In: *Reviews of Modern Physics* 86 (2014), pp. 419–478 (cit. on pp. 16, 19).

[BNV13] N. Brunner, M. Navascués, and T. Vértesi. “Dimension Witnesses and Quantum State Discrimination”. en. In: *Physical Review Letters* 110.15 (2013) (cit. on p. 17).

[Bru02] D. Bruß. “Characterizing entanglement”. In: *Journal of Mathematical Physics* 43.9 (2002), pp. 4237–4251 (cit. on p. 12).

[BF04] H. Bruus and K. Flensburg. *Many-body quantum theory in condensed matter physics*. en. Oxford Graduate Texts. London, England: Oxford University Press, 2004 (cit. on p. 107).

[Buh+10] H. Buhrman, R. Cleve, S. Massar, and R. de Wolf. “Nonlocality and communication complexity”. en. In: *Reviews of Modern Physics* 82.1 (2010), pp. 665–698 (cit. on p. 19).

[CSW14] A. Cabello, S. Severini, and A. Winter. “Graph-theoretic approach to quantum correlations”. In: *Physical review letters* 112.4 (2014), p. 040401 (cit. on p. 20).

[Cai+16] Y. Cai, J.-D. Bancal, J. Romero, and V. Scarani. “A new device-independent dimension witness and its experimental implementation”. In: *Journal of Physics A: Mathematical and Theoretical* 49.30 (2016), p. 305301 (cit. on p. 54).

[CHT20] C. Carmeli, T. Heinosaari, and A. Toigo. “Quantum random access codes and incompatibility of measurements”. en. In: *EPL (Europhysics Letters)* 130.5 (2020), p. 50001 (cit. on p. 17).

[CS16] D. Cavalcanti and P. Skrzypczyk. “Quantum steering: a review with focus on semidefinite programming”. In: *Reports on Progress in Physics* 80.2 (2016), p. 024001 (cit. on pp. 14, 24).

[CFS02] C. M. Caves, C. A. Fuchs, and R. Schack. “Unknown quantum states: the quantum de Finetti representation”. In: *Journal of Mathematical Physics* 43.9 (2002), pp. 4537–4559. eprint: https://doi.org/10.1063/1.1494475 (cit. on pp. 11, 50).

[Cha+15] R. Chaves, R. Kueng, J. B. Brask, and D. Gross. “Unifying Framework for Relaxations of the Causal Assumptions in Bell’s Theorem”. In: *Phys. Rev. Lett.* 114 (2015), p. 140403 (cit. on p. 17).

[CBB15] R. Chaves, J. B. Brask, and N. Brunner. “Device-Independent Tests of Entropy”. In: *Physical Review Letters* 115 (2015), p. 110501 (cit. on p. 56).

[CMG15] R. Chaves, C. Majenz, and D. Gross. “Information–theoretic implications of quantum causal structures”. In: *Nature Communications* 6.1 (2015), pp. 1–8 (cit. on pp. 16, 56).

[Che+22] S. Chen, J. Cotler, H.-Y. Huang, and J. Li. “Exponential Separations Between Learning With and Without Quantum Memory”. In: *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*. 2022, pp. 574–585 (cit. on pp. 31, 92).

[CGY24] S. Chen, W. Gong, and Q. Ye. “Optimal tradeoffs for estimating Pauli observables”. In: *arXiv:2404.19105* (2024) (cit. on pp. 92, 93).

[CR12] M. Christandl and R. Renner. “Reliable Quantum State Tomography”. In: *Physical Review Letters* 109 (2012), p. 120403 (cit. on p. 71).

[Chu+06] M. Chudnovsky, N. Robertson, P. Seymour, and R. Thomas. “The strong perfect graph theorem”. In: *Annals of mathematics* (2006), pp. 51–229 (cit. on p. 36).

[CKV21] C. Coey, L. Kapelevich, and J. P. Vielma. *Solving natural conic formulations with Hypatia.jl*. 2021. arXiv: 2005.01136 [math.OC] (cit. on p. 24).

[CS20] A. Coladangelo and J. Stark. “An inherently infinite-dimensional quantum correlation”. In: *Nature Communications* 11.1 (2020), p. 3335 (cit. on p. 65).

[CD06] C. Colbourn and J. Dinitz. *Handbook of Combinatorial Designs*. Chapman and Hall/CRC, 2006 (cit. on pp. 21, 97, 110).

[Col04] C. J. Colbourn. “Combinatorial aspects of covering arrays”. In: *Le Matematiche* 59.1, 2 (2004), pp. 125–172 (cit. on p. 97).

[Col10] C. J. Colbourn. *CRC Handbook of Combinatorial Designs*. CRC Press, 2010 (cit. on p. 95).

[Col] C. Colbourn. *Covering Array Tables for t=2,3,4,5,6*. https://www.public.asu.edu/~ccolbou/src/tabby/catable.html (cit. on p. 96).

[CK08] R. Cole and Ł. Kowalik. “New linear-time algorithms for edge-coloring planar graphs”. In: *Algorithmica* 50.3 (2008), pp. 351–368 (cit. on p. 109).

[Col63] A. J. Coleman. “Structure of Fermion Density Matrices”. In: *Reviews of Modern Physics* 35 (1963), pp. 668–686 (cit. on pp. 89, 106, 112).

[CGM16] A. Conte, R. Grossi, and A. Marino. “Clique covering of large real-world networks”. In: *Proceedings of the 31st Annual ACM Symposium on Applied Computing*. SAC ’16. Pisa, Italy: Association for Computing Machinery, 2016, pp. 1134–1139 (cit. on p. 100).

[CW20] J. Cotler and F. Wilczek. “Quantum Overlapping Tomography”. In: *Physical Review Letters* 124 (2020), p. 100401 (cit. on pp. 8, 69, 89–91, 94, 101, 106).

[Cra+10] M. Cramer, M. B. Plenio, S. T. Flammia, R. Somma, D. Gross, et al. “Efficient quantum state tomography”. In: *Nature Communications* 1.149 (2010) (cit. on pp. 8, 69).

[DCT25] N. D’Alessandro, C. R. i. Carceller, and A. Tavakoli. “Semidefinite Relaxations for High-Dimensional Entanglement in the Steering Scenario”. In: *Physical Review Letters* 134 (2025), p. 090802 (cit. on p. 53).

[Des22] S. Designolle. “Robust genuine high-dimensional steering with many measurements”. In: *Physical Review A* 105 (2022), p. 032430 (cit. on pp. 40, 45, 48, 54).

[DFK19] S. Designolle, M. Farkas, and J. Kaniewski. “Incompatibility robustness of quantum measurements: a unified framework”. In: *New Journal of Physics* 21.11 (2019), p. 113053 (cit. on p. 19).

[Des+21] S. Designolle, V. Srivastav, R. Uola, N. H. Valencia, W. McCutcheon, et al. “Genuine High-Dimensional Quantum Steering”. In: *Physical Review Letters* 126 (2021), p. 200404 (cit. on pp. 16, 39, 40, 42, 45, 48).

