Doctoral Thesis

Positive and Invariant Tensor Decompositions

Approximations and Computational Complexity

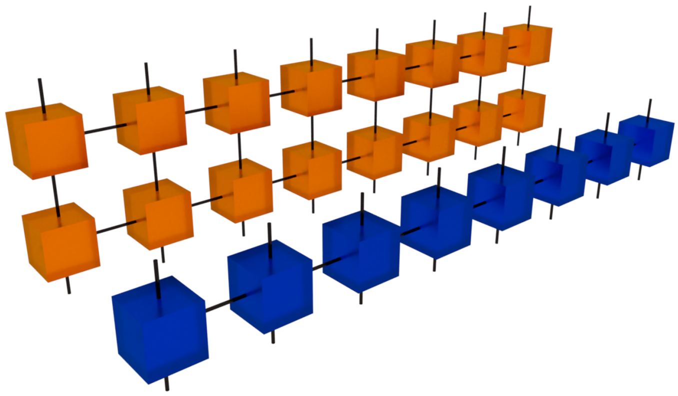

Author: Andreas Klingler
01517439

Date: May 16, 2024

Submitted to the University of Innsbruck, Faculty of Mathematics, Computer Science, and Physics for the academic degree

Doctor of Philosophy (Doktoratsstudium Physik)

Supervisors: assoz.-Prof. Dr. Gemma De les Coves
Institute for Theoretical Physics

Univ.-Prof. Dr. Tim Netzer
Department for Mathematics

Faculty for Mathematics, Computer Science and Physics

Colophon
This document was typeset with the help of KOMA-Script and LaTeX using the kaobook class.

Acknowledgments

I am deeply grateful to my first supervisor, Gemma De les Coves, for giving me the opportunity to do research in this fascinating field and to work on this challenging and exciting project. Under her mentorship, I acquired the essential methodologies for conducting research and learned so many skills in presenting and communicating my findings with clarity. I also appreciate her efforts in introducing me to the academic world, for giving me so much freedom in my research, and, in particular, for encouraging me to follow my goals.

I also extend my most profound appreciation to my second supervisor, Tim Netzer, for his consistent support throughout this journey. His door was always open, offering invaluable patience and guidance whether I encountered complex challenges or simply sought clarification on mathematical concepts. I will always cherish our numerous inspiring discussions, which so often gave me new ideas and kept me optimistic about the success of the projects.

The research presented in this dissertation was partially funded by the Austrian Science Fund (FWF) via the stand-alone project "Positivity structures in quantum many-body systems" (doi: 10.55776/P33122), as well as the Austrian Academy of Sciences (ÖAW) via the DOC fellowship "Decompositions of tensors with invariance, positivity and approximations" (project number 26547).

Throughout my PhD, I had the privilege of getting to know many wonderful friends and colleagues who accompanied me on this journey, which would not have been as enriching without their support.

I want to thank all the members of my research group for fostering a great atmosphere within our team. Special gratitude to Tomáš Gonda for his patience in helping me navigate through Category theory. To Johannes Fankhauser, for all our stimulating discussions about philosophy, physics, and psychology. To Mirte van der Eyden, Sebastian Stengele, and Tobias Reinhart for engaging in countless discussions that deepened my understanding of various concepts in physics and mathematics. I am also grateful to our secretary, Jade Meysami-Hörtnagl, for always keeping track of organizational matters, especially concerning my fellowship.

I am also immensely grateful to Tobias Fritz for introducing me to the world of Categorical Probability, which opened up a new line of research for me. Furthermore, I want to thank Areeb Shah-Mohammed and Antonio Lorenzin for their insightful discussions that enhanced my understanding of Categorical Probability.

I also want to thank Paria Abbasi for the many fruitful discussions and our collaboration during the beginning of my PhD. I am also grateful to Paria and to Martin Berger for their support during my application process for the DOC fellowship.

I extend a heartfelt thanks to all my friends who supported me throughout my entire academic journey. Particularly to Michael Fellner, with whom I shared this journey from the very first week of my Bachelor studies. I also want to thank all my friends outside of academia, especially Caro, Dani, Eva, Steffi, Sabine, Alex and Josi, for their unwavering emotional support, even during the most challenging times.

Finally, I am deeply grateful to my parents for their unwavering support and sacrifices that enabled me to pursue studies in physics and mathematics. Their support allowed me to fully immerse myself in my studies.

Abstract

Many composite systems are described by a tensor product and feature a notion of positivity.

Describing multipartite positive tensors is challenging for two reasons. One is the exponential growth in the number of parameters. The second is the fact that the tensor product interacts with the positivity cones in an intricate way. For example, it may be costly to enforce the positivity in the local terms of the decomposition.

This thesis studies composite systems subject to positivity structures from the perspective of algebraic geometry and computational complexity.

In the first perspective, we present a framework to decompose positive and invariant tensors so that these properties manifest in the local terms and prove under which conditions optimizations over such tensors are stable. We then apply this framework to positive, invariant multivariate polynomials. Finally, we explore implications for the topology of the space of quantum correlation scenarios.

The second perspective concerns computational problems inspired by tensor decompositions. We leverage a relation between tensor decompositions and certain linear recurrence sequences (called moment sequences) to prove the decidability or undecidability of the positivity of such sequences. Finally, we show that many undecidable problems in physics, computer science, and mathematics concerning arbitrary large composite systems have bounded versions that are NP-hard.

Overall, this thesis sheds light on the algebraic, numerical, and computational properties of composite systems, particularly on tensor product spaces, with positivity structures and invariance. It also unveils tensor decompositions in unexpected places, to which a wealth of results can be applied.

List of publications

This thesis is based on the following publications and preprints:

[P1] **Border Ranks of Positive and Invariant Tensor Decompositions: Applications to Correlations**
A. Klingler, T. Netzer, G. De les Coves
arXiv:2304.13478

[P2] **Polynomial decompositions with invariance and positivity inspired by tensors**
G. De las Cuevas, A. Klingler, T. Netzer
arXiv:2109.06680

[P3] **Positive Moments Forever: Undecidable and Decidable Cases**
G. De les Coves, J. Graf, A. Klingler, T. Netzer
arXiv:2404.15053

[P4] **Many bounded versions of undecidable problems are NP-hard**
A. Klingler, M. van der Eyden, S. Stengele, T. Reinhart, G. De las Cuevas
SciPost Physics 14 (6), 173. doi: 10.21468/SciPostPhys.14.6.173

Specifically, Chapters 3-5 and Chapters 7-8 mirror segments of these publications, albeit with modifications and restructurings made for coherence and uniformity. In [P4], the abstract framework was collaboratively developed by all authors, with the first and second authors equally contributing to the elaboration of ideas and the primary writing responsibility falling on the first author. In [P1,P2,P3], the author of this thesis was in charge of writing the manuscript, while all authors equally participated in idea development and provided critical reviews of the content. Note that the articles [P2] and [P3] follow an alphabetical ordering of the authors.

During this research period, I also contributed to articles which are not included in this thesis:

[P5] **Approximate Pythagoras Numbers on $\star$-algebras over $\mathbb{C}$**
P. Abbasi, S. Gribling, A. Klingler, T. Netzer
Journal of Complexity 74, 101698 (2023). doi: 10.1016/j.jco.2022.101698

[P6] **Approximate Completely Positive Semidefinite Factorizations and their Ranks**
P. Abbasi, A. Klingler, T. Netzer
Linear Algebra and its Applications 677, 323-336 (2023). doi: 10.1016/j.laa.2023.08.005

[P7] **The d-Separation Criterion in Categorical Probability**
T. Fritz, A. Klingler
Journal of Machine Learning Research 24(46), 1-49 (2023). doi: 10.48550/arXiv.2207.05740

[P8] **Hidden Markov Models and the Bayes Filter in Categorical Probability**
T. Fritz, A. Klingler, D. McNeely, A. Shah-Mohammed, Y. Wang
arXiv:2401.14669

[P9] **Homotopy Methods for Convex Optimization**
A. Klingler, T. Netzer
arXiv:2403.02095

Contents

1 Introduction

To specify a theory or framework, one needs to describe its basic components and how they are composed, i.e. how they can be combined to give rise to other elements. A prime example lies in the postulates of quantum mechanics, which not only detail the description of individual systems but also their composition into larger, composite systems. It follows that the notion of composition is thus a fundamental and essential part of a theory.

The *tensor product* is a salient instance of composition found in theories like quantum theory and probability theory. It captures the essence of composition in systems where correlations between subsystems are fundamental. For instance, in probabilistic settings, characterizing joint systems requires defining probabilities for all combinations of outcomes across subsystems.

Another crucial feature in these theories is *positivity*, by which the tensor product space is equipped with a *positivity structure*. A certain cone of positive elements defines valid objects in the theory. In discrete probability theory, for instance, each outcome is associated with a nonnegative number — the probability of that specific outcome. Consequently, only tensors with nonnegative entries can describe probability distributions. Similarly, in quantum theory: Open quantum systems are described by mixed states, which are *positive semidefinite matrices*, establishing a cone within the space of matrices.

Yet, tensor product structures pose challenges, notably due to the exponential increase in the dimension of the system. For instance, simulating the dynamics of a small quantum system with more than 100 particles is impossible due to the exponential amount of degrees of freedom. If the tensor product space is equipped with a positivity structure, additional challenges arise due to the difficulty of verifying the positivity constraint in the global tensors.

To address all these challenges, *tensor (network) decompositions* offer a practical and powerful approach, both with analytical and numerical applications. They describe elements in a multipartite tensor product space by breaking them down into elementary components, enabling the simulation of large quantum systems in a tractable way. Prominent examples of tensor (network) decompositions are *matrix product states*, an efficient representation of certain one-dimensional systems, or *projected entangled pair states*, a generalization of matrix product states to higher dimensional grids. The complexity of representing a tensor using such decompositions is determined by the rank of the decomposition, which reflects the number of degrees of freedom needed to represent the original vector.

A tensor product space incorporating an additional positivity structure introduces numerous challenges for tensor decompositions. On the one hand, the global positivity of the tensor may not be clearly reflected in the resulting decomposition. In other words, the positivity of the tensor is not inherent in the tensor decomposition. On the other hand, attempting to reflect positivity in the decomposition can significantly increase its complexity in the local terms representing the tensor [36].

Recently, a framework to decompose positive and invariant tensors was introduced [37]. This framework generates tensor decompositions along with three variations:

- **Decomposition geometry:** The framework offers arbitrary decomposition geometries, each mimicking a structural arrangement resembling tensor networks. The geometry is determined by a *weighted simplicial complex*.
- **Explicit positivity:** The decompositions can be made explicitly positive in various ways. They have to ensure that the resulting tensors have the required positivity constraint.
- **Explicit invariance:** For tensors invariant under permutations of the local systems, we introduce constraints on the local elements in the decompositions that lead to explicitly invariant global tensors.

In this thesis, we investigate this decomposition framework from two perspectives:

- **Applications:** Is there an operational interpretation of positive tensor decompositions? In other words, do positive tensors that admit a certain decomposition have an interpretation beyond their mere mathematical representation? We prove that tensors that attain particular decompositions correspond to specific correlation scenarios in quantum information (Chapter 3). We also introduce a novel framework inspired by positive tensor decompositions to decompose positive, multivariate polynomials into univariate ones. This framework tracks the positivity and invariance of the polynomials in local (univariate) polynomials (Chapter 5).
- **Approximations:** How do tensor decompositions behave under approximations? In particular, is the rank a stable parameter, or can it collapse for small approximations? We prove that positive and invariant tensor decompositions can exhibit instablities when subjected to approximations. We also elucidate the implications of this instability for optimization strategies and correlation scenarios in quantum information theory (Chapter 4).

We also explore the relationship between positivity and large systems, like those present in the tensor product, from the computational complexity perspective. Specifically, multipartite tensor product spaces with positivity constraints share properties and challenges with other large systems. We demonstrate that specific tensor decompositions give rise to what are known as *moment sequences*, and verifying the positivity of these decompositions corresponds to solving the *positivity problem* for these moment sequences. The computational complexity results for such tensor decompositions offer a novel perspective on the positivity problem for arbitrary moment sequences, specifically for *linear recurrence sequences*. Moreover, we show that a specific property of tensor decompositions — the bounded version of an undecidable problem becomes NP-hard — applies to various problems in quantum information, quantum many-body physics, mathematics and computer science.

This thesis is divided into two parts (see also Figure 1.1). In the first part, we provide a comprehensive review of the framework for decom-

Part I:	**Part II:**
Decomposing positive and invariant tensors	**Computational problems motivated by tensor decompositions**
Applications: ▶ Correlations (Chapter 3) ▶ Polynomials (Chapter 5)	**Moment membership**: ▶ Complexity (Chapter 7)
Approximations: ▶ Stability (Chapter 4)	**Undecidable problems**: ▶ Bounding (Chapter 8)

Figure 1.1: Structure of this thesis. In the first part, we study the framework to decompose positive and invariant tensors (introduced in [37]) from the perspectives of approximations and applications. In the second part, we study two questions in computational complexity motivated by known computational results for tensor decompositions.

posing positive and invariant tensors [37]. We study its stability under approximations and explore its applications to correlation scenarios and polynomial decompositions. The second part focuses on demonstrating how computational complexity results for positive tensor decompositions inspire various questions in computational complexity, especially when combining positivity and large systems.

Let us now give a brief overview of the specific questions, results, and methods in the different parts and chapters of this thesis.

Part I: Decompositions of Positive Tensors: Approximations and Applications

In the first part, we review the framework to decompose positive and invariant tensors, introduced in [37] (Chapter 2). Building upon this framework, we present three results based on [74] and [39].

This part intersects two fields: (semi-)algebraic geometry and quantum theory (see Figure 1.2). Specifically, we relate tools from algebraic geometry, like the border rank of a tensor, with concepts in quantum information and quantum many-body physics, like correlation scenarios and tensor network decompositions of mixed states. Conversely, we show that the tensor (network) decompositions initially conceived for quantum many-body systems give rise to a novel family of decompositions for positive polynomials, which are the main characters in semi-algebraic geometry.

Let us now elaborate on the questions and results in each chapter.

Tensor decompositions and correlation scenarios (Chapter 3). What probability distributions can emerge from such shared resources when multiple distant parties share a particular class of quantum states? Entanglement within the shared state determines the strength of correlations from the probability distributions arising from measuring the state locally. For instance, without any shared resources, the resulting probability distributions can only be independent.

We show that positive tensor decompositions relate to particular correlation scenarios. Specifically, we prove that positive tensors with bounded ranks correspond to probability distributions that arise via local measurements from quantum states with a specific entanglement structure.

This correspondence provides us with both an operational interpretation of the decomposition framework and a means to link properties of tensor decompositions. This link will be further explored in Chapter 4.

Instabilities of tensor decompositions (Chapter 4). A crucial aspect of tensor decompositions is their sensitivity to approximations, governed by a defining parameter: the *rank* of the decompositions. When the rank of a tensor is low, fewer degrees of freedom suffice to express the tensor. For this reason, the rank is often used as a parameter to upper bound the cost of representing tensors in numerical simulations. For instance, to make an optimization problem over tensors tractable, one relaxes the problem by only optimizing over the set of tensors with a bounded rank.

Unlike matrices, whose matrix rank remains stable under slight perturbations, the tensor rank can collapse for arbitrarily small approximation errors. This leads to undesirable properties for fixed-rank approximations of tensors, like instability of optimization problems. We show this instability by introducing the border rank of a tensor, a well-known rank notion from algebraic geometry, and enrich it with positivity and invariance constraints. The border rank of a tensor measures the best way to represent a tensor up to arbitrary small approximations. If the border rank of a tensor is strictly smaller than its original rank (indicating that arbitrarily good approximations of the tensor can be represented more efficiently than the exact tensor), then the tensor decomposition is instable.

Finally, we relate this instability with the correlation scenarios presented in Chapter 3. The instabilities on the tensor side lead to constraints on the feasibility of testing resources from finite samples.

Polynomial decompositions inspired by tensors (Chapter 5). Motivated by the tensor decomposition framework, we introduce a novel approach to decompose multivariate polynomials, which explicitly showcases their invariance and positivity. The symmetry with respect to permutations of variables specifies the invariance and the notion of sum-of-square polynomials specify the positivity in the space of multivariate polynomials. We show that these polynomial decompositions behave very similarly to the original tensor decomposition framework. Specifically, we prove that they parametrize the entire space of positive and invariant polynomials in certain situations. Moreover, we show that separations between ranks appear as well.

Part II: Computational Aspects of Tensor Decompositions and Beyond

Many tensor (network) decomposition problems are very hard to solve on a computer. Even worse, some of these problems are even undecidable, i.e. there is no algorithm that solves them. In the second part of this

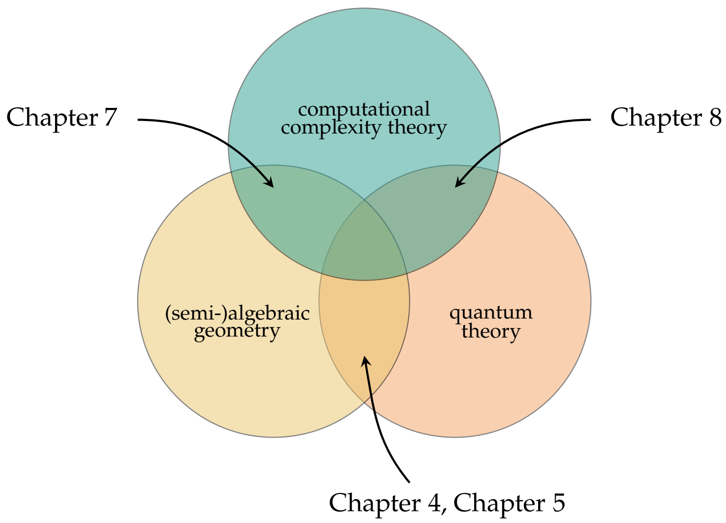

Figure 1.2: This thesis applies tools from computational complexity theory, semi-algebraic geometry, and quantum theory. In Chapter 4, we study tensor (network) decompositions from the perspective of algebraic geometry by computing their *border ranks*. In Chapter 5, we apply tensor decompositions to introduce a novel type of polynomial decompositions that inherits positivity and invariance. In Chapter 7, we apply tools from computational complexity theory and semi-algebraic geometry to prove that the moment membership problem is decidable, and in Chapter 8, we show that many bounded versions of undecidable problems in quantum information theory are NP-hard.

thesis, we introduce problems and questions motivated by computational aspects of tensor decompositions.

First, we show that certain tensor decompositions give rise to so-called matrix moments. Second, we show that the computational behavior of many large systems is already known for tensor decompositions, namely undecidable problems give rise to NP-hard *bounded versions*.

Positivity of matrix moments (Chapter 7). Matrix moment sequences are sequences of the form

$$n \mapsto \operatorname{tr}(A^n)$$

where A is a matrix and tr is the trace of a matrix. While these matrix moment sequences are usually considered over matrices with real or complex entries, they also generalize to matrices with entries that live in a ring (i.e. a structure that allows only for multiplication and addition but not inversion). This generalization also includes, for example, a particular class of tensor decompositions by choosing a specific ring.

Moment sequences are particular instances of so-called *linear recurrences sequences* which find applications in many different contexts. We show that the undecidability of a particular tensor decomposition problem gives rise to an undecidable problem for moment sequences and, therefore, also for linear recurrence sequences. Moreover, we prove that specific problems for these moment sequences remain decidable.

In this chapter, we use tools from semi-algebraic geometry to prove the decidability of specific moment membership problems.

Bounded versions of undecidable problems (Chapter 8). Many problems in physics, mathematics and computer science have been proven undecidable. All these problems share a common theme: There is a parameter in the problem statement that can be arbitrarily large. In the example of tensor network decompositions, this parameter is the number of tensor product spaces; these problems ask for properties of tensor decompositions of arbitrary size.

What happens to the complexity of the problem if this parameter is bounded? In Chapter 8, we show that many bounded versions of undecidable problems that arise from bounding the parameter become NP-hard. Specifically, we elucidate how the proof for undecidable problems can be leveraged to prove the NP-hardness of the bounded versions. While this was already known for several tensor network problems [36, 72, 108], we extended this principle to many other problems in mathematics and physics. For this reason, the tools used in this part are at the intersection of computational complexity theory and quantum theory.

Part I

Decompositions of Positive Tensors: Approximations and Applications

A framework to decompose positive and invariant tensors 2

The tensor product is a mathematical construct that models the composition of single systems into a joint system in many theories, including quantum theory and probability theory. In simple scenarios, a tensor product can be understood as a collection of scalar values (for example, real or complex numbers) in a multi-dimensional array. The simplest examples of tensors are one-dimensional arrays, often referred as vectors

$$\left(a_i\right)_{i=1,\ldots,d} = \left(a_1\; a_2\; \ldots\; a_d\right).$$

Expanding this concept to the two-dimensional realm yields bipartite tensors, known as matrices

$$\left(a_{ij}\right)_{i,j=1,\ldots,d} = \begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1d} \\ a_{21} & a_{22} & & \\ \vdots & & \ddots & \\ a_{d1} & & & a_{dd} \end{pmatrix}.$$

Analogously, n-partite tensors are represented by n-dimensional arrays

$$\left(a_{i_1,\ldots,i_n}\right)_{i_1,\ldots,i_n}$$

where every entry is a scalar (see Figure 2.1).

Multipartite tensors encounter a significant challenge: their degrees of freedom grow exponentially with the system size. While a single d-dimensional vector is uniquely defined by d scalar values, an n-partite tensor with each part having a local dimension of d necessitates d^n distinct parameters. Tensor decompositions offer a strategies to represent specific tensors more parsimoniously.

Take for example the n-partite tensor of the following form:

$$a_{i_1} \cdot b_{i_2} \cdots c_{i_n}$$

This particular tensor requires only $n \cdot d$ distinct scalar values to fully parametrize the n-partite tensor. Tensors structured in this manner are termed *elementary*. The conventional tensor decomposition involves breaking down tensors into a combination of elementary tensors (as depicted in Figure 2.2). The count of elementary tensors needed to

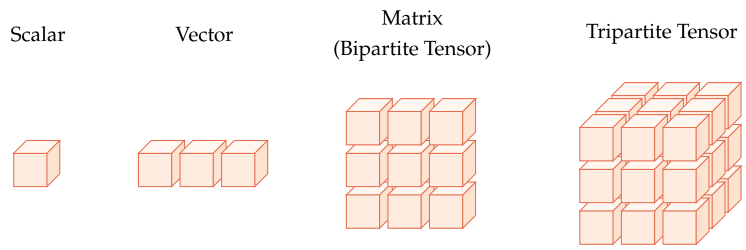

Figure 2.1: From scalars and arrays to bipartite and tripartite tensors. A scalar is a single number, represented by a box in the figure. A vector in a finite-dimensional vector space can be understood as one-dimensional array of scalars. A bipartite tensor is a two-dimensional array, a tripartite tensor is a three dimensional array. Similarly, a n-partite tensor is a n-dimensional array. Of course this is only possible if a basis is chosen in the vectorspaces.

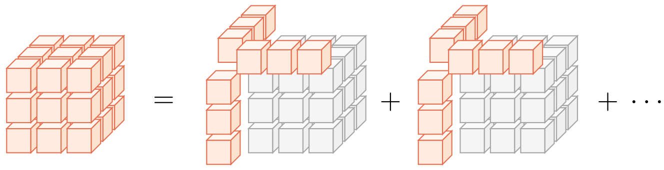

Figure 2.2: A tensor admits a decomposition into a sum of elementary tensors, where each elementary tensor (gray) is specified by a vector for every axis. For tripartite tensors, an elementary tensor is specified by three vectors.

decompose a tensor is known as the tensor rank. This parameter in particular serves as a complexity measure for the tensor.

Tensors frequently possess an inherent *positivity structure*. Positive tensors form a subset of tensors that exhibit additional properties typical of positive elements. For instance, positive tensors retain their positivity when multiplied by a positive factor or when combined with another positive tensor. Note that the global positivity of a tensor typically does not manifest in the individual local elements, i.e., the elementary tensors. Consequently, this absence of local reflection does not guarantee the overall positivity of the global tensor.

In this chapter, we present a framework to decompose positive and invariant tensors. We will utilize this framework to show the following results in the following chapters of this part:

- Positive tensors describe certain correlation scenarios within quantum information theory. Particularly, the count of elementary tensors in a positive tensor decomposition serves as a measure of correlation inherent in multipartite probability distributions.
- When subjected to approximations, positive and invariant tensor decomposition methods are susceptible to instability. Specifically, the count of elementary tensors may diminish when tolerating minor approximation errors. This phenomenon entails instabilities in numerical optimization processes for tensors.
- The tensor decomposition framework gives rise to a novel framework to decompose multipartite, positive polynomials.

2.1 Basic definitions

In the following, we introduce the basic definitions of the tensor product in vector spaces and the notion of a positivity structure on tensors.

Throughout this thesis, we make use of the braket notation for vectors. In particular, we denote an element in a $\mathbb{C}$-vector space[1] $\mathcal{V}$ by

1: We will also often make use of $\mathbb{R}$-vector spaces. The construction of the tensor product space does not rely on the specific choice of the ground field.

$$|v\rangle \in \mathcal{V}$$

with its dual element is given by $\langle v| \in \mathcal{V}^*$. Applying a linear operator $A : \mathcal{V} \to \mathcal{W}$ to $|v\rangle$ is denoted by $A\,|v\rangle$. Applying an element in the dual vector space $\langle w| \in \mathcal{V}^*$ to an element $|v\rangle \in \mathcal{V}$ is denoted by $\langle w \,|\, v\rangle$.

We denote the standard basis of $\mathbb{C}^d$ by $|1\rangle, \ldots, |d\rangle$.[2] If $\mathcal{V} = \mathrm{Mat}_d(\mathbb{C})$, then the standard basis is given by elements $|i\rangle\langle j|$ for $i, j = 1, \ldots, d$. Moreover, for a matrix $A \in \mathrm{Mat}_d(\mathbb{C})$ the entry at position (i, j) is determined by

2: In quantum information, it is often customary to begin counting from 0. Here, we adopt the convention of counting from 1 to d for the sake of readability.

$$\mathrm{tr}\left(A\,|j\rangle\langle i|\,\right) = \langle i|\,A\,|j\rangle\,.$$

2.1.1 The tensor product

Following the intuition of tensors depicted in Figure 2.1, the tensor product of the finite-dimensional vector spaces $\mathbb{C}^d \otimes \mathbb{C}^d$ is the vector space spanned by the basis vectors

$$|j_1, j_2\rangle \quad \text{for } j_1, j_2 = 1, \ldots, d$$

and the n-partite tensor product space $\left(\mathbb{C}^d\right)^{\otimes n}$ is spanned by the basis vectors

$$|j_1, \ldots, j_n\rangle \quad \text{for } j_1, \ldots, j_n = 1, \ldots, d$$

We now give a brief overview of the construction of the tensor product for arbitrary vector spaces. Let $\mathcal{V}$, $\mathcal{W}$ be two vector spaces. Consider the vector space $\mathcal{Q}$ spanned by the basis vectors

There are also alternative constructions of the tensor product using the universal property. For details on the different approaches, we refer to [79, Chapter 16].

$$(|v\rangle, |w\rangle) \in \mathcal{V} \times \mathcal{W}.$$

Note that $\mathcal{Q}$ is always infinite-dimensional[3] as for $|v_1\rangle \neq |v_2\rangle$, the elements $(|v_1\rangle, |w\rangle)$ and $(|v_2\rangle, |w\rangle)$ are linearly independent, even if $|v_1\rangle = \lambda |v_2\rangle$. Consequently, we need to consider a specific subspace of $\mathcal{Q}$ to define the tensor product $\mathcal{V} \otimes \mathcal{W}$—a subspace where examples like the one above are linearly dependent.

3: Of course assuming that $\mathcal{V}, \mathcal{W}$ are nontrivial.

Let $\mathcal{L} \subseteq \mathcal{Q}$ be the subspace spanned by the following elements

$$\begin{aligned} (|v_1\rangle + |v_2\rangle, |w\rangle) - (|v_1\rangle, |w\rangle) - (|v_2\rangle, |w\rangle) \\ (|v\rangle, |w_1\rangle + |w_2\rangle) - (|v\rangle, |w_1\rangle) - (|v\rangle, |w_2\rangle) \\ (\lambda |v\rangle, |w\rangle) - \lambda(|v\rangle, |w\rangle) \\ (|v\rangle, \lambda |w\rangle) - \lambda(|v\rangle, |w\rangle) \end{aligned} \tag{2.1}$$

for every $\lambda \in \mathbb{C}$, $|v\rangle, |v_1\rangle, |v_2\rangle \in \mathcal{V}$, $|w\rangle, |w_1\rangle, |w_2\rangle \in \mathcal{W}$. This space allows us to construct the tensor product space of $\mathcal{V}$ and $\mathcal{W}$.

Definition 2.1.1 (The tensor product space)
The tensor product space of $\mathcal{V}$ and $\mathcal{W}$ is defined by

$$\mathcal{V} \otimes \mathcal{W} := \mathcal{Q}/\mathcal{L}$$

where $\mathcal{Q}/\mathcal{L}$ is the quotient of $\mathcal{Q}$ by $\mathcal{L}$. The representatives in $\mathcal{Q}/\mathcal{L}$ of elements $(|v\rangle, |w\rangle) \in \mathcal{Q}$ are denoted by

$$|v\rangle \otimes |w\rangle .$$

The quotient space is defined as follows: Every subspace $\mathcal{U} \subseteq \mathcal{V}$ of a vector space gives rise to an equivalence relation

$$x \sim y \iff x - y \in \mathcal{U}$$

The set $\mathcal{V}/\mathcal{U}$ is defined by all equivalence classes induced by $\sim$. These equivalence classes define themselves a vector space. Intuitively, the quotient space arises by identifying all elements in $\mathcal{U}$ to be zero.

Note that according to Equation (2.1), the tensor product of vectors is bilinear, i.e.

$$|v\rangle \otimes \big(|w_1\rangle + \lambda |w_2\rangle \big) = |v\rangle \otimes |w_1\rangle + \lambda |v\rangle \otimes |w_2\rangle .$$

This holds true for every $|v\rangle, |w_1\rangle, |w_2\rangle$ and λ, and similarly for the first component.

While this construction of $\mathcal{V} \otimes \mathcal{W}$ is very abstract and non-constructive, the following proposition elucidates the behavior of the tensor product, particularly showcasing the intuitive properties of tensor products of $\mathbb{C}^d$:

Proposition 2.1.1
Let $\mathcal{V}$ and $\mathcal{W}$ be vector spaces with bases $\{|v_i\rangle\}_{i\in\mathcal{I}}$ and $\{|w_j\rangle\}_{j\in\mathcal{J}}$ respectively. The tensor product space

$$\mathcal{V}\otimes\mathcal{W}$$

is the vector space spanned by the basis vectors:

$$\big\{\, |v_i\rangle \otimes |w_j\rangle \,:\, (i,j)\in\mathcal{I}\times\mathcal{J}\big\}.$$

Proposition 2.1.1 makes apparent that the vector space dimension is multiplicative, i.e.

$$\dim(\mathcal{V}\otimes\mathcal{W}) = \dim(\mathcal{V})\cdot\dim(\mathcal{W}).$$

If $\mathcal{V},\mathcal{W}=\mathbb{C}^d$, then the standard basis vector define a basis on $\mathbb{C}^d\otimes\mathbb{C}^d$ given by $|i\rangle\otimes|j\rangle$ for $i,j\in\{1,\dots,d\}$. We will denote these vectors by

$$|i,j\rangle := |i\rangle\otimes|j\rangle$$

for simplicity. Similarly, the basis of an n-partite tensor product space is spanned by all combinations of basis vectors of the local vector spaces.

In the case $\mathcal{V},\mathcal{W}=\mathbb{C}^d$, the bipartite tensor product can also be realized using the matrix space $\mathrm{Mat}_d(\mathbb{C})$. Specifically, identifying

$$|j_1,j_2\rangle := |j_1\rangle\langle j_2|\,,$$

every matrix $T\in\mathrm{Mat}_d(\mathbb{C})$ corresponds to a tensor $|T\rangle\in\mathbb{C}^d\otimes\mathbb{C}^d$ as follows:

$$|T\rangle = \sum_{j_1,j_2=1}^{d} \langle j_1|\,T\,|j_2\rangle\,|j_1,j_2\rangle$$

and vice versa. This correspondence reflects the representation of a bipartite tensor as a matrix, illustrated in Figure 2.1.

2.1.2 Positivity structures on tensor product spaces

Vector spaces often come equipped with a positivity structure. In the following, we give three examples of tensor product structures with positivity constraints.

For a vector space $\mathcal{V}$, we call $\mathcal{C}\subseteq\mathcal{V}$ a *positivity structure* if it satisfies

- If $|v\rangle\in\mathcal{C}$, then $\lambda\,|v\rangle\in\mathcal{C}$ for every $\lambda\geqslant 0$.
- If $|v\rangle\,,|w\rangle\in\mathcal{C}$, then $|v\rangle+|w\rangle\in\mathcal{C}$.

In other words, $\mathcal{C}$ is a convex cone, i.e. positive combinations of elements in $\mathcal{C}$ are again contained in $\mathcal{C}$.

In this thesis, we will study three concrete examples of positivity structures in tensor product vector spaces arising from different applications:

Multipartite probability distributions: Providing a discrete probability distribution merely requires specifying the probabilities of the outcomes. Specifically, if X is a random variable taking values $1, \dots, d$, we can associate the probabilities $P(X = j)$ with a vector $|T\rangle \in \mathbb{R}^d$ such that

$$P(X = j) = \langle j \mid T \rangle .$$

Extending this concept to probability distributions involving multiple random variables $X_1, \dots, X_n$, each ranging from 1 to d, the correspondence expands to

$$P(X_1 = j_1, \dots, X_n = j_n) = \langle j_1, \dots, j_n \mid T \rangle$$

for a tensor

$$|T\rangle \in \mathbb{R}^d \otimes \cdots \otimes \mathbb{R}^d \cong \mathbb{R}^{d^n}.$$

Tensors that represent probability distribution are *entrywise nonnegative*, i.e.

$$\langle j_1, \dots, j_n \mid T \rangle \geqslant 0.$$

This establishes a positivity structure within the multipartite tensor product space.

Multipartite mixed quantum states: Following the axioms of quantum mechanics, physical degrees of freedom are described by a pure quantum state — a vector $|\psi\rangle$ in a Hilbert space $\mathcal{H}$. The concrete choice of the Hilbert space $\mathcal{H}$ depends on the system; for instance, a fixed-in-space spin-$\frac{1}{2}$ particle is modeled by $\mathcal{H} = \mathbb{C}^2$.

The composition of multiple quantum systems is captured by the tensor product. For example, the joint system of n spin-$\frac{1}{2}$ particles is described via a state

$$|\psi\rangle \in \mathbb{C}^2 \otimes \cdots \otimes \mathbb{C}^2 \cong \mathbb{C}^{2^n}.$$

In practice, one often has access only to a part of the entire physical system. Instead of defining a wave function for the entire system, a complete description of the reduced system is given by *mixed states*, also called *density matrices*. These are described by a positive semidefinite (psd) operator $\rho \in \mathcal{B}(\mathcal{H})$ with $\operatorname{tr}(\rho) = 1$. As in the pure state picture, combining open quantum systems is accomplished through the tensor product. For instance, a single spin-$\frac{1}{2}$ particle is described by a psd 2×2 matrix $\rho \in \mathrm{Psd}_2(\mathbb{C})$, while an open system of n spin-$\frac{1}{2}$ particles is described by a state

$$\rho \in \mathrm{Her}_2(\mathbb{C}) \otimes \cdots \otimes \mathrm{Her}_2(\mathbb{C}) \cong \mathrm{Her}_{2^n}(\mathbb{C}) \;\text{ with }\; \rho \succcurlyeq 0.$$

Density matrices represent a positivity structure on the multipartite matrix tensor product space, as the space of psd matrices forms a convex cone.

Multivariate polynomials: Multivariate polynomials emerge as a tensor product structure of univariate polynomials. A univariate polynomial (in the variable x) is a linear combination of monomials x^k. In essence, the space of polynomials in x, $\mathbb{R}[x]$, is the vector space generated by the monomial basis $\{x^k : k \in \mathbb{N}\}$. When we combine monomial bases in two variables, we once again acquire a monomial basis:

$$\{x^k y^\ell : k, \ell \in \mathbb{N}\}.$$

These monomials span the space of bivariate polynomials $\mathbb{R}[x, y]$ representing polynomials of the form

$$p = \sum_{k,\ell=1}^{n} c_{k,\ell} x^k y^\ell.$$

where $c_{k,\ell} \in \mathbb{R}$. With Proposition 2.1.1 this demonstrates that

$$\mathbb{R}[x, y] = \mathbb{R}[x] \otimes \mathbb{R}[y].$$

Multivariate polynomials embody multiple positivity structures. One example are *nonnegative polynomials*, which are polynomials satisfying

$$p(x, y) \geqslant 0 \quad \text{for all } x, y \in \mathbb{R}.$$

Another example is the cone of sum-of-square polynomials, i.e. polynomials of the form

$$p = \sum_{t=1}^{r} q_t^2.$$

For a detailed discussion of positivity structures on multipartite polynomials, we refer to Chapter 5.

2.2 The building blocks to decompose tensors

Multipartite tensors are, in general, very costly to represent. This follows from the exponential increase of the vector space dimension with respect to the number of local spaces.

Elementary tensors are specific elements that are easy to represent. An elementary tensor in $\left(\mathbb{C}^d\right)^{\otimes n}$

$$|T\rangle = |v^{[1]}\rangle \otimes \cdots \otimes |v^{[n]}\rangle$$

is specified by n vectors in $\mathbb{C}^d$. Therefore, we only need $n \cdot d$ scalars to describe an elementary tensor. Every tensor admits a decomposition into elementary tensors, called a *tensor decomposition*, i.e.

$$|T\rangle = \sum_{\alpha=1}^{r} |v_\alpha^{[1]}\rangle \otimes \cdots \otimes |v_\alpha^{[n]}\rangle$$

The minimal parameter r realizing such a decomposition, called the *tensor rank* of $|T\rangle$, is a measure of the cost of describing the tensor, as it enables representing the tensor using only $r \cdot n \cdot d$ scalars. Consequently, tensors with a low tensor rank can be efficiently represented.

There are many variants of tensor decompositions with other summation geometries, local positivity constraints, or local invariance constraints.

For example, one can decompose a tensor via a cyclic arrangement of indices

$$|T\rangle = \sum_{\alpha_1,\ldots,\alpha_n=1}^{r} |v^{[1]}_{\alpha_1,\alpha_2}\rangle \otimes |v^{[2]}_{\alpha_2,\alpha_3}\rangle \cdots \otimes |v^{[n]}_{\alpha_n,\alpha_1}\rangle\,.$$

This decomposition is known as the matrix product state (MPS) decomposition, and is widely applied in quantum many-body physics [89, 90, 30, 3].

Introducing *local symmetry constraints* such as

$$|v^{[i]}_{\alpha,\beta}\rangle = |v^{[j]}_{\alpha,\beta}\rangle \qquad \text{for every } i,j \in \{1,\ldots,n\} \tag{2.2}$$

gives rise to a *global symmetry* constraint. Concretely, Equation (2.2) leads to *translational invariance* of the $|T\rangle$, i.e. invariance under translations of the local tensor factors.

Introducing *local positivity constraints* such as

$$\langle j\,|\,v^{[i]}_{\alpha,\beta}\rangle \geqslant 0 \qquad \text{for all } \alpha, i, j \tag{2.3}$$

gives rise to a *global positivity constraint*. Concretely, Equation (2.3) guarantees that the global tensor $|T\rangle$ is entrywise nonnegative.

In the following we review a framework for decomposing positive and invariant tensors based on weighted simplicial complexes. Weighted simplicial complexes specify a geometry in the decomposition (i.e. a specific arrangement of summation indices). Equipping this geometry with an additional symmetry constraint via a group action on the weighted simplicial complex (WSC) will be the basic building block to define invariant and positive decompositions. The idea of this framework is based on [37] and has been used since then in [38, 39, 74].

In Section 2.2.1 and Section 2.2.2 we introduce the basic machinery of weighted simplicial complexes and group actions. Finally, in Section 2.2.3, we will present numerous examples of WSC that will appear throughout this thesis.