[DB16] S. Diamond and S. Boyd. “CVXPY: A Python-embedded modeling language for convex optimization”. In: *Journal of Machine Learning Research* 17.83 (2016), pp. 1–5 (cit. on p. 127).

[Die82] D. Dieks. “Communication by EPR devices”. In: *Physics Letters A* 92.6 (1982), pp. 271–272 (cit. on p. 6).

[Die25] R. Diestel. *Graph Theory*. Springer Berlin Heidelberg, 2025 (cit. on p. 19).

[DPS02] A. C. Doherty, P. A. Parrilo, and F. M. Spedalieri. “Distinguishing Separable and Entangled States”. In: *Physical Review Letters* 88 (2002), p. 187904 (cit. on pp. 11, 44, 50, 128).

[DPS04] A. C. Doherty, P. A. Parrilo, and F. M. Spedalieri. “Complete family of separability criteria”. In: *Physical Review A* 69 (2004), p. 022308 (cit. on pp. 11, 44, 50, 128, 129).

[DC21] B. Doolittle and E. Chitambar. “Certifying the classical simulation cost of a quantum channel”. In: *Physical Review Research* 3 (2021), p. 043073 (cit. on p. 61).

[Dur+10] T. Durt, B.-G. Englert, I. Bengtsson, and K. Życzkowski. “On Mutually Unbiased Bases”. In: *International Journal of Quantum Information* 08.04 (2010), pp. 535–640 (cit. on pp. 46, 48, 73).

[Eck+19] S. Ecker, F. Bouchard, L. Bulla, F. Brandt, O. Kohout, et al. “Overcoming Noise in Entanglement Distribution”. In: *Physical Review X* 9 (2019), p. 041042 (cit. on p. 39).

[Ege+25] S. Egelhaaf, J. Pauwels, M. T. Quintino, and R. Uola. “Certifying measurement incompatibility in prepare-and-measure and Bell scenarios”. In: *Journal of Physics A: Mathematical and Theoretical* (2025) (cit. on pp. 19, 66).

[EPR35] A. Einstein, B. Podolsky, and N. Rosen. “Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?” In: *Physical Review* 47 (1935), pp. 777–780 (cit. on p. 14).

[ES14] C. Eltschka and J. Siewert. “Quantifying entanglement resources”. In: *Journal of Physics A: Mathematical and Theoretical* 47.42 (2014), p. 424005 (cit. on p. 12).

[FR16] P. Faist and R. Renner. “Practical and Reliable Error Bars in Quantum Tomography”. In: *Physical Review Letters* 117 (2016), p. 010404 (cit. on p. 71).

[Fei95] U. Feige. “Randomized graph products, chromatic numbers, and Lovasz j-function”. In: *Proceedings of the twenty-seventh annual ACM symposium on Theory of computing*. 1995, pp. 635–640 (cit. on p. 35).

[FW15] P. E. Frenkel and M. Weiner. “Classical Information Storage in an n-Level Quantum System”. In: *Communications in Mathematical Physics* 340.2 (2015), pp. 563–574 (cit. on p. 64).

[FW22] P. E. Frenkel and M. Weiner. “On entanglement assistance to a noiseless classical channel”. In: *Quantum* 6 (2022), p. 662 (cit. on pp. 19, 56, 57, 59, 66).

[Gal+10] R. Gallego, N. Brunner, C. Hadley, and A. Acín. “Device-Independent Tests of Classical and Quantum Dimensions”. en. In: *Physical Review Letters* 105.23 (2010), p. 230501 (cit. on pp. 16, 56).

[Gar+20] G. García-Pérez, M. A. C. Rossi, B. Sokolov, E.-M. Borrelli, and S. Maniscalco. “Pairwise tomography networks for many-body quantum systems”. In: *Physical Review Research* 2 (2020), p. 023393 (cit. on pp. 8, 69, 90, 91, 94, 101, 106).

[GJ83] M. R. Garey and D. S. Johnson. “Computers and Intractability. A Guide to the Theory of NP-Completeness”. In: *Journal of Symbolic Logic* 48.2 (1983), pp. 498–500 (cit. on p. 109).

[GM12] B. Gärtner and J. Matousek. *Approximation algorithms and semidefinite programming*. Springer Science & Business Media, 2012 (cit. on p. 23).

[GAN14] I. M. Georgescu, S. Ashhab, and F. Nori. “Quantum simulation”. In: *Reviews of Modern Physics* 86 (2014), pp. 153–185 (cit. on p. 105).

[Gho+25] A. Ghoshal, C. de Gois, K. Hansenne, O. Gühne, and H.-C. Nguyen. “Partial tomography of fermionic simulators”. In: *(In preparation)* (2025) (cit. on pp. 105, 109, 139).

[GST07] A. Gil, J. Segura, and N. M. Temme. *Numerical methods for special functions*. SIAM, 2007 (cit. on p. 132).

[GMR19] P. Giorda, L. Maccone, and A. Riccardi. “State-independent uncertainty relations from eigenvalue minimization”. In: *Physical Review A* 99 (2019), p. 052121 (cit. on pp. 31, 32).

[GLM06] V. Giovannetti, S. Lloyd, and L. Maccone. “Quantum Metrology”. In: *Physical Review Letters* 96 (2006), p. 010401 (cit. on p. 39).

[Gis89] N. Gisin. “Stochastic quantum dynamics and relativity”. In: *Helvetica Physica Acta* 62.4 (1989), pp. 363–371 (cit. on p. 5).

[Gis+02] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden. “Quantum cryptography”. In: *Reviews of modern physics* 74.1 (2002), p. 145 (cit. on p. 9).

[Goi21] C. de Gois. “Classicality and dense coding in the prepare and measure scenario”. Available at `https://hdl.handle.net/20.500.12733/1252`. M.Sc. Thesis. Universidade Estadual de Campinas, 2021 (cit. on p. 19).

[Goi22] C. de Gois. *Code to accompany “Interplays between classical and quantum entanglement-assisted communication scenarios”*. 2022 (cit. on pp. 61, 64).

[GG25] C. de Gois and O. Gühne. “A translation of G. Rumer’s *“Zur Theorie der Spinvalenz”*”. In: *(In preparation)* (2025) (cit. on p. 140).

[GHG23] C. de Gois, K. Hansenne, and O. Gühne. “Uncertainty relations from graph theory”. In: *Physical Review A* 107.6 (2023), p. 062211 (cit. on pp. 31, 34, 36, 121, 139).

[GK23] C. de Gois and M. Kleinmann. “User-friendly confidence regions for quantum state tomography”. In: *Physical Review A* 109 (2023), p. 062417 (cit. on pp. 69, 75, 77, 78, 80, 82, 84, 86, 139).

[Goi+21] C. de Gois, G. Moreno, R. Nery, S. Brito, R. Chaves, et al. “General Method for Classicality Certification in the Prepare and Measure Scenario”. In: *PRX Quantum* 2 (2021), p. 030311 (cit. on pp. 19, 56, 140).

[Goi+23] C. de Gois, M. Plávala, R. Schwonnek, and O. Gühne. “Complete Hierarchy for High-Dimensional Steering Certification”. In: *Physical Review Letters* 131 (2023), p. 010201 (cit. on pp. 39, 73, 139).

[Gok+19] P. Gokhale, O. Angiuli, Y. Ding, K. Gui, T. Tomesh, et al. “Minimizing state preparations in variational quantum eigensolver by partitioning into commuting families”. In: *arXiv:1907.13623* (2019) (cit. on p. 114).

[Gon+23] D. González-Cuadra, D. Bluvstein, M. Kalinowski, R. Kaubruegger, N. Maskara, et al. “Fermionic quantum processing with programmable neutral atom arrays”. In: *Proceedings of the National Academy of Sciences* 120.35 (2023), e2304294120 (cit. on pp. 105, 108).

[GB14] M. Grant and S. Boyd. *CVX: Matlab Software for Disciplined Convex Programming, version 2.1*. `https://cvxr.com/cvx`. 2014 (cit. on p. 126).

[GB08] M. Grant and S. Boyd. “Graph implementations for nonsmooth convex programs”. In: *Recent Advances in Learning and Control*. Ed. by V. Blondel, S. Boyd, and H. Kimura. Lecture Notes in Control and Information Sciences. `http://stanford.edu/~boyd/graph_dcp.html`. Springer-Verlag Limited, 2008, pp. 95–110 (cit. on p. 126).

[GPS24] D. Grier, H. Pashayan, and L. Schaeffer. “Sample-optimal classical shadows for pure states”. en-GB. In: *Quantum* 8 (2024). Publisher: Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften, p. 1373 (cit. on p. 93).

[Gro11] D. Gross. “Recovering Low-Rank Matrices From Few Coefficients in Any Basis”. In: *IEEE Transactions in Information Theory* 57.3 (2011), pp. 1548–1566 (cit. on p. 74).

[Gro+10] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert. “Quantum State Tomography via Compressed Sensing”. In: *Physical Review Letters* 105 (2010), p. 150401 (cit. on pp. 8, 69, 74).

[GT09] O. Gühne and G. Tóth. “Entanglement detection”. In: *Physics Reports* 474.1 (2009), pp. 1–75 (cit. on pp. 10, 14, 85).

[Gup78] R. P. Gupta. “An edge-coloration theorem for bipartite graphs with applications”. In: *Discrete Mathematics* 23.3 (1978), pp. 229–233 (cit. on p. 109).

[Gur] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. `http://www.gurobi.com` (cit. on p. 25).

[Guţ+20] M. Guţă, J. Kahn, R. Kueng, and J. A. Tropp. “Fast state tomography with optimal error bounds”. In: *Journal of Physics A: Mathematical and Theoretical* 53.20 (2020), p. 204001 (cit. on pp. 71–74, 78–81, 88).

[Haa+16] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu. “Sample-Optimal Tomography of Quantum States”. In: *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*. STOC ’16. Cambridge, MA, USA: Association for Computing Machinery, 2016, pp. 913–925 (cit. on pp. 8, 79).

[Han24] K. Hansenne. *Exploring quantum information theory through graphs and networks*. 2024 (cit. on pp. 31, 36, 37).

[HG23] K. Hansenne and O. Gühne. “Covariance-Matrix-Based Criteria for Network Entanglement”. In: *Entropy* 25.9 (2023) (cit. on p. 14).

[Han+24] K. Hansenne, R. Qu, L. T. Weinbrenner, C. de Gois, H. Wang, et al. “Optimal overlapping tomography”. In: *arXiv:2408.05730* (2024) (cit. on pp. 69, 89, 101, 106, 114, 121, 139).

[Han+22] K. Hansenne, Z.-P. Xu, T. Kraft, and O. Gühne. “Symmetries in quantum networks lead to no-go theorems for entanglement distribution and to verification techniques”. In: *Nature communications* 13.1 (2022), p. 496 (cit. on pp. 14, 32).

[Har+20] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, et al. “Array programming with NumPy”. en. In: *Nature* 585.7825 (2020), pp. 357–362 (cit. on p. 126).

[HO22] M. B. Hastings and R. O’Donnell. “Optimizing strongly interacting fermionic Hamiltonians”. In: *Proceedings of the 54th annual ACM SIGACT symposium on theory of computing*. 2022, pp. 776–789 (cit. on pp. 32, 37).

[Hei+06] M. Hein, W. Dür, J. Eisert, R. Raussendorf, M. Van den Nest, et al. “Entanglement in graph states and its applications”. In: *Quantum computers, algorithms and chaos*. IOS Press, 2006, pp. 115–218 (cit. on p. 20).

[Hol73] A. S. Holevo. “Bounds for the quantity of information transmitted by a quantum communication channel”. In: *Problemy Peredachi Informatsii* 9.3 (1973), pp. 3–11 (cit. on pp. 19, 60).

[Hol81] I. Holyer. “The NP-Completeness of Some Edge-Partition Problems”. In: *SIAM Journal on Computing* 10.4 (1981), pp. 713–717. eprint: `https://doi.org/10.1137/0210054` (cit. on p. 109).

[HHH96] M. Horodecki, P. Horodecki, and R. Horodecki. “Separability of mixed states: necessary and sufficient conditions”. In: *Physics Letters A* 223.1 (1996), pp. 1–8 (cit. on p. 10).

[Hor+09] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki. “Quantum entanglement”. In: *Reviews of Modern Physics* 81 (2009), pp. 865–942 (cit. on pp. 10, 14, 85).

[Hou+06] M. Hou, P. Berman, L. Zhang, and W. Miller. “Controlling size when aligning multiple genomic sequences with duplications”. In: *Algorithms in Bioinformatics: 6th International Workshop, WABI 2006, Zurich, Switzerland, September 11-13, 2006. Proceedings 6*. Springer. 2006, pp. 138–149 (cit. on p. 95).

[HKP21] H.-Y. Huang, R. Kueng, and J. Preskill. “Efficient Estimation of Pauli Observables by Derandomization”. en. In: *Physical Review Letters* 127.3 (2021), p. 030503 (cit. on p. 94).

[HKP20] H.-Y. Huang, R. Kueng, and J. Preskill. “Predicting many properties of a quantum system from very few measurements”. en. In: *Nature Physics* 16.10 (2020), pp. 1050–1057 (cit. on pp. 8, 70, 92–94).

[HJW93] L. P. Hughston, R. Jozsa, and W. K. Wootters. “A complete classification of quantum ensembles having a given density matrix”. In: *Physics Letters A* 183.1 (1993), pp. 14–18 (cit. on p. 5).