Throughout this thesis, we denote the set $\{1,\ldots,n\}$ by $[n]$.

2.2.1 Weighted simplicial complexes

A weighted simplicial complex is a mathematical structure that models relations between different objects, similar to graphs. More specifically, it consists of *vertices* representing the objects and *facets*, which connect the different vertices.

Definition 2.2.1 (Weighted simplicial complex)
A *weighted simplicial complex* (in short WSC) Ω on the set $[n]$ is a function

$$\Omega : \mathcal{P}([n]) \to \mathbb{N}$$

which satisfies the condition

$$S_1 \subseteq S_2 \subseteq [n] \quad \Longrightarrow \quad \Omega(S_1) \;\text{ divides }\; \Omega(S_2)$$

If $\Omega(S) \in \{0,1\}$ for every $S \subseteq [n]$, we call Ω a *simplicial complex*.

A subset $S \subseteq [n]$ such that $\Omega(S) \neq 0$ is termed a *simplex* of Ω. We assume that for every $i \in [n]$, the set $\{i\}$ is considered a simplex, which we call a *vertex* of Ω.

If S is maximal with respect to inclusion[4], we call it a *facet*. We denote the set of facets by

4: i.e. for every $T \supsetneq S$, we have $\Omega(T) = 0$

$$\mathcal{F} := \{F \subseteq [n] : F \text{ facet of } \Omega\}.$$

Moreover, we define the set of facets on $\{i\}$ by

$$\mathcal{F}_i := \{F \in \mathcal{F} : i \in F\}$$

The sets $\mathcal{F}$ and $\mathcal{F}_i$ will play a central role for defining tensor decomposition.[5]

5: We refer to Section 2.3 for definitions and examples of these tensor decompositions.

Restricting the function Ω to $\mathcal{F}$ and $\mathcal{F}_i$ makes these sets into multisets[6], which we denote by $\widetilde{\mathcal{F}}$ and $\widetilde{\mathcal{F}}_i$. For simplicity, we will treat these multisets analogously to sets. Therefore, for any facet F, the value $\Omega(F)$ represents the multiplicity of F in the WSC.

6: A *multiset* with elements in A is a function $m : A \to \mathbb{N}_+$, where $\mathbb{N}_+$ is the set of positive natural numbers. Each element $a \in A$ is contained in the multiset precisely $m(a)$ times.

Note that a WSC is a special type of multihypergraph [20], where each simplex is contained within a facet, and the multiplicities of the simplices adhere to the condition in Definition 2.2.1. Consequently, a WSC can be understood as a properly structured multihypergraph. We refer to the examples in Section 2.2.3 to elucidate this analogy further.

2.2.2 Group actions on weighted simplicial complexes

In the following, we introduce the concept of a group acting on a WSC Ω. Essentially, a group acting on a WSC consists of a permutation of vertices that is compatible with the structure of the WSC, similar to a graph-automorphism for graphs [20].

We say that a group G *acts on a set* X if there is a map

$$\alpha : G \times X \to X$$

that satisfies the *identity* and the *compatibility axiom*.[7] For convenience we will use the shorthand notation gx for $\alpha(g,x)$.

7: The identity axiom states that $\alpha(e,x) = x$ for all $x \in X$ and compatibility means that $\alpha(g, \alpha(h,x)) = \alpha(gh,x)$ for all $g,h \in G$ and $x \in X$.

In the following we define a some basic notions regarding group actions.

Definition 2.2.2 (G-invariant functions)
Let $f : X \to Y$ be a function and let G act on X. We say that f is *G-invariant* if

$$f(gx) = f(x) \quad \text{for all } x \in X, g \in G.$$

Intuitively, a G-invariant function remains the same under group actions on the X.

Moreover, for a function $f : X \to Y$ and a group action G on X, we define the shifted function

$$\begin{aligned} {}^{g}f \colon X &\to Y \\ x &\mapsto f(g^{-1}x). \end{aligned} \tag{2.4}$$

By definition, we have that ${}^{gh}f = {}^{g}\left({}^{h}f\right)$ as well as ${}^{e}f = f$. Moreover, the map

$${}^{g} : Y^X \to Y^X : f \mapsto {}^{g}f$$

is a bijection.[8]

8: Y^X indicates the space of functions from X to Y. These functions can alternatively be represented as a tuple indexed by X with values in Y, hence the notation.

We now introduce the notion of a group action on a WSC Ω.

Definition 2.2.3 (Group actions on WSC)
A group action of G on a WSC Ω is given by the following two parts:

- An action of G on the set $[n]$, such that Ω is G-invariant with respect to the action of G induced on $\mathcal{P}([n])$, i.e.

$$\Omega(gA) = \Omega(\{ga : a \in A\}) = \Omega(A).$$

  This group action then reduces to a group action on the set $\mathcal{F}$.
- A compatible refinement of the group action G to the multiset $\widetilde{\mathcal{F}}$. In other words, a group action G on $\widetilde{\mathcal{F}}$ such that the collapse map

$$c : \widetilde{\mathcal{F}} \to \mathcal{F}$$

  is G-linear, i.e. $c(gF) = gc(F)$ for all $g \in G$ and $F \in \widetilde{\mathcal{F}}$.

In simple terms, a group action on a WSC outlines how the vertices $[n]$ can be rearranged while preserving the original structure of the WSC. If the WSC contains, in addition, multi-facets, then the action of the group on the multi-facets needs to be refined since the rearrangement of the vertices does not uniquely determine the permutation of the facets anymore.

Now, we introduce two crucial properties of group actions essential for characterizing tensor decompositions based on WSC. Initially, we define these concepts for general group actions and subsequently refine the definitions for actions tailored to WSC.

Definition 2.2.4 (Free and blending group actions)
Let G be a group acting on a set X.

(i) G is *free* if the only stabilizer is the identity, i.e. $\mathrm{Stab}(x) = \{e\}$ for every $x \in X$, where

$$\mathrm{Stab}(x) := \{g \in G \colon gx = x\}.$$

(ii) G is *blending* if for every choice $g_1, \ldots, g_n \in G$ such that

$$\{g_1 1, g_2 2, \ldots, g_n n\} = [n]$$

there exists $g \in G$ such that $gi = g_i i$ for every $i \in [n]$.

Intuitively, a free group action consists of permutations that keep no element fixed. For example, if $X = \mathbb{Z}_k$ is the set of natural numbers $0, \dots, k-1$ with addition modulo k, then addition

$$\mathbb{Z}_k \times \mathbb{Z}_k \to \mathbb{Z}_k \colon (c, a) \mapsto c + a \bmod k$$

is a free group action on $\mathbb{Z}_k$.

Definition 2.2.5

A group action of G on a WSC Ω is called:

(i) *free* if the action of G is free on $\widetilde{\mathcal{F}}$;
(ii) *blending* if the action of G is blending on $\widetilde{\mathcal{F}}$;
(iii) *external* if for all $g \in G$ such that $gi = i$ we have that

$$gF = F \quad \text{for every } \; F \in \widetilde{\mathcal{F}}_i.$$

We now present examples of WSC Ω and group actions G on Ω, which satisfy various such properties.

2.2.3 Examples of weighted simplicial complexes with group actions

We now construct various examples of WSC that play a central role in this part of the thesis:[9]

9: These examples give rise to conceptually distinct types of tensor decompositions (c.f. Section 2.3.1), each exhibiting entirely different characteristics and behaviors (see, for instance, Chapter 4 or Chapter 5).

- The *simplex*
- The *line* with n vertices
- The *cycle* with n vertices
- The *double edge*

Furthermore, we illustrate instances of group actions on these WSC. A summary of which properties apply to the examples is provided in Table 2.1.

Table 2.1: Which properties of Definition 2.2.5 are satisfied for the examples of (Ω, G)? This table shows when the simplex with full symmetry group, the line and the double edge with the cyclic group, and the cycle with the cyclic group are free, blending, or external.

	free	**blending**	**external**
(Σ_n, C_n)	no	yes	yes
(Λ_n, C_2)	yes (n odd) no (n even)	yes ($n \leqslant 3$) no ($n \geqslant 4$)	yes (n even) no (n odd)
(Θ_n, C_n)	yes	no	yes
(Δ, C_2)	yes	yes	yes

Example 2.2.1 (The simplex)

The simplicial complex $\Omega = \Sigma_n$ that maps each subset of $[n]$ to 1 is called the *simplex*. In particular, this WSC contains precisely one facet

$$\mathcal{F} := \{[n]\}.$$

For $n = 5$, the hypergraph corresponding to Σ_n is illustrated in Figure 2.3. It is worth noting that any group action on $[n]$ results in a trivial group action on $\widetilde{\mathcal{F}}$, thereby defining a group action on Σ_n.
Note that the action of the full permutation group S_n on $[n]$ is blending. Moreover, the trivial group action $G = \{e\}$ is the only free group action, and every group action on Σ_n is also external.

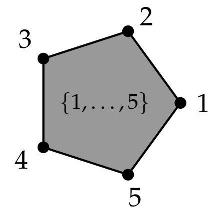

Figure 2.3: The simplex for $n = 5$ with its 5 vertices and its single facet $\{1,\ldots,5\}$ connecting all vertices.

Example 2.2.2 (The line with n vertices)

For $n \geqslant 1$, the *line with n vertices* is the simplicial complex $\Omega = \Lambda_n$ given by the graph shown in Figure 2.4. Specifically, the set of facets is given by

$$\mathcal{F} = \widetilde{\mathcal{F}} := \{\{1,2\},\{2,3\},\ldots,\{n-1,n\}\}$$

and therefore consists of $n-1$ elements. The only non-trivial group action on Λ_n is the cyclic group with two elements $G = C_2$. Here, the generator reverses the order of the vertices, meaning that vertex i is mapped to vertex $n+1-i$. This action is free if and only if n is odd.[10] Moreover, it is blending if and only if $n \leqslant 3$, and it is an external group action if and only if n is even.[11]

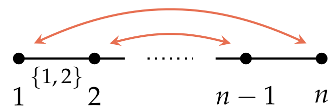

Figure 2.4: The line with n vertices. Every facet connects two neighboring vertices. The arrows in orange illustrate the only non-trivial group action C_2 on Λ_n which reflects the vertices.

10: If n is even, the middle edge is a fixed point of the action.

11: If n is odd, the group action keeps the middle vertex fixed but permutes the two edges connected to it.

Example 2.2.3 (The cycle with n vertices)

For $n \geqslant 3$, the *cycle with n vertices* is the simplicial complex $\Omega = \Theta_n$ corresponding to the graph shown in Figure 2.5. Specifically, the set of n facets is given by

$$\mathcal{F} = \{\{1,2\},\{2,3\},\ldots,\{n-1,n\},\{n,1\}\}.$$

One group action on Θ_n is the cyclic group C_n that is defined by the generator

$$\tau : \; : i \mapsto i+1 \mod (n+1).$$

Translations of vertices induce a group action on $\mathcal{F}$ (see also Figure 2.5). C_n is a free and external group action on Θ_n for every n, but it is not blending.

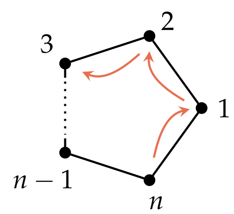

Figure 2.5: The cycle with n vertices. A vertex is characterized by the set of vertices which are contained in it. The arrows in orange illustrate the group action C_n on Λ_n, which is a translation of the facets.

Example 2.2.4 (The double edge)

A WSC can include multiple facets containing the same vertices. A basic example showing this property is the double edge Δ. It comprises two vertices $\{1,2\}$ and its multiset of facets is given by $\widetilde{\mathcal{F}} := \{\mathfrak{a},\mathfrak{b}\}$ where $1,2 \in \mathfrak{a}$ and $1,2 \in \mathfrak{b}$. The double edge is depicted in Figure 2.6.
Note that the single edge corresponds to the WSC $\Lambda_2 = \Sigma_2$. While Σ_2 has no non-trivial free group action, there is a non-trivial free group action on the double edge. Let $C_2 = \{e,s\}$ be the cyclic group with two elements. According to Definition 2.2.3, a group action on Δ is a refinement on the level of multi-sets. If $s\mathfrak{a} = \mathfrak{b}$, i.e. C_2 flips the edges, then C_2 is a free group action on Δ. This action is illustrated in Figure 2.6.

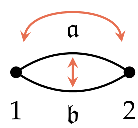

Figure 2.6: The double edge Δ. Its multiset of facets is given by $\{\mathfrak{a},\mathfrak{b}\}$, where both facets contain both vertices.

2.3 Positive and invariant decompositions

In the following, we define the notion of an (Ω, G)-*decomposition*, a tensor decomposition based on a WSC Ω, and a group action G on Ω.

Let $\mathcal{V}_1, \ldots, \mathcal{V}_n$ be vector spaces, where we call $\mathcal{V}_i$ the *local vector space* at site i, and define the *global vector space* as

$$\mathcal{V} := \mathcal{V}_1 \otimes \mathcal{V}_1 \otimes \cdots \otimes \mathcal{V}_n$$

where $\otimes$ denotes the algebraic tensor product. For this reason, every element in $\mathcal{V}$ is a *finite* sum of elementary tensors. Note that these vector spaces do not have to be finite-dimensional in general; in Chapter 5, we study the example of the infinite-dimensional vector space $\mathcal{V}_i = \mathbb{R}[x]$.

Note that any group action G on $[n]$ induces a linear action on $\mathcal{V}$ by permuting the tensor factors, if $\mathcal{V}_{gi} = \mathcal{V}_i$. More precisely, we consider the representation

$$\rho : G \to \mathrm{GL}(\mathcal{V})$$

where the action on the elementary tensors is given by[12]

$$\rho(g)\left[|v^{[1]}\rangle \otimes |v^{[2]}\rangle \otimes \cdots \otimes |v^{[n]}\rangle\right] = |v^{[g1]}\rangle \otimes |v^{[g2]}\rangle \otimes \cdots \otimes |v^{[gn]}\rangle$$

12: This defines a unique action on $\mathcal{V}$ by linearity.

For convenience, we will write $g\,|v\rangle$ as a shorthand notation for $\rho(g)\,|v\rangle$. Further, note that we will assume throughout this work that $\mathcal{V}_i = \mathcal{V}_j$; therefore, every group action G on $[n]$ induces an action on $\mathcal{V}$. All results presented also apply for different $\mathcal{V}_i$ respecting the symmetry constraints.

Note that the set of all functions Y^X can be written up as a tuple if X is finite. For example, if $\mathcal{I}$ is a finite index set, then for a WSC Ω, the set

$$\mathcal{I}^{\widetilde{\mathcal{F}}}$$

can be understood as a set of tuples

$$(\alpha_F)_{F \in \widetilde{\mathcal{F}}}$$

where every entry is indexed by a facet $F \in \widetilde{\mathcal{F}}$ and takes values in $\mathcal{I}$.[13]

13: Each viewpoint has its own advantages. While tuples are often used for specific examples, the functional approach proves advantageous in the general setting due to its greater flexibility and reduced technical complexity.

For $i \in \mathcal{I}$, the set

$$\mathcal{I}^{\widetilde{\mathcal{F}}_i}$$

can be analogously understood as tuples, but now only indexed by facets containing the vertex i. Representing α in the functional way, allows to define for $\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}$ the restriction

$$\alpha_{|_i} := \alpha_{|\widetilde{\mathcal{F}}_i} \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}.$$

2.3.1 Invariant tensor decompositions and ranks

We define the notion of an (Ω, G)-*decomposition*. Afterwards, we provide explicit examples of these decompositions for the simplex, the line, the cycle, and the double edge.

In essence, we introduce a sum of elementary tensors, where the local vectors of these tensors possess multiple summation indices. The organization of these indices is mirrored by WSC, such that each facet of the WSC corresponds to one summation index. Furthermore, the group action from G on the WSC introduces a symmetry within the elementary tensors, according to the arrangement of summation indices.

Definition 2.3.1 ((Ω, G)-decomposition)
Let $|v\rangle \in \mathcal{V}$. An (Ω, G)-*decomposition* of $|v\rangle$ is given by a family of local vectors

$$\left(|v^{[i]}_{\beta}\rangle \right)_{\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}}$$

for every $i \in [n]$ with $|v^{[i]}_{\beta}\rangle \in \mathcal{V}_i$, satisfying the following:

- Decomposing $|v\rangle$, i.e.

$$|v\rangle = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} |v^{[1]}_{\alpha_{|1}}\rangle \otimes \cdots \otimes |v^{[n]}_{\alpha_{|n}}\rangle \tag{2.5}$$

- Invariance: For every $i \in [n]$, $g \in G$, and $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$, we have

$$|v^{[gi]}_{{}^{g}\beta}\rangle = |v^{[i]}_{\beta}\rangle$$

  where ${}^{g}\beta$ is defined in Equation (2.4).

The smallest cardinality of the index set $\mathcal{I}$ among all possible (Ω, G)-decompositions is called the (Ω, G)-*rank* of $|v\rangle$, denoted by

$$\mathrm{rank}_{(\Omega,G)}(|v\rangle).$$

For convenience, we will call (Ω, G)-decompositions for trivial groups $G = \{e\}$ just Ω-decompositions and denote its corresponding Ω-rank by

$$\mathrm{rank}_{\Omega}(|v\rangle).$$

Intuitively, an (Ω, G)-decomposition is a way of decomposing $|v\rangle$ that is explicitly invariant. Specifically, we have that

$$\begin{aligned}
g\,|v\rangle &= \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} |v^{[g1]}_{\alpha_{|g1}}\rangle \otimes \cdots \otimes |v^{[gn]}_{\alpha_{|gn}}\rangle \\
&= \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} |v^{[g1]}_{{}^{g}\left((g^{-1}\alpha)_{|1}\right)}\rangle \otimes \cdots \otimes |v^{[gn]}_{{}^{g}\left((g^{-1}\alpha)_{|n}\right)}\rangle \\
&= \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} |v^{[1]}_{(g^{-1}\alpha)_{|1}}\rangle \otimes \cdots \otimes |v^{[n]}_{(g^{-1}\alpha)_{|n}}\rangle \\
&= \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} |v^{[1]}_{\alpha_{|1}}\rangle \otimes \cdots \otimes |v^{[n]}_{\alpha_{|n}}\rangle
\end{aligned}$$

where we use the invariance condition of Definition 2.3.1 in the third equality and the fact that $\alpha \mapsto {}^{g}\alpha$ is a bijection on $\mathcal{I}^{\widetilde{\mathcal{F}}}$ in the last equality.

We now present examples of (Ω, G)-decompositions by using the running examples of Section 2.2.3. For these specific choices of Ω and G, the decompositions reproduce well-known tensor decompositions and tensor network decompositions. For simplicity, we assume that $\mathcal{V}_i = \mathbb{C}^d$.

Example 2.3.1 (The standard and symmetric tensor decomposition)
Let Σ_n be the simplex with n vertices.[14] The Σ_n-decomposition is given by

$$|v\rangle = \sum_{\alpha=1}^{r} |v_\alpha^{[1]}\rangle \otimes \cdots \otimes |v_\alpha^{[n]}\rangle$$

which corresponds to the *standard tensor decompositions* [78, Section 2.4]. This decomposition consists of one summation index, which is reflected by the single facet in Σ_n.[15]
For the full permutation group S_n, the (Σ_n, S_n)-decomposition is is given by

$$|v\rangle = \sum_{\alpha=1}^{r} |v_\alpha\rangle \otimes \cdots \otimes |v_\alpha\rangle$$

i.e. all local vectors are identical. This follows from the invariance condition of Definition 2.3.1

$$|v_\alpha\rangle := |v_\alpha^{[i]}\rangle = |v_\alpha^{[j]}\rangle \quad \text{for all } i, j \in [n].$$

This decomposition is known as the *symmetric tensor decomposition* [78, Section 2.4]. The corresponding rank is called *symmetric rank*.

14: See Example 2.2.1.

15: In other words, since $|\widetilde{\mathcal{F}}| = 1$, every function $\alpha : \widetilde{\mathcal{F}} \to \mathcal{I}$ is characterized by a value $\alpha \in \mathcal{I}$.

Example 2.3.2 (Matrix Product States I)
For $n \geqslant 3$, let Θ_n be the cycle with n vertices.[16] The Θ_n-decomposition is given by

$$|v\rangle = \sum_{\alpha_1,\ldots,\alpha_n=1}^{r} |v_{\alpha_1,\alpha_2}^{[1]}\rangle \otimes |v_{\alpha_2,\alpha_3}^{[2]}\rangle \otimes \cdots \otimes |v_{\alpha_n,\alpha_1}^{[n]}\rangle \tag{2.6}$$

Here, the index α_i represents the entry of the $(\alpha_F)_{F\in\mathcal{I}^{\widetilde{\mathcal{F}}}}$ tuple indexed by $F = \{i, i+1\}$ where addition is modulo $n+1$. Since all vertices are contained in two facets, we have two summation indices for every local vector.
This decomposition is known as the MPS decomposition. Usually MPS are presented parametrizing the coefficients of the tensor in a fixed basis, i.e. finding a description of $v_{j_1,\ldots,j_n}$ for

$$|v\rangle = \sum_{j_1,\ldots,j_n=1}^{d} v_{j_1,\ldots,j_n} |j_1,\ldots,j_n\rangle .$$

To obtain this representation, let $A_j^{[i]} \in \mathrm{Mat}_r(\mathbb{C})$ with

$$\langle\alpha| A_j^{[i]} |\beta\rangle = \langle j \,|\, v_{\alpha,\beta}^{[i]}\rangle .$$

16: See Example 2.2.3.

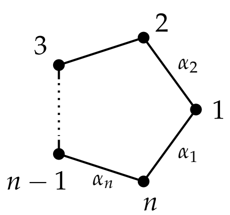

Figure 2.7: The cycle with n vertices and its correspondence to the summation indices in the Θ_n-decomposition. The connecting facets of the WSC represent all summation indices. For example the facets containing vertex 1 represent the summation indices α_1 and α_2 which are associated to the local vectors in the first local space.

Then, we obtain

$$|v\rangle = \sum_{j_1,\dots,j_n=1}^{d} \operatorname{tr}\left(A^{[1]}_{j_1} \cdot A^{[2]}_{j_2} \cdots A^{[n]}_{j_n}\right) |j_1,\dots,j_n\rangle$$

which corresponds to a MPS decomposition with closed boundary conditions. Specifically, $\operatorname{rank}_{\Theta_n}(|v\rangle)$ corresponds to the *bond dimension* of the MPS. Tensor networks are often illustrated using a diagrammatic calculus. We refer to Figure 2.8 for the tensor network diagram of the MPS. For more details on the diagrammatic formalism of tensor networks, we refer to [89, 22, 30].

Let in addition $G = C_n$ be the cyclic group. The (Θ_n, C_n)-decomposition is given by

$$|v\rangle = \sum_{\alpha_1,\dots,\alpha_n=1}^{r} |v_{\alpha_1,\alpha_2}\rangle \otimes |v_{\alpha_2,\alpha_3}\rangle \otimes \cdots \otimes |v_{\alpha_n,\alpha_1}\rangle \tag{2.7}$$

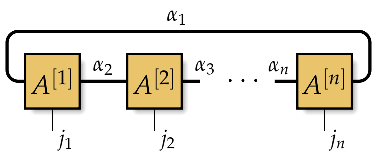

Figure 2.8: The tensor network diagram of the MPS with closed boundary conditions. The thick lines correspond to the matrix contraction, the thin open lines represent the local physical systems of dimension d.

The local vectors in Equation (2.7) are all the same, in contrast to Equation (2.6). This is guaranteed by the invariance condition of Definition 2.3.1. Note that this decomposition is called the translational invariant (ti) MPS, defined as

$$|v\rangle = \sum_{j_1,\dots,j_n=1}^{r} \operatorname{tr}\left(A_{j_1} \cdot A_{j_2} \cdots A_{j_n}\right) |j_1,\dots,j_n\rangle$$

with

$$\langle\alpha|\, A_j\, |\beta\rangle = \langle j\,|\, v_{\alpha,\beta}\rangle\,.$$

Example 2.3.3 (Matrix Product States II)

Let Λ_n be the line with n vertices.[17] The Λ_n-decomposition is given by

17: See Example 2.2.2.

$$|v\rangle = \sum_{\alpha_1,\dots,\alpha_{n-1}=1}^{r} |v^{[1]}_{\alpha_1}\rangle \otimes |v^{[2]}_{\alpha_1,\alpha_2}\rangle \otimes \cdots \otimes |v^{[n-1]}_{\alpha_{n-2},\alpha_{n-1}}\rangle \otimes |v^{[n]}_{\alpha_{n-1}}\rangle\,.$$

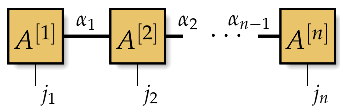

Figure 2.9: The tensor network diagram of the MPS with open boundary conditions. The thick lines correspond to the matrix contraction, the thin open lines represent the local physical systems of dimension d.

In this context, the index α_i represents the entry of the $(\alpha_F)_{F\in\mathcal{I}^{\tilde{\mathcal{F}}}}$ tuple indexed by $F = \{i, i+1\}$. The vertices 1 and n are only included in one facet, reflecting that these local tensors possess only one summation index.

This decomposition corresponds to an MPS decomposition with open boundary conditions, as there is no connection between the last and the first local space. This decomposition can be expressed as a tensor network as

$$|v\rangle = \sum_{j_1,\dots,j_n=1}^{r} A^{[1]}_{j_1} \cdot A^{[2]}_{j_2} \cdots A^{[n-1]}_{j_{n-1}} \cdot A^{[n]}_{j_n} |j_1,\dots,j_n\rangle$$

where $A^{[i]}_j \in \operatorname{Mat}_r(\mathbb{C})$ for $i \in \{2,\dots,n-1\}$, $A^{[1]}_j \in \mathbb{C}^{1\times r}$ and $A^{[n]}_j \in \mathbb{C}^r$. We refer to Figure 2.9 for a representation of this decomposition via tensor network diagrams.

Example 2.3.4 (The double edge decomposition)
Let Δ be the double edge.[18] The corresponding Δ-decomposition is given by

$$|v\rangle = \sum_{\alpha,\beta=1}^{r} |v^{[1]}_{\alpha,\beta}\rangle \otimes |v^{[2]}_{\beta,\alpha}\rangle \,.$$

Moreover, the (Δ, C_2)-decomposition is given by[19]

$$|v\rangle = \sum_{\alpha,\beta=1}^{r} |v_{\alpha,\beta}\rangle \otimes |v_{\beta,\alpha}\rangle \,.$$

Therefore, the the double edge decomposition can be viewed as an MPS decomposition when $n = 2$.

18: See Example 2.2.4.

19: Note that here the order of the summation indices is important. Specifically, the decomposition

$$v = \sum_{\alpha,\beta=1}^{r} |v_{\alpha,\beta}\rangle \otimes |v_{\alpha,\beta}\rangle$$

corresponds to the double edge where the edges are not swapped via the action of C_2. In this situation, the group action is not free (see Definition 2.2.4).

2.3.2 Positive tensor decompositions

We now introduce invariant tensor decompositions tailored for tensors that satisfy a positivity constraint. Specifically, these decompositions inherently maintain positivity, ensuring the global tensor remains positive under local perturbations. This is achieved by imposing additional constraints on the local vectors in the tensor decomposition.

Specifically, we introduce decompositions for tensors in the space

$$\mathcal{V} := \mathbb{R}^d \otimes \cdots \otimes \mathbb{R}^d \cong \mathbb{R}^{d^n},$$

i.e. every local space corresponds to $\mathbb{R}^d$. Moreover, we equip this space with a notion of positivity, namely *entrywise nonnegativity*.[20]

A tensor $|T\rangle \in \mathbb{R}^d \otimes \cdots \otimes \mathbb{R}^d$ is called *entrywise nonnegative*, if

$$\langle j_1, \ldots, j_n \,|\, T\rangle \geqslant 0 \quad \text{for every } j_1, \ldots, j_n \in \{1, \ldots, d\}.$$

Entrywise nonnegative tensors describe, for example, multi-partite probability distributions [80, 101]. For random variables $X_1, X_2, \ldots, X_n$ taking values in $\{1, \ldots, d\}$, the joint probability distribution is represented by a nonnegative tensor, specifically:

$$P(X_1 = j_1, \ldots, X_n = j_n) = \langle j_1, \ldots, j_n \,|\, T\rangle \,.$$

We utilize this correspondence to make statements about correlation scenarios via ranks of positive tensor decompositions.[21]

In the following, we describe two notions of locally positive tensor decompositions:

- the *nonnegative* decomposition
- the *positive semidefinite* decomposition

While the former employs entrywise nonnegative vectors as fundamental components, the latter utilizes psd matrices as local constitutents.

These two decompositions extend well-known matrix factorizations, including the nonnegative matrix factorization [31, 124, 18, 113], the

20: We introduce similar positive decompositions for multipartite psd matrices in Section 2.3.5 and for positive polynomials in Chapter 5.

21: See Chapter 3 for details on this relation.

completely positive decomposition [12], and the positive semidefinite matrix factorization [49, 118, 112, 68]. We refer to Example 2.3.5 and Example 2.3.7 for further details.

The nonnegative tensor decomposition

In the following, we introduce the nonnegative (Ω, G)-decomposition. Intuitively, this decomposition builds upon the unconstrained (Ω, G)-decomposition but restricts the local vectors to be entrywise nonnegative.

Definition 2.3.2 (Nonnegative (Ω, G)-decompositions)
Let $|T\rangle \in \mathbb{R}^d \otimes \cdots \mathbb{R}^d$. A *nonnegative* (Ω, G)-decomposition of $|T\rangle$ consists of an (Ω, G)-decomposition of $|T\rangle$

$$|T\rangle = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} |T^{[1]}_{\alpha_{|1}}\rangle \otimes \cdots \otimes |T^{[n]}_{\alpha_{|n}}\rangle$$

such that[22]

$$|T^{[i]}_{\beta}\rangle \geqslant 0$$

for every $i \in [n]$ and $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$.
The smallest cardinality of the index set $\mathcal{I}$ among all nonnegative (Ω, G)-decompositions is called the *nonnegative* (Ω, G)*-rank* of $|T\rangle$. We denote it by

$$\text{nn-rank}_{(\Omega,G)}(|T\rangle).$$

22: For a vector $|v\rangle \in \mathbb{R}^d$, we write

$$|v\rangle \geqslant 0$$

if $\langle j \,|\, v\rangle \geqslant 0$ for every $j \in [d]$.

For convenience, we call a nonnegative (Ω, G)-decomposition for the trivial group $G = \{e\}$ just nonnegative Ω-decomposition and denote its corresponding nonnegative Ω-rank by

$$\text{nn-rank}_{\Omega}(|T\rangle).$$

Intuitively, a nonnegative (Ω, G)-decomposition ensures explicit invariance[23] and explicit entrywise nonnegativity, since

$$\langle j_1, \ldots, j_n \,|\, T\rangle = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \langle j_1 \,|\, T^{[1]}_{\alpha_{|1}}\rangle \cdots \langle j_n \,|\, T^{[n]}_{\alpha_{|n}}\rangle \geqslant 0.$$

23: i.e. every element that achieves an (Ω, G)-decomposition is automatically G-invariant, i.e. $g\,|T\rangle = |T\rangle$ (see the remarks of Definition 2.3.1).

Let us now review examples of the nonnegative decompositions for specific choices of WSC Ω and group action from G.

Example 2.3.5 (The simplex decomposition)
Let Σ_n be the simplex with n vertices.[24] The nonnegative Σ_n-decomposition is given by

24: See Example 2.2.1 for its definition.

$$|T\rangle = \sum_{\alpha=1}^{r} |T^{[1]}_{\alpha}\rangle \otimes \cdots \otimes |T^{[n]}_{\alpha}\rangle \,.$$

This is commonly referred to as the *nonnegative tensor decomposition*.
For $n = 2$, i.e. the single edge, this yields the *nonnegative matrix fac-*

torization. For an entrywise nonnegative matrix $M \in \mathrm{Mat}_d(\mathbb{R})$, the nonnegative matrix factorization is defined as

$$M = B_1 B_2^t,$$

where B_1, B_2 are $d \times r$ matrices with nonnegative entries and $(\,\cdot\,)^t$ is the matrix transposition. More specifically,

$$\langle j_1 | \, M \, | j_2 \rangle = \sum_{\alpha=1}^{r} \langle j_1 | \, B_1 \, | \alpha \rangle \cdot \langle j_2 | \, B_2 \, | \alpha \rangle$$

which agrees with a Σ_2-decomposition by identifying

$$\langle j_1, j_2 \, | \, T \rangle := \langle j_1 | \, M \, | j_2 \rangle \quad \text{and} \quad \langle j \, | \, T_\alpha^{[i]} \rangle := \langle j_i | \, B_i \, | \alpha \rangle \, .$$

In particular, the column dimension of B_1 and B_2 agree with the number of elementary tensors in the decomposition of $|T\rangle$.
If $G = S_n$, we obtain the *symmetric nonnegative decomposition*

$$|T\rangle = \sum_{\alpha=1}^{r} |T_\alpha\rangle \otimes \cdots \otimes |T_\alpha\rangle \, ,$$

i.e. all local vectors are identical.
For $n = 2$, this decomposition gives rise to the *completely positive (cp) matrix factorization*. For a matrix $M \in \mathrm{Mat}_d(\mathbb{R})$, this is defined as

$$M = BB^t,$$

where B is a $d \times r$ matrix. Again the minimal number of columns of B is precisely the number of elementary tensors in the tensor decomposition.

Example 2.3.6

Let Θ_n be the cycle with n vertices.[25] The nonnegative Θ_n-decomposition is given by

25: See Example 2.2.3 for its definition.

$$|T\rangle = \sum_{\alpha_1,\ldots,\alpha_n=1}^{r} |T^{[1]}_{\alpha_1,\alpha_2}\rangle \otimes |T^{[2]}_{\alpha_2,\alpha_3}\rangle \otimes \cdots \otimes |T^{[n]}_{\alpha_n,\alpha_1}\rangle \, .$$

This is also known as a *nonnegative MPS*, *stochastic MPS* [121], or *nonnegative tensor train decomposition* [58]. Similar to MPS, this is expressed in the computational basis as

$$|T\rangle = \sum_{j_1,\ldots,j_n=1}^{d} \mathrm{tr}\left(A^{[1]}_{j_1} \cdots A^{[n]}_{j_n}\right) |j_1, \ldots, j_n\rangle$$

using the same correspondence as in Example 2.3.2. For nonnegative tensor train decompositions the local matrices $A^{[i]}_j$ are in addition entrywise nonnegative.

The positive semidefinite decomposition

We now define the psd (Ω, G)-decomposition. Intuitively, this replaces the nonnegativity constraint of the local elements by positive semidefiniteness.

Definition 2.3.3 (Positive semidefinite (Ω, G)-decomposition)
A *positive semidefinite* (Ω, G)-decomposition consists of psd matrices

$$A_j^{[i]} \in \mathrm{Mat}_{\mathcal{I}^{\widetilde{\mathcal{F}}_i}}(\mathbb{C}) \quad \text{for every } i \in [n], j \in [d]$$

with the constraint that

$$\left(A_j^{[gi]}\right)_{g\beta, g\beta'} = \left(A_j^{[i]}\right)_{\beta,\beta'}$$

decomposing $|T\rangle$ as

$$\langle j_1, \ldots, j_n \,|\, T\rangle = \sum_{\alpha,\alpha' \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \left(A_{j_1}^{[1]}\right)_{\alpha_{|1}, \alpha'_{|1}} \cdots \left(A_{j_n}^{[n]}\right)_{\alpha_{|n}, \alpha'_{|n}}$$

The minimal cardinality of $\mathcal{I}$ among all psd (Ω, G)-decompositions of $|T\rangle$ is called the psd (Ω, G)-rank, denoted by

$$\text{psd-rank}_{(\Omega,G)}(|T\rangle).$$

If $G = \{e\}$, we call the decomposition again just psd Ω-decomposition and denote its corresponding rank

$$\text{psd-rank}_{\Omega}(|T\rangle).$$

Again, every tensor admitting a positive semidefinite (Ω, G)-decomposition is inherently G-invariant as well as inherently entrywise nonnegative.[26]

26: To see that a psd decomposition only gives rise to nonnegative tensors, we refer to Section 2.3.4.

Example 2.3.7 (The simplex decomposition)
Let Σ_n be the simplex with n vertices. The Σ_n-decomposition is given by

$$\langle j_1, \ldots, j_n \,|\, T\rangle = \sum_{\alpha_1,\alpha_2=1}^{r} \left(A_{j_1}^{[1]}\right)_{\alpha_1,\alpha_2} \cdots \left(A_{j_n}^{[n]}\right)_{\alpha_1,\alpha_2}.$$

This decomposition has been studied before in the context of quantum correlation and quantum communication scenarios [69].
If $n = 2$, the decomposition specializes to the *positive semidefinite matrix factorization* [49], which is defined as

$$\langle j_1, j_2 \,|\, T\rangle = \mathrm{tr}\left(\left(A_{j_1}^{[1]}\right) \cdot \left(A_{j_2}^{[2]}\right)^t\right).$$

If in addition $G = S_2$, this leads to the *completely positive semidefinite transpose (cpsdt)* decomposition [40], defined as

$$\langle j_1, j_2 \,|\, T\rangle = \mathrm{tr}\left(A_{j_1} A_{j_2}^t\right).$$

While the cpsdt decomposition looks similar to the *completely positive semidefinite (cpsd)* decomposition

$$\langle j_1, j_2 \,|\, T\rangle = \mathrm{tr}(A_{j_1} A_{j_2}),$$

it deviates significantly from it in its behavior. For example, the cpsd decomposition cannot be expressed as a tensor decomposition. We refer to [100] for details.

2.3.3 Inequalities of ranks

We now briefly review the relation of the different ranks, shown in [37] for (Ω, G)-decompositions and in [40] for the Σ_n- and the Λ_n-decompositions.

Lemma 2.3.1

Let $|T\rangle$ be a nonnegative tensor. Then, the following inequalities hold:

(i) $\text{rank}_{(\Omega,G)}(|T\rangle) \leqslant \text{nn-rank}_{(\Omega,G)}(|T\rangle)$
(ii) $\text{rank}_{(\Omega,G)}(|T\rangle) \leqslant \text{psd-rank}_{(\Omega,G)}(|T\rangle)^2$
(iii) $\text{psd-rank}_{(\Omega,G)}(|T\rangle) \leqslant \text{nn-rank}_{(\Omega,G)}(|T\rangle)$ if G is a free[27] action on Ω.

27: See Definition 2.2.5 for free group actions.

For the proof of this statement, we refer to [37, Corollary 37].

2.3.4 The structure tensor $|\Omega_r\rangle$

We now introduce, for every WSC Ω with n vertices, a corresponding n-partite tensor $|\Omega_r\rangle$ of (Ω, G)-rank r which inherits the geometry of Ω. This tensor facilitates a concise representation of (Ω, G)-decompositions with (Ω, G)-rank r, which we will use in the proofs of Theorem 3.1.2 and Theorem 3.2.1. Defining tensor (network) decompositions via structure tensors is a common approach in tensor decompositions without positivity constraints [29].

For the vector space $\mathbb{C}^{\mathcal{I}^{\tilde{\mathcal{F}}_i}}$, we consider the standard basis

$$\{\,|\beta\rangle\,\}_{\beta\in\mathcal{I}^{\tilde{\mathcal{F}}_i}}.$$

In other words, the basis vectors in system i are indexed by $|\beta_1, \ldots, \beta_k\rangle$, where k is the number of facets that contain the vertex i and $\beta_\ell \in \{1, \ldots, r\}$ for every $\ell \in [k]$.

Given a WSC Ω, we define

$$|\Omega_r\rangle := \sum_{\alpha\in\mathcal{I}^{\tilde{\mathcal{F}}}} |\alpha_{|_1}\rangle \otimes \cdots \otimes |\alpha_{|_n}\rangle \;\in\; \bigotimes_{i=1}^{n} \mathbb{C}^{r_i}$$

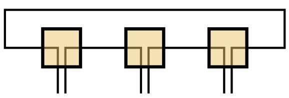

Figure 2.10: Tensor network representation of a matrix multiplication (MaMu)-tensor. The double output correspond to the space $\mathbb{C}^d \otimes \mathbb{C}^d$.

where $\mathcal{I} = \{1, \ldots, r\}$ and $r_i = |\mathcal{I}^{\tilde{\mathcal{F}}_i}|$.

For the n-cycle Θ_n, we obtain the n-fold *MaMu-tensor* (see Figure 2.10)

$$|\Theta_{n,r}\rangle = \sum_{\alpha_1,\ldots,\alpha_n=1}^{r} |\alpha_1, \alpha_2\rangle \otimes |\alpha_2, \alpha_3\rangle \otimes \cdots \otimes |\alpha_n, \alpha_1\rangle\,.$$

For the n-fold simplex Σ_n, we obtain the *unnormalized r-dimensional GHZ-state*

$$|\Sigma_{n,r}\rangle = \sum_{\alpha=1}^{r} |\alpha\rangle^{\otimes n}\,.$$

Note that every (Ω, G)-decomposition of (Ω, G)-rank at most r can be written as

$$|T\rangle = W^{[1]} \otimes \cdots \otimes W^{[n]} |\Omega_r\rangle \tag{2.8}$$

with

$$W^{[i]} := \sum_{\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}} |v_\beta^{[i]}\rangle \langle \beta| \,.$$

In this case, the G-invariance of $|v_\beta^{[i]}\rangle$ translates to

$$W^{[gi]} |{}^g\beta\rangle = W^{[i]} |\beta\rangle \,.$$

The nonnegative (Ω, G)-decomposition translates similarly, except for the additional constraint

$$\langle j| W^{[i]} |\beta\rangle \geqslant 0$$

for every $j \in \{1, \ldots, d\}$ and every $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$.

The psd (Ω, G)-decomposition translates to

$$\langle j_1, \ldots, j_n \,|\, T\rangle = \langle \Omega_r| \, A_{j_1}^{[1]} \otimes \cdots \otimes A_{j_n}^{[n]} \, |\Omega_r\rangle \tag{2.9}$$

where $A_j^{[i]}$ are the matrices of Definition 2.3.3. From Equation (2.9) it is also evident that $|T\rangle$ is a nonnegative tensor, since the matrices $A_j^{[i]}$ are psd.

Note that in all examples, the (Ω, G)-rank is determined by the minimal parameter r in the structure tensor that admits such a decomposition.

2.3.5 Positive matrix tensor decompositions

We now introduce two positive tensor decompositions for multipartite psd matrices, i.e. elements of the space

$$\mathrm{Mat}_d(\mathbb{C}) \otimes \mathrm{Mat}_d(\mathbb{C}) \otimes \cdots \otimes \mathrm{Mat}_d(\mathbb{C}) \cong \mathrm{Mat}_{d^n}(\mathbb{C}),$$

known as *separable decomposition* and *local purification form*. These decompositions can be perceived as generalizations of the positive decompositions for nonnegative tensors. A relation between their ranks shall be presented in Proposition 2.3.2.

We start with the definition of the separable (Ω, G)-decomposition.

Definition 2.3.4 (The separable (Ω, G)-decomposition)
Let $\rho \in (\mathrm{Mat}_d(\mathbb{C}))^{\otimes n}$. The separable (Ω, G)-decomposition is given by a family of matrices

$$\left(\rho_\beta^{[i]}\right)_{\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}}$$

for every $i \in [n]$ with $\rho_\beta \in \mathrm{Mat}_d(\mathbb{C})$, satisfying

$$\rho = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \rho^{[1]}_{\alpha_{|1}} \otimes \cdots \otimes \rho^{[n]}_{\alpha_{|n}}$$

with the additional constraints:

- Symmetry: $\rho^{[gi]}_{g\beta} = \rho^{[i]}_{\beta}$
- Positivity: ρ_β is psd.

The minimal cardinality of $\mathcal{I}$ among all separable (Ω, G)-decompositions is called the (Ω, G)-*rank*, denoted

$$\text{sep-rank}_{(\Omega,G)}(\rho).$$

If $G = \{e\}$ is the trivial group, we call the decomposition just separable Ω-decomposition and its rank just separable Ω-rank, denoted

$$\text{sep-rank}_{\Omega}(\rho).$$

Note that every ρ attaining a separable (Ω, G)-decomposition is inherently G-invariant, i.e. $g\rho = \rho$.[28] Moreover, matrices attaining a separable decomposition are separable by construction.[29]

28: For a proof of $g\rho = \rho$ for unconstrained (Ω, G)-decompositions, we refer to the remark after Definition 2.3.1.

29: A matrix $A \in \text{Mat}_d(\mathbb{C}) \otimes \text{Mat}_d(\mathbb{C})$ is called *separable* if it accepts a decomposition

$$A = \sum_{k=1}^{r} A_k^{[1]} \otimes A_k^{[2]}$$

with $A_k^{[i]} \succcurlyeq 0$ psd. Separable matrices model quantum states without entanglement.

Examples of separable (Ω, G)-decompositions are constructed similarly to those of nonnegative (Ω, G)-decompositions, with the only difference that nonnegative local vectors are replaced by psd local matrices. We also refer to [40, 37] for more examples.

The separable decomposition exclusively parametrizes separable matrices, a strict subset of psd matrices. Let us introduce a positive tensor decomposition covering all multipartite psd matrices: the (Ω, G)-purification form.

Definition 2.3.5 ((Ω, G)-purification)
For $\rho \in \text{Mat}_d(\mathbb{C})^{\otimes n}$, an (Ω, G)-*purification* is a factorization

$$\rho = L^\dagger L$$

where $L \in \text{Mat}_{\ell,d}(\mathbb{C})^{\otimes n}$ for some $\ell \in \mathbb{N}$ together with an (Ω, G)-decomposition of L.
The most efficient (Ω, G)-decompositions among all purifications L defines the *purification rank* of ρ, i.e.

$$\text{puri-rank}_{(\Omega,G)}(\rho) := \min_{\rho = L^\dagger L} \text{rank}_{(\Omega,G)}(L).$$

Again, if $G = \{e\}$ is the trivial group, we call the decomposition just Ω-purification and its corresponding rank Ω-purification rank, denoted

$$\text{puri-rank}_{\Omega}(\rho).$$

If ρ admits a (Ω, G)-purification, then it is automatically G-invariant as well as psd, as every matrix of the form $L^\dagger L$ is psd. Moreover, if G is a free group action on Ω, then every psd G-invariant matrix admits an (Ω, G)-purification[30].

30: We refer to [37, Theorem 20] for the proof.

We now present the example of the unconstrained Θ_n-decomposition and Θ_n-purification for multipartite matrices.

Example 2.3.8 (The Matrix Product Operator form)
Let $\Omega = \Theta_n$ be the cycle with n vertices.[31] The Θ_n-decomposition is given by

31: See Example 2.2.3 for the definition.

$$\rho = \sum_{\alpha_1,\ldots,\alpha_n=1}^{r} \rho^{[1]}_{\alpha_1,\alpha_2} \otimes \rho^{[2]}_{\alpha_2,\alpha_3} \cdots \otimes \rho^{[n]}_{\alpha_n,\alpha_1}$$

where $\rho_{\alpha,\beta} \in \mathrm{Mat}_d(\mathbb{C})$ are arbitrary matrices (not necessarily psd). This is known as the matrix product operator (MPO) form or matrix product density operator (MPDO) form [126, 130], also defined as

$$\rho = \sum_{j_1,k_1,\ldots,j_n,k_n=1}^{d} \mathrm{tr}\left(A^{[1]}_{j_1,k_1} \cdots A^{[n]}_{j_n,k_n}\right) |j_1,\ldots,j_n\rangle \langle k_1,\ldots,k_n|$$

where $A^{[i]}_{j,k} \in \mathrm{Mat}_r(\mathbb{C})$. The correspondence is given by

$$\langle\alpha| A^{[i]}_{j,k} |\beta\rangle = \langle j| \rho^{[i]}_{\alpha,\beta} |k\rangle\,.$$

See Figure 2.11 for a tensor network diagram.

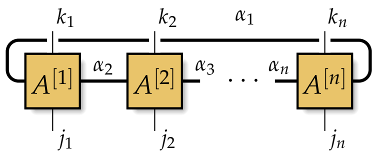

Figure 2.11: Tensor network diagram of the MPO. Thick lines (indexed by $\alpha_1,\ldots,\alpha_n$) correspond to the matrix contractions, the thin open lines (indexed by $j_1,k_1,\ldots,j_n,k_n$) represent the local physical systems of dimension d.

Example 2.3.9 (The locally purified MPDO form)
Let $\Omega = \Theta_n$ be the cycle with n vertices. The Θ_n-purification is given by

$$\rho = L^\dagger L$$

together with a Θ_n-decomposition of L,

$$L = \sum_{j_1,k_1,\ldots,j_n,k_n=1}^{d} \mathrm{tr}\left(B^{[1]}_{j_1,k_1} \cdots B^{[n]}_{j_n,k_n}\right) |j_1,\ldots,j_n\rangle \langle k_1,\ldots,k_n|\,. \qquad (2.10)$$

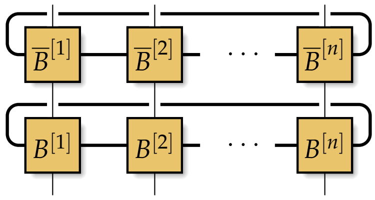

Figure 2.12: The tensor network diagram of the local purification MPDO form. The thick horizontal lines correspond to the matrix contraction of the tensor network and the thin lines to the matrix contractions $L^\dagger L$. The thin open lines represent the local physical systems of dimension d.

The locally purified form also admits a tensor network representation (see Figure 2.12).
This decomposition is known as the local purification form [41, 36] or the locally purified density operator (LPDO) form.

The separable (Ω, G)-decomposition and the (Ω, G)-purification are generalizations of the nonnegative and the psd (Ω, G)-decomposition, correspondingly, in the following way: For a tensor $|T\rangle \in \mathbb{C}^d \otimes \cdots \otimes \mathbb{C}^d$, let

$$\rho_{|T\rangle} := \sum_{j_1,\ldots,j_n=1}^{d} \langle j_1,\ldots j_n \mid T\rangle\, |j_1,\ldots,j_n\rangle \langle j_1,\ldots,j_n| \qquad (2.11)$$

be the diagonal embedding of $|T\rangle$ into a diagonal matrix. $\rho_{|T\rangle}$ is psd if and only if $|T\rangle$ is entrywise nonnegative.

We now state the correspondence between the ranks of nonnegative tensors and multipartite psd matrices. For a proof we refer to [37, Theorem 43].

Proposition 2.3.2 (Decompositions of diagonal matrices)
Let $|T\rangle$ be a nonnegative tensor.

- $\text{nn-rank}_{(\Omega,G)}(|T\rangle) = \text{sep-rank}_{(\Omega,G)}(\rho_{|T\rangle})$
- $\text{psd-rank}_{(\Omega,G)}(|T\rangle) = \text{puri-rank}_{(\Omega,G)}(\rho_{|T\rangle})$

Tensor decompositions and correlation scenarios 3

This chapter is based on Section 4 and Appendix E, F of [74].

In physics, we frequently encounter situations where we have access only to a limited set of observable quantities whose behavior depends on a hidden entity. In quantum physics, the wavefunction serves as an example where our access is restricted. In this context, we only possess access through measurements conducted on the quantum system, which effectively projects the wavefunction onto a probability distribution of measurement outcomes. (see Figure 3.1).

This begs the question: Can we deduce properties about the hidden quantity—the wavefunction—from the outcomes of measurements?

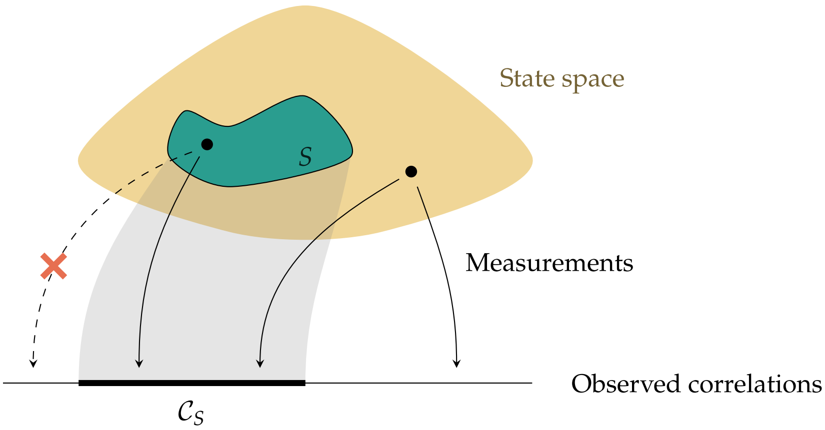

Figure 3.1: Applying measurements on a given set of states S gives rise to a subset of probability distributions $\mathcal{C}_S$. Therefore, observing a correlation outside of $\mathcal{C}_S$ witnesses that the state of the system is not contained in S.

Bell's theorem [9] addresses this question in a specific setting. It demonstrates that bipartite conditional probability distributions whose correlation arises from a particular classical causal structure satisfy the so-called Bell inequality. Consequently, a probability distribution that violates this inequality cannot emerge from this particular causal structure.

In this chapter, we show a correspondence of similar flavor between positive tensor decompositions and certain quantum correlation scenarios. More specifically, we show the following:

> *Applying local measurements on multipartite quantum states that obey a particular entanglement structure gives rise to probabilitiy distributions with a bounded positive tensor rank.*

Therefore, if a nonnegative tensor violates the tensor rank inequality, it cannot arise from the specific measurement scenario. Specifically, we show that if the subset of states S is given by states with a bounded (Ω, G)-rank, then the set $\mathcal{C}_S$ of arising correlations are characterized by a bounded (Ω, G) psd rank. Furthermore, we will prove a generalized correspondence replacing the observed correlations by density matrices and measurements by quantum channels (see Section 3.2).

We will leverage this correspondence in Chapter 4 to prove that these sets of correlations $\mathcal{C}_S$ are not topologically closed by showing this result on the level of tensors with bounded rank.

Probabilistic structures and their connection to tensor- and matrix-ranks have been previously explored:

- The nonnegative matrix decomposition is equivalent to a bipartite classical correlation scenario [31]. Since then, this correspondence has appeared in many different contexts and has also been generalized to the nonnegative tensor decomposition with one summation index [80, 101].
- A similar relation between the positive semidefinite matrix rank and bipartite quantum correlation scenarios has appeared in several works, see [49, 68, 69, 50].
- Nondeterministic quantum communication also relates to a notion of tensor rank, called the support rank [23].
- Nonnegative tensor network decompositions share a duality to *undirected graphical models*[1] as shown in [103, 58, 57, 84].

1: An undirected graphical model is a probabilistic model, where a graph expresses the conditional independence structure of the probability distribution. For details on undirected graphical models, we refer to [76].

To the best of our knowledge, there is no existing relation between the psd-rank for tensor networks and correlation scenarios.

This chapter is organized as follows: In Section 3.1 we introduce two correlation scenarios arising from classical hidden variables, as well as quantum states with a particular entanglement structure. We relate the sets of these correlation scenarios with nonnegative tensors of bounded rank. In Section 3.2 we extend these findings to mixed state correlation scenarios.

3.1 Classical correlations

Multipartite, finite probability distributions can be associated with nonnegative tensors. In particular, if $X_1, \dots, X_n$ are random variables taking values in $\{1, \dots, d\}$, then the tensor $|T\rangle$, defined via

$$\langle j_1, \dots, j_n \,|\, T\rangle := P(X_1 = j_1, \dots, X_n = j_n) \tag{3.1}$$

is a nonnegative tensor which is in addition normalized, i.e.

$$\sum_{j_1,\dots,j_n=1}^{d} \langle j_1, \dots, j_n \,|\, T\rangle = 1$$

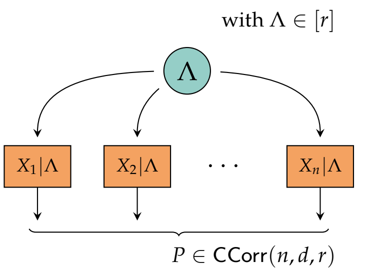

Figure 3.2: The classical correlation scenario defined in Equation (3.2). The joint probability distribution arises from a joint hidden variable that is shared between n parties.

Conversely, every normalized, nonnegative tensor gives rise to a probability distribution via Equation (3.1).

In the following, we use both notations probability distributions P and corresponding tensors $|T\rangle$ interchangeably. Specifically, we define specific correlation scenarios for probability distributions P and link them with the positive ranks for the corresponding nonnegative tensors $|T\rangle$.

3.1.1 Classical correlations from (Ω, G)-structures

We now define two correlation scenario sets that can be characterized via positive ranks.

First, we define the set

$$\mathsf{CCorr}(n, d, r)$$

as the set of probability distributions on n parties with local dimension d arising from local distributions conditioned on a shared hidden variable taking values in $\{1, \ldots, r\}$ (see Figure 3.2), i.e.

$$P(X_1 = j_1, \ldots, X_n = j_n) = \sum_{\alpha=1}^{r} P(\Lambda = \alpha) \prod_{i=1}^{n} P(X_i = j_i \mid \Lambda = \alpha) \quad (3.2)$$

where $X_1, \ldots, X_n$ are random variables taking values in $\{1, \ldots, d\}$.

Second, we define the set

$$\mathsf{CQCorr}_{(\Omega,G)}(n, d, r)$$

for a given WSC Ω and a group action G on Ω as the set of all n-partite probability distributions P arising as

$$P(X_1 = j_1, \ldots, X_n = j_n) = \langle\psi| A_{j_1}^{[1]} \otimes \cdots \otimes A_{j_n}^{[n]} |\psi\rangle$$

where

$$\left(A_j^{[i]}\right)_{j=1}^{d}$$

are POVMs[2] that are G-symmetric, i.e. the measurement on position i coincides with the measurement on gi for every $g \in G$. In other words, we have that $A_j^{[gi]} = A_j^{[i]}$ for every $g \in G$. Moreover, the state $|\psi\rangle$ satisfies the constraint that

$$\mathrm{rank}_{(\Omega,G)}(|\psi\rangle) \leqslant r.$$

We refer to Figure 3.3 for an illustration of this scenario.

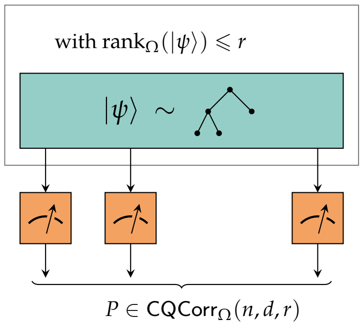

Figure 3.3: The quantum-classical correlation scenario for a trivial group action G. The state $|\psi\rangle$ admits an Ω-decomposition with $\mathrm{rank}_\Omega(|\psi\rangle) \leqslant r$. Each of the n measurements is performed locally and outputs a d-dimensional random variable.

2: A positive operator-valued measurement (POVM) is defined by a family of psd matrices E_j that satisfy the normalization condition

$$\sum_{j=1}^{k} E_j = \mathbb{1}.$$

This describes a measurement on a state ρ with probability distribution

$$P(X = j) = \mathrm{tr}(E_j \rho).$$

If, for example, $\Omega = \Theta_n$ is a cycle with n vertices, then

$$\mathsf{CQCorr}_{\Theta_n}(n, d, r)$$

is the set of all n-partite probability distributions obtained from an MPS $|\psi\rangle$ with bond dimension at most r via measurements on each local space. For the cyclic group $G = C_n$,

$$\mathsf{CQCorr}_{(\Theta_n, C_n)}(n, d, r)$$

is the set of probability distributions obtained from a ti MPS $|\psi\rangle$ with bond dimension at most r via identical measurements on each local space.

3.1.2 A correspondence to positive tensor ranks

In the following, we show that the sets

$$\mathsf{CCorr}(n, d, r) \quad \text{and} \quad \mathsf{CQCorr}_{(\Omega,G)}(n, d, r)$$

are characterized by the positive tensor ranks introduced in Section 2.3.

For this purpose, let $|T\rangle$ be the corresponding tensor to the probability distribution P as defined in Equation (3.1).

First, we show the correspondence for classical probability distributions. Note that a similar result has also been proven in [80].

Theorem 3.1.1 (The nonnegative rank and classical correlations) The following statements are equivalent:

(i) $\text{nn-rank}_{\Sigma_n}(|T\rangle) \leqslant r$
(ii) $P \in \mathsf{CCorr}(n, d, r)$.

The same equivalence holds for $\text{nn-rank}_{(\Sigma_n, S_n)}$ with the additional constraint in (ii) that the conditional probability distributions

$$P(X_i = - \mid Z = \alpha)$$

are identical for every $i \in \{1, \ldots, n\}$.

Proof. We show the equivalence only for $\text{nn-rank}_{(\Sigma_n, S_n)}$ as the other follows analogously.

(i) $\Longrightarrow$ (ii): Since $\text{rank}_{(\Sigma_n, S_n)}(|T\rangle) \leqslant r$ there is a nonnegative decomposition

$$|T\rangle = \sum_{\alpha=1}^{r} |v_\alpha\rangle \otimes \cdots \otimes |v_\alpha\rangle\,. \tag{3.3}$$

Define

$$P(X_i = j \mid Z = \alpha) := \frac{\langle j \mid v_\alpha\rangle}{\sum_{j=1}^{d} \langle j \mid v_\alpha\rangle}$$

and

$$P(Z = \alpha) = \left(\sum_{j=1}^{d} \langle j \mid v_\alpha\rangle\right)^{n}\,.$$

By definition, $P(X_i = - \mid Z = \alpha)$ is a probability distribution. Moreover, $P(Z = -)$ is a probability distribution since

$$\begin{aligned}
\sum_{\alpha=1}^{r} P(Z = \alpha) &= \sum_{\alpha=1}^{r} \left(\sum_{j=1}^{d} \langle j \mid v_\alpha\rangle\right)^{n} \\
&= \sum_{\alpha=1}^{r} \sum_{j_1,\ldots,j_n=1}^{d} \langle j_1, \ldots, j_n| \Big(|v_\alpha\rangle\Big)^{\otimes n} \\
&= \sum_{j_1,\ldots,j_n=1}^{d} \langle j_1, \ldots, j_n| \left(\sum_{\alpha=1}^{r} |v_\alpha\rangle^{\otimes n}\right) \\
&= \sum_{j_1,\ldots,j_n=1}^{d} P(X_1 = j_1, \ldots, X_n = j_n) = 1
\end{aligned}$$

where we have used the correspondence between P and $|T\rangle$ in the last step. Finally, $P(X_i = - \mid Z = \alpha)$ and $P(Z = -)$ give rise to the probability distribution P.

(ii) $\implies$ (i): Let

$$P(X_1 = j_1, \dots, X_n = j_n) = \sum_{\alpha=1}^{r} P(Z = \alpha) \prod_{i=1}^{n} P(X_i = j_i \mid Z = \alpha).$$

Defining

$$|v_\alpha^{[i]}\rangle := \sum_{j=1}^{r} P(X_i = j | Z = \alpha) \cdot P(Z = \alpha)^{\frac{1}{n}} \, |j\rangle$$

gives rise to nonnegative vectors in the computational basis. Since all conditional distributions $P(X_i = - | Z = \alpha)$ are identical, we have that $|v_\alpha^{[i]}\rangle = |v_\alpha^{[j]}\rangle =: |v_\alpha\rangle$ for every $i, j \in \{1, \dots, n\}$. It is immediate that Equation (3.3) holds. □

We now prove that elements of $\mathsf{CQCorr}_{(\Omega,G)}(n,d,r)$ are precisely these tensors with $\text{psd-rank}_{(\Omega,G)}(|T\rangle) \leqslant r$ if G is a external group action on Ω.[3] The special case $\Omega = \Sigma_n$ and $G = \{e\}$ is proven in [69, Theorem 13].

3: See Definition 2.2.5 for external group actions.

Theorem 3.1.2
Let Ω be a WSC and G an external group action on Ω. The following statements are equivalent:

(i) $P \in \mathsf{CQCorr}_{(\Omega,G)}(n,d,r)$.
(ii) $\text{psd-rank}_{(\Omega,G)}(|T\rangle) \leqslant r$.

We first need a preparatory lemma about the joint diagonalizability of G-invariant families of matrices.

Lemma 3.1.3 (G-symmetric matrix diagonalization)
Let Ω be a wsc and G an external group action on Ω. Moreover, let $K^{[i]} \in \mathrm{Her}_{\mathcal{I}^{\tilde{\mathcal{F}}_i}}(\mathbb{C})$ for $i \in [n]$ be Hermitian matrices such that

$$\langle {}^g\beta | \, K^{[gi]} \, | {}^g\beta' \rangle = \langle \beta | \, K^{[i]} \, | \beta' \rangle \qquad \text{for all } \beta, \beta' \in \mathcal{I}^{\tilde{\mathcal{F}}_i}$$

Then, there exists a joint eigendecomposition of all matrices $K^{[i]}$

$$K^{[i]} = \sum_{\ell=1}^{m} \lambda_\ell^{[i]} \, |w_\ell^{[i]}\rangle \langle w_\ell^{[i]}|$$

such that

$$\langle {}^g\beta \, | \, w_\ell^{[gi]} \rangle = \langle \beta \, | \, w_\ell^{[i]} \rangle \quad \text{and} \quad \lambda_\ell^{[gi]} = \lambda_\ell^{[i]}$$

Proof. Choose $i_1, \dots, i_m \in [n]$ representatives of the m orbits of the group action G on $[n]$. Computing the eigenvectors and eigenvalues of $K^{[i_1]}, \dots, K^{[i_m]}$ we obtain a generating set of eigendecompositions for every matrix $K^{[i]}$ by setting

$$\lambda_\ell^{[i]} = \lambda_\ell^{[gi_k]} \quad \text{and} \quad |w_\ell^{[i]}\rangle = \sum_{\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_{i_k}}} |{}^g\beta\rangle \langle \beta \,|\, w_\ell^{[i_k]}\rangle$$

for $g \in G$ and a representative i_k such that $i = gi_k$. Since the action is external, this is independent of the choice of g, which shows the statement. □

We are now ready to prove the main statement of this section.

Proof of Theorem 3.1.2. (i) $\Longrightarrow$ (ii): Let $P \in \mathsf{CQCorr}_{(\Omega,G)}(n,d,r)$. By definition, there exists a state

$$|\psi\rangle = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} |v_{\alpha_{|1}}^{[1]}\rangle \otimes \cdots \otimes |v_{\alpha_{|n}}^{[n]}\rangle$$

with $|\mathcal{I}| \leqslant r$ and G-invariant POVMs $\left(A_j^{[i]}\right)_{j=1}^{d}$ such that

$$\langle j_1, \ldots, j_n \,|\, T\rangle = \operatorname{tr}\left(A_{j_1}^{[1]} \otimes \cdots \otimes A_{j_n}^{[n]} \,|\psi\rangle\langle\psi|\right).$$

Define

$$B_j^{[i]} := \left(X^{[i]}\right)^{\dagger} A_j^{[i]} X^{[i]} \quad \text{with} \quad X^{[i]} = \sum_{\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}} |v_\beta^{[i]}\rangle\langle\beta|.$$

Note that $B_j^{[i]} \in \mathrm{Psd}_{\mathcal{I}^{\widetilde{\mathcal{F}}_i}}(\mathbb{C})$. Moreover, we have

$$\langle {}^g\beta|\, B_j^{[gi]}\, |{}^g\beta'\rangle = \langle v_{{}^g\beta}^{[gi]}|\, A_j^{[gi]}\, |v_{{}^g\beta'}^{[gi]}\rangle = \langle v_\beta^{[i]}|\, A_j^{[i]}\, |v_{\beta'}^{[i]}\rangle = \langle\beta|\, B_j^{[i]}\, |\beta'\rangle$$

where we have used that $|v_\beta^{[i]}\rangle$ forms a (Ω, G)-decomposition and that $A_j^{[i]}$ are G-invariant. Moreover,

$$\begin{aligned}
&\sum_{\alpha,\alpha' \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \left(B_{j_1}^{[1]}\right)_{\alpha_{|1},\alpha'_{|1}} \cdots \left(B_{j_n}^{[n]}\right)_{\alpha_{|n},\alpha'_{|n}} \\
&= \langle\psi|\, A_{j_1}^{[1]} \otimes \cdots \otimes A_{j_n}^{[n]}\, |\psi\rangle = P_{j_1,\ldots,j_n}
\end{aligned}$$

which proves that $\text{psd-rank}_{(\Omega,G)}(P) \leqslant r$.

(ii) $\Longrightarrow$ (i): Let

$$\begin{aligned}
\langle j_1, \ldots, j_n \,|\, T\rangle &= \sum_{\alpha,\alpha' \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \left(B_{j_1}^{[1]}\right)_{\alpha_{|1},\alpha'_{|1}} \cdots \left(B_{j_n}^{[n]}\right)_{\alpha_{|n},\alpha'_{|n}} \\
&= \langle\Omega_r|\, B_{j_1}^{[1]} \otimes \cdots \otimes B_{j_n}^{[n]}\, |\Omega_r\rangle
\end{aligned} \tag{3.4}$$

be a psd (Ω, G)-decomposition of P with $\text{psd-rank}_{(\Omega,G)}(P) \leqslant r = |\mathcal{I}|$. As the last expression in Equation (3.4) suggests, we use $B_j^{[i]}$ to construct a POVM and $|\Omega_r\rangle$ to construct a state whose combination leads to P.

While the matrices $B_j^{[i]}$ are psd, they need not form a POVM since

$$\sum_{j=1}^{k} B_j^{[i]} \neq \mathbb{1}_{r_i}$$

with $r_i = |\mathcal{I}^{\widetilde{\mathcal{F}}_i}|$. To this end, define

$$S^{[i]} := \sum_{j=1}^{d} B_j^{[i]} = \sum_{\ell=1}^{m_i} \lambda_\ell^{[i]} |w_\ell^{[i]}\rangle \langle w_\ell^{[i]}|$$

with $\lambda_\ell^{[i]} > 0$ being only the positive eigenvalues of $S^{[i]}$ and $|w_\ell^{[i]}\rangle$ being the G-invariant eigenvectors of the family $S^{[1]}, \ldots, S^{[n]}$ according to Lemma 3.1.3. Define

$$T^{[i]} = \sum_{\ell=1}^{m_i} \left(\lambda_\ell^{[i]}\right)^{-1/2} |w_\ell^{[i]}\rangle \langle \ell| \quad \text{and} \quad W^{[i]} = \sum_{\ell=1}^{m_i} \left(\lambda_\ell^{[i]}\right)^{1/2} |\ell\rangle \langle w_\ell^{[i]}| .$$

Note that $T^{[i]} \cdot W^{[i]}$ is a projector on

$$\operatorname{span}(\{|w_1^{[i]}\rangle, \ldots, |w_{m_i}^{[i]}\rangle\}).$$

Therefore, we have that

$$B_j^{[i]} = \left(T^{[i]} \cdot W^{[i]}\right)^\dagger \cdot B_j^{[i]} \cdot \left(T^{[i]} \cdot W^{[i]}\right). \tag{3.5}$$

Moreover, we have that

$$\langle {}^g\beta| T^{[gi]} = \langle \beta| T^{[i]} \quad \text{and} \quad W^{[gi]} |{}^g\beta\rangle = W^{[i]} |\beta\rangle \tag{3.6}$$

since the vectors $|w_\ell^{[i]}\rangle$ are G-invariant. We now define a POVM $(A_j^{[i]})_{j=1}^d$ via

$$A_j^{[i]} = \left(T^{[i]}\right)^\dagger \cdot B_j^{[i]} \cdot T^{[i]}.$$

We have that $A_j^{[i]}$ is psd and

$$\sum_{j=1}^{d} A_j^{[i]} = \mathbb{1}_{m_i}$$

which shows that $A^{[i]} := \left(A_j^{[i]}\right)_{j=1,\ldots,d}$ is indeed a POVM for each $i \in [n]$.

Moreover, $\left(A^{[i]}\right)_{i=1}^{n}$ is a G-invariant family since

$$\begin{aligned}
A_j^{[gi]} &= \left(T^{[gi]}\right)^{\dagger} \cdot B_j^{[gi]} \cdot T^{[gi]} \\
&= \sum_{\beta,\beta' \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}} \left(\langle\beta|\, T^{[gi]}\right)^{\dagger} \langle\beta|\, B_j^{[gi]}\, |\beta'\rangle\, \langle\beta'|\, T^{[gi]} \\
&= \sum_{\beta,\beta' \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}} \left(\langle{}^g\beta|\, T^{[gi]}\right)^{\dagger} \langle{}^g\beta|\, B_j^{[gi]}\, |{}^g\beta'\rangle\, \langle{}^g\beta'|\, T^{[gi]} \\
&= \sum_{\beta,\beta' \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}} \left(\langle\beta|\, T^{[i]}\right)^{\dagger} \langle\beta|\, B_j^{[i]}\, |\beta'\rangle\, \langle\beta'|\, T^{[i]} = A_j^{[i]}
\end{aligned}$$

where we have used that $\beta \mapsto {}^g\beta$ is a bijection between $\mathcal{I}^{\widetilde{\mathcal{F}}_i}$ and $\mathcal{I}^{\widetilde{\mathcal{F}}_{gi}}$ in the third step, and Equation (3.6) in the fourth step. Moreover,

$$|\psi\rangle := W^{[1]} \otimes \cdots \otimes W^{[n]}\, |\Omega_r\rangle$$

is a state with $\mathrm{rank}_{(\Omega,G)}(|\psi\rangle) \leqslant r$ since

$$\begin{aligned}
\langle\psi \,|\, \psi\rangle &= \langle\Omega_r|\, \left(W^{[1]}\right)^{\dagger} W^{[1]} \otimes \cdots \otimes \left(W^{[n]}\right)^{\dagger} W^{[n]}\, |\Omega_r\rangle \\
&= \langle\Omega_r|\, S^{[1]} \otimes \cdots \otimes S^{[n]}\, |\Omega_r\rangle \\
&= \sum_{j_1,\ldots,j_n=1}^{d} \sum_{\alpha,\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \left(B_{j_1}^{[1]}\right)_{\alpha_{|_1},\alpha'_{|_1}} \cdots \left(B_{j_n}^{[n]}\right)_{\alpha_{|_n},\alpha'_{|_n}} \\
&= \sum_{j_1,\ldots,j_n=1}^{d} \langle j_1,\ldots,j_n \,|\, T\rangle = 1
\end{aligned}$$

where we have used that the tensor $|T\rangle$ represents a probability distribution in the last step. Finally, the defined POVMs $\left(A_j^{[i]}\right)_{j=1}^{d}$ and the state $|\psi\rangle$ generate the probability distribution P, since

$$\begin{aligned}
\langle\psi|\, A_{j_1}^{[1]} \otimes \cdots \otimes A_{j_n}^{[n]}\, |\psi\rangle &= \sum_{\alpha,\alpha' \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \left(B_{j_1}^{[1]}\right)_{\alpha_{|_1},\alpha'_{|_1}} \cdots \left(B_{j_n}^{[n]}\right)_{\alpha_{|_n},\alpha'_{|_n}} \\
&= \langle j_1,\ldots,j_n \,|\, T\rangle
\end{aligned}$$

where we have used Equation (3.5) in the first step and Equation (3.4) in the second step. □

3.2 Mixed state correlation scenarios

In the following, we consider correlation scenarios where the output is a density matrix instead of a probability distribution. We will generalize the set $\mathsf{CQCorr}_{(\Omega,G)}(n,d,r)$ to this setting and show that the puri-rank of the output density matrix characterizes these correlations.

We define the set $\mathsf{QQCorr}_{(\Omega,G)}(n,d,r)$ as the set of all density matrices arising as

$$\rho = (\mathcal{E}_1 \otimes \cdots \otimes \mathcal{E}_n)\,(|\psi\rangle\,\langle\psi|)$$

where $(\mathcal{E}_i)_{i=1}^n$ is a family of quantum channels[4] which is G-invariant, i.e.

$$\mathcal{E}_i = \mathcal{E}_{gi}.$$

Moreover, $|\psi\rangle$ satisfies the condition $\text{rank}_{(\Omega,G)}(|\psi\rangle) \leqslant r$. We refer to Figure 3.4 for a sketch.

So, for example, when $\Omega = \Theta_n$ is a cycle graph of length n, then $\mathsf{QQCorr}_{\Theta_n}(n,d,r)$ is the set of all n-partite density matrices obtained from a MPS $|\psi\rangle$ with $\text{rank}_{\Theta_n}(|\psi\rangle) \leqslant r$ and applying local quantum channels on each local space. If additionally $G = C_n$ is the cyclic group, then $\mathsf{QQCorr}_{(\Theta_n,C_n)}(n,d,r)$ is the set of density matrices obtained from a MPS $|\psi\rangle$ with $\text{rank}_{(\Theta_n,C_n)}(|\psi\rangle) \leqslant r$ and applying identical quantum channels on each local space.

We now prove the quantum version of Theorem 3.1.2, namely that elements of $\mathsf{QQCorr}_{(\Omega,G)}(n,d,r)$ are precisely psd matrices ρ with $\text{tr}(\rho) = 1$ and $\text{puri-rank}_{(\Omega,G)}(\rho) \leqslant r$.

Theorem 3.2.1 (The puri-rank and quantum scenarios)
Let Ω be a WSC, G an external group action, and ρ an n-partite density matrix. The following statements are equivalent:

(i) $\rho \in \mathsf{QQCorr}_{(\Omega,G)}(n,d,r)$.
(ii) $\text{puri-rank}_{(\Omega,G)}(\rho) \leqslant r$.

The proof of this statement is similar to that of Theorem 3.1.2. The proof idea of (ii) $\Longrightarrow$ (i) is depicted in Figure 3.5 for one-dimensional purification forms, i.e. a Λ_n-purification.

Proof. (i) $\Longrightarrow$ (ii): Let ρ be a density matrix in $\mathsf{QQCorr}_{(\Omega,G)}(n,d,r)$. By definition, there exists a state

$$|\psi\rangle = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} |v^{[1]}_{\alpha_{|1}}\rangle \otimes \cdots \otimes |v^{[n]}_{\alpha_{|n}}\rangle$$

such that $\text{rank}_{(\Omega,G)}(|\psi\rangle) \leqslant r = |\mathcal{I}|$ and G-invariant family of quantum channels

$$\mathcal{E}_i(-) := \sum_{k=1}^{d_i} \left(A_k^{[i]}\right) \cdot - \cdot \left(A_k^{[i]}\right)^\dagger \tag{3.7}$$

with the condition that $A_k^{[i]} = A_k^{[gi]}$. We now define $L \in \text{Mat}_{d,d_1}(\mathbb{C}) \otimes \cdots \otimes \text{Mat}_{d,d_n}(\mathbb{C})$ such that

(a) $\rho = LL^\dagger$
(b) $\text{rank}_{(\Omega,G)}(L) \leqslant r$

which proves (ii). For $i \in [n]$ and $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$ let

$$L_\beta^{[i]} := \sum_{k=1}^{d_i} A_k^{[i]} |v_\beta^{[i]}\rangle \langle k| . \tag{3.8}$$

4: A quantum channel describes the most general transformation of a quantum state. It is a map

$$\mathcal{E} : \text{Mat}_{d_1}(\mathbb{C}) \to \text{Mat}_{d_1}(\mathbb{C})$$

that is *completely positive*, i.e.

$$(\mathbb{1}_n \otimes \mathcal{E})(\rho) \text{ is psd}$$

for every $n \in \mathbb{N}$ and every psd matrix ρ, and *trace preserving*, i.e.

$$\text{tr}(\mathcal{E}(\rho)) = \text{tr}(\rho).$$

A quantum channel is thus called completely positive trace preserving (cptp) map.

Further, set
$$L = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} L^{[1]}_{\alpha_{|1}} \otimes \cdots \otimes L^{[n]}_{\alpha_{|n}}.$$

By definition, we have that $\operatorname{rank}_{(\Omega,G)}(L) \leqslant r$. It remains to prove (a). But this follows from

$$\begin{aligned} LL^\dagger &= \sum_{k_1,\ldots,k_n=1}^{d} \left(A^{[1]}_{k_1} \otimes \cdots \otimes A^{[n]}_{k_n}\right) |\psi\rangle \langle\psi| \left(A^{[1]}_{k_1} \otimes \cdots \otimes A^{[n]}_{k_n}\right)^\dagger \\ &= (\mathcal{E}_1 \otimes \cdots \otimes \mathcal{E}_n)(|\psi\rangle \langle\psi|) = \rho \end{aligned}$$

where we have used Equation (3.8) in the first step and Equation (3.7) in the second.