[Hul+04] F. Hulpke, D. Bruß, M. Lewenstein, and A. Sanpera. “Simplifying Schmidt number witnesses via higher-dimensional embeddings”. In: *Quantum Information and Computation* 4 (2004), pp. 207–221 (cit. on pp. 13, 39, 43).

[Inn+23] L. Innocenti, S. Lorenzo, I. Palmisano, F. Albarelli, A. Ferraro, et al. “Shadow Tomography on General Measurement Frames”. en. In: *PRX Quantum* 4.4 (2023), p. 040328 (cit. on p. 94).

[IR21] M. Ioannou and D. Rosset. “Noncommutative polynomial optimization under symmetry”. In: *arXiv:2112.10803* (2021) (cit. on p. 53).

[Ioa+22] M. Ioannou, P. Sekatski, S. Designolle, B. D. Jones, R. Uola, et al. “Simulability of high-dimensional quantum measurements”. In: *arXiv:2202.12980* (2022) (cit. on p. 53).

[Ji+21] Z. Ji, A. Natarajan, T. Vidick, J. Wright, and H. Yuen. “MIP* = RE”. In: *Commun. ACM* 64.11 (2021), pp. 131–138 (cit. on p. 137).

[Jia+20] Z. Jiang, A. Kalev, W. Mruczkiewicz, and H. Neven. “Optimal fermion-to-qubit mapping via ternary trees with applications to reduced quantum states learning”. In: *Quantum* 4 (2020), p. 276 (cit. on p. 106).

[Joh16] N. Johnston. *QETLAB: A MATLAB toolbox for quantum entanglement, version 0.9*. `https://qetlab.com`. 2016 (cit. on pp. 125, 126).

[Jon+22] B. D. Jones, R. Uola, T. Cope, M. Ioannou, S. Designolle, et al. “Equivalence between simulability of high-dimensional measurements and high-dimensional steering”. In: *arXiv:2207.04080* (2022) (cit. on p. 53).

[JL05] N. S. Jones and N. Linden. “Parts of quantum states”. In: *Physical Review A* 71.1 (2005), p. 012324 (cit. on p. 89).

[JW28] P. Jordan and E. Wigner. “Über das Paulische Äquivalenzverbot”. In: *Zeitschrift für Physik* 47.9 (1928), pp. 631–651 (cit. on p. 105).

[Jul24] JuliaCon.org. *JuliaCon 2024*. `https://juliacon.org/2024/`. [Accessed on 29-01-2025]. 2024 (cit. on p. 126).

[JMG11] B. Jungnitsch, T. Moroder, and O. Gühne. “Taming Multiparticle Entanglement”. In: *Physical Review Letters* 106 (2011), p. 190502 (cit. on pp. 85, 86, 133).

[Kar72] R. M. Karp. “Reducibility among Combinatorial Problems”. In: *Complexity of Computer Computations: Proceedings of a symposium on the Complexity of Computer Computations, held March 20–22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, and sponsored by the Office of Naval Research, Mathematics Program, IBM World Trade Corporation, and the IBM Research Mathematical Sciences Department*. Ed. by R. E. Miller, J. W. Thatcher, and J. D. Bohlinger. Boston, MA: Springer US, 1972, pp. 85–103 (cit. on p. 20).

[Kin+25] R. King, D. Gosset, R. Kothari, and R. Babbush. “Triply efficient shadow tomography”. In: *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2025, pp. 914–946. arXiv: 2404.19211 (cit. on pp. 37, 92, 94, 106).

[KR21] M. Kliesch and I. Roth. “Theory of Quantum System Certification”. In: *PRX Quantum* 2 (2021), p. 010201 (cit. on p. 137).

[Knu94] D. E. Knuth. “The Sandwich Theorem”. In: *The Electronic Journal of Combinatorics* 1.1 (1994) (cit. on p. 34).

[Kok+20] J. I. Kokkala, K. Meagher, R. Naserasr, K. J. Nurmela, P. R. Östergård, et al. “On the structure of small strength-2 covering arrays”. In: *Journal of Combinatorial Designs* 28.1 (2020), pp. 5–24 (cit. on p. 101).

[Kra+21] T. Kraft, S. Designolle, C. Ritz, N. Brunner, O. Gühne, et al. “Quantum entanglement in the triangle network”. In: *Phys. Rev. A* 103 (2021), p. L060401 (cit. on p. 14).

[Kra+18] T. Kraft, C. Ritz, N. Brunner, M. Huber, and O. Gühne. “Characterizing Genuine Multilevel Entanglement”. In: *Physical Review Letters* 120 (2018), p. 060502 (cit. on p. 42).

[KRT17] R. Kueng, H. Rauhut, and U. Terstiege. “Low rank matrix recovery from rank one measurements”. In: *Applied and Computational Harmonic Analysis* 42.1 (2017), pp. 88–116 (cit. on p. 79).

[KKL+10] D. R. Kuhn, R. N. Kacker, Y. Lei, et al. “Practical combinatorial testing”. In: *NIST special Publication* 800.142 (2010), p. 142 (cit. on pp. 21, 95).

[Kur+11] P. Kurzyński, T. Paterek, R. Ramanathan, W. Laskowski, and D. Kaszlikowski. “Correlation complementarity yields Bell monogamy relations”. In: *Physical Review Letters* 106.18 (2011), p. 180402 (cit. on pp. 32, 35).

[LPW02] N. Linden, S. Popescu, and W. K. Wootters. “Almost Every Pure State of Three Qubits Is Completely Determined by Its Two-Particle Reduced Density Matrices”. In: *Physical Review Letters* 89 (2002), p. 207901 (cit. on p. 89).

[Liu+22] P. Liu, Z. Liu, S. Chen, and X. Ma. “Fundamental Limitation on the Detectability of Entanglement”. In: *Physical Review Letters* 129 (2022), p. 230503 (cit. on pp. 40, 85).

[Llo96] S. Lloyd. “Universal Quantum Simulators”. In: *Science* 273 (1996), pp. 1073–1078 (cit. on p. 105).

[Lod10] A. Lodi. “Mixed integer programming computation”. In: *50 Years of Integer Programming 1958-2008: From the Early Years to the State-of-the-Art* (2010), pp. 619–645 (cit. on p. 25).

[Lof04] J. Lofberg. “YALMIP: A toolbox for modeling and optimization in MATLAB”. In: *2004 IEEE international conference on robotics and automation (IEEE Cat. No. 04CH37508)*. IEEE. 2004, pp. 284–289 (cit. on p. 126).

[LR15] S. Lörwald and G. Reinelt. “PANDA: a software for polyhedral transformations”. In: *EURO Journal on Computational Optimization* 3 (2015), pp. 297–308 (cit. on p. 61).

[Lov79] L. Lovász. “On the Shannon capacity of a graph”. In: *IEEE Transactions on Information Theory* 25.1 (1979), pp. 1–7 (cit. on pp. 32, 34).

[Lov86] L. Lovász. *An algorithmic theory of numbers, graphs and convexity*. SIAM, 1986 (cit. on p. 34).

[Löw55] P.-O. Löwdin. “Quantum Theory of Many-Particle Systems. I. Physical Interpretations by Means of Density Matrices, Natural Spin-Orbitals, and Convergence Problems in the Method of Configurational Interaction”. In: *Physical Review* 97 (1955), pp. 1474–1489 (cit. on pp. 106, 112).

[Lub+23] M. Lubin, O. Dowson, J. Dias Garcia, J. Huchette, B. Legat, et al. “JuMP 1.0: Recent improvements to a modeling language for mathematical optimization”. In: *Mathematical Programming Computation* 15 (2023), pp. 581–589 (cit. on pp. 24, 25, 125).

[MP14] L. Maccone and A. K. Pati. “Stronger Uncertainty Relations for All Incompatible Observables”. In: *Physical Review Letters* 113 (2014), p. 260401 (cit. on p. 31).

[Mal09] E. Maltais. “Covering arrays avoiding forbidden edges and edge clique covers”. PhD thesis. University of Ottawa (Canada), 2009 (cit. on p. 95).

[Mar+15] M. Marciniak, A. Rutkowski, Z. Yin, M. Horodecki, and R. Horodecki. “Unbounded violation of quantum steering inequalities”. In: *Physical Review Letters* 115.17 (2015), p. 170401 (cit. on p. 39).

[Maz12] D. A. Mazziotti. “Structure of Fermionic Density Matrices: Complete N-Representability Conditions”. In: *Physical Review Letters* 108 (2012), p. 263002 (cit. on p. 89).

[Mik+22] N. Miklin, A. Chaturvedi, M. Bourennane, M. Pawłowski, and A. Cabello. “Exponentially Decreasing Critical Detection Efficiency for Any Bell Inequality”. In: *Physical Review Letters* 129 (2022), p. 230403 (cit. on p. 39).

[MO21] N. Miklin and M. Oszmaniec. “A universal scheme for robust self-testing in the prepare-and-measure scenario”. In: *Quantum* 5 (2021), p. 424 (cit. on p. 16).

[Min17] S. Minsker. “On some extensions of Bernstein’s inequality for self-adjoint operators”. In: *Statistics & Probability Letters* 127 (2017), pp. 111–119 (cit. on pp. 74, 75).

[MG92] J. Misra and D. Gries. “A constructive proof of Vizing’s theorem”. In: *Information Processing Letters* 41.3 (1992), pp. 131–133 (cit. on p. 109).

[MH24] M. B. Morán and F. Huber. “Uncertainty Relations from State Polynomial Optimization”. In: *Physical Review Letters* 132 (2024), p. 200202 (cit. on pp. 36, 37).

[Mor+21] G. Moreno, R. Nery, C. de Gois, R. Rabelo, and R. Chaves. “Semi-device-independent certification of entanglement in superdense coding”. In: *Physical Review A* 103 (2021), p. 022426 (cit. on pp. 19, 56, 57, 140).

[Mor+12] T. Moroder, P. Hyllus, G. Tóth, C. Schwemmer, A. Niggebaum, et al. “Permutationally invariant state reconstruction”. In: *New Journal of Physics* 14.10 (2012), p. 105001 (cit. on pp. 8, 69).

[MOS] MOSEK ApS. *MOSEK optimizer documentation*. `https://www.mosek.com/documentation/` (cit. on p. 24).

[NOP09a] M. Navascués, M. Owari, and M. B. Plenio. “Complete Criterion for Separability Detection”. In: *Physical Review Letters* 103 (2009), p. 160404 (cit. on pp. 11, 54, 131, 132).

[NOP09b] M. Navascués, M. Owari, and M. B. Plenio. “Power of symmetric extensions for entanglement detection”. In: *Physical Review A* 80 (2009), p. 052306 (cit. on pp. 48, 54, 130, 131).

[NPA08] M. Navascués, S. Pironio, and A. Acín. “A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations”. In: *New Journal of Physics* 10.7 (2008), p. 073013 (cit. on p. 24).

[NPA07] M. Navascués, S. Pironio, and A. Acín. “Bounding the Set of Quantum Correlations”. In: *Physical Review Letters* 98 (2007), p. 010401 (cit. on p. 24).

[Nav+20] M. Navascués, E. Wolfe, D. Rosset, and A. Pozas-Kerstjens. “Genuine Network Multipartite Entanglement”. In: *Phys. Rev. Lett.* 125 (2020), p. 240505 (cit. on p. 14).

[NL25] A. Nayak and A. Lowe. “Lower bounds for learning quantum states with single-copy measurements”. In: *ACM Transactions on Computation Theory* (2025) (cit. on pp. 8, 79).

[NS02] A. Nayak and J. Salzman. “On communication over an entanglement-assisted quantum channel”. In: *Proceedings of the Thirty-fourth Annual ACM Symposium on Theory of Computing*. 2002, pp. 698–704 (cit. on p. 60).

[Neu+25] J. Neumann, T. V. Kondra, K. Hansenne, L. T. Weinbrenner, H. Kampermann, et al. “No quantum advantage without classical communication: fundamental limitations of quantum networks”. In: *arXiv:2503.09473* (2025) (cit. on p. 14).

[Ngu+20] H. C. Nguyen, S. Designolle, M. Barakat, and O. Gühne. “Symmetries between measurements in quantum mechanics”. In: *arXiv:2003.12553* (2020) (cit. on pp. 48, 49, 53, 136).

[Ngu+22] H. C. Nguyen, J. L. Bönsel, J. Steinberg, and O. Gühne. “Optimizing Shadow Tomography with Generalized Measurements”. en. In: *Physical Review Letters* 129.22 (2022), p. 220502 (cit. on p. 94).

[Nie99] M. A. Nielsen. “Conditions for a Class of Entanglement Transformations”. In: *Physical Review Letters* 83 (1999), pp. 436–439 (cit. on p. 13).

[NC10] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, 2010 (cit. on pp. 2, 12).

[ODo21] B. O’Donoghue. “Operator Splitting for a Homogeneous Embedding of the Linear Complementarity Problem”. In: *SIAM Journal on Optimization* 31 (2021), pp. 1999–2023 (cit. on p. 24).

[ODo+16] B. O’Donoghue, E. Chu, N. Parikh, and S. Boyd. “Conic Optimization via Operator Splitting and Homogeneous Self-Dual Embedding”. In: *Journal of Optimization Theory and Applications* 169.3 (2016), pp. 1042–1068 (cit. on p. 24).

[OEI] OEIS Foundation Inc. (2024). *Entry A000438 in The On-Line Encyclopedia of Integer Sequences*. `https://oeis.org/A000438` (cit. on p. 111).