(ii) $\Longrightarrow$ (i): Let $\rho = LL^\dagger$ where

$$L = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} L^{[1]}_{\alpha_{|1}} \otimes \cdots \otimes L^{[n]}_{\alpha_{|n}}$$

be an (Ω, G)-purificiation with $\text{puri-rank}_{(\Omega,G)}(\rho) \leqslant r = |\mathcal{I}|$.

Defining the completely positive maps

$$\mathcal{N}_i(-) := \sum_{k=1}^{d'} \left(B^{[i]}_k\right) \cdot - \cdot \left(B^{[i]}_k\right)^\dagger \quad \text{with} \quad \left(B^{[i]}_k\right)_{\ell,\beta} = \left(L^{[i]}_\beta\right)_{\ell,k}$$

we have that
$$\rho = (\mathcal{N}_1 \otimes \cdots \otimes \mathcal{N}_n)(|\Omega_r\rangle \langle\Omega_r|) \tag{3.9}$$

where $|\Omega_r\rangle$ is the structure tensor defined in Subsection 2.3.4. However, $\mathcal{N}_i$ is neither trace-preserving nor a G-invariant family, and $|\Omega_r\rangle$ is not normalized. For this reason, define

$$S^{[i]} := \sum_{k=1}^{d'} \left(B^{[i]}_k\right)^\dagger \cdot \left(B^{[i]}_k\right) = \sum_{\ell=1}^{m_i} \lambda^{[i]}_\ell |w^{[i]}_\ell\rangle \langle w^{[i]}_\ell|$$

where $|w^{[i]}_\ell\rangle$ is a G-invariant eigendecomposition of the family $S^{[1]}, \ldots, S^{[n]}$ according to Lemma 3.1.3. Similarly to the proof of Theorem 3.1.2 we define

$$\begin{aligned} T^{[i]} &:= \sum_{\ell=1}^{m_i} \left(\lambda^{[i]}_\ell\right)^{-1/2} |w^{[i]}_\ell\rangle \langle\ell| \\ W^{[i]} &:= \sum_{\ell=1}^{m_i} \left(\lambda^{[i]}_\ell\right)^{1/2} |\ell\rangle \langle w^{[i]}_\ell| \end{aligned} \tag{3.10}$$

and completely positive maps

$$\mathcal{E}_i(\rho) := \sum_{k=1}^{d'} \left(A^{[i]}_k\right) \cdot \rho \cdot \left(A^{[i]}_k\right)^\dagger \ \text{ with } \ A^{[i]}_k := B^{[i]}_k \cdot T^{[i]}. \tag{3.11}$$

Note that $(\mathcal{E}_i)_{i=1,\ldots,n}$ is by definition a G-invariant family of quantum channels. Moreover, by the reasoning of the proof of Theorem 3.1.2,

$$|\psi\rangle = W^{[1]} \otimes \cdots \otimes W^{[n]} |\Omega_r\rangle \tag{3.12}$$

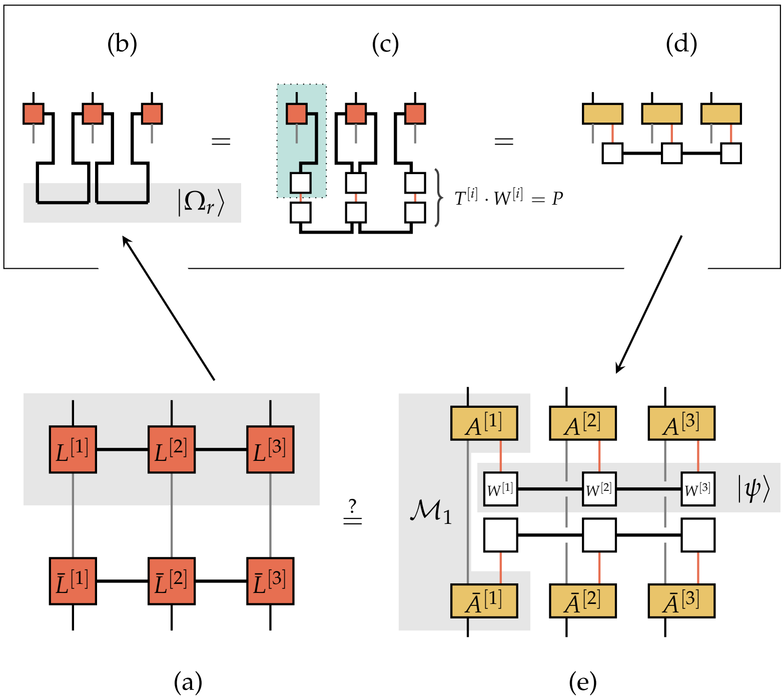

Figure 3.5: Proof of Theorem 3.2.1 (ii) $\Longrightarrow$ (i) on a 1d chain, i.e. proving the equality of expressions (a) and (e). (a) is the local purification form with puri-rank$_{\Lambda_n}(\rho) \leqslant r$. (b) When rearranging the wires, we obtain the definition of a Ω-decomposition with the structure-tensor $|\Omega_r\rangle$ according to Equation (2.8). This decomposition can also be understood as applying a completely positive map to $|\Omega_r\rangle$ according to Equation (3.9). In (c), we insert a projector $P^{[i]}$ of the space where the tensor $L^{[i]}$ acts non-trivially and factorize it into a product $T^{[i]} \cdot W^{[i]}$ according to Equation (3.10). To obtain (d) we merge the upper box ($T^{[i]}$) with the red box (Equation (3.11)). This gives rise to a normalized state (Equation (3.12)) together with local quantum channels (e).

defines a normalized state with $\text{rank}_{(\Omega,G)}(|\psi\rangle) \leqslant r$. Moreover,

$$(\mathcal{E}_1 \otimes \cdots \otimes \mathcal{E}_n)\,(|\psi\rangle\,\langle\psi|) = (\mathcal{N}_1 \otimes \cdots \otimes \mathcal{N}_n)\,(|\Omega_r\rangle\,\langle\Omega_r|) = \rho$$

which proves the statement. □

Note that Theorem 3.2.1 implies Theorem 3.1.2 when restricting to diagonal density matrices. This follows from the fact that every quantum channel that outputs only classical states corresponds to a POVM.

More specifically, every POVM $E_1, \ldots, E_k$ gives rise to a quantum channel

$$\begin{aligned} \mathcal{E}\colon\ \text{Mat}_d(\mathbb{C}) &\to \text{Mat}_k(\mathbb{C}) \\ \rho &\to \sum_{i=1}^{k} |i\rangle\,\langle i|\,\text{tr}(E_i\rho) \end{aligned}$$

Conversely, every quantum channel that maps into the space of diagonal matrices can be specified by

$$\begin{aligned} \mathcal{E}\colon\ \text{Mat}_d(\mathbb{C}) &\to \text{Mat}_k(\mathbb{C}) \\ \rho &\mapsto \sum_{i=1}^{r} |i\rangle\,\langle i|\,\text{tr}(A_i\rho) \end{aligned}$$

since $|i\rangle\,\langle i|$ for $i \in [k]$ is a basis of the space of diagonal matrices in $\text{Mat}_k(\mathbb{C})$. Since $\mathcal{E}$ is positive, we have that $\text{tr}(A_i\rho) \geqslant 0$ for all psd matrices ρ. This implies that A_i is psd. Moreover, since $\mathcal{E}$ is trace preserving,

$$\text{tr}\left(\sum_{i=1}^{k} A_i\rho\right) = \text{tr}(\rho)$$

for every ρ. But this shows that[5]

$$\sum_{i=1}^{k} A_i = \mathbb{1}$$

and hence $(A_i)_{i=1,\dots,k}$ is a POVM.

To summarize, Theorem 3.1.2 and Theorem 3.2.1 reveal that bounding psd-rank and puri-rank leads to an information-theoretic interpretation, elucidating correlations arising from quantum states with particular entanglement structures. We will revisit these correlation scenarios in Section 4.3, demonstrating that the sets of correlation scenarios do not exhibit topological closure for certain configurations of Ω and G.

5: Given $A \in \mathrm{Her}_k(\mathbb{C})$, the condition $\mathrm{tr}(AB) = \mathrm{tr}(B)$ for every $B \in \mathrm{Her}_k(\mathbb{C})$ implies that $A = \mathbb{1}$. This follows since

$$\mathrm{tr}((A - \mathbb{1})B) = 0$$

for every B. Since tr is an inner product on $\mathrm{Her}_k(\mathbb{C})$, we can conclude that $A - \mathbb{1} = 0$.

4 Border ranks of positive tensor decompositions

It is well-known that low-rank approximations of matrices exhibit desirable properties: For every matrix, there is a best low-rank approximation with a fixed error, and any element closer to the original matrix must have a larger rank. In other words, the approximate rank

$$\mathrm{rank}^{\varepsilon}(|T\rangle) := \min_{\||W\rangle - |T\rangle\| \leqslant \varepsilon} \mathrm{rank}(|W\rangle) \tag{4.1}$$

coincides with the exact rank when ε is small enough.

This chapter is based on Section 1, 3, 4, and 5 of [74].

The multipartite tensor rank behaves very differently: There exist tensors $|T\rangle$ where the *border rank*

$$\underline{\mathrm{rank}}(|T\rangle) := \lim_{\varepsilon \to 0} \mathrm{rank}^{\varepsilon}(|T\rangle)$$

is *strictly smaller* than the rank of $|T\rangle$ (see Figure 4.1). For the mathematician, this means that the rank is not lower semi-continuous. This is equivalent to the statement that the set of tensors whose rank is upper bounded by a constant r

$$\mathcal{T} := \{\, |T\rangle \in \mathcal{V}^{\otimes n} : \mathrm{rank}(|T\rangle) \leqslant r\}$$

is topologically not closed since there are sequences in $\mathcal{T}$ whose limit is not in $\mathcal{T}$. As a consequence, optimization problems over such sets, such as computing an optimal low-rank approximation, are generally ill-posed [114]. It is known that tensor decompositions with three or more local spaces exhibit a gap between rank and border rank [78], and so do tensor network decompositions containing loops [77, 29, 5], where some of these results concern symmetric decompositions of invariant tensors.

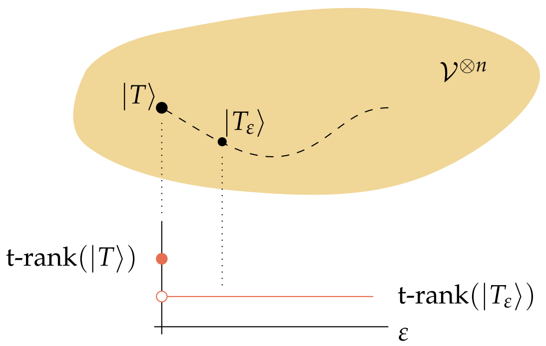

Figure 4.1: Border rank. Given a tensor $|T\rangle$ in an n-fold tensor product space and a certain type of rank t-rank, if there exists a family of tensors $(|T_\varepsilon\rangle)_{\varepsilon>0}$ such that $|T_\varepsilon\rangle \to |T\rangle$ for $\varepsilon \to 0$ and t-rank$(|T_\varepsilon\rangle) <$ t-rank$(|T\rangle)$ for all $\varepsilon > 0$, we say that t-rank exhibits a gap between rank and border rank.

In this chapter, we prove that several locally positive and invariant decompositions exhibit a gap between rank and border rank, as summarized in Figure 4.2. This includes positive and/or symmetric versions of Matrix Product States (MPS) and Matrix Product Operators (MPO), as well as the multipartite generalizations of the psd-rank.

We leverage the gaps between border ranks and ranks together with the connection to quantum correlations presented in Chapter 3 to show that:

Imposing stronger positivity constraints ↓

Decomposition type / Type of rank	Standard 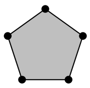	Symmetric 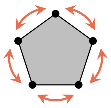	Cyclic 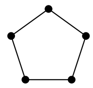	Transl. invariant 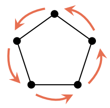	Tree 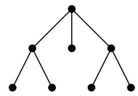
rank	Yes ($n \geqslant 3$) [13, 78]	Yes ($n \geqslant 3$) [78]	Yes ($n \geqslant 3$) [5, 29, 77]	Yes ($n \geqslant 3$) [5, 29]	No [5, 77]
psd-rank puri-rank	**Yes** ($n \geqslant 5$) (Sec. 4.1.1)	**Yes** ($n \geqslant 3$) (Sec. 4.1.1)	**?**	**Yes** ($n \geqslant 17$) (Sec. 4.1.2)	**No** (Thm. 4.2.6)
nn-rank sep-rank	**No** (Thm. 4.2.2)	**No** (Thm. 4.2.2)	**Yes** ($n \geqslant 3$) (Sec. 4.1.3)	**Yes** ($n \geqslant 5$) (Sec. 4.1.2)	**No** (Thm. 4.2.6)

Figure 4.2: Is there a gap between rank and border rank in an n-fold tensor product space? This table summarizes known results and the contributions of this paper (marked boldface): We prove that gaps persist when imposing positivity constrains corresponding to quantum correlation scenarios (second row), and that certain gaps disappear for stronger positivity constrains corresponding to classical correlation scenarios (third row). The types of ranks and of decompositions are defined in Chapter 2.

- If a tensor network geometry (i.e. the WSC) contains a loop, computing the best approximation with a fixed positive rank is ill-posed. Specifically, given a mixed state ρ, there is typically no mixed state σ which is the best approximation among all decompositions with a positive rank bounded by r, because for any $\varepsilon > 0$ there is an ε-close mixed state of rank r, while the rank of ρ is strictly greater than r.
- The set of probability distributions generated by a multipartite state with local measurements (Figure 3.3) is not closed. Consequently, it is impossible to verify the necessity of a certain resource state from sampling the distribution, even in arbitrarily many rounds. The same applies to generating multipartite mixed states from local quantum channels (Figure 3.4).
- We provide correlation scenarios where the quantum case is fragile with respect to approximations, while the classical case is robust. This shows a novel type of separation between these two scenarios.

4.1 Gaps between ranks and border ranks

Here we provide examples of tensor decompositions with gaps summarized in Figure 4.2. Throughout, the gaps between ranks and border ranks are established by giving explicit examples of tensors exhibiting them.

Note that for every (Ω, G)-rank, we define the corresponding border (Ω, G)-rank as the minimal (Ω, G)-rank of a sequence approach to the original element. More precisely

$$\underline{\mathrm{rank}}_{(\Omega,G)}(|T\rangle) \leqslant r \iff \begin{array}{rl} \exists (|T_n\rangle)_{n\in\mathbb{N}}: & |T_n\rangle \to |T\rangle \\ \text{and} & \mathrm{rank}_{(\Omega,G)}(|T_n\rangle) \leqslant r. \end{array}$$

We define $\underline{\text{psd-rank}}_{(\Omega,G)}$, $\underline{\text{nn-rank}}_{(\Omega,G)}$, $\underline{\text{puri-rank}}_{(\Omega,G)}$, and $\underline{\text{sep-rank}}_{(\Omega,G)}$ analogously.

4.1.1 Standard tensor decomposition

Since the matrix rank does not exhibit a gap between border rank and rank, systems of size $n = 3$ are the smallest examples with a gap between border rank and rank. While this has been extensively studied for the standard and symmetric tensor rank[1], we extend these investigations in this subsection to psd matrices. The nonnegative standard decomposition is treated in Section 4.2.1.

1: See for example [78] and references therein.

For the standard (unconstrained) tensor decomposition, the unnormalized n-partite W-state

$$|W_n\rangle := |0,\dots,0,1\rangle + |0,\dots,1,0\rangle + \dots + |1,0,\dots,0\rangle$$

exhibits a gap between border rank and rank as well as between symmetric border rank and rank for system sizes $n \geqslant 3$. Specifically, for $\varepsilon > 0$, the family of tensors

$$|W_n^\varepsilon\rangle = \frac{1}{\varepsilon}\big(\,|0\rangle + \varepsilon\,|1\rangle\,\big)^{\otimes n} - \frac{1}{\varepsilon}\,|0,\dots,0\rangle \tag{4.2}$$

implies that

$$\underline{\text{rank}}_{(\Sigma_n,S_n)}(|W_n\rangle) = \underline{\text{rank}}_{\Sigma_n}(|W_n\rangle) = 2 \tag{4.3}$$

since $|W_n^\varepsilon\rangle \to |W_n\rangle$ as $\varepsilon \to 0$. On the other hand, we obtain the following statement:

Proposition 4.1.1
For $n \geqslant 2$, we have that $\text{rank}_{\Sigma_n}(|W_n\rangle) = n$.

Proof. That $\text{rank}_{\Sigma_n}(|W_n\rangle) \leqslant n$ is clear by the definition of $|W_n\rangle$. We prove that $\text{rank}_{\Sigma_n}(|W_n\rangle) \geqslant n$ by induction. The case $n = 2$ is clear, since $|W_2\rangle \in \mathbb{C}^2 \otimes \mathbb{C}^2$ corresponds to the matrix

$$W_2 = |0\rangle\langle 1| + |1\rangle\langle 0| = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

Therefore $|W_2\rangle$ has Σ_2-rank[2] 2.

2: This is precisely the matrix rank.

For the induction step $n \to n+1$, suppose that $|W_{n+1}\rangle$ has

$$\text{rank}_{\Sigma_{n+1}}(|W_{n+1}\rangle) \leqslant n$$

with a decomposition

$$|W_{n+1}\rangle = \sum_{\alpha=1}^{n} |v_\alpha^{[1]}\rangle \otimes \cdots \otimes |v_\alpha^{[n]}\rangle\,.$$

For the first local system, we will prove that

(a) The vectors $\{|v_\alpha^{[1]}\rangle\}_{\alpha=1,\dots,n}$ span $\mathbb{C}^2$.

(b) $|v_\beta^{[1]}\rangle = c_\beta\,|0\rangle$ for every $\beta \in \{1,\dots,n\}$.

These two conditions contradict each other, hence proving the statement of the proposition.

To prove (a) assume that the family $\{|v_\alpha^{[1]}\rangle\}_{\alpha=1,\dots,n}$ does not span $\mathbb{C}^2$. Then there exists a non-zero vector $|x\rangle \in \mathbb{C}^2$ such that $\langle x\,|\,v_\alpha^{[1]}\rangle = 0$ for every α. Appyling $\langle x|$ to the first tensor factor of $|W_{n+1}\rangle$ leads to

$$0 = \langle x\,|\,0\rangle\,|W_n\rangle + \langle x\,|\,1\rangle\,|0,0,0,\dots,0\rangle\,.$$

Since $|W_n\rangle$ and $|0,\dots,0\rangle$ are linearly independent this implies that $|x\rangle = 0$, which is a contradiction.

To prove (b), note that

$$\operatorname{rank}_{\Sigma_n}(|W_n\rangle + b\,|0,\dots,0\rangle) \geqslant \operatorname{rank}_{\Sigma_n}(|W_n\rangle) \geqslant n$$

for every $b \in \mathbb{R}$ since

$$|W_n\rangle = A^{\otimes n}\Big(\,|W_n\rangle + b\,|0,0,0,\dots,0\rangle\,\Big)$$

with

$$A : |0\rangle \mapsto |0\rangle\,, \quad |1\rangle \mapsto |1\rangle - \frac{b}{n}\,|0\rangle\,.$$

This shows that

$$\operatorname{rank}_{\Sigma_n}\Big(\,|W_n\rangle + b\,|0,\dots,0\rangle\,\Big) \geqslant \operatorname{rank}_{\Sigma_n}(|W_n\rangle)$$

since the rank is non-increasing under local operations. Now let $\beta \in \{1,\dots,r\}$ be fixed and choose $|x\rangle \in \mathbb{C}^2$ such that $\langle x\,|\,v_\beta^{[1]}\rangle = 0$. Applying $\langle x|$ to the first tensor factor of $|W_{n+1}\rangle$ we obtain

$$\sum_{\substack{\alpha=1\\ \alpha\neq\beta}}^{n} \langle x\,|\,v_\alpha^{[1]}\rangle\,|v_\alpha^{[2]}\rangle \otimes \dots \otimes |v_\alpha^{[n]}\rangle = \langle x\,|\,0\rangle\,W_n + \langle x\,|\,1\rangle\,|0,0,0,\dots,0\rangle\,.$$

Since the sum on the left hand side contains $n-1$ elementary tensors and the right hand side has rank at least n, if $\langle x\,|\,0\rangle \neq 0$, it follows that $\langle x\,|\,0\rangle = 0$. But this implies that

$$|v_\beta^{[1]}\rangle = c_\beta\,|0\rangle\,.$$

□

Equation (4.3) and Proposition 4.1.1 imply the following corollary:

Corollary 4.1.2

For $n \geqslant 3$, the standard and the symmetric tensor rank exhibit a gap. More specifically

$$\begin{aligned}\underline{\operatorname{rank}}_{(\Sigma_n,S_n)}(|W_n\rangle) &= \underline{\operatorname{rank}}_{\Sigma_n}(|W_n\rangle) = 2\\ &< n = \operatorname{rank}_{\Sigma_n}(|W_n\rangle) \leqslant \operatorname{rank}_{(\Sigma_n,S_n)}(|W_n\rangle)\end{aligned}$$

We now show that the $|W_n\rangle$ also exhibits a gap between rank and border rank for the psd Σ_n-rank. Since $\text{rank}_{\Sigma_n}(|W_n\rangle) = n$ and

$$\text{rank}_{\Sigma_n}(|T\rangle) \leqslant \text{psd-rank}_{\Sigma_n}(|T\rangle)^2$$

(see Lemma 2.3.1) we have $\text{psd-rank}_{\Sigma_n}(W_5) \geqslant 3$ and $\text{psd-rank}_{\Sigma_n}(W_n) = \Omega(\sqrt{n})$.[3] It is not known if this lower bound is tight.

3: This means that $\text{psd-rank}_{\Sigma_n}$ is asymptotically lower bounded by $D \cdot \sqrt{n}$ for some constant D.

On the other hand, for $\varepsilon > 0$, the family of tensors $|\widetilde{W}_n^\varepsilon\rangle$ defined by psd matrices

$$A_0^\varepsilon = \frac{C}{\sqrt[n-1]{\varepsilon}} \begin{pmatrix} 1 & e^{\frac{i\pi}{n}} \\ e^{-\frac{i\pi}{n}} & 1 \end{pmatrix}, \quad A_1^\varepsilon = \varepsilon \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} \tag{4.4}$$

for a suitable constant $C \in \mathbb{R}$ provide an arbitrarily close approximation of $|W_n\rangle$ which implies that

$$\underline{\text{psd-rank}}_{\Sigma_n}(|W_n\rangle) = \underline{\text{psd-rank}}_{(\Sigma_n,S_n)}(|W_n\rangle) = 2.$$

In other words, there is a border rank separation for n-partite psd-decompositions with $n \geqslant 5$.

For the symmetric psd rank of $|W_3\rangle$ we obtain a tighter lower bound.

Proposition 4.1.3
We have that

$$3 \leqslant \text{psd-rank}_{(\Sigma_n,S_n)}(W_3).$$

Proof. Assume that $\text{psd-rank}_{(\Sigma_n,S_n)}(W_3) = 2$. Then there exists a symmetric psd-decomposition

$$\langle j_1, j_2, j_3 \,|\, W_3\rangle = \sum_{\alpha,\beta=1}^{2} \left(A_{j_1}\right)_{\alpha,\beta} \cdot \left(A_{j_2}\right)_{\alpha,\beta} \cdot \left(A_{j_3}\right)_{\alpha,\beta}.$$

This can be expressed equivalently as

$$\langle j_1, j_2, j_3 \,|\, W_3\rangle = \langle M|\, A_{j_1} \star A_{j_2} \star A_{j_3} \,|M\rangle$$

where $|M\rangle = (1, \ldots, 1)^t$ and $\star$ is the Hadamard product.[4] We claim that A_0 and A_1 in the decomposition have rank 1. Assume for example that A_0 has full rank; it is positive definite, therefore $A_0 \star A_0 \star A_0$ is positive definite by Schur's product theorem (see [65, Theorem 7.5.3.]). But this implies that

4: The Hadamard Product $\star$ of two matrices is defined as

$$(X \star Y)_{\alpha,\beta} = X_{\alpha,\beta} \cdot Y_{\alpha,\beta}.$$

$$0 = \langle 0,0,0 \,|\, W_3\rangle = \langle M|\, A_0 \star A_0 \star A_0 \,|M\rangle > 0.$$

The same argument applies to A_1.

Since A_0, A_1 have rank 1, we can parametrize them as

$$A_j = \begin{pmatrix} a_{j,0} & \sqrt{a_{j,0}a_{j,1}}\exp(i2\pi\varphi_j) \\ \sqrt{a_{j,0}a_{j,1}}\exp(-i2\pi\varphi_j) & a_{j,1} \end{pmatrix}$$

where $a_{j,0}, a_{j,1} \geqslant 0$. Since $\langle 0,0,0 \,|\, W_3\rangle = \langle 1,1,1 \,|\, W_3\rangle = 0$, we have that $a_{j,0} = a_{j,1}$ for $j = 0,1$ as well as $\varphi_j = 1/2$ which implies that

$\langle j_1, j_2, j_3 \,|\, W_3 \rangle = 0$ for all $j_1, j_2, j_3 \in \{0, 1\}$. □

For the non-symmetric case in the tripartite scenario the existence of a gap between border rank and rank is still open. We summarize these observations in the following corollary:

Corollary 4.1.4
For $n \geqslant 5$, there is a gap between $\underline{\text{psd-rank}}_{\Sigma_n}$ and $\text{psd-rank}_{\Sigma_n}$. More specifically

$$\underline{\text{psd-rank}}_{\Sigma_n}(|W_n\rangle) = \underline{\text{psd-rank}}_{(\Sigma_n, S_n)}(|W_n\rangle) = 2$$

and

$$\sqrt{n} \leqslant \text{psd-rank}_{\Sigma_n}(|W_3\rangle).$$

For the symmetric psd rank, the gap is already present for $n = 3$, since

$$\text{psd-rank}_{\Sigma_n}(|W_n\rangle) = 3.$$

In contrast to the psd-decomposition, the nonnegative (and subsequently also the separable) decomposition exhibit no gap between border rank and rank in the n-partite case for arbitrary n, as we will see Section 4.2.

4.1.2 Cyclic translational invariant decomposition

We now prove the existence of gaps between border rank and rank for ti cyclic decompositions. We obtain border rank separations for all types of decompositions. Similar to Section 4.1.1 the n-partite W-state is can be used as an example showing the gaps.

We start with the unconstrained decomposition.

Proposition 4.1.5
For the n-partite W-state we have that

$$\underline{\text{rank}}_{(\Theta_n, C_n)}(|W_n\rangle) = 2 < \sqrt{n} \leqslant \text{rank}_{(\Theta_n, C_n)}(|W_n\rangle)$$

Therefore, there is a gap for $n \geqslant 5$.

Proof. For $\underline{\text{rank}}_{(\Theta_n, C_n)}(|W_n\rangle) = 2$, we use the construction by Christandl et al. [29]. We define the approximate decomposition using $|v_{12}^{\varepsilon}\rangle = |v_{21}^{\varepsilon}\rangle = 0$ and

$$|v_{11}^{\varepsilon}\rangle = \frac{1}{\varepsilon^{1/n}} \begin{pmatrix} 1 \\ \varepsilon \end{pmatrix} \quad \text{and} \quad |v_{22}^{\varepsilon}\rangle = \frac{1}{\varepsilon^{1/n}} \begin{pmatrix} (-1)^{\frac{1}{n}} \\ 0 \end{pmatrix}$$

for arbitrary $\varepsilon > 0$.

For the lower bound $\sqrt{n} \leqslant \text{rank}_{(\Theta_n, C_n)}(|W_n\rangle)$ we refer to [40, Proposition 23] which relies on the irreducible form of MPS [35].[5] □

5: A weaker lower bound

$$\text{rank}_{(\Theta_n, C_n)}(|W_n\rangle) \geqslant \Omega(n^{1/3})$$

was shown by Pérez-García et al. [97] using Wieland's inequality [106].

For the ti psd-decomposition, we obtain the following result:

Proposition 4.1.6
We have that

$$\underline{\text{psd-rank}}_{(\Theta_n,C_n)}(|W_n\rangle) = 2$$

and

$$\Omega(n^{1/4}) \leqslant \text{psd-rank}_{(\Theta_n,C_n)}(|W_n\rangle)$$

In particular, there is a gap for $n \geqslant 17$.

Proof. To show $\underline{\text{psd-rank}}_{(\Theta_n,C_n)}(|W_n\rangle) = 2$, we define the psd matrices

$$\left(B_j^\varepsilon\right)_{\alpha,\alpha';\beta,\beta'} = \delta_{\alpha,\alpha'} \cdot \delta_{\beta,\beta'} \cdot \left(A_j^\varepsilon\right)_{\alpha,\beta}$$

where A_j^ε is defined in Equation (4.4). We obtain

$$\begin{aligned} &\sum_{\alpha_i,\beta_i=1}^{2} \left(B_{j_1}^\varepsilon\right)_{\alpha_1,\alpha_2;\beta_1,\beta_2} \cdots \left(B_{j_n}^\varepsilon\right)_{\alpha_n,\alpha_1;\beta_n,\beta_1} \\ &= \langle j_1,\dots,j_n \mid W_n\rangle + \mathcal{O}\left(\varepsilon^{1+\frac{1}{n-1}}\right). \end{aligned}$$

Moreover, using Lemma 2.3.1 together with Proposition 4.1.5 we obtain that

$$\text{psd-rank}_{(\Theta_n,C_n)}(|W_n\rangle) \geqslant \Omega\left(n^{1/4}\right) \tag{4.5}$$

and in particular $\text{psd-rank}_{(\Theta_n,C_n)}(|W_n\rangle) \geqslant 3$ as soon as $n \geqslant 17$. This proves the separation between border rank and rank for the t.i. cyclic psd-decomposition. □

For the ti nonnegative decomposition we construct a tensor with a separation between border rank and rank for every odd $n \geqslant 5$. Consider again the tensor $|W_n\rangle$. By the previous discussion, we have

$$\text{nn-rank}_{(\Theta_n,C_n)}(|W_n\rangle) \geqslant \text{rank}_{(\Theta_n,C_n)}(|W_n\rangle) \geqslant \sqrt{n}.$$

In order to prove an upper bound for $\underline{\text{nn-rank}}_{(\Theta_n,C_n)}$, we use the following representation of a nonnegative cyclic decomposition

$$\langle j_1,\dots,j_n \mid T\rangle = \text{tr}(A_{j_1}\cdots A_{j_n}),$$

where $A_j \in \mathcal{M}_r(\mathbb{C})$ and $(A_j)_{\alpha,\beta} \geqslant 0$. It follows that the rank of the decomposition is specified by the size of the matrices A_j.[6]

6: For details, we refer to Example 2.3.6.

Proposition 4.1.7
We have that

$$\underline{\text{nn-rank}}_{(\Theta_n,C_n)}(|W_n\rangle) \leqslant 2$$

if n is odd.

Proof. Let

$$A_0^\varepsilon = \frac{1}{\sqrt[n-1]{\varepsilon}}\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = \frac{1}{\sqrt[n-1]{\varepsilon}} P_\tau \qquad A_1^\varepsilon = \varepsilon I_2$$

be multiples of a nonnegative representation of the cyclic group on $\{1,2\}$, where τ is the permutation $1 \mapsto 2$ and $2 \mapsto 1$ and P_τ the corresponding permutation matrix. We have

$$\begin{aligned}\langle j_1, \ldots, j_n \mid \widehat{W}_n^\varepsilon \rangle &:= \frac{1}{2} \operatorname{tr}\left(A_{j_1}^\varepsilon \cdots A_{j_n}^\varepsilon\right) \\ &= \frac{1}{2} \begin{cases} 0 & : j_1 + \ldots + j_n \text{ even} \\ \varepsilon^{k-1+\frac{k-1}{n-1}} & : j_1 + \ldots + j_n \text{ odd} \end{cases}\end{aligned}$$

where $k := j_1 + \cdots + j_n$. This implies that $|\widehat{W}_n^\varepsilon\rangle = \frac{1}{2} |W_n\rangle + \mathcal{O}(\varepsilon^2)$. □

Note that this construction generalizes to every n and $p \mid (n-1)$ by replacing $\{1,2\}$ with $\{1,\ldots,p\}$, and τ by the translation on $\{1,\ldots,p\}$. Since the corresponding permutation matrices A_0^ε and A_1^ε are of size $p \times p$, it follows that $\underline{\text{nn-rank}}_{(\Theta_n, C_n)}(|W_n\rangle) \leqslant p$.

Corollary 4.1.8
If n is odd, we have that

$$\begin{aligned}\underline{\text{nn-rank}}_{(\Theta_n,C_n)}(|W_n\rangle) = 2 < \sqrt{n} \leqslant \operatorname{rank}_{(\Theta_n,C_n)}(|W_n\rangle) \\ \leqslant \text{nn-rank}_{(\Theta_n,C_n)}(|W_n\rangle).\end{aligned}$$

This implies that there is a gap for $n \geqslant 5$.

4.1.3 Cyclic decompositions

In the following, we consider the cyclic decomposition without translational invariance. In contrast to Section 4.1.2, the n-partite W-state is not an appropriate example to show a gap. This is because

$$\underline{\operatorname{rank}}_{\Theta_n}(|W_n\rangle) = \operatorname{rank}_{\Theta_n}(|W_n\rangle) = 2$$

since

$$|W_n\rangle = \sum_{\alpha_1,\ldots,\alpha_n=1}^{2} |v_{\alpha_1,\alpha_2}\rangle \otimes |w_{\alpha_2,\alpha_3}\rangle \otimes \cdots \otimes |w_{\alpha_n,\alpha_1}\rangle$$

where

$$|v_{\alpha,\beta}\rangle = \delta_{\alpha,2}\delta_{\beta,1} |0\rangle + \delta_{\alpha,2}\delta_{\beta,2} |1\rangle$$

and

$$|w_{\alpha,\beta}\rangle = \delta_{\alpha,\beta} |0\rangle + \delta_{\alpha,1}\delta_{\beta,2} |1\rangle .$$

Regarding unconstrained decompositions, Barthel et al. [5] prove that for the Θ_n-rank, there is a gap between border rank and rank for the *two-domain state*, given by

$$\begin{aligned}|\tau\rangle :=& \sum_{\alpha=1}^{k} |\alpha,\alpha\rangle^{\otimes n} \\ &+ \sum_{i=0}^{n-1} \sum_{\alpha\neq\beta=1}^{k} |\alpha,\alpha\rangle^{\otimes i} \otimes |\alpha,\beta\rangle \otimes |\beta,\beta\rangle^{\otimes(n-i)} \otimes |\beta,\alpha\rangle .\end{aligned}$$

In particular, they prove that $\underline{\text{rank}}_{\Theta_n}(|\tau\rangle) \leqslant k < \text{rank}_{\Theta_n}(|\tau\rangle)$.

The construction in [5] also leads to a gap between border rank and rank for nonnegative cyclic decompositions, which we briefly review now. Let $\varepsilon > 0$ and define for every $\alpha, \beta \in \{1, \ldots, k\}$ the nonnegative vectors

$$|v^{\varepsilon}_{\alpha,\beta}\rangle = \varepsilon\,|\alpha,\beta\rangle + (1-\varepsilon)\delta_{\alpha,\beta}\,|\alpha,\beta\rangle$$

where $\delta_{\alpha,\beta}$ is the Kronecker-delta, as well as

$$|w^{\varepsilon}_{\alpha,\beta}\rangle = \delta_{\alpha,\beta}\,|\alpha,\beta\rangle + \frac{1}{\varepsilon}(1-\delta_{\alpha,\beta})\,|\alpha,\beta\rangle\,.$$

Setting

$$|\tau^{\varepsilon}\rangle = \sum_{\alpha_i=1}^{k} |v^{\varepsilon}_{\alpha_1,\alpha_2}\rangle \otimes |v^{\varepsilon}_{\alpha_2,\alpha_3}\rangle \otimes \cdots \otimes |v^{\varepsilon}_{\alpha_{n-1},\alpha_n}\rangle \otimes |w^{\varepsilon}_{\alpha_n,\alpha_1}\rangle$$

we obtain $|\tau_{\varepsilon}\rangle = |\tau\rangle + \mathcal{O}(\varepsilon)$ and therefore $\underline{\text{nn-rank}}_{\Theta_n}(|\tau\rangle) \leqslant k$. This implies the following chain of inequalities

$$\underline{\text{rank}}_{\Theta_n}(|\tau\rangle) \leqslant \underline{\text{nn-rank}}_{\Theta_n}(|\tau\rangle) \leqslant k < \text{rank}_{\Theta_n}(|\tau\rangle) \leqslant \text{nn-rank}_{\Theta_n}(|\tau\rangle),$$

where the strict inequality is shown in [5, Proposition 5] and the inequalities between rank_{Θ_n} and $\text{nn-rank}_{\Theta_n}$ hold because the latter is a constrained version of the former.

Lemma 2.3.1 cannot be employed to prove a gap for the $\text{psd-rank}_{\Theta_n}$. The existence of an example for the ti cyclic psd decomposition, motivates us to conjecture that:

Conjecture 4.1.9
There is a nonnegative tensor $|T\rangle$ such that

$$\underline{\text{psd-rank}}_{\Theta_n}(|T\rangle) < \text{psd-rank}_{\Theta_n}(|T\rangle).$$

4.1.4 Multipartite positive semidefinite matrices

The three types of positive decompositions for nonnegative tensors are related to the three positive decompositions for multipartite psd matrices (see Proposition 2.3.2). This enables us to translate gaps between border ranks and ranks for positive tensor decompositions to gaps between border rank and rank for multipartite psd matrices. Given a tensor $|T\rangle$ such that $\underline{\text{psd-rank}}_{\Sigma_n}(|T\rangle) < \text{psd-rank}_{\Sigma_n}(|T\rangle)$, the diagonal matrix $\rho_{|T\rangle}$ (Equation (2.11)) satisfies

$$\begin{aligned}\underline{\text{puri-rank}}_{\Sigma_n}(\rho_{|T\rangle}) &\leqslant \underline{\text{psd-rank}}_{\Sigma_n}(|T\rangle)\\ &< \text{psd-rank}_{\Sigma_n}(|T\rangle)\\ &= \text{puri-rank}_{\Sigma_n}(\rho_{|T\rangle}),\end{aligned}$$

and thereby exhibits a gap between border rank and rank for $\text{puri-rank}_{\Sigma_n}$. Analogously one obtains gaps for matrix tensor decompositions whenever there is a gap in the corresponding tensor decomposition. This strategy results in gaps between border rank and rank for puri-rank$_{(\Sigma_n,S_n)}$, puri-rank$_{(\Theta_n,C_n)}$, sep-rank$_{\Theta_n}$, and sep-rank$_{(\Theta_n,C_n)}$.

4.2 Absence of gaps

In the following, we provide the two remaining cases where no gaps between border rank and rank appear. First, we establish that for standard tensor decompositions (i.e. only containing one summation index), the nn-rank$_{\Sigma_n}$, nn-rank$_{(\Sigma_n,S_n)}$, sep-rank$_{\Sigma_n}$, and the sep-rank$_{(\Sigma_n,S_n)}$ do not exhibit a gap. Second, we prove that Ω-decompositions arising from a tree Ω do not exhibit gaps between rank and border rank regardless of the local positivity constraints.

The proof strategy is similar in all cases. When considering a sequence of tensors $|T_k\rangle$ converging to a tensor $|T\rangle$ and their decompositions

$$|T_k\rangle = \sum_{\alpha=1}^{r} |v_{\alpha,k}\rangle \otimes \cdots \otimes |v_{\alpha,k}\rangle ,$$

the local vectors $|v_{\alpha,k}\rangle$ do usually not converge when $k \to \infty$;[7] however, we show that in the specific cases below, every decomposition can be reduced to a normalized version[8]. Then we apply the Bolzano–Weierstraß Theorem to the local elements to guarantee that every sequence of decompositions obtained from a converging sequence of global elements converges to a decomposition of the same rank.

7: This is for example the case for all examples exhibiting a gap. There, the local vectors diverge when approaching the limit.

8: This means that every local element satisfies a normalization constraint

Let us now state the version of the Bolzano–Weierstraß Theorem for finite dimensional normed vector spaces.

Theorem 4.2.1 (Bolzano–Weierstraß)

Let $\mathcal{S} \subseteq \mathcal{V}$ be a compact set in a finite dimensional normed vector space. Then *every* sequence $(s_i)_{i\in\mathbb{N}} \in \mathcal{S}^{\mathbb{N}}$ has a convergent subsequence, i.e. there is a strictly increasing sequence $(k_\ell)_{\ell\in\mathbb{N}}$ in $\mathbb{N}$ such that

$$\lim_{\ell\to\infty} s_{k_\ell} = s \in \mathcal{S}.$$

A set $\mathcal{S}$ in a finite dimensional normed vector space is compact if it is

- closed, i.e. every converging sequence $(s_n)_{n\in\mathbb{N}}$ with sequence elements $s_n \in \mathcal{S}$ has its limit in $\mathcal{S}$, and
- bounded, i.e. there is a $C \in \mathbb{R}$ such that $\|s\| \leqslant C$ for all $s \in \mathcal{S}$.

Note that the choice of the vector space norm in Theorem 4.2.1 does not matter, as all norms that define a finite dimensional vector space are equivalent[9]. For this reason, we will equip the multipartite tensor product space with the most convenient norm to prove the statements.

9: More specifically, if $\|-\|_1$, $\|-\|_2$ are two norms on $\mathcal{V}$, there exist constants $c_1, c_2 > 0$ such that

$$c_1\|v\|_1 \leqslant \|v\|_2 \leqslant c_2\|v\|_1.$$

for every $v \in \mathcal{V}$.

4.2.1 Standard tensor decomposition

Let us now show that nn-rank$_{\Sigma_n}$, nn-rank$_{(\Sigma_n,S_n)}$, sep-rank$_{\Sigma_n}$, and the sep-rank$_{(\Sigma_n,S_n)}$ do not exhibit a gap between rank and border rank.

Theorem 4.2.2

Let $(\rho_k)_{k\in\mathbb{N}}$ be a sequence of n-partite separable matrices with limit

$\rho_k \to \rho$ and $\text{sep-rank}_{\Sigma_n}(\rho_k) \leqslant r$ for every k. Then,

$$\text{sep-rank}_{\Sigma_n}(\rho) \leqslant r$$

The same statement holds for $\text{sep-rank}_{(\Sigma_n, S_n)}$. It also holds for sequences of nonnegative tensors together with $\text{nn-rank}_{\Sigma_n}$, and $\text{nn-rank}_{(\Sigma_n, S_n)}$.

Since the nonnegative decomposition corresponds to the separable decomposition of a diagonal matrix, it suffices to prove the statement for separable decompositions. This generalizes the result in [102], by which the multipartite nonnegative standard tensor decomposition does not exhibit a gap between rank and border rank.

To prove Theorem 4.2.2 we need the following preparatory lemma.

Lemma 4.2.3
Let $A, B \in \text{Psd}_d(\mathbb{C})$. Then,

$$\max\left\{\lambda_{\max}(A), \lambda_{\max}(B)\right\} \leqslant \lambda_{\max}(A+B)$$

Proof. Let

$$\mathcal{R}_X(x) := \frac{\langle x|\, X\, |x\rangle}{\langle x\,|\,x\rangle}$$

for $|x\rangle \in \mathbb{C}^d$. We have that $\mathcal{R}_A(x) + \mathcal{R}_B(x) = \mathcal{R}_{A+B}(x)$ and since A, B are psd, we have that $\mathcal{R}_A(x), \mathcal{R}_B(x) \geqslant 0$ for every x. This implies that

$$\max\left\{\mathcal{R}_A(x), \mathcal{R}_B(x)\right\} \leqslant \mathcal{R}_{A+B}(x).$$

Since

$$\lambda_{\max}(X) = \max_{|x\rangle \in \mathbb{C}^d} \mathcal{R}_X(x),$$

the result follows. □

Proof of Theorem 4.2.2. We prove it for $\text{sep-rank}_{(\Sigma_n, S_n)}$. The proof for $\text{sep-rank}_{\Sigma_n}$ is analogous, and the proof for $\text{nn-rank}_{\Sigma_n}$ and $\text{nn-rank}_{(\Sigma_n, S_n)}$ follows from restricting to diagonal matrices and the fact that[10]

$$\text{nn-rank}_{\Sigma_n}(|T\rangle) = \text{sep-rank}_{\Sigma_n}(\rho_{|T\rangle}).$$

10: we refer to Proposition 2.3.2 for this correspondence.

Let $(\rho_k)_{k\in\mathbb{N}}$ be a sequence of separable matrices with

$$\text{sep-rank}_{(\Sigma_n, S_n)}(\rho_k) \leqslant r,$$

i.e. with a separable decomposition

$$\rho_k = \sum_{\alpha=1}^{r} \rho_{\alpha,k} \otimes \cdots \otimes \rho_{\alpha,k}$$

with $\rho_{\alpha,k}$ psd. Since all elementary tensors are themselves psd, we have that for all α and all k

$$\|\rho_{\alpha,k}\|_\infty^n = \|\rho_{\alpha,k}^{\otimes n}\|_\infty \leqslant \|\rho_k\|_\infty \leqslant \|\rho\|_\infty + C$$

for some constant $C \in \mathbb{N}$, where the first equality is true since

$$\lambda_{\max}(\rho^{\otimes n}) = \lambda_{\max}(\rho)^n,$$

the first inequality follows by Lemma 4.2.3, and the last inequality follows from the convergence of ρ_k to ρ.

This implies that $(\rho_{\alpha,k})_{k\in\mathbb{N}}$ is a bounded sequence. By Bolzano–Weierstraß (Theorem 4.2.1) there is a subsequence $(k_\ell)_{\ell\in\mathbb{N}}$ such that ρ_{α,k_ℓ} converges to a limiting point ρ_α, which is again psd. Since $\rho_k \to \rho$ by assumption, we have that

$$\rho = \sum_{\alpha=1}^{r} \rho_\alpha \otimes \cdots \otimes \rho_\alpha,$$

i.e. $\text{sep-rank}_{(\Sigma_n, S_n)}(\rho) \leqslant r$, which proves the statement. □

4.2.2 Tree tensor networks

Tensor networks without local positivity exhibit border rank phenomena if and only if they contain loops in the hypergraph Ω that specifies the decomposition structure [5]. In particular, if a hypergraph Ω is a tree, the corresponding unconstrained tensor network decomposition exhibits no gap between rank and border rank. In the following we will prove that the same is the case for positive tensor networks. We show the following:

If Ω is a tree[11], then all positive Ω-ranks do not exhibit a gap between border rank and rank.

11: i.e. it corresponds to a graph with $|F| = 2$ for every facet $F \in \mathcal{F}$ and contains no loops of facets, i.e. there is no choice of distinct vertices $i_1, \ldots, i_k \in [n]$ such that

$$\{i_1, i_2\}, \ldots, \{i_{k-1}, i_k\}, \{i_k, i_1\} \in \mathcal{F}.$$

The proof idea is similar to the proof of Theorem 4.2.2. So we first show that every tensor decomposition can be transformed to a normalized version without increasing the rank. Second, we show that applying the limit with respect to the elementary tensors yields a tensor decomposition of the limit element.

The unconstrained decomposition

In this part, we review the result that unconstrained Ω-decompositions on trees Ω do not exhibit a gap between border-rank and rank, i.e.

$$\underline{\text{rank}}_\Omega(|T\rangle) = \text{rank}_\Omega(|T\rangle).$$

The idea is as follows. A tensor decomposition where an index only joins two local spaces, such as

$$|T\rangle = \sum_{\alpha=1}^{r} |v_\alpha\rangle \otimes |w_\alpha\rangle$$

is equivalent to a matrix factorization of the corresponding matrix $T = A \cdot B$ with $A \in \text{Mat}_{d,r}(\mathbb{C})$ and $B \in \text{Mat}_{r,d}(\mathbb{C})$, where each column of A is given by a vector $|v_\alpha\rangle$ and each row of B is given by a vector $|w_\alpha\rangle$. Note that there is a "gauge freedom" in these decompositions, as for every $X \in \text{Mat}_{r,r}(\mathbb{C})$ invertible, $\widetilde{A} = A \cdot X^{-1}$ and $\widetilde{B} = X \cdot B$ give rise to a new decomposition of T of the same rank. Computing a thin (or reduced) QR-decomposition of A [59, Chapter 5], we obtain $A = Q \cdot R$

with Q an isometry in $\mathrm{Mat}_{d,r}(\mathbb{C})$ and $R \in \mathrm{Mat}_r(\mathbb{C})$ an invertible matrix. Hence,

$$\widetilde{A} := Q \quad \text{and} \quad \widetilde{B} := R \cdot B$$

give rise to a decomposition where all tensor factors in the first part form an orthonormal basis, and the local vectors satisfy normalization conditions with respect to the Hilbert–Schmidt norm

$$\|X\|_2 := \sqrt{\operatorname{tr}(X^\dagger X)} = \sqrt{\sum_{i,j=1}^{d} |X_{i,j}|^2},$$

namely $\|\widetilde{A}\|_2 = \sqrt{r}$ and

$$\|T\|_2 = \|\widetilde{A}\widetilde{B}\|_2 = \sqrt{\operatorname{tr}\left(\widetilde{B}^\dagger Q^\dagger Q \widetilde{B}\right)} = \sqrt{\operatorname{tr}\left(\widetilde{B}^\dagger \widetilde{B}\right)} = \|\widetilde{B}\|_2.$$

Similarly, for any tree Ω there exists a normalized Ω-decomposition. Such decompositions are known as *canonical* forms in the tensor network literature[12]

12: We refer to [97] for the left- and right-canonical form on the line, and to [110] for the canonical form on trees. See also [89] for a detailed treatment.

Lemma 4.2.4
Let Ω be a tree and $|\psi\rangle \in \left(\mathbb{C}^d\right)^{\otimes n}$ with $\operatorname{rank}_\Omega(|\psi\rangle) \leqslant r$. There exists a decomposition[13]

13: see Section 2.3.4 for the relation between the structure tensor $|\Omega_r\rangle$ and Ω-decompositions.

$$|\psi\rangle = W^{[1]} \otimes \cdots \otimes W^{[n]} \, |\Omega_r\rangle$$

such that

$$\|W^{[i]}\|_2 = \sqrt{r} \quad \text{for } i = 1, \ldots, n-1, \text{ and} \quad \|W^{[n]}\|_2 = \sqrt{\langle \psi \mid \psi \rangle}$$

Proof. Follows directly from the proof in [5, Proposition 1]. □

Lemma 4.2.4 entails that there is no gap between border rank and rank for unconstrained Ω-decompositions whenever Ω is a tree.

Theorem 4.2.5
If Ω is a tree, then $\operatorname{rank}_\Omega = \underline{\operatorname{rank}}_\Omega$.

Proof. Let $|\psi_k\rangle$ be a sequence of states with $|\psi_k\rangle \to |\psi\rangle$ such that $\operatorname{rank}_\Omega(|\psi_k\rangle) \leqslant r$. We show that $\operatorname{rank}_\Omega(|\psi\rangle) \leqslant r$. By Lemma 4.2.4 there exists tensor decomposition

$$|\psi_k\rangle = W_k^{[1]} \otimes \cdots \otimes W_k^{[n]} \, |\Omega_r\rangle$$

sucht that $\|W_k^{[i]}\|_2 = \sqrt{r}$ for $i = 1, \ldots, n-1$ and $\|W_k^{[n]}\|_2 = \sqrt{\langle \psi_k \mid \psi_k \rangle}$. Since $|\psi_k\rangle \to |\psi\rangle$ there exists a constant C such that

$$\sqrt{\langle \psi_k \mid \psi_k \rangle} \leqslant \sqrt{\langle \psi \mid \psi \rangle} + C$$

which implies that $(W_k^{[i]})_{k \in \mathbb{N}}$ is a bounded sequence for every $i \in [n]$. By the Bolzano–Weierstraß Theorem (Theorem 4.2.1), there exists a subsequence $(W^{[i]}_{k_\ell})_{\ell\in\mathbb{N}}$ converging to a matrix $W^{[i]}$ for every $i \in \{1,\dots,n\}$ which implies that

$$|\psi\rangle = W^{[1]} \otimes \cdots \otimes W^{[n]} \, |\Omega_r\rangle \,,$$

i.e. $\mathrm{rank}_\Omega(|\psi\rangle) \leqslant r$. □

Note that the same results hold for unconstrained Ω-decompositions of multipartite matrices.

The nonnegative and the separable decomposition

Theorem 4.2.6
Let Ω be a tree, $|T\rangle$ a nonnegative tensor and ρ an n-partite separable matrix. Then, the following holds:

(i) $\underline{\text{nn-rank}_\Omega}(|T\rangle) = \text{nn-rank}_\Omega(|T\rangle)$
(ii) $\underline{\text{sep-rank}_\Omega}(\rho) = \text{sep-rank}_\Omega(\rho)$

Similar to the proof of Theorem 4.2.2, we first prove a lemma on the existence of normalized decompositions.

Lemma 4.2.7
Let Ω be a tree and $\rho \in \mathrm{Mat}_d(\mathbb{C})^{\otimes n}$ be a separable matrix with $\text{sep-rank}_\Omega(\rho) \leqslant r$. There exists a separable Ω-decomposition with $|\mathcal{I}| \leqslant r$

$$\rho = \sum_{\alpha\in\mathcal{I}^{\widetilde{\mathcal{F}}}} \rho^{[1]}_{\alpha_{|1}} \otimes \cdots \otimes \rho^{[n]}_{\alpha_{|n}}$$

such that

- $\mathrm{tr}(\rho^{[i]}_\beta) \leqslant 1$ for every $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$ and $i \in [n-1]$
- $\sum_{\beta\in\mathcal{I}^{\widetilde{\mathcal{F}}_i}} \mathrm{tr}(\rho^{[n]}_\beta) = \mathrm{tr}(\rho)$.

We first give an idea of the normalization procedure when Ω is a tree of three vertices according to Figure 4.3. In this case, the separable decomposition of a state ρ is given by

$$\rho = \sum_{\alpha,\gamma=1}^{r} \rho^{[1]}_\alpha \otimes \rho^{[2]}_\gamma \otimes \rho^{[3]}_{\alpha,\gamma} \tag{4.6}$$

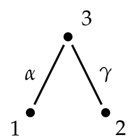

Figure 4.3: A tree with 3 vertices giving rise to the decomposition in Equation (4.6).

Note that none of the local matrices in the decomposition is normalized, except the global one by $\mathrm{tr}(\rho) = 1$. Replacing the first two local families of matrices by

$$\sigma^{[i]}_\beta = \frac{\rho^{[i]}_\beta}{\mathrm{tr}(\rho^{[i]}_\beta)}$$

for $i \in \{1, 2\}$ and the third family by

$$\sigma^{[3]}_{\beta_1,\beta_2} = \operatorname{tr}\left(\rho^{[1]}_{\beta_1}\right) \cdot \operatorname{tr}\left(\rho^{[2]}_{\beta_2}\right) \rho^{[3]}_{\beta_1,\beta_2}$$

we again obtain a separable decomposition

$$\rho = \sum_{\alpha,\gamma=1}^{r} \sigma^{[1]}_{\alpha} \otimes \sigma^{[2]}_{\gamma} \otimes \sigma^{[3]}_{\alpha,\gamma}$$

that satisfies the properties in the lemma, since

$$\operatorname{tr}\left(\sigma^{[1]}_{\beta}\right) = \operatorname{tr}\left(\sigma^{[2]}_{\beta}\right) = 1$$

for $\beta \in \{1, \dots, r\}$ and

$$\begin{aligned}\sum_{\alpha,\beta=1}^{r} \operatorname{tr}\left(\sigma^{[3]}_{\alpha,\gamma}\right) &= \sum_{\alpha,\beta=1}^{r} \operatorname{tr}\left(\rho^{[1]}_{\alpha}\right) \operatorname{tr}\left(\rho^{[2]}_{\beta}\right) \operatorname{tr}\left(\rho^{[3]}_{\alpha,\gamma}\right) \\ &= \sum_{\alpha,\beta=1}^{r} \operatorname{tr}\left(\rho^{[1]}_{\alpha} \otimes \rho^{[2]}_{\beta} \otimes \rho^{[3]}_{\alpha,\gamma}\right) = \operatorname{tr}(\rho)\end{aligned}$$

where we have used the multiplicativity of the trace with respect to the tensor product[14]. Note that a similar normalization procedure can be done for every other arrangement of local spaces.

14: i.e. for two square matrices A, B, we have that

$$\operatorname{tr}(A \otimes B) = \operatorname{tr}(A) \cdot \operatorname{tr}(B).$$

Proof of Lemma 4.2.7. We prove a stronger statement by induction over the number of vertices n. Specifically, we show that for every family $(\rho_\delta)_{\delta \in \mathcal{I}}$ with a joint Ω-decomposition

$$\rho_\delta = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \rho^{[1]}_{\alpha_{|1}} \otimes \rho^{[2]}_{\alpha_{|2}} \otimes \cdots \otimes \rho^{[n-1]}_{\alpha_{|n-1}} \otimes \rho^{[n]}_{\alpha_{|n},\delta} \tag{4.7}$$

the local tensors can be chosen such that $\operatorname{tr}(\rho^{[i]}_{\beta}) = 1$ for $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$ and $i \in \{1, \dots, n-1\}$, and

$$\sum_{\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_n}} \operatorname{tr}(\rho^{[n]}_{\beta,\delta}) = \operatorname{tr}(\rho_\delta).$$

Setting $\delta = 1$ proves the claim. The idea of the induction step is shown in Figure 4.4.

For $n = 1$ (i.e. a single vertex) the statement is trivial.

For the induction step $n-1 \to n$, choose a joint Ω-decomposition according to Equation (4.7) without normalization constraints. We assume without loss of generality that vertex n is connected to precisely two other vertices.[15] We denote the vertices of the first subtree Ω_1 by $\{1, \dots, k_1\}$, and the vertices on the second subtree Ω_2 by $\{k_1+1, \dots, n-1\}$. Moreover, vertices k_1 and $n-1$ are connected to vertex n (Figure 4.4). For this reason, we can rewrite the separable Ω-decomposition ρ_δ as

15: If it is connected to more or less vertices the proof works analogously.

$$\rho_\delta = \sum_{\gamma,\eta \in \mathcal{I}} \rho^{[1,\dots,k_1]}_{\gamma} \otimes \rho^{[k_1+1,\dots,n-1]}_{\eta} \otimes \rho^{[n]}_{\gamma,\eta,\delta}$$

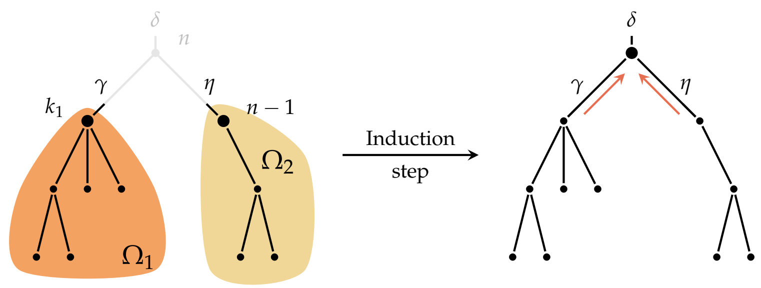

Figure 4.4: Sketch of the induction step in the proof of Lemma 4.2.7. We assume that a normalized decomposition on every subtree Ω_1, Ω_2 exists. This implies that all local elements at the small nodes have trace 1. Large nodes represent local elements whose normalization is given by the global element. In the induction step, we shift the global normalization constraint of node k_1 and $n-1$ to node n.

with

$$\rho_\gamma^{[1,\dots,k_1]} = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{G}}}} \rho_{\alpha_{|1}}^{[1]} \otimes \cdots \otimes \rho_{\alpha_{|k_1},\gamma}^{[k_1]}$$

and

$$\rho_\eta^{[k_1+1,\dots,n-1]} = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{H}}}} \rho_{\alpha_{|k_1+1}}^{[k_1+1]} \otimes \cdots \otimes \rho_{\alpha_{|n-1},\eta}^{[n-1]}$$

where $\widetilde{\mathcal{G}}$ and $\widetilde{\mathcal{H}}$ are the sets of facets of Ω_1 and Ω_2 respectively. Applying the induction hypothesis to $\rho_\gamma^{[1,\dots,k_1]}$ and $\rho_\eta^{[k_1+1,\dots,n-1]}$, we obtain that all tensor factors have trace one, except the tensor factors at position k_1 and $n-1$. There, we have

$$\sum_{\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_{k_1}}} \mathrm{tr}(\rho_{\beta,\gamma}^{[k_1]}) = \mathrm{tr}(\rho_\gamma^{[1,\dots,k_1]})$$

and

$$\sum_{\beta' \in \mathcal{I}^{\widetilde{\mathcal{F}}_{n-1}}} \mathrm{tr}(\rho_{\beta',\eta}^{[n-1]}) = \mathrm{tr}(\rho_\eta^{[k_1+1,\dots,n-1]}).$$

Defining

$$\widetilde{\rho}_{\beta,\gamma}^{[k_1]} := \frac{1}{\mathrm{tr}(\rho_\gamma^{[1,\dots,k_1]})} \rho_{\beta,\gamma}^{[k_1]},$$

$$\widetilde{\rho}_{\beta',\eta}^{[n-1]} := \frac{1}{\mathrm{tr}(\rho_\eta^{[k_1+1,\dots,n-1]})} \rho_{\beta,\eta}^{[n-1]},$$

and

$$\widetilde{\rho}_{\gamma,\eta,\delta}^{[n]} := \mathrm{tr}(\rho_\gamma^{[1\dots k_1]}) \cdot \mathrm{tr}(\rho_\eta^{[k_1+1\dots n-1]}) \cdot \rho_{\gamma,\eta,\delta}^{[n]}$$

we obtain a joint Ω-decomposition

$$\rho_\delta = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \rho_{\alpha_{|1}}^{[1]} \otimes \cdots \otimes \widetilde{\rho}_{\alpha_{|k_1}}^{[k_1]} \otimes \rho_{\alpha_{|k_1+1}}^{[k_1+1]} \otimes \cdots \otimes \widetilde{\rho}_{\alpha_{|n-1}}^{[n-1]} \otimes \widetilde{\rho}_{\alpha_{|n},\delta}^{[n]}$$

that satisfies the desired properties. Since every tree arises by sequentially attaching vertices in the described way, this proves the statement. □

We are now ready to prove the absence of gaps for separable and nonnegative tree tensor decompositions.

Proof of Theorem 4.2.6. The proof is analogous to Theorem 4.2.2. We prove it again only for separable decompositions; the statement for nonnegative decompositions follows by considering separable decompositions of diagonal matrices. Let $(\rho_k)_{k\in\mathbb{N}}$ be a sequence of separable matrices such that $\text{sep-rank}_\Omega(\rho_k) \leqslant r$ and $\rho_k \to \rho$. We show that $\text{sep-rank}_\Omega(\rho) \leqslant r$. To this end, let

$$\rho_k = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \rho^{[1]}_{\alpha_{|1},k} \otimes \cdots \otimes \rho^{[n]}_{\alpha_{|n},k}$$

be a normalized decomposition according to Lemma 4.2.7. We have that $\text{tr}(\rho^{[i]}_{\beta,k}) = 1$ for every $i \in \{1, \ldots, n-1\}$ and $\text{tr}(\rho^{[n]}_{\beta,k}) \leqslant \text{tr}(\rho) + C$ for a suitable choice of C due to the convergence $\rho_k \to \rho$. Hence, every tensor factor is a bounded sequence which has a convergent subsequence $\rho^{[i]}_{\beta,k_\ell} \to \rho^{[i]}_\beta$ for $\ell \to \infty$ due to Theorem 4.2.1. Since $\rho_k \to \rho$, we have that

$$\rho = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \rho^{[1]}_{\alpha_{|1}} \otimes \cdots \otimes \rho^{[n]}_{\alpha_{|n}}$$

which shows that $\text{sep-rank}_\Omega(\rho) \leqslant r$. □

The psd decomposition and the local purification form

We now prove that for every tree Ω, neither psd Ω-decompositions nor Ω-purifications exhibit a gap between rank and border rank. The proof strategy is similar to other cases without gaps: We use that there is a bounded decomposition with the same expressiveness and then apply the Bolzano–Weierstraß Theorem. In this case, we additionally use the correspondence to correlation scenarios (Theorem 3.2.1) and the absence of gaps for unconstrained decompositions (Theorem 4.2.5).

Theorem 4.2.8
Let Ω be a tree, $|T\rangle$ a nonnegative tensor and ρ a psd matrix. Then,

(i) $\underline{\text{psd-rank}}_\Omega(|T\rangle) = \text{psd-rank}_\Omega(|T\rangle)$
(ii) $\underline{\text{puri-rank}}_\Omega(\rho) = \text{puri-rank}_\Omega(\rho)$

To prove the theorem, we need the following preparatory lemma:

Lemma 4.2.9
For every sequence of quantum channels

$$\Big(\mathcal{E}_k : \text{Mat}_{d_1}(\mathbb{C}) \to \text{Mat}_{d_1}(\mathbb{C})\Big)_{k \in \mathbb{N}},$$

there exists a convergent subsequence.

Proof. Let $\mathsf{Lin}(d_1, d_2)$ be the set of all linear maps

$$L\colon \text{Mat}_{d_1}(\mathbb{C}) \to \text{Mat}_{d_2}(\mathbb{C}).$$

We prove that the set

$$\mathsf{CPTP}(d_1, d_2) := \{\mathcal{E} \in \mathsf{Lin}(d_1, d_2) \colon \mathcal{E} \text{ is cptp}\}$$

is compact in $\mathsf{Lin}(d_1, d_2)$. The statement follows then by Bolzano–Weierstraß (Theorem 4.2.1).

Equipping the space $\mathsf{Lin}(d_1, d_2)$ with the norm

$$\|\mathcal{E}\| := \max_{\|\rho\|_1 \leqslant 1} \|\mathcal{E}(\rho)\|_1$$

where $\|\cdot\|_1$ is the trace-norm on $\mathrm{Mat}_{d_i}(\mathbb{C})$, we obtain that $\|\mathcal{E}\| \leqslant 1$ for every $\mathcal{E} \in \mathsf{CPTP}(d_1, d_2)$, which shows the boundedness.

Moreover, $\mathsf{CPTP}(d_1, d_2)$ is closed since it can be characterized by the closed conditions $\mathrm{id}_n \otimes \mathcal{E}(A) \geqslant 0$ for every psd $A \in \mathrm{Mat}_{d_1 \cdot n}(\mathbb{C})$ and $\mathrm{tr}(\mathcal{E}(\rho)) = \mathrm{tr}(\rho)$ for every $\rho \in \mathrm{Mat}_{d_1}(\mathbb{C})$. Since intersections of closed sets are closed, we obtain compactness of $\mathsf{CPTP}(d_1, d_2)$. ☐

Proof of Theorem 4.2.8. We prove the statement only for puri-rank$_\Omega$ as the case of psd-rank$_\Omega$ works similarly. Let $(\rho_k)_{k\in\mathbb{N}}$ be a sequence of psd matrices such that puri-rank$_\Omega(\rho_k) \leqslant r$ and $\rho_k \to \rho$. We need to prove that puri-rank$_\Omega(\rho) \leqslant r$.

By Theorem 3.2.1 there exists a sequence of states $|\psi_k\rangle$ with $\mathrm{rank}_\Omega(|\psi_k\rangle) \leqslant r$ and a sequence of quantum channels $\mathcal{E}_i^{(k)}$ for every $i \in [n]$ such that

$$\rho_k = \left(\mathcal{E}_1^{(k)} \otimes \cdots \otimes \mathcal{E}_n^{(k)}\right) \left(|\psi_k\rangle \langle\psi_k| \right).$$

Since the space of quantum states is compact (we have that $\langle\psi \mid \psi\rangle = 1$ for every $|\psi\rangle$), and by Lemma 4.2.9, there exists a joint subsequence k_ℓ such that

$$\mathcal{E}_i := \lim_{\ell\to\infty} \mathcal{E}_i^{(k_\ell)} \quad \text{and} \quad |\psi\rangle := \lim_{\ell\to\infty} |\psi_{k_\ell}\rangle,$$

which implies that

$$\rho = (\mathcal{E}_1 \otimes \cdots \otimes \mathcal{E}_n)(|\psi\rangle\langle\psi|).$$

Since $\mathrm{rank}_\Omega = \underline{\mathrm{rank}}_\Omega$ (see Theorem 4.2.5), we have that $\mathrm{rank}_\Omega(|\psi\rangle) \leqslant r$, which proves that puri-rank$_\Omega(\rho) \leqslant r$.

The proof for the psd-rank similarly uses Theorem 3.1.2 and the fact that every sequence of a POVM has a convergent subsequence that converges to a POVM by the Bolzano–Weierstraß Theorem. ☐

4.3 Applications

Let us now present three implications of the existence and absence of gaps between ranks and border ranks:

- In Section 4.3.1 we show that the existence of gaps leads to instabilities for optimization problems over tensor network manifolds.
- In Section 4.3.2 we prove a correspondence between postive tensor decompositions and quantum correlation sets. The gaps between border ranks and ranks then imply that certain sets of quantum correlations are not closed.
- In Section 4.3.3 we prove that gaps also lead to new types of separations between positive tensor ranks.

4.3.1 Instability in optimization

Tensors are in general very costly to represent. For this reason, one often restricts to approximate representations with a restriction on the rank of the approximation. In this context, one wants to find for a given n-partite tensor $|T\rangle \in \left(\mathbb{C}^d\right)^{\otimes n}$ the best rank r approximation of $|T\rangle$, i.e.

$$\begin{aligned} &\text{minimize } \| \, |T\rangle - |W\rangle \, \| \\ &\text{subject to } \operatorname{rank}(|W\rangle) \leqslant r. \end{aligned} \tag{4.8}$$

For the case of bipartite tensors (i.e. matrices), this minimization problem has an analytic solution by the Eckart–Young–Mirsky theorem.[16] Specifically, if $\|\cdot\|$ is an unitarily invariant norm[17], then for every matrix $A \in \operatorname{Mat}_d(\mathbb{C})$ with singluar value decomposition

$$A = \sum_{k=1}^{d} \sigma_k \, |u_k\rangle \, \langle v_k|$$

and singular values $\sigma_1 \geqslant \sigma_2 \geqslant \ldots \geqslant \sigma_d \geqslant 0$, the solution of Equation (4.8) is given by

$$A_r := \sum_{k=1}^{r} \sigma_k \, |u_k\rangle \, \langle v_k|$$

i.e. considering the largest r singular values.

For other norms, no analytic formula is given; however, Equation (4.8) has a solution since the set of feasible points

$$\mathcal{T} := \{A \in \operatorname{Mat}_d(\mathbb{C}) \colon \operatorname{rank}(A) \leqslant r\}$$

is topologically closed. This is equivalent to the matrix-rank being *lower semi-continuous*, i.e. for every sequence $A_k \to A$ for $k \to \infty$, we have

$$\operatorname{rank}(A_k) \leqslant r \quad \Longrightarrow \quad \operatorname{rank}(A) \leqslant r. \tag{4.9}$$

Also for positive matrix ranks, Equation (4.8) has a solution. This again follows from the fact that the nonnegative and the psd matrix rank[18] are lower-semi continuous. For a direct proof of these results we refer to [18] for the nonnegative matrix rank and to [49, Theorem 2.12] for the positive semidefinite matrix rank. Note that these results are a special case of Theorem 4.2.6, Theorem 4.2.5, and Theorem 4.2.8 considering the tree with two vertices and one edge.

By Equation (4.9), we have shown that the multipartite tensor ranks are lower-semicontinuous for tree structures, which implies that the best rank r approximation problem has a solution in these cases.[19]

However, the gaps between border rank and rank exemplify that Equation (4.8) does not have a solution for arbitrary tensor decompositions. For example the problem

$$\begin{aligned} &\text{minimize } \| \, |T\rangle - |W_n\rangle \, \| \\ &\text{subject to } \operatorname{rank}_{\Sigma_n}(|T\rangle) \leqslant 2 \end{aligned} \tag{4.10}$$

16: This result goes back to Eckart and Young [45] for the Frobenius norm, and to Mirsky [85] for arbitrary unitarily invariant norms.

17: A norm $\|\cdot\|$ is called *unitarily invariant* if

$$\|UA\| = \|A\|$$

for every matrix A and unitary matrix U. Examples of unitarily invariant norms are the Frobenius norm, the spectral norm, or more generally every Schatten p-norm with parameter $p \geqslant 1$.

18: See Example 2.3.5 and Example 2.3.7 for the definition of these ranks.

19: Similar to the positive matrix factorizations, this result does not say anything about the efficiency of computing these approximations.

where $|W_n\rangle$ is the n-partite W-state does not have a solution because we can find a rank-2 approximation of $|W_n\rangle$ for every approximation error $\varepsilon > 0$ (see Corollary 4.1.2). In other words, the set of feasible tensors

$$\mathcal{T} := \left\{ |T\rangle \in \left(\mathbb{C}^d\right)^{\otimes n} : \operatorname{rank}_{\Sigma_n}(|T\rangle) \leqslant 2 \right\}$$

is not closed for $n \geqslant 3$.

In summary, we have the following statement:

Observation 4.3.1

For any type of t-rank, the minimization problem of Equation (4.8) has no solution if and only if there is a gap between border rank and rank for this t-rank.

4.3.2 Quantum correlation scenarios

In Chapter 3 we proved a correspondence between positive tensor decompositions and correlation scenarios. We now show that these correspondences together with the gaps between ranks and border ranks imply that the sets of correlations are not closed. It follows that it is generally impossible to test membership of a probability distribution in these sets with a finite number of measurements.

We prove non-closedness for $\mathsf{CQCorr}_{(\Theta_n,C_n)}(n,d,r)$, for other sets, the argument is analogous. Let $(|P_k\rangle)_{k\in\mathbb{N}}$ be a sequence of tensors representing a probability distribution with $\lim_{k\to\infty}|P_k\rangle = |P\rangle$ and exhibiting a gap between rank and border rank (see Proposition 4.1.6), i.e.

$$\text{psd-rank}_{(\Theta_n,C_n)}(|P_k\rangle) \leqslant r < \text{psd-rank}_{(\Theta_n,C_n)}(|P\rangle).$$

Then $P_k \in \mathsf{CQCorr}_{(\Theta_n,C_n)}(n,d,r)$ for all $k \in \mathbb{N}$ while

$$P \notin \mathsf{CQCorr}_{(\Theta_n,C_n)}(n,d,r),$$

i.e. $\mathsf{CQCorr}_{(\Omega,G)}(n,d,r)$ is not closed.

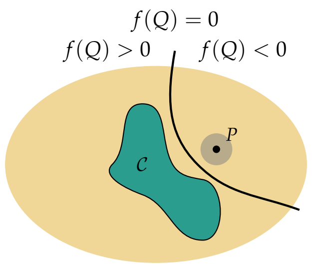

Figure 4.5: A witness function f for a given probability distribution P outside of a subset $\mathcal{C}$. f separates P from $\mathcal{C}$. Since $f(P) < 0$ and f is continuous, it remains negative for a small neighborhood of P. This is only possible if $\mathcal{C}$ is closed.

The closedness of correlation sets is essential to test membership. Certifying that a probability distribution P does not arise from a certain correlation scenario is based on constructing a continuous witness function

$$f : \left(\mathbb{R}^d\right)^{\otimes n} \to \mathbb{R}$$

that satisfies the following properties:

- $f(Q) > 0$ for every $Q \in \mathsf{CQCorr}_{(\Omega,G)}(n,d,r)$
- $f(P) < 0$

Guessing P from finitely many samples results in an approximation $\widetilde{P}$ that is close to P with high probability. Therefore, if the guess $\widetilde{P}$ satisfies $f(\widetilde{P}) < 0$, we can infer that P does not arise from the correlation scenarios with high probability. This follows from the fact that if $f(P) < 0$, then also $f(\widetilde{P}) < 0$, if $\widetilde{P}$ is in some neighborhood of P (see Figure 4.5).

But such witness functions only exist if $\mathsf{CQCorr}_{(\Omega,G)}(n,d,r)$ is closed. If $P \notin \mathsf{CQCorr}_{(\Omega,G)}(n,d,r)$ lies on the boundary, a potential witness function must 'jump' in P which contradicts its continuity. Thus it is impossible to witness $P \notin \mathsf{CQCorr}_{(\Omega,G)}(n,d,r)$ from finitely many samples of the probability distribution.

According to the gaps between ranks and border ranks (see Figure 4.2) the same behavior appears in the following cases:

- Testing the rank_{Σ_n} for $n \geqslant 5$.
- Symmetrically testing $\text{rank}_{(\Sigma_n,S_n)}$ for $n \geqslant 3$.
- Symmetrically testing $\text{rank}_{(\Theta_n,C_n)}$ for $n \geqslant 17$.

Analogously, one can show that $\mathsf{QQCorr}_{(\Omega,G)}(n,d,r)$ is not closed in the above situations.

In contrast, the set of classical correlations $\mathsf{CCorr}(n,d,r)$ is closed for every choice of $n,d,r \in \mathbb{N}$. This follows from the fact that $\text{nn-rank}_{\Sigma_n}$ does not exhibit a gap between border rank and rank, and hence for every converging sequence of nonnegative tensors $|P_k\rangle \to |P\rangle$ with $\text{nn-rank}_{\Sigma_n}(P_k) \leqslant r$ we also have $\text{nn-rank}_{\Sigma_n}(|P\rangle) \leqslant r$. For every $P \notin \mathsf{CCorr}(n,d,r)$ there exists a separating witness since the distance between $\mathsf{CCorr}(n,d,r)$ and P is strictly positive. Moreover, the sets of quantum correlations $\mathsf{CQCorr}_\Omega(n,d,r)$ and $\mathsf{QQCorr}_\Omega(n,d,r)$ are closed if Ω is a tree.

4.3.3 Separations for approximate tensor decompositions

Various notions of positive tensor decompositions exhibit separations [49, 68], meaning that there exist families of bipartite tensors $(|T_d\rangle)_{d\in\mathbb{N}}$ where $|T_d\rangle \in \mathbb{C}^d \otimes \mathbb{C}^d$ such that

$$\text{rank}(|T_d\rangle) = \text{const.} \quad \text{and} \quad \text{psd-rank}(|T_d\rangle) \to \infty$$

as $d \to \infty$. Moreover, there is also a family of bipartite tensors $(|S_d\rangle)_{d\in\mathbb{N}}$ such that

$$\text{psd-rank}(|S_d\rangle) = \text{const.} \quad \text{and} \quad \text{nn-rank}(|S_d\rangle) \to \infty.$$

Are these separations robust with respect to approximations? In [38] it is proven that for fixed approximation error $\varepsilon > 0$ and a fixed norm, the separations between rank_Ω, psd-rank_Ω and nn-rank_Ω disappear. More precisely, $\text{rank}^\varepsilon(T)$, $\text{psd-rank}^\varepsilon(T)$, $\text{nn-rank}^\varepsilon(T)$ (see Equation (4.1)) can be upper bounded by a function depending only on ε and $\|T\|$, independent of the dimension of the tensor product space. However, if the choice of $\varepsilon > 0$ and vector space dimension is too small, this upper bound exceeds trivial dimension-dependent upper bounds. So the bounds are only meaningful when the dimension of the tensor product space is large.

We will now prove a "dual" statement. If the dimension of the tensor product space is fixed, there exists an error $\varepsilon > 0$ such that the separation between rank and nn-rank persists.

Theorem 4.3.2
There exists a family of nonnegative tensors $(|T_n\rangle)_{n\in\mathbb{N}}$ with

$$|T_n\rangle \in \left(\mathbb{C}^d\right)^{\otimes n}$$

and a family of approximation errors $\varepsilon_n > 0$ such that

$$\text{nn-rank}^{\varepsilon_n}(|T_n\rangle) = n.$$

We have also that

$$\text{rank}^{\varepsilon}_{\Sigma_n}(|T_n\rangle) = \text{psd-rank}^{\varepsilon}_{\Sigma_n}(|T_n\rangle) = 2$$

for every $\varepsilon > 0$ independent of n.