[Orl77] J. Orlin. “Contentment in graph theory: Covering graphs with cliques”. In: *Indagationes Mathematicae (Proceedings)*. Vol. 80. 1977, pp. 406–424 (cit. on pp. 21, 95).

[Osz+17] M. Oszmaniec, L. Guerini, P. Wittek, and A. Acín. “Simulating Positive-Operator-Valued Measures with Projective Measurements”. In: *Physical Review Letters* 119 (2017), p. 190501 (cit. on p. 5).

[PV10] K. F. Pál and T. Vértesi. “Maximal violation of a bipartite three-setting, two-outcome Bell inequality using infinite-dimensional quantum systems”. In: *Physical Review A* 82 (2010), p. 022116 (cit. on p. 65).

[PS98] C. H. Papadimitriou and K. Steiglitz. *Combinatorial optimization: algorithms and complexity*. Courier Corporation, 1998 (cit. on p. 24).

[Pas+15] E. Passaro, D. Cavalcanti, P. Skrzypczyk, and A. Acín. “Optimal randomness certification in the quantum steering and prepare-and-measure scenarios”. In: *New Journal of Physics* 17.11 (2015), p. 113010 (cit. on p. 16).

[Pau+22a] J. Pauwels, S. Pironio, E. Z. Cruzeiro, and A. Tavakoli. “Adaptive Advantage in Entanglement-Assisted Communications”. In: *Physical Review Letters* 129 (2022), p. 120504 (cit. on pp. 57, 59).

[PPT25] J. Pauwels, S. Pironio, and A. Tavakoli. “Information capacity of quantum communication under natural physical assumptions”. In: *Quantum* 9 (2025), p. 1637 (cit. on p. 17).

[Pau+22b] J. Pauwels, A. Tavakoli, E. Woodhead, and S. Pironio. “Entanglement in prepare-and-measure scenarios: many questions, a few answers”. In: *New Journal of Physics* 24.6 (2022), p. 063015 (cit. on pp. 19, 56–59, 61, 65).

[PB11] M. Pawłowski and N. Brunner. “Semi-device-independent security of one-way quantum key distribution”. In: *Physical Review A* 84.1 (2011), p. 010302 (cit. on p. 16).

[Paw+09] M. Pawłowski, T. Paterek, D. Kaszlikowski, V. Scarani, A. Winter, et al. “Information causality as a physical principle”. en. In: *Nature* 461.7267 (2009), pp. 1101–1104 (cit. on pp. 16, 56).

[PŻ10] M. Pawłowski and M. Żukowski. “Entanglement-assisted random access codes”. en. In: *Physical Review A* 81.4 (2010), p. 042326 (cit. on p. 19).

[Per96] A. Peres. “Separability Criterion for Density Matrices”. In: *Physical Review Letters* 77 (1996), pp. 1413–1415 (cit. on p. 10).

[PD24] L. Perron and F. Didier. *CP-SAT*. Version v9.11. Google, 2024 (cit. on p. 99).

[Per+10] S. Perseguers, M. Lewenstein, A. Acín, and J. I. Cirac. “Quantum random networks”. In: *Nature Physics* 6.7 (2010), pp. 539–543 (cit. on p. 20).

[PW15] M. Piani and J. Watrous. “Necessary and Sufficient Quantum Information Characterization of Einstein-Podolsky-Rosen Steering”. In: *Physical Review Letters* 114 (2015), p. 060404 (cit. on p. 49).

[Pod+20] D. Poderini, S. Brito, R. Nery, F. Sciarrino, and R. Chaves. “Criteria for nonclassicality in the prepare-and-measure scenario”. en. In: *Physical Review Research* 2.4 (2020) (cit. on p. 56).

[Qu+22a] R. Qu, Y. Wang, M. An, F. Wang, Q. Quan, et al. “Retrieving High-Dimensional Quantum Steering from a Noisy Environment with N Measurement Settings”. In: *Physical Review Letters* 128 (2022), p. 240402 (cit. on p. 39).

[Qu+22b] R. Qu, Y. Wang, X. Zhang, S. Ru, F. Wang, et al. “Robust method for certifying genuine high-dimensional quantum steering with multimeasurement settings”. In: *Optica* 9.5 (2022), pp. 473–478 (cit. on pp. 40, 45).

[QVB14] M. T. Quintino, T. Vértesi, and N. Brunner. “Joint Measurability, Einstein-Podolsky-Rosen Steering, and Bell Nonlocality”. In: *Physical Review Letters* 113 (2014), p. 160402 (cit. on pp. 16, 53).

[Ren+04] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves. “Symmetric informationally complete quantum measurements”. In: *Journal of Mathematical Physics* 45.6 (2004), pp. 2171–2180 (cit. on pp. 73, 86).

[Ren+21] M.-O. Renou, D. Trillo, M. Weilenmann, T. P. Le, A. Tavakoli, et al. “Quantum theory based on real numbers can be experimentally falsified”. In: *Nature* 600.7890 (2021), pp. 625–629 (cit. on p. 137).

[RC09] A. H. Ronneseth and C. J. Colbourn. “Merging covering arrays and compressing multiple sequence alignments”. In: *Discrete Applied Mathematics* 157.9 (2009), pp. 2177–2190 (cit. on p. 95).

[Rus20] V. Russo. *Theory of quantum information toolkit: A Python package for studying quantum information*. `https://zenodo.org/records/4743211`. 2020 (cit. on p. 127).

[SS22] G. Sagnol and M. Stahlberg. “PICOS: A Python interface to conic optimization solvers”. In: *Journal of Open Source Software* 7.70 (2022), p. 3915 (cit. on p. 127).

[SBL01] A. Sanpera, D. Bruß, and M. Lewenstein. “Schmidt-number witnesses and bound entanglement”. In: *Physical Review A* 63 (2001), p. 050301 (cit. on pp. 13, 39).

[Sca19] V. Scarani. *Bell Nonlocality*. en. 1st ed. Oxford University Press, 2019 (cit. on p. 19).

[Sch36] E. Schrödinger. “Probability relations between separated systems”. In: *Mathematical Proceedings of the Cambridge Philosophical Society* 32.3 (1936), pp. 446–452 (cit. on p. 5).

[Sch35] E. Schrödinger. “Discussion of probability relations between separated systems”. In: *Mathematical Proceedings of the Cambridge Philosophical Society*. Vol. 31–4. Cambridge University Press. 1935, pp. 555–563 (cit. on p. 9).

[Sch+15] C. Schwemmer, L. Knips, D. Richart, H. Weinfurter, T. Moroder, et al. “Systematic Errors in Current Quantum State Tomography Tools”. In: *Physical Review Letters* 114 (2015), p. 080403 (cit. on p. 73).

[SDW17] R. Schwonnek, L. Dammeier, and R. F. Werner. “State-Independent Uncertainty Relations and Entanglement Detection in Noisy Systems”. In: *Physical Review Letters* 119 (2017), p. 170404 (cit. on p. 31).