Proof. Let $|T_n\rangle := |W_n\rangle$ the family of n-partite W-states. For fixed $n \in \mathbb{N}$, we know that

$$\underline{\text{nn-rank}}_{\Sigma_n}(|W_n\rangle) = \text{nn-rank}_{\Sigma_n}(|W_n\rangle) = n.$$

Therefore there exists a $\varepsilon_n > 0$ such that

$$\text{nn-rank}^{\varepsilon_n}_{\Sigma_n}(|W_n\rangle) = n.$$

For the second statement, recall that

$$\underline{\text{rank}}_{\Sigma_n}(|W_n\rangle) = \underline{\text{psd-rank}}_{\Sigma_n}(|W_n\rangle) = 2.$$

Since

$$\text{rank}^{\varepsilon}_{\Sigma_n}(|W_n\rangle) \leqslant \underline{\text{rank}}_{\Sigma_n}(|W_n\rangle) = 2$$

and

$$\text{psd-rank}^{\varepsilon}_{\Sigma_n}(|W_n\rangle) \leqslant \underline{\text{psd-rank}}_{\Sigma_n}(|W_n\rangle) = 2$$

for every $\varepsilon > 0$, this proves the statement. □

4.4 Conclusions and outlook

In this chapter, we have shown that many gaps between ranks and border ranks persist when introducing positivity and invariance constraints for tensor decompositions, and explored its consequences. More precisely, we have proven that:

- The standard and symmetric tensor decompositions exhibit gaps between border rank and rank for the psd-decomposition and local purifications (Subsection 4.1.1), and the gaps disappear for the nonnegative and separable decomposition (Theorem 4.2.2);
- Most of the gaps persist for cyclic and translational invariant decompositions (Subsection 4.1.3 and Subsection 4.1.2);
- There are no gaps for tree tensor decompositions, regardless of positivity constraints (Theorem 4.2.6);

Many of the examples exhibiting a separation are n-partite tensor decompositions with $n > 3$. This leaves open the question whether gaps between border ranks and ranks exist for positive and invariant 3-partite decompositions.

Other surprising properties of tensor decompositions appearing already at $n = 3$ include the fact that tensor rank and border rank are non-additive with respect to the direct sum [109, 111, 27], and that they are also non-multiplicative with respect to tensor products [28, 26]. Do these properties also hold for positive and invariant decompositions? And what are their implications for correlation scenarios?

Polynomial decompositions inspired by tensors 5

This chapter is based on Section 1, 3, 4, and 7 in [39].

In Chapter 2, we introduced two classes of objects: nonnegative tensors and multipartite psd matrices. These two classes encompass the two central structures studied in this part, namely a tensor product structure and a notion of global positivity.

In this chapter, we introduce yet another vector space structure encompassing these two elements: *real multivariate polynomials*. These are objects in the tensor product space of polynomials in each of their variables,

$$\mathcal{P} := \mathbb{R}[\mathbf{x}^{[1]}, \mathbf{x}^{[2]}, \dots, \mathbf{x}^{[n]}] \cong \mathbb{R}[\mathbf{x}^{[1]}] \otimes \mathbb{R}[\mathbf{x}^{[2]}] \otimes \cdots \otimes \mathbb{R}[\mathbf{x}^{[n]}],$$

where $\otimes$ denotes the algebraic tensor product and $\mathbf{x}^{[i]}$ a collection of variables $x_1^{[i]}, \dots x_{m_i}^{[i]}$. In other words, every polynomial $p \in \mathcal{P}$ can be expressed as a finite sum of "elementary constituents"

$$p^{[1]}(\mathbf{x}^{[1]}) \cdot p^{[2]}(\mathbf{x}^{[2]}) \cdots p^{[n]}(\mathbf{x}^{[n]}),$$

where every $p^{[i]}$ is itself a polynomial that only depends on the variables $\mathbf{x}^{[i]}$. We consider two questions:

- If p is symmetric under the exchange of, say, systems i and j, can this symmetry be reflected in the decomposition?
- If p is positive (for some notion of positivity), can this positivity be reflected in the decomposition?

Our framework addresses these two questions as follows, when applied to polynomials:

(a) The summation structure is described by a weighted simplicial complex Ω, so that every system i is associated to a vertex of Ω, and every summation index to a facet of Ω.
(b) By definition, an (Ω, G)-decomposition of a polynomial contains a certificate of invariance under the group G. We characterize which G-invariant polynomials admit an (Ω, G)-decomposition.
(c) By definition, a separable or sum-of-squares (sum-of-squares (sos)) (Ω, G)-decomposition contains a certificate of invariance and of membership in the separable or sos cone, respectively. We characterize which separable or sos polynomials admit such decompositions.

Our framework models symmetries as follows: we have a group G acting on the set $\{1, \dots, n\}$, and the induced action on the polynomial space $\mathcal{P}$ is obtained by permuting system $[i]$ to $[gi]$,

$$g : \mathbf{x}^{[i]} \mapsto g\mathbf{x}^{[i]} := \mathbf{x}^{[gi]}. \tag{5.1}$$

A polynomial is G-invariant if it is invariant with respect to all such permutations $g \in G$, and we want to make this invariance explicit in the decomposition of p. For example, the decomposition

$$p = \sum_{\alpha_1,\dots,\alpha_n=1}^{r} p_{\alpha_1,\alpha_2}(\mathbf{x}^{[1]}) \cdot p_{\alpha_2,\alpha_3}(\mathbf{x}^{[2]}) \cdots p_{\alpha_n,\alpha_1}(\mathbf{x}^{[n]}) \tag{5.2}$$

Note that there are no superscripts $[i]$ in the polynomials in the invariant decompositions.

makes explicit that p is invariant under the cyclic group, $\mathbf{x}^{[i]} \mapsto \mathbf{x}^{[i+1]}$. And

$$p = \sum_{\alpha=1}^{r} p_\alpha(\mathbf{x}^{[1]}) \cdot p_\alpha(\mathbf{x}^{[2]}) \cdots p_\alpha(\mathbf{x}^{[n]}) \tag{5.3}$$

makes explicit that p is invariant under the full symmetry group.

Finally, if p is in a cone[1], we want a certificate of this fact (cf. (c)). In quantum physics, a mixed quantum state is represented by a psd matrix and the certificate is called a purification. In probabilistic modelling, the certificate of a probability distribution is a nonnegative decomposition. In real algebraic geometry, the natural certificate of positivity of a polynomial is being sum of squares. In all of these cases, witnessing the positivity of a global element is a central problem with many ramifications.

1: Examples of cones are the sum-of-squares (sos) polynomials, the cone of nonnegative polynomials, or the cone of polynomials with nonnegative coefficients.

Note that decompositions of tensors and polynomials have been studied from different perspectives. Also symmetries and positivity have been considered together, but the arising decompositions are by far not as clean as the corresponding separate decompositions. To give a short overview, and also motivate our combined approach, let us explain some of the existing decompositions, and point out why they are not directly related to our approach.

- The Waring decomposition is a decomposition of polynomials, also inspired by tensors. Let $p \in \mathbb{R}[x_1, \ldots, x_n]$ of degree d. The *Waring rank* of p is defined as the minimum $r \in \mathbb{N}$ such that

$$p = \sum_{\alpha=1}^{r} c_\alpha \ell_\alpha(x_1, \ldots, x_n)^d$$

where $\ell_\alpha(x_1, \ldots, x_n) = a_{\alpha,1}x_1 + \ldots + a_{\alpha,n}x_n$ is a linear form. The Waring rank is equivalent to the symmetric tensor rank via the correspondence

$$p = \sum_{j_1,\ldots,j_n=1}^{d} \langle j_1, \ldots, j_n \mid T \rangle \, x_{j_1} \cdots x_{j_n}$$

between symmetric tensors in $T \in \left(\mathbb{C}^d\right)^{\otimes n}$ and homogeneous polynomials of degree n. Yet, the Waring decomposition cannot exhibit any additional symmetry of the polynomial, since the corresponding tensor is already fully symmetric for any polynomial. For generalizations of the Waring problem to polynomials instead of linear forms, we refer to [52]. Another related decomposition is the completely decomposable decomposition [1].
- For symmetric polynomials, the decomposition into power-sum polynomials is an example of an explicitly invariant decomposition. Every symmetric polynomial p can be written as $p = q(p_1, \ldots, p_n)$, where

$$p_\alpha = \sum_{i=1}^{n} x_i^\alpha .$$

In other words, the ring of symmetric polynomials with real coefficients corresponds to the ring $\mathbb{R}[p_1, \ldots, p_n]$ generated by power-sum polynomials. The same statement is true by replacing the set of power-sum polynomials by elementary symmetric polynomials.

- Also, the combination of symmetry and positivity is well-studied. It is, for example, known that symmetric sum-of-squares polynomials do, in general, not decompose into a sum of symmetric squares, to fully characterize the set of symmetric sum-of-squares polynomials, one has to introduce a more general notion of symmetric sum-of-square decomposition [43].

In this chapter, we do the following:

- We define invariant decompositions of polynomials (Definition 5.1.1). We show that every invariant polynomial admits an invariant decomposition if the group action is free on the weighted simplicial complex (Theorem 5.1.2). In addition, every invariant polynomial can be written as the difference of two invariant decompositions if the group action is blending (Theorem 5.1.7).
- We define the invariant separable decomposition (Definition 5.1.2), and the invariant sos decomposition (Definition 5.1.3), and show that every invariant separable/sos polynomial admits an invariant separable/sos decomposition if the group action is free (Theorem 5.1.8 and Corollary 5.1.12, respectively). These decompositions combine positivity and symmetry in a clean way.
- We provide inequalities and separations between the ranks of three invariant decompositions (Proposition 5.2.2 and Corollary 5.2.6, respectively).

5.1 Invariant polynomial decompositions

In this section we define invariant polynomial decompositions and their ranks. To this end we first set the stage (Section 5.1.1), define and study the invariant decomposition (Section 5.1.2), the invariant separable decomposition (Section 5.1.3), and finally the invariant sum-of-squares decomposition (Section 5.1.4).

5.1.1 Setting the stage

Throughout this section we consider polynomials in the space

$$\mathcal{P} := \mathbb{R}[\mathbf{x}^{[1]}, \mathbf{x}^{[2]}, \dots, \mathbf{x}^{[n]}] \cong \mathbb{R}[\mathbf{x}^{[1]}] \otimes \mathbb{R}[\mathbf{x}^{[2]}] \otimes \cdots \otimes \mathbb{R}[\mathbf{x}^{[n]}]$$

where $\mathbb{R}[\mathbf{x}^{[i]}] := \mathbb{R}[x_1^{[i]}, \dots, x_{m_i}^{[i]}]$ is the space of real polynomials in m_i variables, and $\otimes$ denotes the algebraic tensor product. These polynomials use collections of local variables, denoted $\mathbf{x}^{[i]}$, for each local site $i = 1, \dots, n$. The case where all $m_i = 1$ is already very interesting, as it describes how the multivariate polynomial ring is decomposed into a tensor product of univariate polynomial rings.

In particular,

$$\mathbb{R}[x^{[1]}, \dots, x^{[n]}] \cong \mathbb{R}[x^{[1]}] \otimes \mathbb{R}[x^{[2]}] \otimes \cdots \otimes \mathbb{R}[x^{[n]}],$$

where $x^{[i]}$ is a single variable, means that every multivariate polynomial can be expressed as a sum of products of univariate polynomials, i.e.

$$p = \sum_{\alpha=1}^{r} p_\alpha^{[1]}(x^{[1]}) \cdots p_\alpha^{[n]}(x^{[n]}).$$

We define the local degree of $p \in \mathcal{P}$, denoted $\deg_{\text{loc}}(p)$, as the smallest positive integer $d \in \mathbb{N}$ such that

$$p \in \mathcal{P}_d := \mathbb{R}[\mathbf{x}^{[1]}]_d \otimes \mathbb{R}[\mathbf{x}^{[2]}]_d \otimes \cdots \otimes \mathbb{R}[\mathbf{x}^{[n]}]_d$$

where $\mathbb{R}[\mathbf{x}]_d$ is the space of real polynomials in $\mathbf{x}$ of degree at most d. A polynomial with $\deg_{\text{loc}}(p) \leqslant d$ contains monomials consisting of variables in $\mathbf{x}^{[i]}$ with degree at most d, for each i. Note that the local degree can be related to the (global) degree of the polynomial by

$$\deg_{\text{loc}}(p) \leqslant \deg(p) \leqslant n \cdot \deg_{\text{loc}}(p).$$

A group action G on $[n]$ induces a group action on the space $\mathcal{P}$, defined for $g \in G$ and $p \in \mathcal{P}$ by

$$(gp)(\mathbf{x}^{[1]}, \dots, \mathbf{x}^{[n]}) := p(\mathbf{x}^{[g1]}, \dots, \mathbf{x}^{[gn]}). \tag{5.4}$$

Note that this definition only makes sense if the local polynomial spaces $\mathbb{R}[\mathbf{x}^{[i]}]$ and $\mathbb{R}[\mathbf{x}^{[j]}]$ are isomorphic whenever $i, j \in [n]$ are in the same orbit of G (i.e. $gi = j$ for some $g \in G$), i.e. the number of local variables needs to coincide for i, j, namely $m_i = m_j$. The canonical isomorphism between elements in $\mathbb{R}[\mathbf{x}^{[i]}]$ and $\mathbb{R}[\mathbf{x}^{[j]}]$ is given by replacing the variables $\mathbf{x}^{[i]}$ with $\mathbf{x}^{[j]}$ in every polynomial and vice versa. We will frequently use this isomorphism in an implicit way, as for a polynomial $p^{[i]} \in \mathbb{R}[\mathbf{x}^{[i]}]$ we will denote its corresponding element in $\mathbb{R}[\mathbf{x}^{[j]}]$ as $p^{[i]}(\mathbf{x}^{[j]})$.

We say that $p \in \mathcal{P}$ is *G-invariant* if for each $g \in G$ we have $gp = p$, or equivalently

$$p(\mathbf{x}^{[g1]}, \dots, \mathbf{x}^{[gn]}) = p(\mathbf{x}^{[1]}, \dots, \mathbf{x}^{[n]}) \quad \text{for every } g \in G.$$

For example, if $m_i = 1$ and G is the full permutation group on $[n]$, then a polynomial p is invariant if

$$p(x^{[1]}, \dots, x^{[n]}) = p(x^{[\sigma(1)]}, \dots, x^{[\sigma(n)]})$$

for every permutation $\sigma : [n] \to [n]$, which means that p is invariant with respect to arbitrary permutations of variables.

Similar to the tensor decompositions in Chapter 2, we consider $\mathcal{I}$ to be a finite index set, and write a map $\alpha \colon \widetilde{\mathcal{F}} \to \mathcal{I}$ as a tuple $\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}$ with entries from $\mathcal{I}$ indexed by the facets in $\widetilde{\mathcal{F}}$. If we have a function $\alpha \colon \widetilde{\mathcal{F}} \to \mathcal{I}$ and want to restrict its domain to $\widetilde{\mathcal{F}}_i$ (for some index $i \in [n]$), in the tuple notation we again write

$$\alpha_{|_i} := \alpha_{|_{\widetilde{\mathcal{F}}_i}} \in \mathcal{I}^{\widetilde{\mathcal{F}}_i},$$

which means that we delete all entries which are indexed by a facet not containing i. We will in general stick to the functional notation except for the examples, where we will switch to the tuple notation. Their connection will be made explicit in the examples.

5.1.2 The invariant decomposition

We now define the basic invariant decomposition similar to Definition 2.3.1, called (Ω, G)-decomposition. Afterwards we will study the existence of decompositions without invariance, the existence of invariant decompositions with free group actions and with blending group actions.

The idea of the invariant decomposition is to consider finite sums of elementary polynomials (i.e. polynomials written as a product of local polynomials depending on one collection of variables $\mathbf{x}^{[i]}$), where each local polynomial is associated to a vertex of Ω, and the summation indices are described as functions $\alpha_{|_i}$ on the facets. The following definition is illustrated in Example 5.1.1, Example 5.1.2, Example 5.1.3, Example 5.1.4, and Example 5.1.5.

Definition 5.1.1 ((Ω, G)-decomposition of polynomials)
Let $p \in \mathcal{P}$. An (Ω, G)-*decomposition* of p consists of a finite index set $\mathcal{I}$ and families of polynomials

$$\mathcal{P}^{[i]} := \left(p_\beta^{[i]}\right)_{\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}}$$

where $p_\beta^{[i]} \in \mathbb{R}\left[\mathbf{x}^{[i]}\right]$ for all $i \in [n]$, such that

(a) p can be written as

$$p = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} p_{\alpha_{|_1}}^{[1]}(\mathbf{x}^{[1]}) \cdots p_{\alpha_{|_n}}^{[n]}(\mathbf{x}^{[n]})$$

(b) For all $i \in [n]$, $g \in G$ and $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$ we have

$$p_\beta^{[i]}(\mathbf{x}^{[i]}) = p_{{}^g\beta}^{[gi]}(\mathbf{x}^{[i]})$$

where ${}^g\beta$ is defined in Equation (2.4).

The minimal cardinality of $\mathcal{I}$ among all (Ω, G)-decomposition of p is the (Ω, G)-*rank of* p, denoted $\operatorname{rank}_{(\Omega,G)}(p)$. If p does not admit an (Ω, G)-decomposition, we set $\operatorname{rank}_{(\Omega,G)}(p) = \infty$.

Also, if G is the trivial group action, we call the (Ω, G)-decomposition just Ω-*decomposition* and denote its rank by $\operatorname{rank}_\Omega$.

Condition (a) provides an arrangement of the summation indices encoded in the functions α, and Condition (b) ensures that the decomposition has the desired symmetry by requiring that the coefficients of local polynomials in different local spaces coincide. Note again that this equality only makes sense if the collections $\mathbf{x}^{[i]}$ and $\mathbf{x}^{[gi]}$ have the same cardinality (i.e. $m_i = m_{gi}$).

If a polynomial has a (Ω, G)-decomposition then it is G-invariant since

$$\begin{aligned} gp = p(\mathbf{x}^{[g1]}, \dots, \mathbf{x}^{[gn]}) &= \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} p^{[1]}_{\alpha_{|_1}}(\mathbf{x}^{[g1]}) \cdots p^{[n]}_{\alpha_{|_n}}(\mathbf{x}^{[gn]}) \\ &= \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} p^{[g1]}_{g(\alpha_{|_1})}(\mathbf{x}^{[g1]}) \cdots p^{[gn]}_{g(\alpha_{|_n})}(\mathbf{x}^{[gn]}) \\ &= \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} p^{[g1]}_{({}^g\alpha)_{|_{g1}}}(\mathbf{x}^{[g1]}) \cdots p^{[gn]}_{({}^g\alpha)_{|_{gn}}}(\mathbf{x}^{[gn]}) \\ &= \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} p^{[1]}_{\alpha_{|_1}}(\mathbf{x}^{[1]}) \cdots p^{[n]}_{\alpha_{|_n}}(\mathbf{x}^{[n]}) = p, \end{aligned}$$

where we have used Definition 5.1.1 (b) in the third equality, and the fact that $\alpha \mapsto {}^g\alpha$ is a bijection on $\mathcal{I}^{\widetilde{\mathcal{F}}}$ and that $i \mapsto gi$ is a bijection on $[n]$ in the fifth equality.

In the converse direction, the following holds: If a polynomial is G-invariant, then it has an (Ω, G)-decomposition if G acts freely on Ω (see Theorem 5.1.2).

The existence of an (Ω, G)-decomposition might imply an even stronger symmetry than G-invariance. As we will see in Example 5.1.4, the existence of a (Σ_n, G)-decomposition for any transitive group action of some group G already implies S_n-invariance. This is closely related to the action not being free.

Let us now revisit our running examples — the simple and double edge — in the light of invariant polynomial decompositions.

Example 5.1.1 (The simple edge with invariance)

On the simple edge Λ_2, the elements in $\mathcal{I}^{\widetilde{\mathcal{F}}}$ are just single values, and thus the corresponding decomposition is given by

$$p = \sum_{\alpha=1}^{r} p^{[1]}_{\alpha}(\mathbf{x}^{[1]}) \cdot p^{[2]}_{\alpha}(\mathbf{x}^{[2]}).$$

The C_2-invariant decomposition is given by

$$p = \sum_{\alpha=1}^{r} p_{\alpha}(\mathbf{x}^{[1]}) \cdot p_{\alpha}(\mathbf{x}^{[2]}).$$

Example 5.1.2 (The double edge with invariance)

For the double edge[2] Δ we have two facets and thus the Δ-decomposition reads

$$p = \sum_{\alpha,\beta=1}^{r} p^{[1]}_{\alpha,\beta}(\mathbf{x}^{[1]}) \cdot p^{[2]}_{\beta,\alpha}(\mathbf{x}^{[2]}).$$

Note that the order of the indiced α, β does not matter here, since there is no connection between the local polynomials at site 1 and 2. But for

2: See Example 2.2.4 for its definition.

We refer to Example 2.3.4 for the analogous example concerning tensors.

the non-trival C_2 action, Definition 5.1.1 (b) specifies that

$$p^{[1]}_{\alpha,\beta} = p^{[2]}_{\alpha,\beta},$$

so an (Δ, C_2)-decomposition is of the form

$$p = \sum_{\alpha,\beta=1}^{r} p_{\alpha,\beta}(\mathbf{x}^{[1]}) \cdot p_{\beta,\alpha}(\mathbf{x}^{[2]}). \tag{5.5}$$

Let us now consider an invariant polynomial on the double edge which we will revisit in Example 5.1.11 in the light of sum-of-squares invariant decompositions.

Example 5.1.3 (Invariant polynomial on the double edge)
Consider the polynomial

$$\begin{aligned} p &= x^2 + y^2 + 4(1+xy)^2 \\ &= 4 + 8xy + x^2 + y^2 + 4x^2y^2 \in \mathbb{R}[x] \otimes \mathbb{R}[y] \end{aligned}$$

which is invariant with respect to the permutation of x and y. A (Δ, C_2)-decomposition of p has the form

$$p = \sum_{\alpha,\beta=1}^{2} p_{\alpha,\beta}(x) p_{\beta,\alpha}(y),$$

with

$$p_{1,1}(t) = \frac{1}{2} + 2t^2, \quad p_{1,2}(t) = p_{2,1}(t) = \sqrt{\frac{15}{8}}, \quad p_{2,2}(t) = \sqrt{8}t.$$

It is easy to see that a decomposition of rank 1 does not exist, showing that the (Δ, C_2)-rank is indeed 2.

Let us now see more standard examples of (Ω, G)-decompositions based off the weighted simplicial complexes presented in Section 2.2.1.

Example 5.1.4 (The simplex decomposition)
For $n \geqslant 2$ consider an n-simplex Σ_n, whose facets are given by $\widetilde{\mathcal{F}} = \{[n]\}$. Since $\widetilde{\mathcal{F}}$ only contains one facet encompassing all vertices, the corresponding Σ_n-decomposition is given by

$$p = \sum_{\alpha=1}^{r} p^{[1]}_{\alpha}(\mathbf{x}^{[1]}) \cdot p^{[2]}_{\alpha}(\mathbf{x}^{[2]}) \cdots p^{[n]}_{\alpha}(\mathbf{x}^{[n]}).$$

The minimal integer r among all such decompositions is the $\mathrm{rank}_{\Sigma_n}(p)$. Now assume there is a group action G on $[n]$ which is transitive, i.e. it generates only one orbit, namely $Gi = [n]$ for all $i \in [n]$. Then Definition 5.1.1 (b) requires $p^{[i]}_{\alpha} = p^{[j]}_{\alpha}$ for all i, j, α, and hence the corresponding

(Σ_n, G)-decomposition reads

$$p = \sum_{\alpha=1}^{r} p_\alpha(\mathbf{x}^{[1]}) \cdot p_\alpha(\mathbf{x}^{[2]}) \cdots p_\alpha(\mathbf{x}^{[n]}).$$

This decomposition is manifestly fully symmetric with respect to every permutation of $\mathbf{x}^{[i]}$ with $\mathbf{x}^{[j]}$. The minimal such r is the $\mathrm{rank}_{(\Sigma_n,G)}(p)$.

Example 5.1.5 (The cyclic decomposition)

For $n \geqslant 3$ consider the circle Θ_n.[3] The Θ_n-decomposition of p reads

3: See Example 2.2.3 for its definition

$$p = \sum_{\alpha_1,\ldots,\alpha_n=1}^{r} p^{[1]}_{\alpha_1,\alpha_2}(\mathbf{x}^{[1]}) \cdot p^{[2]}_{\alpha_2,\alpha_3}(\mathbf{x}^{[2]}) \cdots p^{[n]}_{\alpha_n,\alpha_1}(\mathbf{x}^{[n]}).$$

The minimal such r is the $\mathrm{rank}_{\Theta_n}(p)$.
Since the cyclic group C_n acts freely on Θ_n, we obtain the (Θ_n, C_n)-decomposition

$$p = \sum_{\alpha_1,\ldots,\alpha_n=1}^{r} p_{\alpha_1,\alpha_2}(\mathbf{x}^{[1]}) \cdot p_{\alpha_2,\alpha_3}(\mathbf{x}^{[2]}) \cdots p_{\alpha_n,\alpha_1}(\mathbf{x}^{[n]}).$$

This decomposition is manifestly ti, that is, invariant with respect to permutations $\mathbf{x}^{[i]} \mapsto \mathbf{x}^{[a+i]}$ for $a \in \mathbb{N}$ where addition is modulo $n+1$. Note that polynomials with such a decomposition are generally not S_n-invariant. The minimal such r is called the $\mathrm{rank}_{(\Theta_n,C_n)}(p)$.

Decompositions without invariance

The first result on the existence of polynomial decompositions does not involve any invariance. It is an adaption of the result for tensor decompositions [37, Theorem 11], which we will prove here for completeness.

Theorem 5.1.1 (Existence of Ω-decompositions)
For every connected WSC Ω and every $p \in \mathcal{P}$ there exists an Ω-decomposition of p. Moreover, the Ω-decomposition can be obtained by using nonnegative multiples of the elementary decomposition

$$p = \sum_{j\in\mathcal{I}} p_j^{[1]}(\mathbf{x}^{[1]}) \cdot p_j^{[2]}(\mathbf{x}^{[2]}) \cdots p_j^{[n]}(\mathbf{x}^{[n]}) \tag{5.6}$$

where $\mathcal{I}$ is a finite index set and $p_j^{[i]} \in \mathbb{R}[\mathbf{x}^{[i]}]$ for all $j \in \mathcal{I}$.

Proof. We start with an elementary polynomial decomposition of Equation (5.6). This will show that $\mathrm{rank}_{\Sigma_n}(p) < \infty$. For $i \in [n]$ and $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$ we define

$$p_\beta^{[i]} := \begin{cases} p_j^{[i]} & : \beta \text{ takes the constant value } j \in \mathcal{I} \\ 0 & : \text{else.} \end{cases} \tag{5.7}$$

Since Ω is connected, for $\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}$ the restricted functions $\alpha_{|_i}$ are all constant if and only if α is constant. It follows that

$$\sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} p^{[1]}_{\alpha_{|_1}}(\mathbf{x}^{[1]}) \cdots p^{[n]}_{\alpha_{|_n}}(\mathbf{x}^{[n]}) = \sum_{j \in \mathcal{I}} p^{[1]}_j(\mathbf{x}^{[1]}) \cdots p^{[n]}_j(\mathbf{x}^{[n]})$$
$$= p(\mathbf{x}^{[1]}, \ldots, \mathbf{x}^{[n]})$$

is an Ω-decomposition of p. □

Note that the Ω-decomposition obtained by reusing the polynomials of Equation (5.6) may not be optimal, i.e. it may need more terms than its rank.

Invariant decompositions with free group actions

We now show that if G acts freely on Ω, then every G-invariant polynomial admits an (Ω, G)-decomposition.[4] The proof is similar to that of [37, Theorem 13] for tensors. We will illustrate the idea of the proof on the double edge Δ in Example 5.1.6.

4: Recall that free was defined in Definition 2.2.5.

Theorem 5.1.2 ((Ω, G)-decompositions with free group actions)
Let Ω be a connected weighted simplicial complex, G a free group action on Ω, and $p \in \mathcal{P}$ a G-invariant polynomial. Then:

- The polynomial p admits an (Ω, G)-decomposition.
- Given a Σ_n-decomposition, an (Ω, G)-decomposition of p can be obtained by using only nonnegative multiples of the local polynomials in the Σ_n-decomposition.

As in Theorem 5.1.1, the (Ω, G)-decomposition obtained by "reusing" the polynomials of Equation (5.6) will generally not be optimal.

The idea of the proof is simple. Starting from the decomposition in Equation (5.6), we essentially build

$$\frac{1}{|G|} \sum_{g \in G} gp = p$$

where gp is defined in Equation (5.4), and let g act on each of the local terms in the decomposition. The latter can then be transformed into an (Ω, G)-decomposition of p.

For the proof of this theorem we need a preparatory lemma:

Lemma 5.1.3
A group action G on the WSC Ω is free if and only if there exists a G-linear map

$$\mathbf{z} \colon \widetilde{\mathcal{F}} \to G$$

where G acts on itself via left-multiplication (which is obviously free).

Proof. To construct $\mathbf{z}$ for a free action, choose for each orbit an element F and map gF to g. The reverse implication is immediate. □

Proof of Theorem 5.1.2. Since G acts freely, by Lemma 5.1.3, there exists a G-linear map $\mathbf{z} : \widetilde{\mathcal{F}} \to G$, where G acts on itself by left-multiplication. In the following, we fix one such mapping. For the polynomial p we first obtain by Theorem 5.1.1 an Ω-decomposition and denote the local elements by

$$Q^{[i]} := \left(q_\beta^{[i]}(\mathbf{x}^{[i]})\right)_{\beta\in\mathcal{I}^{\widetilde{\mathcal{F}}_i}}$$

where $q_\beta^{[i]}(\mathbf{x}^{[i]}) \in \mathbb{R}[\mathbf{x}^{[i]}]$ for every $i \in [n]$. We define a new index set

$$\hat{\mathcal{I}} := \mathcal{I} \times G$$

together with the projection maps $\pi_1 : \hat{\mathcal{I}} \to \mathcal{I}$ and $\pi_2 : \hat{\mathcal{I}} \to G$. For each $i \in [n]$ and $\beta \in \hat{\mathcal{I}}^{\widetilde{\mathcal{F}}_i}$ we define the following local polynomials:

$$p_\beta^{[i]} := \begin{cases} q^{[gi]}_{g(\pi_1\circ\beta)}(\mathbf{x}^{[i]}) & : \pi_2\circ\beta = ({}^{g^{-1}}\mathbf{z})_{|i} \\ 0 & : \text{else.} \end{cases}$$

Note that $p_\beta^{[i]}(\mathbf{x}^{[i]})$ is well-defined since g is uniquely determined by the relation $\pi_2\circ\beta = ({}^{g^{-1}}\mathbf{z})_{|i}$ if such a g exists. This is due to the fact that if $({}^{g_1^{-1}}\mathbf{z})_{|i} = ({}^{g_2^{-1}}\mathbf{z})_{|i}$ we have $g_1\cdot\mathbf{z}(F) = g_2\cdot\mathbf{z}(F)$ for any $F\in\widetilde{\mathcal{F}}_i$ by G-linearity of $\mathbf{z}$. But this implies that $g_1 = g_2$. In addition, the defined local polynomials satisfy Definition 5.1.1 (b) since for $g, h \in G$ we obtain

$$p^{[hi]}_{{}^h\beta}(\mathbf{x}^{[i]}) = q^{[ghi]}_{g(\pi_1\circ{}^h\beta)}(\mathbf{x}^{[i]}) = q^{[ghi]}_{gh(\pi_1\circ\beta)}(\mathbf{x}^{[i]}) = p_\beta^{[i]}(\mathbf{x}^{[i]})$$

using the fact that

$$\pi_2\circ{}^h\beta = \left({}^{g^{-1}}\mathbf{z}\right)_{|hi} \iff \pi_2\circ\beta = \left({}^{(gh)^{-1}}\mathbf{z}\right)_{|i}.$$

It only remains to show that the local polynomials form an (Ω, G)-decomposition of p. To this end we compute

$$\begin{aligned}&\sum_{\hat{\alpha}\in\hat{\mathcal{I}}^{\widetilde{\mathcal{F}}}} p^{[1]}_{\hat{\alpha}_{|1}}(\mathbf{x}^{[1]})\cdots p^{[n]}_{\hat{\alpha}_{|n}}(\mathbf{x}^{[n]}) \\ &= \sum_{\substack{z\in G^{\widetilde{\mathcal{F}}} \\ \forall i\exists g_i : z_{|i} = \left({}^{g_i^{-1}}\mathbf{z}\right)_{|i}}} \sum_{\alpha\in\mathcal{I}^{\widetilde{\mathcal{F}}}} q^{[g_1 1]}_{g_1(\alpha_{|1})}(\mathbf{x}^{[1]})\cdots q^{[g_n n]}_{g_n(\alpha_{|n})}(\mathbf{x}^{[n]}).\end{aligned}$$

Using that Ω is connected and $\mathbf{z}$ is G-linear, for each z fulfilling the conditions from the outer sum on the right, we obtain $g_i = g_j =: g$ for all $i, j \in [n]$. So the corresponding inner sum becomes

$$\begin{aligned}\sum_{\alpha\in\mathcal{I}^{\widetilde{\mathcal{F}}}} q^{[g1]}_{g(\alpha_{|1})}(\mathbf{x}^{[1]})\cdots q^{[gn]}_{g(\alpha_{|n})}(\mathbf{x}^{[n]}) &= p(\mathbf{x}^{[g^{-1}1]},\ldots,\mathbf{x}^{[g^{-1}n]}) \\ &= p(\mathbf{x}^{[1]},\ldots,\mathbf{x}^{[n]}),\end{aligned}$$

using G-invariance of p. Hence the total sum equals a positive multiple of

p, where the factor is the number of all z which fulfill the above conditions. In fact, this number is just $|G|$, since the ${}^{g^{-1}}\mathbf{z}$ for $g \in G$ are precisely the different choices for z. So dividing by $|G|$ and absorbing its positive n^{th} root into the local polynomials yields an (Ω, G)-decomposition of p. The last statement is immediate by construction. □

The following are some immediate relations between the various notions of ranks, based on the proofs of Theorem 5.1.2 and Theorem 5.1.2.

Corollary 5.1.4 (Relations among ranks)
Let Ω be connected and G a free group action on Ω, and Σ_n the simplex.[5] Then for every G-invariant $p \in \mathcal{P}$ we have

$$\operatorname{rank}_{(\Omega,G)}(p) \leqslant |G| \cdot \operatorname{rank}_{\Omega}(p) \leqslant |G| \cdot \operatorname{rank}_{\Sigma_n}(p).$$

5: Defined in Example 2.2.1.

In words, the first inequality says that one can impose invariance by increasing the rank by a factor of at most $|G|$. The second inequality says that the standard tensor rank is always the most expensive rank, i.e. having one joint index is the most costly decomposition.

Proof. The first inequality is immediate from the construction in the proof of Theorem 5.1.2, and the second inequality follows from the construction in the proof of Theorem 5.1.1. □

Let us now illustrate the proof of Theorem 5.1.2 for the double edge.

Example 5.1.6 (Invariant decomposition on the double edge)
The cyclic group C_2 provides a free group action on the double edge Δ, so every C_2-invariant polynomial admits a (Δ, C_2)-decomposition, given by Equation (5.5). Let us now construct this decomposition.
For the group action of $C_2 = \{e, c\}$ on $\widetilde{\mathcal{F}} = \{\mathfrak{a}, \mathfrak{b}\}$ (with $c\mathfrak{a} = \mathfrak{b}$) there exists a G-linear map $\mathbf{z} : \widetilde{\mathcal{F}} \to G$, which can be chosen as[6]

$$\mathbf{z} \colon \mathfrak{a} \mapsto e, \; \mathfrak{b} \mapsto c.$$

6: There is exactly one other choice, namely exchanging the two outcomes of $\mathbf{z}$.

We start with a Δ-decomposition of p, namely

$$p = \sum_{\alpha,\beta=1}^{r} q^{[1]}_{\alpha,\beta}(\mathbf{x}^{[1]}) \cdot q^{[2]}_{\beta,\alpha}(\mathbf{x}^{[2]}),$$

where we associate index α with $\mathfrak{a}$ and β with $\mathfrak{b}$. To construct a (Δ, C_2)-decomposition, we extend the indices α, β to tuples $(\alpha, g_1), (\beta, g_2)$ where $g_1, g_2 \in C_2$. We define the local polynomials as

$$p^{[1]}_{(\alpha,g_1),(\beta,g_2)}(\mathbf{x}^{[1]}) := \begin{cases} q^{[1]}_{\alpha,\beta}(\mathbf{x}^{[1]}) & \text{if } (g_1, g_2) = (e, c) \\ q^{[2]}_{\beta,\alpha}(\mathbf{x}^{[1]}) & \text{if } (g_1, g_2) = (c, e) \\ 0 & \text{else} \end{cases}$$

and

$$p^{[2]}_{(\alpha,g_1),(\beta,g_2)}(\mathbf{x}^{[2]}) := \begin{cases} q^{[2]}_{\alpha,\beta}(\mathbf{x}^{[2]}) & \text{if } (g_1,g_2) = (e,c) \\ q^{[1]}_{\beta,\alpha}(\mathbf{x}^{[2]}) & \text{if } (g_1,g_2) = (c,e) \\ 0 & \text{else.} \end{cases}$$

For $\alpha, \beta \in \{1,\ldots,r\}$ and $g_1, g_2 \in C_2$, the symmetry condition gives rise to the definition

$$p^{[c1]}_{c((\alpha,g_1),(\beta,g_2))} = p^{[2]}_{(\beta,g_2),(\alpha,g_1)} = p^{[1]}_{(\alpha,g_1),(\beta,g_2)} =: p_{(\alpha,g_1),(\beta,g_2)}.$$

In addition, it is easy to verify that

$$\begin{aligned} &\sum_{g_1,g_2\in C_2}\sum_{\alpha,\beta=1}^{r} p_{(\alpha,g_1),(\beta,g_2)}(\mathbf{x}^{[2]}) \cdot p_{(\beta,g_2),(\alpha,g_1)}(\mathbf{x}^{[2]}) \\ =&p(\mathbf{x}^{[1]},\mathbf{x}^{[2]}) + p(\mathbf{x}^{[2]},\mathbf{x}^{[1]}) = 2p(\mathbf{x}^{[1]},\mathbf{x}^{[2]}) \end{aligned}$$

which shows that the local polynomials $\frac{1}{\sqrt{2}} \cdot p_{(\alpha,g_1),(\beta,g_2)}$ form a (Δ, C_2)-decomposition of p. This also implies $\operatorname{rank}_{(\Delta,C_2)}(p) \leqslant 2 \cdot r$.

Invariant decompositions with blending group actions

Since the full symmetry group S_n is not free on the simplex Σ_n, Theorem 5.1.2 does not say anything about the existence of (Σ_n, S_n)-decompositions. In fact, for real polynomials, such decompositions may not exist (see Example 5.1.7). Nonetheless, we can prove another, weaker existence result for polynomial decompositions with a blending[7] group action G (Theorem 5.1.7). In preparation for this result we need the following two lemmas. The first lemma introduces a "negative part" in the symmetric decomposition, which can be omitted if n is odd:

7: See Definition 2.2.5 for free group actions.

Lemma 5.1.5 (Symmetric decompositions for tensors [32])
Let $|T\rangle \in \mathbb{R}^d \otimes \cdots \otimes \mathbb{R}^d \cong \mathbb{R}^{nd}$ be S_n-invariant.[8] Then there exist $r_1, r_2 \in \mathbb{N}$ and $|v_1\rangle, \ldots, |v_{r_1}\rangle, |v_{r_1+1}\rangle, \ldots, |v_{r_1+r_2}\rangle \in \mathbb{R}^d$ such that

$$|T\rangle = \sum_{\ell=1}^{r_1} |v_\ell\rangle^{\otimes n} - \sum_{\ell=r_1+1}^{r_1+r_2} |v_\ell\rangle^{\otimes n} \tag{5.8}$$

If n is odd, there exists a decomposition

$$|T\rangle = \sum_{\ell=1}^{r_1} |v_\ell\rangle^{\otimes n}.$$

8: i.e. for every $i_1, \ldots, i_n \in \{1, \ldots, d\}$ and permutation $\sigma \in S_n$ we have

$$\langle i_1, \ldots, i_n \,|\, T\rangle = \langle \sigma(i_1), \ldots, \sigma(i_n) \,|\, T\rangle .$$

The last statement is not given in [32], but it is obvious, since the minus sign can be absorbed into the odd number of terms n.[9]

9: This is because $(-1)^n = -1$ for odd n.