[SU08] M. Seevinck and J. Uffink. “Partial separability and entanglement criteria for multiqubit quantum states”. In: *Physical Review A* 78 (2008), p. 032101 (cit. on pp. 14, 85).

[Sha+13] J. Shang, H. K. Ng, A. Sehrawat, X. Li, and B.-G. Englert. “Optimal error regions for quantum state estimation”. In: *New Journal of Physics* 15.12 (2013), p. 123026 (cit. on pp. 8, 69).

[Sha56] C. Shannon. “The zero error capacity of a noisy channel”. In: *IRE Transactions on Information Theory* 2.3 (1956), pp. 8–19 (cit. on p. 33).

[SC15] P. Skrzypczyk and D. Cavalcanti. “Loss-tolerant Einstein-Podolsky-Rosen steering for arbitrary-dimensional states: Joint measurability and unbounded violations under losses”. In: *Physical Review A* 92.2 (2015), p. 022354 (cit. on p. 39).

[SC23] P. Skrzypczyk and D. Cavalcanti. *Semidefinite programming in quantum information science*. IOP Publishing, 2023 (cit. on p. 24).

[Slo19] W. Slofstra. “The set of quantum correlations is not closed”. In: *Forum of Mathematics, Pi* 7 (2019), e1 (cit. on p. 65).

[Sou23] C. H. de Souza Vieira. “Exploring the role of entanglement in sets of behaviors from prepare-and-measure scenarios”. Available at `https://hdl.handle.net/20.500.12733/8536`. Ph.D. Thesis. Universidade Estadual de Campinas, 2023 (cit. on pp. 19, 55).

[Str+22] R. Stricker, M. Meth, L. Postler, C. Edmunds, C. Ferrie, et al. “Experimental Single-Setting Quantum State Tomography”. In: *PRX Quantum* 3 (2022), p. 040310 (cit. on pp. 6, 86, 94).

[STM13] T. Sugiyama, P. S. Turner, and M. Murao. “Precision-Guaranteed Quantum Tomography”. In: *Physical Review Letters* 111 (2013), p. 160406 (cit. on p. 71).

[SŻ19] K. Szymański and K. Życzkowski. “Geometric and algebraic origins of additive uncertainty relations”. In: *Journal of Physics A: Mathematical and Theoretical* 53.1 (2019), p. 015302 (cit. on pp. 31, 32).

[TO23] S. T. Flammia and R. O’Donnell. “Quantum chi-squared tomography and mutual information testing”. In: *arXiv.2305.18519* (2023) (cit. on p. 79).

[Tac+20] F. Tacchino, A. Chiesa, S. Carretta, and D. Gerace. “Quantum Computers as Universal Quantum Simulators: State-of-the-Art and Perspectives”. In: *Advanced Quantum Technologies* 3.3 (2020), p. 1900052 (cit. on p. 105).

[Tav+18] A. Tavakoli, J. Kaniewski, T. Vértesi, D. Rosset, and N. Brunner. “Self-testing quantum states and measurements in the prepare-and-measure scenario”. en. In: *Physical Review A* 98.6 (2018), p. 062307 (cit. on p. 16).

[Tav+21] A. Tavakoli, J. Pauwels, E. Woodhead, and S. Pironio. “Correlations in Entanglement-Assisted Prepare-and-Measure Scenarios”. In: *PRX Quantum* 2 (2021). Preprint at arXiv:2103.10748, p. 040357 (cit. on pp. 19, 56–61, 65).

[Tav+23] A. Tavakoli, A. Pozas-Kerstjens, P. Brown, and M. Araújo. “Semidefinite programming relaxations for quantum correlations”. In: *arXiv:2307.02551* (2023) (cit. on p. 86).

[Tav+24] A. Tavakoli, A. Pozas-Kerstjens, P. Brown, and M. Araújo. “Semidefinite programming relaxations for quantum correlations”. In: *Reviews of Modern Physics* 96.4 (2024), p. 045006 (cit. on p. 24).

[TRR19] A. Tavakoli, D. Rosset, and M.-O. Renou. “Enabling Computation of Correlation Bounds for Finite-Dimensional Quantum Systems via Symmetrization”. In: *Physical Review Letters* 122.7 (2019), p. 070501 (cit. on pp. 53, 136).

[TŻB20] A. Tavakoli, M. Żukowski, and Č. Brukner. “Does violation of a Bell inequality always imply quantum advantage in a communication complexity problem?” In: *Quantum* 4 (2020), p. 316 (cit. on p. 19).

[Ter00] B. M. Terhal. “Bell inequalities and the separability criterion”. In: *Physics Letters A* 271.5 (2000), pp. 319–326 (cit. on p. 85).

[TH00] B. M. Terhal and P. Horodecki. “Schmidt number for density matrices”. In: *Physical Review A* 61 (2000), p. 040301 (cit. on pp. 12, 13, 39, 42).

[Til+21] J. Tilly, P. V. Sriluckshmy, A. Patel, E. Fontana, I. Rungger, et al. “Reduced density matrix sampling: Self-consistent embedding and multiscale electronic structure on current generation quantum computers”. In: *Physical Review Research* 3 (2021), p. 033230 (cit. on p. 106).

[Tor+18] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, et al. “Neural-network quantum state tomography”. In: *Nature Physics* 14.5 (2018), pp. 447–450 (cit. on p. 137).

[Tor] J. Torres-Jimenez. *Uniform CA*. `https://www.tamps.cinvestav.mx/~oc/` (cit. on p. 96).

[TI13] J. Torres-Jimenez and I. Izquierdo-Marquez. “Survey of covering arrays”. In: *2013 15th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing*. IEEE. 2013, pp. 20–27 (cit. on pp. 95, 97).

[TA14] G. Tóth and I. Apellaniz. “Quantum metrology from a quantum information science perspective”. In: *Journal of Physics A: Mathematical and Theoretical* 47.42 (2014), p. 424006 (cit. on p. 39).

[TG05] G. Tóth and O. Gühne. “Entanglement detection in the stabilizer formalism”. In: *Physical Review A* 72 (2005), p. 022340 (cit. on pp. 31, 34).

[Tro15] J. A. Tropp. “An Introduction to Matrix Concentration Inequalities”. In: *Foundations and Trends in Machine Learning* 8.1-2 (2015), pp. 1–230 (cit. on pp. 74, 75).

[Tro12] J. A. Tropp. “User-Friendly Tail Bounds for Sums of Random Matrices”. In: *Foundations of Computational Mathematics* 12.4 (2012), pp. 389–434 (cit. on p. 74).

[Uol+15] R. Uola, C. Budroni, O. Gühne, and J.-P. Pellonpää. “One-to-One Mapping between Steering and Joint Measurability Problems”. In: *Physical Review Letters* 115 (2015), p. 230402 (cit. on pp. 16, 53).

[Uol+20] R. Uola, A. C. S. Costa, H. C. Nguyen, and O. Gühne. “Quantum steering”. In: *Reviews of Modern Physics* 92 (2020), p. 015001 (cit. on pp. 14, 16, 40).