The minus sign in Equation (5.8) is necessary. Consider for example the the simple case of real matrices, namely when the corresponding tensor $|T\rangle$ lives in the space $\mathbb{R}^d \otimes \mathbb{R}^d \cong \mathrm{Mat}_d(\mathbb{R})$. Without a minus sign, Equation (5.8) in the matrix picture would read[10]

$$T = \sum_{\ell=1}^{r_1} |v_\ell\rangle \langle v_\ell| \succcurlyeq 0$$

implying that every symmetric matrix is psd which is false. (see also Example 5.1.7).

10: by using the correspondence between $|v\rangle\,|w\rangle$ and $|v\rangle\,\langle w|$.

In the next lemma we show subadditivity and submultiplicativity of the (Ω, G)-rank. For a proof, we refer to [37, Proposition 16].

Lemma 5.1.6 (Subadditivity and submultiplicativity [37])
Let $p_1, p_2 \in \mathcal{P}$.

(i) $\operatorname{rank}_{(\Omega,G)}(p_1 + p_2) \leqslant \operatorname{rank}_{(\Omega,G)}(p_1) + \operatorname{rank}_{(\Omega,G)}(p_2)$
(ii) $\operatorname{rank}_{(\Omega,G)}(p_1 \cdot p_2) \leqslant \operatorname{rank}_{(\Omega,G)}(p_1) \cdot \operatorname{rank}_{(\Omega,G)}(p_2)$

We are now ready to prove the existence of invariant decompositions with blending group actions.

Theorem 5.1.7 (Invariant decompositions, blending actions)
Let Ω be a connected WSC and G a blending group action on Ω. For any G-invariant $p \in \mathcal{P}$ there exist two polynomials $q_1, q_2 \in \mathcal{P}$ with

$$p = q_1 - q_2,$$

where q_1, q_2 attain an (Ω, G)-decomposition. If n is odd we can set $q_2 = 0$.

Proof. We start with a non-invariant decomposition of p, as given in Equation (5.6), where $\mathcal{I}$ is a finite index set. Now we choose real numbers $d_\ell^{[i]} \in \mathbb{R}$ for $i \in [n]$ and $\ell \in \{1, \ldots, r_1 + r_2\}$, such that the following holds:

$$\sum_{\ell=1}^{r_1} d_\ell^{[i_1]} \cdots d_\ell^{[i_n]} - \sum_{\ell=r_1+1}^{r_2} d_\ell^{[i_1]} \cdots d_\ell^{[i_n]} = \begin{cases} 1 & : \{i_1, \ldots, i_n\} = [n] \\ 0 & : \text{else} \end{cases}$$

This is possible because the tensor on the right hand side is real and symmetric, hence the existence follows by Lemma 5.1.5. For $i \in [n]$, $\ell \in \{1, \ldots, r_1 + r_2\}$ and $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$ we define

$$p_{\ell,\beta}^{[i]}(\mathbf{x}^{[i]}) := \begin{cases} \sum_{g \in G} d_\ell^{[gi]} p_j^{[gi]}(\mathbf{x}^{[i]}) & : \beta \text{ is the constant } j \in \mathcal{I} \\ 0 & : \text{else} \end{cases}.$$

For fixed ℓ, the polynomials $p_{\ell,\beta}^{[i]}$ satisfy Definition 5.1.1 (b) and hence give rise to (Ω, G)-decompositions of polynomials

$$p_1, \ldots, p_{r_1}, p_{r_1+1}, \ldots, p_{r_1+r_2}.$$

We now define q_1 as

$$q_1 := \sum_{\ell=1}^{r_1} p_\ell = \sum_{\ell=1}^{r_1} \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} p^{[1]}_{\ell,\alpha_{|_1}}(\mathbf{x}^{[1]}) \cdots p^{[n]}_{\ell,\alpha_{|_n}}(\mathbf{x}^{[n]})$$
$$= \sum_{g_1,\dots,g_n \in G} \sum_{\ell=1}^{r_1} d_\ell^{[g_1 1]} \cdots d_\ell^{[g_n n]} \sum_{j \in \mathcal{I}} p_j^{[g_1 1]}(\mathbf{x}^{[1]}) \cdots p_j^{[g_n n]}(\mathbf{x}^{[n]})$$

where we have used that Ω is connected in the third equality, and thus $\alpha_{|_i}$ constant for all i if and only if α is constant. Note that q_1 has an (Ω, G)-decomposition by Lemma 5.1.6, since all p_ℓ do. We define q_2 similarly as

$$q_2 := \sum_{\ell=r_1+1}^{r_2} p_\ell.$$

Because of the definition of $d_\ell^{[i]}$, and the fact that the action of G is blending, the difference $q_1 - q_2$ simplifies to

$$q_1 - q_2 = \sum_{\substack{g_1,\dots,g_n \in G \\ \{g_1 1,\dots,g_n n\}=[n]}} \sum_{j \in \mathcal{I}} p_j^{[g_1 1]}(\mathbf{x}^{[1]}) \cdots p_j^{[g_n n]}(\mathbf{x}^{[n]})$$
$$\sim \sum_{g \in G} \sum_{j \in \mathcal{I}} p_j^{[g1]}(\mathbf{x}^{[1]}) \cdots p_j^{[gn]}(\mathbf{x}^{[n]}) = |G| \cdot p$$

where $\sim$ stands for positive multiple of. Note that we have used that p is G-invariant in the last equality. Dividing by $|G|$ and the positive scaling factor proves the statement, since the scaling can be absorbed in the local polynomials. The last statement of the theorem follows from the statement in Lemma 5.1.5 for even n. □

Example 5.1.7 (The minus sign in the single and double edge)

The minus sign in the decomposition of Theorem 5.1.7 is necessary (as long as we do not switch to complex coefficients). For example, the polynomial $p = x^2 + y^2$ is C_2-invariant, and since C_2 is blending on the single edge Λ_2, there exists an (Λ_2, C_2)-decomposition for p with this additional minus sign (by Theorem 5.1.7):

$$p = x^2 + y^2 = p_1(x) \cdot p_1(y) - p_2(x) \cdot p_2(y)$$

where

$$p_1(t) = \frac{1}{\sqrt{2}}(1 + t^2) \text{ and } p_2(t) = \frac{1}{\sqrt{2}}(1 - t^2).$$

But for degree reasons there cannot exist an actual (Λ_2, C_2)-decomposition for p, i.e. an invariant decomposition without the additional minus sign.

On the other hand, the refinement of Λ_2 to the double edge Δ allows for a free group action of C_2. Hence there exists a (Δ, C_2)-decomposition of p (by Theorem 5.1.2), given for example by

$$x^2 + y^2 = \sum_{\alpha,\beta=1}^{2} p_{\alpha,\beta}(x) \cdot p_{\beta,\alpha}(y)$$

where $p_{1,1}(t) = 0, p_{1,2}(t) = t^2, p_{2,1}(t) = 1$ and $p_{2,2}(t) = 0$. This shows that $\mathrm{rank}_{(\Delta,C_2)}(p) = 2$.

5.1.3 The invariant separable decomposition

In this section we assume that every local space of polynomials is equipped with a convex cone $\mathcal{C}^{[i]} \subseteq \mathbb{R}[\mathbf{x}^{[i]}]$, i.e. a set which fulfills $\alpha p + \beta q \in \mathcal{C}$ for all $p, q \in \mathcal{C}$ and $\alpha, \beta \geqslant 0$. Important examples of such cones are the cone of sum-of-squares (sos) polynomials

$$\mathcal{C}_{\mathrm{sos}} := \left\{ p \in \mathbb{R}[\mathbf{x}] : p = \sum_{k=1}^{N} q_k^2 \text{ for some } q_k \in \mathbb{R}[\mathbf{x}], N \in \mathbb{N} \right\},$$

the cone of nonnegative polynomials

$$\mathcal{C}_{\mathrm{nn}} := \{ p \in \mathbb{R}[\mathbf{x}] : p(a) \geqslant 0 \text{ for all } a \in \mathbb{R}^m \},$$

and the cone of polynomials with nonnegative coefficients[11]

$$\mathcal{C}_{\text{nn-coeff}} := \left\{ p \in \mathbb{R}[\mathbf{x}] : p = \sum_{\alpha \in \{1,\ldots,d\}^n} c_\alpha \mathbf{x}^\alpha \text{ with all } c_\alpha \geqslant 0 \right\}.$$

11: In this definition, we use the notation where $\alpha := (\alpha_1, \ldots, \alpha_n)$ is an n-tuple, with $\mathbf{x}^\alpha := x_1^{\alpha_1} \cdots x_n^{\alpha_n}$.

For a given set of local cones $\mathcal{C}^{[1]}, \ldots, \mathcal{C}^{[n]}$ we define the global separable cone

$$\begin{aligned} \mathcal{C}_{\mathrm{sep}} &:= \mathcal{C}^{[1]} \otimes \mathcal{C}^{[2]} \otimes \cdots \otimes \mathcal{C}^{[n]} \\ &:= \left\{ \sum_{j=1}^{r} p_j^{[1]} \cdots p_j^{[n]} : r \in \mathbb{N},\ p_j^{[i]} \in \mathcal{C}^{[i]} \right\} \subseteq \mathcal{P}. \end{aligned}$$

This is the smallest global convex cone generated by the elementary tensors formed from the local cones. For a given group action of G on Ω, we further assume that $\mathcal{C}^{[i]} = \mathcal{C}^{[gi]}$ for all $g \in G$.[12]

12: We suppress again the canonical isomorphism between the local polynomial spaces in the notation.

We now define and study the invariant separable decomposition of polynomials, i.e. decompositions which are inherently G-invariant and where the containment in the separable cone is explicit — i.e. a positive combination of elementary polynomials where each factor is in the local cone.

Definition 5.1.2 (Invariant separable decomposition)
Let $p \in \mathcal{C}_{\mathrm{sep}}$. A *separable* (Ω, G)*-decomposition* of p is an (Ω, G)-decomposition

$$\mathcal{P}^{[i]} := \left(p_\beta^{[i]} \right)_{\beta \in \mathcal{I}^{\tilde{\mathcal{F}}_i}}$$

with the restriction that

$$p_\beta^{[i]} \in \mathcal{C}^{[i]}$$

for all $i \in [n]$ and $\beta \in \mathcal{I}^{\tilde{\mathcal{F}}_i}$.

The minimal cardinality of $\mathcal{I}$ among all separable (Ω, G)-decomposition of p is called the *separable* (Ω, G)*-rank of* p, denoted $\text{sep-rank}_{(\Omega,G)}(p)$. If p does not admit an (Ω, G)-decomposition, we set

$$\text{sep-rank}_{(\Omega,G)}(p) = \infty.$$

If G is the trivial group action, we call the separable (Ω, G)-decomposition just *separable* Ω*-decomposition*, and its minimal number terms the *separable rank*, denoted sep-rank_{Ω}.

We now show the existence of invariant separable decompositions with free group actions. This follows from Theorem 5.1.2, as it can be constructed via positive multiples of the initial decomposition.

Theorem 5.1.8 (Invariant separable decompositions)
Let Ω be a connected WSC with a free group action G. Every G-invariant polynomial $p \in \mathcal{C}_{\text{sep}}$ admits a separable (Ω, G)-decomposition.

Proof. Let p be decomposed as in Equation (5.6) with $p_j^{[i]} \in \mathcal{C}^{[i]}$, which is a separable decomposition of p. Applying the construction of the proof of Theorem 5.1.2 we obtain a separable (Ω, G)-decomposition, since all local polynomials $p_\beta^{[i]}$ are positive multiples of $p_j^{[gi]}$ for $g \in G$. Since the local cones coincide on the orbits of G, this guarantees that $p_\beta^{[i]} \in \mathcal{C}^{[i]}$. □

Example 5.1.8 (Separable decomposition on the double edge)
The (Δ, C_2)-decomposition of $p = x^2 + y^2$ given in Example 5.1.7 is in fact an invariant separable decomposition with respect to the local sos cones, proving that $\text{sep-rank}_{(\Delta,C_2)}(p) = \text{rank}_{(\Delta,C_2)}(p) = 2$.

We can now easily promote the results of Corollary 5.1.4 to the (invariant) separable ranks. The proof is analogous.

Corollary 5.1.9 (Relation between separable ranks)
Let Ω be connected and G a free group action on Ω. Then for every G-invariant $p \in \mathcal{P}$ we have

$$\text{sep-rank}_{(\Omega,G)}(p) \leqslant |G| \cdot \text{sep-rank}_{\Omega}(p) \leqslant |G| \cdot \text{sep-rank}_{\Sigma_n}(p).$$

Note that an analogue of Theorem 5.1.8 for blending group actions is not true. One reason is that, if the action is blending, we cannot construct a decomposition using the local polynomials from the initial tensor decomposition. This is visible in the simplest case, namely for (Λ_2, C_2)-decompositions, illustrated in Example 5.1.7. Another reason is that Theorem 5.1.7 (with blending group actions) uses a *difference* of two (Ω, G)-decompositions, and a difference of separable elements is in general not separable.

Finally we show that the global cone of sos polynomials $\mathcal{C}_{\text{sos}}$ is strictly larger than the cone of separable polynomials over local sos polynomials

$\mathcal{C}_{\text{sep}} = \mathcal{C}_{\text{sos}}^{[1]} \otimes \cdots \otimes \mathcal{C}_{\text{sos}}^{[n]}$, i.e.

$$\mathcal{C}_{\text{sos}}^{[1]} \otimes \cdots \otimes \mathcal{C}_{\text{sos}}^{[n]} \subsetneq \mathcal{C}_{\text{sos}}.$$

In other words, there exist polynomials which admit a sos decomposition over all variables, but cannot be written as tensor decomposition where every term is a sos polynomial. This is even true for polynomials in two variables x and y, as the following example shows. The example relies on the *Gram map*, which will be the cornerstone of invariant sos decompositions (Section 5.1.4). Moreover, it relies on the standard result that the set of separable matrices is strictly smaller than the set of psd matrices.

Example 5.1.9 (sos polynomials which are not separable)

We consider the following Gram map $\mathcal{G}$ between real-valued matrices $M \in \text{Mat}_2(\mathbb{R}) \otimes \text{Mat}_2(\mathbb{R})$ and polynomials $p \in \mathbb{R}[x, y]$:

$$\mathcal{G} : M \mapsto p := \mathfrak{m}_1(x)^t \otimes \mathfrak{m}_1(y)^t \cdot M \cdot \mathfrak{m}_1(x) \otimes \mathfrak{m}_1(y)$$

where $\mathfrak{m}_1(x) := (1, x)^t$ is the monomial basis in x of degree at most 1. It is well-known (and easy to see) that for $\deg_{\text{loc}}(p) \leqslant 2$ we have $p \in \mathcal{C}_{\text{sos}}$ if and only if there exists a psd $M \in \text{Mat}_2(\mathbb{R}) \otimes \text{Mat}_2(\mathbb{R})$ with $\mathcal{G}(M) = p$. Further, $p \in \mathcal{C}_{\text{sep}}$ if and only if there exists an $M \in \text{Mat}_2(\mathbb{R}) \otimes \text{Mat}_2(\mathbb{R})$ that is separable[13] and $\mathcal{G}(M) = p$.

For example, consider the matrix

$$M = \sum_{ij=0}^{1} |i\rangle \langle j| \otimes |i\rangle \langle j| = |\Phi^+\rangle \langle \Phi^+| = \begin{pmatrix} 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 \end{pmatrix}$$

where $|\Phi^+\rangle = |0,0\rangle + |1,1\rangle \in \mathbb{R}^2 \otimes \mathbb{R}^2$ is known as an (unnormalized) Bell state. Note that M is psd but not separable, which can easily be seen with the positive partial transposition criterion [96, 66].[14] Furthermore, M is the *only* psd matrix representing the polynomial

$$p = 1 + 2xy + x^2y^2 = (1 + xy)^2 = \mathcal{G}(M),$$

since the matrix

$$M_\alpha = \begin{pmatrix} 1 & 0 & 0 & 1-\alpha \\ 0 & 0 & \alpha & 0 \\ 0 & \alpha & 0 & 0 \\ 1-\alpha & 0 & 0 & 1 \end{pmatrix}$$

is not psd for any $\alpha \in \mathbb{R} \setminus \{0\}$, and

$$\mathcal{G}^{-1}(\{p\}) = \{M_\alpha : \alpha \in \mathbb{R}\}.$$

This implies that $p = (1 + xy)^2$ is sos but not separable with respect to the local sos cones.

More generally, in order to show that a polynomial is sos but not separable, one needs to show that every psd matrix M with $\mathcal{G}(M) = p$ is not separable. This is generally a hard problem.

13: i.e. there exists a decomposition

$$M = \sum_{j=1}^{r} M_j^{[1]} \otimes M_j^{[2]}$$

where all $M_j^{[i]}$ are psd.

14: The positive partial transpose criterion is a necessary criterion for bipartite states to be separable. If $\rho \in \text{Mat}_d(\mathbb{R}) \otimes \text{Mat}_d(\mathbb{R})$ is separable, then ρ^{t_2} is psd, where

$$\rho^{t_2} := \sum_{j=1}^{r} A_j \otimes B_j^t$$

for

$$\rho = \sum_{j=1}^{r} A_j \otimes B_j.$$

5.1.4 The invariant sum-of-squares decomposition

In this section we introduce a sum-of-squares (sos) decomposition in the (Ω, G)-framework. First, notice that not every G-invariant sos polynomial p can be decomposed into G-invariant polynomials q_k via $p = \sum_{k=1}^{N} q_k^2$, as the following example shows.

Example 5.1.10 (Absence of stringent invariant sos decomposition) Consider again $p = x^2 + y^2$, which is obviously sos and C_2-invariant, i.e. invariant with respect to permuting x and y. Yet, there does *not* exist a decomposition

$$p = \sum_{k=1}^{N} q_k^2 \qquad \text{where all } q_k \text{ are } C_2\text{-invariant.}$$

To see this, assume the contrary. Since $\deg(q_k) \leqslant \frac{1}{2}\deg(p)$, each polynomial can be written as $q_k = a_k x + a_k y + b_k$. Further, since p has no constant term, we must have $b_k = 0$. But this is impossible, since the xy coefficient of p is zero.

We term the previous definition of an invariant sos decomposition *stringent*, and now introduce a more 'relaxed' one, which allows for permutations among elements of the family $\{q_k\}$, and which is the correct notion as far the existence results are concerned, as we will later show. So let G act on $[n]$, and equip the finite index set

$$\mathcal{S} = \mathcal{S}_1 \times \ldots \times \mathcal{S}_n$$

with the induced group action

$$g\mathbf{k} := (k_{g^{-1}1}, \ldots, k_{g^{-1}n})$$

for every $\mathbf{k} = (k_1 \ldots, k_n) \in \mathcal{S}$ and $g \in G$. We say that the family of polynomials $\mathsf{q} = (q_\mathbf{k})_{\mathbf{k}\in\mathcal{S}}$ is G-*invariant* if

$$q_{g\mathbf{k}} = g q_\mathbf{k}$$

for all $g \in G$ and $\mathbf{k} \in \mathcal{S}$. This equation can be spelled out as

$$q_{g\mathbf{k}}(\mathbf{x}^{[1]}, \ldots, \mathbf{x}^{[n]}) = q_\mathbf{k}(\mathbf{x}^{[g1]}, \ldots, \mathbf{x}^{[gn]}).$$

Now, if q is G-invariant, the resulting sos polynomial

$$p = \sum_{\mathbf{k}\in\mathcal{S}} q_\mathbf{k}^2$$

is also G-invariant (since $\mathbf{k} \mapsto g\mathbf{k}$ is a bijection on $\mathcal{S}$). In Theorem 5.1.11 (i), we will prove the reverse direction, namely that every G-invariant sos polynomial p has a G-invariant family of polynomials q.

To prove this result, we leverage a correspondence between matrices and polynomials given by the *Gram map* $\mathcal{G}$ (similarly to Example 5.1.9). For simplicity, we assume for the rest of this section that every local polynomial space uses the same number of variables, i.e.

$$\mathcal{P} = \mathbb{R}[\mathbf{x}^{[1]}] \otimes \cdots \otimes \mathbb{R}[\mathbf{x}^{[n]}]$$

where $\mathbf{x}^{[i]} = (x_1^{[i]}, \ldots, x_m^{[i]})$ for each $i \in [n]$. Now consider a polynomial $p \in \mathcal{P}$ with $\deg_{\mathrm{loc}}(p) \leqslant 2d$. We can represent p via the Gram map

$$\begin{aligned} \mathcal{G} \colon \operatorname{Mat}_D(\mathbb{R})^{\otimes n} &\to \mathcal{P} \\ M &\mapsto \langle \mathfrak{m}_{n,d} | \, M \, | \mathfrak{m}_{n,d} \rangle \end{aligned}$$

where $|\mathfrak{m}_{n,d}\rangle = |\mathfrak{m}_d(\mathbf{x}^{[1]})\rangle \otimes \cdots \otimes |\mathfrak{m}_d(\mathbf{x}^{[n]})\rangle$ and we define $|\mathfrak{m}_d(\mathbf{x})\rangle$ to be the monomial basis in $\mathbf{x}$ consisting of all monomials of degree at most d. In other words, for indices $i_1, \ldots, i_n \in \{0, \ldots, d\}$ such that $i_1 + \ldots + i_n \leqslant d$, we have

$$\langle i_1, \ldots, i_n \, | \, \mathfrak{m}_d(\mathbf{x}) \rangle = x_1^{i_1} \cdots x_n^{i_n}.$$

In addition, $\operatorname{Mat}_D(\mathbb{R})$ is the space of real matrices of size $D \times D$, where $D = \binom{m+d}{d}$. Note that D is also the number of monomials in m variables of degree at most d. We say that the matrix $M = \sum_{j=1}^N M_j^{[1]} \otimes \cdots \otimes M_j^{[n]}$ is G-invariant if

$$gM := \sum_{j=1}^{N} M_j^{[g^{-1}1]} \otimes \cdots \otimes M_j^{[g^{-1}n]} = M$$

for every $g \in G$, that is, if M is invariant with respect to all permutations of the tensor factors induced by the group action of G on $[n]$. This generalizes the Gram map for multivariate polynomials without invariance [25].

Lemma 5.1.10 (Gram matrix of invariant sos polynomials)
Let $p \in \mathcal{P}$ with $\deg_{\mathrm{loc}}(p) \leqslant 2d$. The following are equivalent:

(i) p is sos and G-invariant.
(ii) There exists an $M \in \operatorname{Mat}_D(\mathbb{R})^{\otimes n}$ that is psd and G-invariant such that $\mathcal{G}(M) = p$.

Proof. (ii) $\Longrightarrow$ (i). If there exists such an M, since it is psd, it has a rank decomposition $M = \sum_k |v_k\rangle \langle v_k|$ where $|v_k\rangle \in \left(\mathbb{R}^D\right)^{\otimes n}$. This gives rise to a sos decomposition of p via $\mathcal{G}$. Furthermore, since $gM = M$ for all $g \in G$, we obtain

$$\begin{aligned} gp &= p(\mathbf{x}^{[g1]}, \ldots, \mathbf{x}^{[gn]}) \\ &= \langle \mathfrak{m}_d(\mathbf{x}^{[g1]}) | \cdots \langle \mathfrak{m}_d(\mathbf{x}^{[gn]}) | \, M \, | \mathfrak{m}_d(\mathbf{x}^{[g1]}) \rangle \cdots | \mathfrak{m}_d(\mathbf{x}^{[gn]}) \rangle \\ &= \langle \mathfrak{m}_d(\mathbf{x}^{[g1]}) | \cdots \langle \mathfrak{m}_d(\mathbf{x}^{[gn]}) | \, g^{-1}M \, | \mathfrak{m}_d(\mathbf{x}^{[g1]}) \rangle \cdots | \mathfrak{m}_d(\mathbf{x}^{[gn]}) \rangle \\ &= p \end{aligned}$$

where the second equality holds by the G-invariance of M, and the last equality by the commutativity of polynomial multiplication.

(i) $\Longrightarrow$ (ii). Assume that $p = \sum_{k=1}^N q_k^2$ is G-invariant. Define $|v_k\rangle \in \left(\mathbb{R}^D\right)^{\otimes n}$ such that $q_k = \langle v_k \, | \, \mathfrak{m}_{n,d} \rangle$ defines a psd matrix $M' = \sum_{k=1}^N v_k v_k^t$ with $\mathcal{G}(M') = p$, where M' need not be G-invariant. By the G-invariance of p, we additionally have that $\mathcal{G}(gM') = p$ for every $g \in G$. Defining M as the average

$$M = \frac{1}{|G|} \sum_{g \in G} gM'$$

we obtain a G-invariant and psd matrix M. By linearity of the Gram map, we have that $\mathcal{G}(M) = p$. □

Remark 5.1.1 (Gram matrix of invariant separable polynomials)
A similar version of Lemma 5.1.10 relates invariant separable polynomials

$$p \in \mathcal{C}_{\text{sep}} = \mathcal{C}_{\text{sos}}^{[1]} \otimes \cdots \otimes \mathcal{C}_{\text{sos}}^{[n]}$$

with invariant separable matrices M. The only difference is that the vectors $|v_k\rangle$ should be elementary tensors factors.

In order to state and prove the main result of this section (Theorem 5.1.11), it only remains to define sos (Ω, G)-decompositions—this is the non-stringent version advocated above.

Definition 5.1.3 (Invariant sos decompositions)
Let G be a group action on the WSC Ω, and let

$$\mathfrak{q} = (q_{\mathbf{k}})_{\mathbf{k} \in \mathcal{S}}$$

be a family of polynomials.

(i) An (Ω, G)-*decomposition of the family* $\mathfrak{q}$ is a decomposition

$$q_{\mathbf{k}} = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} q^{[1]}_{k_1, \alpha_{|1}}(\mathbf{x}^{[1]}) \cdots q^{[n]}_{k_n, \alpha_{|n}}(\mathbf{x}^{[n]})$$

for every $\mathbf{k} \in \mathcal{S}$, where

$$q^{[i]}_{k_i, \beta} \in \mathbb{R}[\mathbf{x}^{[i]}]$$

and

$$q^{[i]}_{k_i, \beta}(\mathbf{x}^{[i]}) = q^{[gi]}_{k_i, g\beta}(\mathbf{x}^{[i]})$$

for every $i \in [n]$, $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$, $g \in G$ and $\mathbf{k} \in \mathcal{S}$. The smallest cardinality of $\mathcal{I}$ among all (Ω, G)-decompositions is called the (Ω, G)-*rank* of $\mathfrak{q}$, denoted $\text{rank}_{(\Omega,G)}(\mathfrak{q})$.

(ii) An *sos* (Ω, G)-*decomposition* of $p \in \mathcal{P}$ is given by a sos decomposition into a family $\mathfrak{q}$

$$p = \sum_{\mathbf{k} \in \mathcal{S}} q_{\mathbf{k}}^2,$$

together with an (Ω, G)-decomposition of $\mathfrak{q}$. The minimal (Ω, G)-rank among all such sos decompositions is called the *sos* (Ω, G)-*rank* of p, denoted $\text{sos-rank}_{(\Omega,G)}(p)$. If G is the trivial group action, we call the sos (Ω, G)-decomposition just *sos* Ω-*decomposition* and denote its rank by sos-rank_{Ω}.

We are now ready to prove the main result on the existence of invariant sos polynomials:

Every G-invariant sos polynomial p has a G-invariant family $\mathfrak{q}$ (Theorem 5.1.11 (i)), and $\mathfrak{q}$ has an (Ω, G)-decomposition if G is a free group action on Ω (Theorem 5.1.11 (ii)). The idea of the proof of Theorem 5.1.11 (i) is to define $\mathfrak{q}$ as the square root of p, and show that this square root is also G-invariant. Some ideas of the proof are illustrated in Example 5.1.11.

Theorem 5.1.11 (Invariant sos decompositions)
Let Ω be a connected WSC with a free group action G. Furthermore, let $p \in \mathcal{P}$ be a G-invariant sos polynomial.

(i) There exists a G-invariant family of polynomials $\mathfrak{q} = (q_{\mathbf{k}})_{\mathbf{k}\in\mathcal{S}}$ such that

$$p = \sum_{\mathbf{k}\in\mathcal{S}} q_{\mathbf{k}}^2.$$

Moreover, every element $q_{\mathbf{k}}$ admits a decomposition in which the local polynomials at site i only depend on k_i, namely

$$q_{\mathbf{k}} = \sum_{j\in\mathcal{I}} q^{[1]}_{k_1,j}(\mathbf{x}^{[1]}) \cdots q^{[n]}_{k_n,j}(\mathbf{x}^{[n]}).$$

(ii) The invariant family $\mathfrak{q}$ admits an (Ω, G)-decomposition.

Note that in [55, Theorem 5.3], the authors prove the existence of so-called semi-symmetric sos decompositions for general representations of finite groups, by using Schur's lemma on the Gram matrix. Theorem 5.1.11 (i) is weaker than this statement, as it only considers group actions that permute the tensor product spaces, but gives an elementary proof.

Proof. (i) We denote the monomial $\mathbf{x} = (x_1, \dots, x_m)$ with exponent $\alpha = (\alpha_1, \dots, \alpha_m)$ by $\mathbf{x}^\alpha = x_1^{\alpha_1} \cdot x_2^{\alpha_2} \cdots x_m^{\alpha_m}$. Without loss of generality we can assume that $\deg_{\mathrm{loc}}(p) \leqslant 2d$. Define

$$\mathcal{S}_i = \{k \in \mathbb{N}^m : |k| \leqslant d\}$$

and $\mathcal{S} = \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$. Note that $\mathcal{S}$ can be identified with the set of monomials in $\mathcal{P}$ of local degree at most d via the correspondence

$$\mathcal{S} \to \mathcal{P}_d\colon \quad \mathbf{k} \mapsto \mathbf{x}^{\mathbf{k}} := \left(\mathbf{x}^{[1]}\right)^{k_1} \cdots \left(\mathbf{x}^{[n]}\right)^{k_n}.$$

Note also that the permutations of variables $\mathbf{x}^{[i]} \mapsto \mathbf{x}^{[gi]}$ coincide with the group action of G on $\mathcal{S}$, since

$$\left(\mathbf{x}^{[g1]}\right)^{k_1} \cdots \left(\mathbf{x}^{[gn]}\right)^{k_n} = \left(\mathbf{x}^{[1]}\right)^{k_{g^{-1}1}} \cdots \left(\mathbf{x}^{[n]}\right)^{k_{g^{-1}n}}. \tag{5.9}$$

Since p is G-invariant and sos, by Lemma 5.1.10 there exists a psd and G-invariant matrix M such that $\mathcal{G}(M) = p$. Now let B be the (unique) psd square root of M, i.e. $M = B^2$. Since M is a psd matrix, B admits a polynomial expression in M and hence B is also G-invariant. Define the polynomials $q_{\mathbf{k}}$ as

$$q_{\mathbf{k}} = \sum_{\mathbf{k}'\in\mathcal{S}} B_{\mathbf{k},\mathbf{k}'} \left(\mathbf{x}^{[1]}\right)^{k'_1} \cdots \left(\mathbf{x}^{[n]}\right)^{k'_n}$$

for $\mathbf{k} \in \mathcal{S}$. The family $\mathfrak{q} = (q_{\mathbf{k}})_{\mathbf{k}\in\mathcal{S}}$ is G-invariant, since

$$gq_k = \sum_{\mathbf{k}'\in\mathcal{S}} B_{g\mathbf{k},g\mathbf{k}'} \left(\mathbf{x}^{[1]}\right)^{k'_{g^{-1}1}} \cdots \left(\mathbf{x}^{[n]}\right)^{k'_{g^{-1}n}}$$
$$= \sum_{\mathbf{k}'\in\mathcal{S}} B_{g\mathbf{k},\mathbf{k}'} \left(\mathbf{x}^{[1]}\right)^{k'_1} \cdots \left(\mathbf{x}^{[n]}\right)^{k'_n} = q_{g\mathbf{k}}$$

where we have used the fact that $B_{\mathbf{k},\mathbf{k}'} = B_{g\mathbf{k},g\mathbf{k}'}$ for every $g \in G$ (which is just the G-invariance of B), together with Equation (5.9) and bijectivity of the map $\mathbf{k}' \mapsto g\mathbf{k}'$. In addition,

$$\sum_{\mathbf{k}\in\mathcal{S}} q_{\mathbf{k}}^2 = \langle \mathfrak{m}_{n,d} |\, B^t B \,| \mathfrak{m}_{n,d} \rangle = \mathcal{G}(M) = p$$

since $B^t B = B^2 = M$. Moreover, B admits a tensor decomposition

$$B_{\mathbf{k},\mathbf{k}'} = \sum_{j\in\mathcal{I}} \left(B_j^{[1]}\right)_{k_1,k'_1} \cdots \left(B_j^{[n]}\right)_{k_n,k'_n}.$$

Using the definition of $q_{\mathbf{k}}$ leads to the last statement of (i).

(ii) The proof is similar to that of Theorem 5.1.2. Start with decompositions

$$q_{\mathbf{k}} = \sum_{j\in\mathcal{I}} q_{k_1,j}^{[1]}(\mathbf{x}^{[1]}) \cdots q_{k_n,j}^{[n]}(\mathbf{x}^{[n]})$$

for every $\mathbf{k} = (k_1, \ldots, k_n) \in \mathcal{S}$. From the construction of Theorem 5.1.1 it follows that every polynomial $q_{\mathbf{k}}$ has a decomposition of the form

$$q_{\mathbf{k}} = \sum_{\alpha\in\mathcal{I}^{\widetilde{\mathcal{F}}}} p_{k_1,\alpha_{|1}}^{[1]}(\mathbf{x}^{[1]}) \cdots p_{k_n,\alpha_{|n}}^{[n]}(\mathbf{x}^{[n]})$$

where $\widetilde{\mathcal{F}}$ is the set of facets of Ω. We now construct a decomposition for every $q_{\mathbf{k}}$ which additionally satisfies the symmetry conditions of Definition 5.1.3 (i). Since G is free, by Lemma 5.1.3, there exists a G-linear map $\mathbf{z} : \widetilde{\mathcal{F}} \to G$. We consider the new index set $\hat{\mathcal{I}} := \mathcal{I} \times G$, together with the projection maps $\pi_1 : \hat{\mathcal{I}} \to \mathcal{I}$ and $\pi_2 : \hat{\mathcal{I}} \to G$. For each $i \in [n]$ and $\beta \in \hat{\mathcal{I}}^{\widetilde{\mathcal{F}}_i}$ we define the following local polynomials

$$q_{k_i,\beta}^{[i]}(\mathbf{x}^{[i]}) := \begin{cases} p_{k_i,g(\pi_1\circ\beta)}^{[gi]}(\mathbf{x}^{[i]}) & : \pi_2 \circ \beta = ({}^{g^{-1}}\mathbf{z})_{|i} \\ 0 & : \text{else.} \end{cases}$$

Similarly to the discussion in the proof of Theorem 5.1.2 we see that

$$q_{k_i,g\beta}^{[gi]}(\mathbf{x}^{[i]}) = q_{k_i,\beta}^{[i]}(\mathbf{x}^{[i]})$$

and

$$|G| \cdot q_{\mathbf{k}} = \sum_{\hat{\alpha}\in\hat{\mathcal{I}}^{\widetilde{\mathcal{F}}}} q_{k_1,\hat{\alpha}_{|1}}^{[1]}(\mathbf{x}^{[1]}) \cdots q_{k_n,\hat{\alpha}_{|n}}^{[n]}(\mathbf{x}^{[n]})$$

holds for every $\mathbf{k} \in \mathcal{S}$. But this implies the existence of an (Ω, G)-decomposition of $\mathfrak{q}$. □

From Theorem 5.1.11 the following statement immediately follows:

Corollary 5.1.12 (sos polynomials with free group action)
Let Ω be a connected WSC with a free group action from G. Then every sos and G-invariant $p \in \mathcal{P}$ admits an sos (Ω, G)-decomposition.

We end this section with an explicit example of an invariant sos decomposition.

Example 5.1.11 (Invariant sos decompositions)
Consider again the polynomial from Example 5.1.3,

$$p = x^2 + y^2 + 4(1 + xy)^2,$$

which is sos and invariant with respect to the permutation of x and y. We have already seen that $\mathrm{rank}_{(\Delta, C_2)}(p) = 2$.
To obtain a sos (Δ, C_2)-decomposition, we follow the proof of Theorem 5.1.11. We obtain $\mathcal{S} = \{0,1\} \times \{0,1\}$ with $G = C_2$ permuting the entries of the tuples, and obtain a C_2-invariant sos decomposition of p via the following family of polynomials:

By a similar argument as in Example 5.1.9, it can be shown that p is not separable with respect to the local sos cones.

$$q_{(0,0)} = q_{(1,1)} = \sqrt{2}(1 + xy), \quad q_{(0,1)} = y, \quad q_{(1,0)} = x.$$

On the double edge Δ we obtain an (Δ, C_2)-decomposition of the family via the following family of polynomials

$$q_0^{[1]} = \begin{pmatrix} \sqrt[4]{2}t & \frac{1}{\sqrt{2}} & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \quad q_0^{[2]} = q_0^{[1]^t}$$
$$q_1^{[1]} = \begin{pmatrix} 0 & 0 & 0 \\ \sqrt{2}t & \sqrt[4]{2}t & 0 \\ 0 & 0 & \sqrt[4]{2} \end{pmatrix}, \quad q_1^{[2]} = q_1^{[1]^t}.$$

where the matrix notation denotes that the rows are indexed by $\alpha = 1, 2, 3$ and the columns by $\beta = 1, 2, 3$. This shows that

$$\text{sos-rank}_{(\Delta, C_2)}(p) \leqslant \text{sos-rank}_{(\Delta, C_2)}(\mathfrak{q}) \leqslant 3.$$

On the single edge Σ_1, a decomposition of $\mathfrak{q}$ requires vectors $|a\rangle, |b\rangle, |c\rangle, |d\rangle \in \mathbb{R}^d$ of length $\sqrt[4]{2}$, with $|a\rangle, |b\rangle, |c\rangle$ pairwise orthogonal, $|d\rangle$ orthogonal to $|b\rangle$ and $|c\rangle$, and $\langle a \,|\, d\rangle = 1$. This is provided by

$$q_0^{[1]} = q_0^{[2]} = \big(\langle \alpha \,|\, a\rangle + \langle \alpha \,|\, b\rangle\, t \big)_{\alpha = 1, \ldots, d}$$
$$q_1^{[1]} = q_1^{[2]} = \big(\langle \alpha \,|\, c\rangle + \langle \alpha \,|\, d\rangle\, t \big)_{\alpha = 1, \ldots, d}$$

where $(\ )_\alpha$ denotes a vector indexed by α. Since such vectors can only be found in dimension $d \geqslant 4$, we obtain

$$\text{sos-rank}_{(\Lambda_2, C_2)}(p) \leqslant \text{sos-rank}_{(\Lambda_2, C_2)}(\mathfrak{q}) = 4.$$

We can also write p as a sum of symmetric squares:

$$p = \left(2 + \frac{3}{2}xy\right)^2 + (x+y)^2 + \left(\sqrt{\frac{7}{4}}xy\right)^2.$$

We now reset the variables $\mathcal{S}_1 = \mathcal{S}_2 = \{1,2,3\}, \mathcal{S} = \mathcal{S}_1 \times \mathcal{S}_2$, as well as

$$q_{(1,1)} = 2 + \frac{3}{2}xy,\ q_{(2,2)} = x + y,\ q_{(3,3)} = \sqrt{\frac{7}{4}}xy,$$

and all other $q_{\mathbf{k}} = 0$. This gives rise to the C_2-invariant family $\mathfrak{q} = (q_{\mathbf{k}})_{\mathbf{k}\in\mathcal{S}}$ that provides an sos decomposition of p with

$$\text{sos-rank}_{(\Delta,C_2)}(\mathfrak{q}) \leqslant 3.$$

But for the single edge, there does not exist a decomposition for the family $\mathfrak{q}$. This is because already $q_{(2,2)} = x + y$ does not admit an (Λ_2, C_2)-decomposition (without a minus sign). So

$$\text{sos-rank}_{(\Lambda_2,C_2)}(\mathfrak{q}) = \infty.$$

5.2 Inequalities and separations between the ranks

In this section, we study rank inequalities (Section 5.2.1), provide an upper bound for the separable rank (Section 5.2.2), and show separations between ranks (Section 5.2.3).