[UMG14] R. Uola, T. Moroder, and O. Gühne. “Joint Measurability of Generalized Measurements Implies Classicality”. In: *Physical Review Letters* 113 (2014), p. 160403 (cit. on p. 16).

[VB96] L. Vandenberghe and S. Boyd. “Semidefinite Programming”. In: *SIAM Review* 38.1 (1996), pp. 49–95 (cit. on p. 23).

[VK24] O. Veltheim and E. Keski-Vakkuri. “Multiset tomography: Optimizing quantum measurements by partitioning multisets of observables”. In: (2024). arXiv: `2403.07068` (cit. on pp. 89, 96, 102, 114, 121).

[Ver18] R. Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018 (cit. on p. 74).

[VYI20] V. Verteletskyi, T.-C. Yen, and A. F. Izmaylov. "Measurement optimization in the variational quantum eigensolver using a minimum clique cover". In: *The Journal of Chemical Physics* 152.12 (2020) (cit. on pp. 96, 102, 106, 114, 121).

[VPB10] T. Vértesi, S. Pironio, and N. Brunner. "Closing the Detection Loophole in Bell Experiments Using Qudits". In: *Physical Review Letters* 104 (2010), p. 060401 (cit. on p. 39).

[VT99] G. Vidal and R. Tarrach. "Robustness of entanglement". In: *Physical Review A* 59 (1999), pp. 141–155 (cit. on p. 49).

[Vie+25] C. Vieira, C. de Gois, S. Designolle, P. Lauand, and M. T. Quintino. "Outcome communication cannot explain quantum correlations". In: *(In preparation)* (2025) (cit. on p. 139).

[Vie+23] C. Vieira, C. de Gois, L. Pollyceno, and R. Rabelo. "Interplays between classical and quantum entanglement-assisted communication scenarios". In: *New Journal of Physics* 25.11 (2023), p. 113004 (cit. on pp. 19, 55, 59, 139).

[Vin14] C. Vinzant. "WHAT IS...a Spectrahedron?" In: *Notices of the American Mathematical Society* 61.5 (2014), p. 1 (cit. on p. 23).

[Vir+20] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, et al. "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python". In: *Nature Methods* 17 (2020), pp. 261–272 (cit. on p. 126).

[Viz65] V. G. Vizing. "The chromatic class of a multigraph". In: *Cybernetics* 1.3 (1965), pp. 32–41 (cit. on p. 109).

[Wai19] M. J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019 (cit. on pp. 74, 75).

[WGE16] M. Walter, D. Gross, and J. Eisert. "Multipartite Entanglement". In: *Quantum Information*. John Wiley & Sons, Ltd, 2016. Chap. 14, pp. 293–330. eprint: `https://onlinelibrary.wiley.com/doi/pdf/10.1002/9783527805785.ch14` (cit. on p. 14).

[WSR19] J. Wang, V. B. Scholz, and R. Renner. "Confidence Polytopes in Quantum State Tomography". In: *Physical Review Letters* 122 (2019), p. 190401 (cit. on p. 71).

[Wan+19] Y. Wang, I. W. Primaatmaja, E. Lavie, A. Varvitsiotis, and C. C. W. Lim. "Characterising the correlations of prepare-and-measure quantum networks". en. In: *npj Quantum Information* 5.1 (2019), pp. 1–6 (cit. on p. 16).

[Wat18] J. Watrous. *The Theory of Quantum Information*. Cambridge University Press, 2018 (cit. on pp. 11, 47, 129).

[WCD08] S. Wehner, M. Christandl, and A. C. Doherty. "Lower bound on the dimension of a quantum system given measured data". In: *Physical Review A* 78.6 (2008), p. 062112 (cit. on p. 16).

[WW08] S. Wehner and A. Winter. "Higher entropic uncertainty relations for anti-commuting observables". In: *Journal of Mathematical Physics* 49.6 (2008) (cit. on pp. 31, 34).

[Wei+20] M. Weilenmann, B. Dive, D. Trillo, E. A. Aguilar, and M. Navascués. "Entanglement Detection beyond Measuring Fidelities". In: *Physical Review Letters* 124 (2020), p. 200502 (cit. on pp. 13, 39).

[Wer01] R. F. Werner. "All teleportation and dense coding schemes". In: *Journal of Physics A: Mathematical and General* 34.35 (2001), pp. 7081–7094 (cit. on p. 59).

[Wer89a] R. F. Werner. "An application of Bell's inequalities to a quantum state extension problem". In: *Letters in Mathematical Physics* 17.4 (1989), pp. 359–363 (cit. on p. 11).

[Wer89b] R. F. Werner. "Quantum states with Einstein-Podolsky-Rosen correlations admitting a hidden-variable model". In: *Physical Review A* 40 (1989), pp. 4277–4281 (cit. on p. 9).

[Wie83] S. Wiesner. "Conjugate coding". In: *ACM SIGACT News* 15.1 (1983), pp. 78–88 (cit. on p. 55).

[WJD07] H. M. Wiseman, S. J. Jones, and A. C. Doherty. "Steering, Entanglement, Nonlocality, and the Einstein-Podolsky-Rosen Paradox". In: *Physical Review Letters* 98 (2007), p. 140402 (cit. on p. 14).

[WZ82] W. K. Wootters and W. H. Zurek. "A single quantum cannot be cloned". In: *Nature* 299.5886 (1982), pp. 802–803 (cit. on p. 6).

[XSW24] Z.-P. Xu, R. Schwonnek, and A. Winter. "Bounding the Joint Numerical Range of Pauli Strings by Graph Parameters". In: *PRX Quantum* 5 (2024), p. 020318 (cit. on p. 36).

[YW23] N. Yu and T.-C. Wei. "Learning marginals suffices!" In: *arXiv:2303.08938* (2023) (cit. on p. 89).

[ZRM21] A. Zhao, N. C. Rubin, and A. Miyake. "Fermionic Partial Tomography via Classical Shadows". In: *Physical Review Letters* 127 (2021), p. 110504 (cit. on p. 106).

[Zha+20] A. Zhao, A. Tranter, W. M. Kirby, S. F. Ung, A. Miyake, et al. "Measurement reduction in variational quantum algorithms". In: *Physical Review A* 101 (2020), p. 062322 (cit. on p. 114).

[Zhu+21] F. Zhu, M. Tyler, N. H. Valencia, M. Malik, and J. Leach. "Is high-dimensional photonic entanglement robust to noise?" In: *AVS Quantum Science* 3.1 (2021), p. 011401 (cit. on p. 39).

[Žni+07] M. Žnidarič, T. Prosen, G. Benenti, and G. Casati. “Detecting entanglement of random states with an entanglement witness”. In: *Journal of Physics A: Mathematical and Theoretical* 40.45 (2007), p. 13787 (cit. on p. 85).