5.2.1 Inequalities between ranks

In this section, we show three relations between the introduced ranks (Proposition 5.2.2), which are similar to the statements established for tensor decompositions in [37, Proposition 29]. For the inequality between sos and separable decompositions we will need to assume that (Ω, G) is *factorizable*:

Note that Equation (5.10) can be seen as a system of linear equations by taking the logarithm on the left and the right hand side.

Definition 5.2.1 (Factorizable)
Let Ω be a WSC with a group action from G. We say that (Ω, G) is *factorizable* if for each finite index set $\mathcal{I}$ the following system of equations admits a solution with all $C^{[i]}_{\beta} > 0$ and $C^{[gi]}_{g\beta} = C^{[i]}_{\beta}$ for all $i \in [n]$, $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$, and $g \in G$:

$$C^{[1]}_{\alpha_{|_1}} \cdot C^{[2]}_{\alpha_{|_2}} \cdots C^{[n]}_{\alpha_{|_n}} = K_\alpha^{-1} \quad \text{for all } \alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}, \tag{5.10}$$

where

$$K_\alpha := \left| \left\{ \gamma \in \mathcal{I}^{\widetilde{\mathcal{F}}} : \begin{array}{l} \exists g_1, \ldots, g_n \in G \text{ with } g_i i = i \text{ and} \\ ({}^{g_i}\gamma)_{|_i} = \alpha_{|_i} \text{ for all } i \in [n] \end{array} \right\} \right|.$$

All examples of group actions on a weighted simplicial complex Ω considered in this paper are factorizable, as the following example shows.

Example 5.2.1 (Factorizable group actions)

Le us now present some examples of factorizable group actions:

(i) If $K_\alpha = 1$ for every $\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}$, then $C^{[i]}_\beta = 1$ solves Equation (5.10). This in particular shows that (Ω, G) is factorizable whenever the action of G on the vertices $[n]$ is free. In addition, this also implies that (Σ_n, S_n) is factorizable.

(ii) Let $\Omega = \Delta$ be the double edge and let $G = \{e, g\}$ act by keeping the vertices fixed (i.e. $ei = gi = i$) and flipping the facets (i.e. $g\mathfrak{a} = \mathfrak{b}$, $g\mathfrak{b} = \mathfrak{a}$)[15]. In this situation, we have

$$K_{\alpha_1,\alpha_2} = \begin{cases} 1 & : \text{if } \alpha_1 = \alpha_2 \\ 2 & : \text{if } \alpha_1 \neq \alpha_2. \end{cases}$$

A solution of Equation (5.10) is given by

$$C^{[i]}_{\alpha_1,\alpha_2} = \begin{cases} 1 & : \text{if } \alpha_1 = \alpha_2 \\ 1/\sqrt{2} & : \text{if } \alpha_1 \neq \alpha_2. \end{cases}$$

Hence, (Δ, G) is also factorizable.

15: This is different from the symmetric double edge of Example 2.2.4, since here the vertices remain fixed. The usual action on the double edge is free, and hence factorizable as well.

In fact, we are not aware of any non-factorizable (Ω, G) structures, leading to the following open question.

Question 5.2.1

Are there non-factorizable (Ω, G) structures?

We are now ready to present the rank inequalities.

Proposition 5.2.2 (Rank inequalities)

Let $p \in \mathcal{P}$.

(i) $\text{rank}_{(\Omega,G)}(p) \leqslant \text{sep-rank}_{(\Omega,G)}(p)$ for any separable cone.

(ii) $\text{rank}_{(\Omega,G)}(p) \leqslant \text{sos-rank}_{(\Omega,G)}(p)^2$.

(iii) If (Ω, G) is factorizable, then

$$\text{sos-rank}_{(\Omega,G)}(p) \leqslant \text{sep-rank}_{(\Omega,G)}(p)$$

for the separable cone over local sos polynomials.

Proof. (i) Every separable decomposition is an unconstrained decomposition. (ii) Let $\mathfrak{q} = (q_{\mathbf{k}})_{\mathbf{k}\in\mathcal{S}}$ be a G-invariant sos-decomposition of p, with an (Ω, G)-decomposition

$$q_{\mathbf{k}} = \sum_{\alpha \in \mathcal{I}^{\widetilde{\mathcal{F}}}} q^{[1]}_{k_1,\alpha_{|_1}}(\mathbf{x}^{[1]}) \cdots q^{[n]}_{k_n,\alpha_{|_n}}(\mathbf{x}^{[n]})$$

for each $\mathbf{k} \in \mathcal{S} = \mathcal{S}_1 \times \cdots \times \mathcal{S}_n$. Defining $\hat{\mathcal{I}} := \mathcal{I} \times \mathcal{I}$ and

$$p^{[i]}_{\beta,\beta'} := \sum_{k \in \mathcal{S}_i} q^{[i]}_{k,\beta}(\mathbf{x}^{[i]}) \cdot q^{[i]}_{k,\beta'}(\mathbf{x}^{[i]})$$

we obtain a valid (Ω, G)-decomposition of p, with

$$\operatorname{rank}_{(\Omega,G)}(p) \leqslant |\hat{\mathcal{I}}| = |\mathcal{I}|^2,$$

namely

$$\begin{aligned} p &= \sum_{\mathbf{k}\in\mathcal{S}} \sum_{\alpha,\alpha' \in \mathcal{I}^{\widetilde{\mathcal{F}}}} q^{[1]}_{k_1,\alpha_{|_1}}(\mathbf{x}^{[1]}) \cdot q^{[1]}_{k_1,\alpha'_{|_1}}(\mathbf{x}^{[1]}) \cdots q^{[n]}_{k_n,\alpha_{|_n}}(\mathbf{x}^{[n]}) \cdot q^{[n]}_{k_n,\alpha'_{|_n}}(\mathbf{x}^{[n]}) \\ &= \sum_{(\alpha,\alpha')\in\hat{\mathcal{I}}^{\widetilde{\mathcal{F}}}} p^{[1]}_{\alpha_{|_1},\alpha'_{|_1}}(\mathbf{x}^{[1]}) \cdots p^{[n]}_{\alpha_{|_n},\alpha'_{|_n}}(\mathbf{x}^{[n]}). \end{aligned}$$

(iii). Let $p^{[i]}_\beta \in \mathcal{C}^{[i]}_{\mathrm{sos}}$ for $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$ and $i \in [n]$ be local polynomials from a separable (Ω, G)-decomposition of p. So there exist sos decompositions

$$p^{[i]}_\beta = \sum_{k=1}^{N} \left(\tau^{[i]}_{k,\beta}\right)^2$$

with $\tau^{[i]}_{k,\beta} \in \mathbb{R}[x^{[i]}]$ (and we can clearly use the same sum length N for all i, β). We can in addition assume without loss of generality that

$$\tau^{[gi]}_{k,{}^g\beta}(\mathbf{x}^{[i]}) = \tau^{[i]}_{k,\beta}(\mathbf{x}^{[i]})$$

holds for all i, β, k and g. Indeed, just consider the action of G on

$$\bigcup_{i\in[n]} \{i\} \times \mathcal{I}^{\widetilde{\mathcal{F}}_i}$$

given by $g \cdot (i, \beta) := (gi, {}^g\beta)$, and fix for every orbit precisely one representative $(i_1, \beta_1), \ldots, (i_M, \beta_M)$. Then choose one sos decomposition for each $p^{[i_\ell]}_{\beta_\ell}$ and use the same along its orbit. This works since we have $p^{[gi]}_{{}^g\beta}(x^{[i]}) = p^{[i]}_\beta(x^{[i]})$ for all i, β by assumption.

Now since (Ω, G) is factorizable, we can choose some positive and G-invariant solution $\left(C^{[i]}_\beta\right)_{\beta,i}$ of Equation (5.10). Using the above representatives (i_ℓ, β_ℓ) again, we now define

$$q^{[i]}_{(\ell,k),\beta} := \begin{cases} \sqrt{C^{[i]}_\beta} \cdot \tau^{[i]}_{k,\beta}(\mathbf{x}^{[i]}) & : \text{if } \exists g \in G : (i,\beta) = (gi_\ell, {}^g\beta_\ell) \\ 0 & : \text{else} \end{cases}$$

where $\ell \in \{1, \ldots, M\}$, $k \in \{1, \ldots, N\}$ and $\beta \in \mathcal{I}^{\widetilde{\mathcal{F}}_i}$. By definition, we have

$$q^{[gi]}_{(\ell,k),{}^g\beta}(\mathbf{x}^{[i]}) = q^{[i]}_{(\ell,k),\beta}(\mathbf{x}^{[i]}),$$

and hence

$$q_{((\ell_1,k_1),\dots,(\ell_n,k_n))} := \sum_{\alpha\in\mathcal{I}^{\widetilde{\mathcal{F}}}} q^{[1]}_{(\ell_1,k_1),\alpha_{|_1}}(\mathbf{x}^{[1]})\cdots q^{[n]}_{(\ell_n,k_n),\alpha_{|_n}}(\mathbf{x}^{[n]})$$

is a valid (Ω, G)-decomposition of the G-invariant family

$$\mathfrak{q} := \left(q_{((\ell_1,k_1),\dots,(\ell_n,k_n))}\right)_{(\ell_i,k_i)\in\mathcal{S}_i}$$

where $\mathcal{S}_i = \{1,\dots,M\}\times\{1,\dots,N\}$. This family is also an sos decomposition of p, since

$$\begin{aligned}&\sum_{\forall i\colon(\ell_i,k_i)\in\mathcal{S}_i} q^2_{((\ell_1,k_1),\dots,(\ell_n,k_n))}\\ &=\sum_{\alpha\in\mathcal{I}^{\widetilde{\mathcal{F}}}} K_\alpha\cdot C^{[1]}_{\alpha_{|_1}}\cdots C^{[n]}_{\alpha_{|_n}}\cdot p^{[1]}_{\alpha_{|_1}}(\mathbf{x}^{[1]})\cdots p^{[n]}_{\alpha_{|_n}}(\mathbf{x}^{[n]}) = p.\end{aligned}$$

Here we have used Equation (5.10), as well as G-invariance of the $C^{[i]}_\beta$ and the $\tau^{[i]}_{k,\beta}$. □

5.2.2 An upper bound for the separable rank

We now provide an upper bound for the separable (Ω, G)-rank with respect to the number of local variables m_i and the polynomial's local degree. For simplicity, we again assume that all local polynomial spaces use the same number of variables, $m := m_i = m_j$ for $i, j \in [n]$. For $p \in \mathcal{P}$ recall that the local degree of p, denoted $\deg_{\text{loc}}(p)$, is the smallest integer $d \in \mathbb{N}$ such that

$$p \in \mathbb{R}[\mathbf{x}^{[1]}]_d \otimes\cdots\otimes \mathbb{R}[\mathbf{x}^{[n]}]_d$$

where $\mathbb{R}[\mathbf{x}]_d$ is the space of polynomials in variables $\mathbf{x}$ of degree at most d.

Proposition 5.2.3 (Upper bound for separable rank)
Let $p \in \mathcal{P}$ be separable and G-invariant, and let Ω be a connected WSC with a free group action G. Then

$$\text{sep-rank}_{(\Omega,G)}(p) \leqslant |G|\cdot\binom{\deg_{\text{loc}}(p)+m}{\deg_{\text{loc}}(p)}^n$$

for any separable cone.

Proof. Let $d = \deg_{\text{loc}}(p)$. Then $p \in \mathbb{R}[\mathbf{x}^{[1]}]_d \otimes\cdots\otimes \mathbb{R}[\mathbf{x}^{[n]}]_d$. Since

$$\dim\left(\mathbb{R}[\mathbf{x}^{[i]}]_d\right) = \binom{d+m}{d}$$

for all $i \in [n]$, p is a conic combination of at most $\binom{d+m}{d}^n$ elementary products with factors from the local cones by Carathéodory's Theorem[16].

16: See for example [6, Theorem 2.3].

From the proof of Theorem 5.1.1, we obtain

$$\text{sep-rank}_{\Omega}(p) \leqslant \binom{d+m}{d}^n.$$

The result now follows from Corollary 5.1.9. □

5.2.3 Separations

Here we will show *separations* between the polynomial ranks, which we will define shortly. Throughout this section we will consider separable decompositions only with respect to the local sos cones.

We know from Proposition 5.2.2 that the separable rank upper bounds both the rank and sos-rank. Here we will show that a reverse inequality is impossible: there are no functions $f, g\colon \mathbb{N} \to \mathbb{N}$ such that

$$\text{sep-rank}_{\Lambda_2}(p) \leqslant f\big(\text{sos-rank}_{\Lambda_2}(p)\big)$$

and

$$\text{sos-rank}_{\Lambda_2}(p) \leqslant g\big(\text{rank}_{\Lambda_2}(p)\big)$$

for all $m \in \mathbb{N}$ and polynomials $p \in \mathbb{R}[\mathbf{x}^{[1]}, \mathbf{x}^{[2]}]$ with $\mathbf{x}^{[i]} := (x_1^{[i]}, \ldots, x_m^{[i]})$. This is called a *separation* between sos-rank and sep-rank, or rank and sos-rank, respectively. We prove the separations by a reduction to matrix factorizations of entrywise nonnegative matrices, which themselves exhibit separations [49, 60].

For this reason, we focus on the subspace of n-quadratic forms in $\mathcal{P}$ and relate it with tensors. For $|T\rangle \in \mathbb{R}^m \otimes \cdots \otimes \mathbb{R}^m$ we define the polynomial

$$p_T := \sum_{j_1,\ldots,j_n=1}^{m} \langle j_1, \ldots, j_n \,|\, T\rangle \left(x_{j_1}^{[1]}\right)^2 \cdots \left(x_{j_n}^{[n]}\right)^2 \in \mathcal{P}. \tag{5.11}$$

There is a one-to-one correspondence between the tensor $|T\rangle$ and the polynomial p_T. In addition, entrywise nonnegativity of $|T\rangle$ fully characterizes the nonnegativity and the sos property of p_T:

Lemma 5.2.4 (Correspondence between tensors and polynomials) The map

$$\begin{aligned} \mathbb{R}^m \otimes \cdots \otimes \mathbb{R}^m &\to \mathcal{P} \\ |T\rangle &\mapsto p_T \end{aligned}$$

(where p_T is given in Equation (5.11)) is linear and injective. In addition, the following statements are equivalent:

(i) $|T\rangle$ is entrywise nonnegative.
(ii) p_T is a sos.
(iii) p_T is globally nonnegative[17].

17: That is,

$$p_T(\mathbf{x}^{[1]}, \ldots, \mathbf{x}^{[n]}) \geqslant 0$$

for $\mathbf{x}^{[i]} \in \mathbb{R}^m$.

Proof. Linearity and injectivity are immediate (each entry of $|T\rangle$ clearly gives rise to a different monomial).

The implications (i) $\Longrightarrow$ (ii) $\Longrightarrow$ (iii) are clear, since a nonnegative tensor T generates a sum of squares, since every sum of squares is globally nonnegative. For (iii) $\Longrightarrow$ (i) assume that $|T\rangle$ is not nonnegative, so there exist $j_1, \dots, j_n$ such that $\langle j_1, \dots, j_n \mid T\rangle < 0$. Then

$$p(\mathbf{e}_{j_1}, \dots, \mathbf{e}_{j_n}) = \langle j_1, \dots, j_n \mid T\rangle < 0,$$

where $\mathbf{e}_j$ is the jth standard vector. This shows that p is not nonnegative. □

In order to "borrow" the separations of tensor decompositions to derive separations of polynomial decompositions, we now show that the different notions of positive ranks for tensors correspond to the polynomial ranks.

The psd (Ω, G)-rank and the nonnegative (Ω, G)-rank are defined in Section 2.3.

Proposition 5.2.5 (Rank correspondence between tensors and polynomials)
Let $|T\rangle \in \mathbb{R}^m \otimes \cdots \otimes \mathbb{R}^m$ and the polynomial p_T be given by Equation (5.11).

(i) $\text{rank}_{(\Omega,G)}(T) = \text{rank}_{(\Omega,G)}(p_T)$.
(ii) $\text{nn-rank}_{(\Omega,G)}(T) = \text{sep-rank}_{(\Omega,G)}(p_T)$.
(iii) $\text{psd-rank}_{(\Omega,G)}(T) \leqslant \text{sos-rank}_{(\Omega,G)}(p_T)$
with equality if G acts freely on $[n]$.

Proof. (i). Let the families $\left(|T_\beta^{[i]}\rangle\right)_{\beta\in\mathcal{I}^{\widetilde{\mathcal{F}}_i}}$ provide an (Ω, G)-decomposition of $|T\rangle$ as in Definition 2.3.1. Now consider the families

$$\mathcal{P}^{[i]} := \left(\Psi_{T_\beta^{[i]}}(\mathbf{x}^{[i]})\right)_{\beta\in\mathcal{I}^{\widetilde{\mathcal{F}}_i}}$$

where for a vector $|V\rangle \in \mathbb{R}^m$ the Ψ notation indicates

$$\Psi_V(\mathbf{x}) := \sum_{j=1}^{m} \langle j \mid V\rangle\, x_j^2.$$

It is immediate to see that these families provide an (Ω, G)-decomposition of p_T, using the same index set $\mathcal{I}$.

Conversely, observe that every (Ω, G)-decomposition of p_T consists without loss of generality of local polynomials of the form

$$p_\beta^{[i]} = \sum_{j=1}^{m} \langle j \mid T_\beta^{[i]}\rangle \left(x_j^{[i]}\right)^2$$

for certain $|T_\beta^{[i]}\rangle \in \mathbb{R}^m$. All other possible monomials will have to cancel out in the total product and sum, and can therefore be omitted. Thus the $|T_\beta^{[i]}\rangle$ give rise to an (Ω, G)-decomposition of $|T\rangle$, again with the same index set $\mathcal{I}$.

Statement (ii) is proven exactly as (i), and using the fact that the local polynomials of an sos (Ω, G)-decomposition of p_T must all be of degree 2, and thus have nonnegative coefficients at all the $\left(x_j^{[i]}\right)^2$.

For (iii) we start with an sos (Ω, G)-decomposition of p_T, where every local polynomial $q_{k,\beta}^{[i]}$ can (for degree reasons) be assumed to be of the form

$$q_{k,\beta}^{[i]} = \sum_{j=1}^{m} \left(B_j^{[i]}\right)_{k,\beta} x_j^{[i]}.$$

Now the matrices

$$E_j^{[i]} := \left(B_j^{[i]}\right)^t \left(B_j^{[i]}\right) \geqslant 0$$

give rise to a psd (Ω, G)-decomposition of $|T\rangle$ of the same rank as the initial decomposition. This can easily be seen by computing the coefficient of p_T at each monomial $(x_{j_1}^{[1]})^2 \cdots (x_{j_n}^{[n]})^2$, and checking that it arises from the sos (Ω, G)-decomposition.

For the reverse inequality, we assume that G acts freely on $[n]$. We start with a psd (Ω, G)-decomposition of T, i.e.

$$\langle j_1, \ldots, j_n \mid T\rangle = \sum_{\alpha,\alpha' \in \mathcal{I}^{\widetilde{\mathcal{F}}}} \left(E_{j_1}^{[1]}\right)_{\alpha_{|_1},\alpha'_{|_1}} \cdots \left(E_{j_n}^{[n]}\right)_{\alpha_{|_n},\alpha'_{|_n}}$$

where all $E_j^{[i]}$ are psd. Decompose $E_j^{[i]} = \left(B_j^{[i]}\right)^t \left(B_j^{[i]}\right)$ with the additional constraint that

$$\left(B_j^{[gi]}\right)_{k,g\beta} = \left(B_j^{[i]}\right)_{k,\beta}.$$

Since G acts freely on $[n]$, we can just choose certain $B_j^{[i]}$ and define the $B_j^{[gi]}$ along the orbit by that formula. Now defining

$$q_{(j,k),\beta}^{[i]} := \left(\widetilde{B}_j^{[i]}\right)_{k,\beta} x_j^{[i]}$$

leads to a sos (Ω, G)-decomposition with sos-rank$_{(\Omega,G)}(p_T) \leqslant |\mathcal{I}|$. □

The proof of Proposition 5.2.5 (iii) does not work in reverse direction if we do not assume that G acts freely on $[n]$. Assume there exists $e \neq g \in G$ and $i \in [n]$ such that $gi = i$. Then, the construction into a symmetric factorization $\widetilde{B}_j^{[i]}$ implies that

$$\begin{aligned}\left(E_j^{[gi]}\right)_{g\beta,\beta'} &= \left(\widetilde{B}_j^{[gi]}\right)^t_{g\beta,-} \left(\widetilde{B}_j^{[gi]}\right)_{-,\beta'} \\ &= \left(\widetilde{B}_j^{[i]}\right)^t_{\beta,-} \left(\widetilde{B}_j^{[i]}\right)_{-,\beta'} = \left(E_j^{[i]}\right)_{\beta,\beta'}\end{aligned}$$

which is stronger than the symmetry of $E_j^{[i]}$ given in a psd (Ω, G)-decomposition.

We now show that there is a separation between the ranks already for decompositions on the single edge.[18]

18: Note that in the following corollary p_m is a polynomial on the single edge.

Corollary 5.2.6 (Rank separations on the single edge)
Let $p_m \in \mathbb{R}[x_1^{[1]}, \dots, x_m^{[1]}, x_1^{[2]}, \dots, x_m^{[2]}]$.

(i) There exists a sequence of polynomials $(p_m)_{m\in\mathbb{N}}$ such that

$$\text{rank}_{\Lambda_2}(p_m) = 3, \quad \text{sos-rank}_{\Lambda_2}(p_m) = 2$$

and

$$\log_2(m) \leqslant \text{sep-rank}_{\Lambda_2}(p_m) < \infty$$

(ii) There exists a sequence of polynomials $(p_m)_{m\in\mathbb{N}}$ such that $\text{rank}_{\Lambda_2}(p_m) = 3$ and[19]

$$\lim_{m\to\infty} \text{sos-rank}_{\Lambda_2}(p_m) = \infty$$

19: Of course we have that

$$\text{sos-rank}_{\Lambda_2}(p_m) < \infty.$$

Proof. (i). The Euclidean distance matrix $M_m \in \text{Mat}_m(\mathbb{R}) \cong \mathbb{R}^m \otimes \mathbb{R}^m$ which is defined as

$$(M_m)_{i,j} = (i-j)^2$$

satisfies[20]

20: See [49, Example 5.17] for details.

$$\text{rank}_{\Lambda_2}(M_m) = 3, \quad \text{psd-rank}_{\Lambda_2}(M_m) = 2,$$

and

$$\text{nn-rank}_{\Lambda_2}(M_m) \geqslant \log_2(m)$$

since all explicit examples are given as a real matrix factorization. Defining $p_m := p_{M_m}$ and using Proposition 5.2.5 shows the statement.

(ii) is similar to (i), this time using the slack matrix of an m-gon for every $m \in \mathbb{N}$.[21] □

21: We refer to [49, Example 5.14] for the definition of a Slack matrix of a polyhedron.

These statements imply that there cannot exist functions $f, g\colon \mathbb{N} \to \mathbb{N}$ such that

$$\text{sep-rank}_{\Lambda_2}(p) \leqslant f\big(\text{sos-rank}_{\Lambda_2}(p)\big)$$

and

$$\text{sos-rank}_{\Lambda_2}(p) \leqslant g\big(\text{rank}_{\Lambda_2}(p)\big)$$

holds for all $m \in \mathbb{N}$ and all polynomials $p \in \mathbb{R}[\mathbf{x}^{[1]}, \mathbf{x}^{[2]}]$ with $\mathbf{x}^{[i]} := (x_1^{[i]}, \dots, x_m^{[i]})$. This also holds true for polynomials of bounded degree, since $\deg(p_m) = 4$ in the above construction.

This immediately leads to the question of whether there are separations between the ranks of polynomials with a bounded number of variables and no bound on the degree. In this setting there does not exist a one-to-one correspondence between polynomials and Gram matrices (as that of Example 5.1.9). We believe that separations will again appear in the simplest setting and leave this question as a conjecture.

Conjecture 5.2.7
There do not exist functions $f, g, h : \mathbb{N} \to \mathbb{N}$ such that for all $p \in \mathbb{R}[x, y]$ (in particular, independently of the degree of p)

(i) $\text{sep-rank}_{\Lambda_2}(p) \leqslant f(\text{rank}_{\Lambda_2}(p))$

(ii) $\text{sep-rank}_{\Lambda_2}(p) \leqslant g(\text{sos-rank}_{\Lambda_2}(p))$
(iii) $\text{sos-rank}_{\Lambda_2}(p) \leqslant h(\text{rank}_{\Lambda_2}(p))$

where p is separable in (i) and (ii), and a sum of squares in (iii). The separable rank is again meant with respect to the local sos-cones.

5.3 Conclusions and outlook

In summary, we have defined and studied several decompositions of multivariate polynomials into local polynomials, each containing only a subset of variables. The variables are divided into blocks, and each local polynomial uses only one block. We describe a decomposition with WSC Ω, whose vertices describe the individual blocks, and facets the summation indices. For polynomials invariant under the permutation of blocks of variables, we have defined and studied an invariant decomposition. We have also defined an invariant decomposition with local positivity conditions, specifically, with the separable and sum of squares condition.

Specifically, we have defined invariant polynomial decompositions (Definition 5.1.1) and shown that every G-invariant polynomial admits an (Ω, G)-decomposition if G acts freely on Ω (Theorem 5.1.2). Moreover, if G is a blending group action, every G-invariant polynomial can be written as a difference of two (Ω, G)-decompositions (Theorem 5.1.7). We have also defined the separable (Ω, G)-decomposition (Definition 5.1.2), and sum of squares (Ω, G)-decomposition (Definition 5.1.3), and have shown that they exist if G acts freely on Ω (Theorem 5.1.8 and Corollary 5.1.12, respectively).

In addition, we have shown that the (Ω, G)-rank of a polynomial can be upper bounded in terms of its separable and sos rank, and that the sos rank can often be upper bounded by its separable rank (Proposition 5.2.2). In the reverse direction such inequalities cannot exist, since there exists a sequence of polynomials with constant (Ω, G)-rank and a diverging sos or separable rank (Corollary 5.2.6).

This work has left two open questions:

- Do the rank separations also hold with respect to a bounded number of variables but unbounded degree (Conjecture 5.2.7), and
- Does there exist non-factorizable (Ω, G) structures (Question 5.2.1)?

A more general open question concerns the full characterization of the existence of invariant polynomial decompositions, as freeness of the group action only provides a sufficient condition. Our investigations indicate that it may also be necessary, but we were not able to prove it.

Part II

Computational Aspects of Tensor Decompositions and Beyond

Computational complexity in semi-algebraic geometry

6

Computation is a concept that has existed in some form for a long period of time. In its usual interpretation, this term refers to the process of producing an output from a set of inputs after applying a finite number of standard operations, for example, addition or multiplication of numbers. In the early 20$^{\text{th}}$ century, many models of computation were formally introduced, leading to the birth of computational complexity theory.

In computational complexity theory, computational procedures are modeled via *Turing machines*. These machines reflect our intuitive notion of computation, namely a fixed number of basic operations performed on an input with the possibility of writing down intermediate results on a scratchpad. The basic operations are modeled by a finite table of transitions and the scratch pad by an infinitely long tape. Despite their simplicity, Turing machines embody the entirety of computational capabilities that are achievable by nature.

Turing machines have been useful to classify the resource usage of different computational problems. This includes the following distinctions of problems:

- Determining whether problems are decidable or undecidable — i.e., whether a given problem can be solved within a finite amount of time.
- Identifying whether a problem admits an *efficient* solution by a Turing machine — i.e., whether the computation time scales reasonably with the size of the input.

In Section 6.1, we introduce the basics of computational complexity, providing a rigorous framework to answer these questions. In this section, we introduce the concept of (non-deterministic) Turing machines alongside the notions of (un-)decidability. Moreover, we survey well-known examples of computational complexity classes such as polynomial-time problems, non-deterministic polynomial-time problems as well as recursively enumerable problems. We also review the concept of *hardness* in computational complexity to lower bound computational complexities.

In Section 6.2, we present tools from (semi-)algebraic geometry that give rise to algorithms for problems in quantum information. Many problems in quantum information involve an infinite amount of polynomial equations or a search over an uncountable amount of values. This includes, for example, checking membership in the set of separable states or block-positive matrices. Semi-algebraic geometry (i.e. the study of systems of polynomial inequalities) provides an algorithmic approach to solve these problems in finite time. Specifically, the *Tarksi–Seidenberg theorem* and *Hilbert's basis theorem* allow us to construct algorithms to solve problems that seem naively not decidable in finite time.

These tools will then be applied in the two remaining chapters of this part:

- In Chapter 7, we prove that the moment membership problem can be decidable or undecidable for certain instance sets. Specifically, we consider the following question: Given a matrix A, check whether
$$\mathrm{tr}(A^n) \in \mathcal{P}$$
for every $n \in \mathbb{N}$ for specific sets $\mathcal{P}$. This problem entails verifying $\mathrm{tr}(A^n) \in \mathcal{P}$ for the countably many cases $n \in \mathbb{N}$. Consequently, no finite decision procedure follows directly from the problem's definition. Nonetheless, leveraging tools from semi-algebraic geometry, we establish that this problem is decidable for certain classes of matrices A, such as unitary matrices. Conversely, we prove that the same problem becomes undecidable when A is a matrix over a ring, such as the ring of commutative or non-commutative polynomials.
- In Chapter 8, we introduce the notion of a bounded version of a decision problem. Many undecidable problems in physics, mathematics, and computer science share a common feature: They consist of infinitely many statements over an unbounded parameter (similar to the moment membership problem). We demonstrate that bounding this parameter makes the problems decidable; however, they remain NP-hard in most situations.

6.1 Basics in computational complexity

Computational complexity provides a formal framework to understand the computational resources required to solve problems. At its core are *Turing machines*, which embody our intuitive notion of computation: Performing basic operations while utilizing a scratch pad to record intermediate results. The adoption of Turing machines as a model of computations stems from the *Church–Turing* thesis;

> *Every physically realizable computation can be executed by a Turing machine.*

This thesis is motivated by the equivalence to many other models of computation, namely λ-calculus, RAM machines, or cellular automata. All of these models were found to be equivalent, i.e. if a function is computable within one model, then it is also computable within every other of the mentioned models [2].

Turing machines serve as representatives for real-world computation. Furthermore, this model gives rise to several fundamental concepts, including *efficient computation* and problems that are efficiently verifiable but not necessarily efficiently disprovable.

We shall present the notion of a *Turing machine*. Moreover, we will review the concept of decision problems and their computational complexity. Specifically, we will present several families of complexity classes, including the class of *recursively enumerable* problems, and the class of *non-deterministic polynomial time* problems.

For a more detailed introduction to computational complexity theory, we refer to the textbooks by Arora and Barak [2], by Papadimitriou [94], by Sipser [115], or by Widgerson [128].

6.1.1 Turing machines

In the following, we introduce the notion of a *Turing machine*, which is the most commonly used model of computation. Turing machines reflect the longheld intuition of computation: Certain mechanical rules are applied to manipulate numbers, and it is allowed to use a notebook for intermediate results. Although introduced at the beginning of the 20^{th} century, Turing machines can be understood as a model of modern computers, with the difference that Turing machines have no built-in upper bound in the memory size.

A Turing machine consists of the following three parts (illustrated in Figure 6.1):

- A *tape* divided into individual cells arranged adjacently. Each cell holds a symbol from a finite set Σ, the tape alphabet. The tape is assumed to be infinitely extendable in both the right and left direction, serving as the computational scratchpad.
- A *head* that can read and write on the tape cells. It can move to the right or to the left, one step at a time.
- A *finite program* equipped with an internal state register comprising finitely many states $\mathcal{Q}$. This program can interact with the head. Conceptually, it is a finite set of instructions that depend on the internal state and on the tape entry. The instructions involve changing the internal state, writing on the tape and moving the head to the left or to the right. For this reason, the set of instructions defines a function

$$\delta : \mathcal{Q} \times \Sigma \to \mathcal{Q} \times \Sigma \times \{L, R\}, \tag{6.1}$$

where Σ is the tape alphabet and $\mathcal{Q}$ is the set of states.[1] For instance, $\delta(q_1, s_1) = (q_2, s_2, L)$ indicates that if the head reads symbol s_1 while in state q_1, the state transitions to q_2, the head overwrites s_1 with s_2, and moves leftward.

1: Both, Σ and $\mathcal{Q}$ are finite sets. This implies that δ can be represented in a finite way.

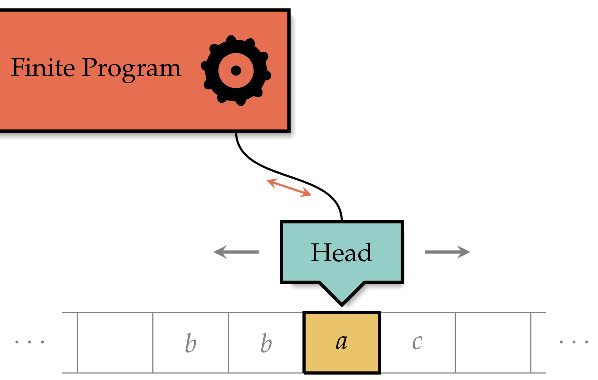

Figure 6.1: Illustration of a Turing machine consisting of an infinite tape, a head, and a finite program. The head can read and write on the tape and move (depending on the instructions of the program). The finite program is modeled by a finite set of states and a transition function introduced in Equation (6.1).

A configuration of a Turing machine, defined by a state and a tape entry, is called an *instantaneous description* of the Turing machine. Consequently, the transition function δ can be viewed as a mapping between various configurations of the Turing machine.

A *computation step* of the Turing machine consists of the head reading the current tape cell entry, resulting in a configuration (q, s), and then applying δ to (q, s). In the new configuration obtained from δ, the Turing machine updates its internal state, inscribes the new symbol onto the tape, and shifts its head one step left or right. Thus, each computation step corresponds to obtaining one instantaneous description.

Turing machines encapsulate the full power of computation. In other words, adding features does not increase the computational power. For example, adding a second tape, a random access memory, letting δ be non-deterministic, or even a quantum device does not change the set of problems that are computable with Turing machines (see [2, Section 1.2.2]).

Accepting and rejecting inputs

All Turing machines considered three states of particular importance: the initial state q_i, the accept state q_a, and the reject state q_r. The Turing machine starts with the initial state, and the accept/reject states are meant to partition inputs into two classes. For this purpose, we assume that the Turing machine does nothing after reaching q_a or q_r.

We define a Turing machine T to *halt on input x* — a string initially written on the tape — if T reaches either the state q_a or the state q_r after a finite number of computation steps, starting from the initial state q_i. If T arrives at q_a, we say that T accepts x (denoted $T(x) = 1$), whereas if it reaches q_r, we say that T rejects x (denoted $T(x) = 0$).

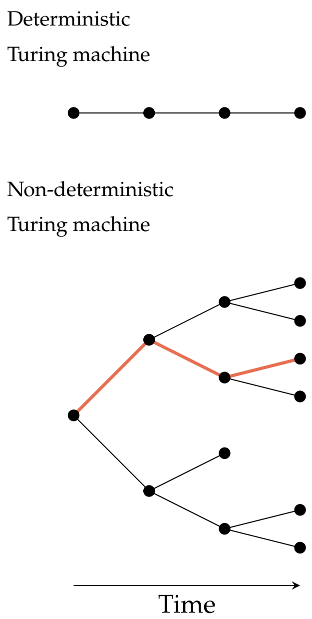

Figure 6.2: (Deterministic) Turing machines vs. non-deterministic Turing machines. Every vertex represents one instantaneous description and every edge a computational step. While the deterministic one has only one computation path (since δ is a function), the computation paths of a non-deterministic one form a tree. In this example every transition has precisely two outcomes, and one example of a computation path is highlighted in orange. The number of distinct computation paths increases exponentially in the number of computation steps. Note that the computation can halt on certain paths earlier than on others.

Non-deterministic Turing machines

There are extensions of Turing machines that are computationally more efficient than standard Turing machines, yet their implementation is not physical. One example is the so-called *non-deterministic Turing machine*. This allow for non-deterministic transitions. Specifically, for a non-deterministic Turing machine, the transition

$$\delta : \mathcal{Q} \times \Sigma \to \mathcal{Q} \times \Sigma \times \{L, R\}$$

can be multi-valued, meaning that $\delta(q, x)$ can have multiple outcomes. While deterministic Turing machines follow a single computational path, non-deterministic Turing machines can explore a tree of computational paths (refer to Figure 6.2), due to the multiple outcomes of each transition. Non-deterministic Turing machines are believed to be more efficient than ordinary deterministic ones; however, non-deterministic transitions cannot be implemented physically. Another extension with a similar behavior is the Turing machine with an additional quantum device (see [2, Chapter 20]). These machines are also believed to be more efficient than standard Turing machines, however, the set of computable problems remains the same for all these models.

6.1.2 Decision problems and computability

In the following, we use the concept of Turing machines to present the notion of *computable* functions and *decision problems*.

For a given finite alphabet Σ, we define the Kleene star on Σ as

$$\Sigma^* := \{c_1 c_2 \dots c_n : \quad n \in \mathbb{N}, c_i \in \Sigma\},$$

i.e. the set of strings generated by characters from Σ. Boolean functions

$$f : \Sigma^* \to \{0, 1\}$$

give rise to the notion of a *decision problem*. Mapping a string to 1 is interpreted as the string being accepted, and mapping it to 0 means the string being rejected. In essence, a decision problem divides the set of all strings Σ^* into two categories: *yes-instances*, where $f(x) = 1$, and *no-instances*, where $f(x) = 0$.

Alternatively, decision problems can also be defined via a *language* that is defined as

$$L := \{x \in \Sigma^* : \ f(x) = 1\} \subseteq \Sigma^*.$$

Throughout this thesis, we will use the terms language and (decision) problem interchangeably.

In practice, we often encounter functions whose domain is not inherently defined by a set of strings, like over the natural numbers $f : \mathbb{N} \to \{0, 1\}$. These can still be understood as decision problems by choosing a proper encoding of the domain. For example, the set $\mathbb{N}$ can be encoded with a binary encoding $\{0, 1\}^*$ by associating every string $(s_0, \dots, s_n) \in \{0, 1\}^*$ with the natural number

$$a = \sum_{k=0}^{n} s_k 2^k.$$

Similar encodings exist for $\mathbb{Z}$, $\mathbb{Q}$, or $\mathrm{Mat}_s(\mathbb{Q})$; however, sets like $\mathbb{R}$ or $\mathbb{C}$ do not admit such a finite encoding, as they are uncountable. For this purpose, instance sets are always restricted to sets that admit a finite encoding.[2]

2: A further example of a set that attains a finite description, are the *algebraic numbers* in $\mathbb{C}$, i.e. elements that arise as roots of a polynomial. A possible finite representation can be constructed by using a polynomial whose root is the element of interest.

A further encodable set is the set of all Turing machines $\mathcal{T}$. Every Turing machine can be represented in a finite way via its finite state set, its finite tape alphabet and its transition function δ which consists of a finite number of instructions.

Let us now introduce one famous decision problem on this instance set, the *halting problem*. We denote its language by $\textsc{Halt}$.

Example 6.1.1 (The halting problem)
The halting problem $\textsc{Halt}$ is a decision problem on $\mathcal{T} \times \Sigma^*$ defined as

$$\langle T, x\rangle \in \textsc{Halt} \quad \Longleftrightarrow \quad T \text{ halts on input } x.$$

One of the central questions in computational complexity is whether decision problems are computable or not. In essence: does there exist a procedure to compute a function $f : \Sigma^* \to \{0, 1\}$? A decision problem given by a language L is called *decidable* (in short $L \in \mathsf{R}$ for recursive), if there exists a Turing machine T such that

$$x \in L \quad \Longleftrightarrow \quad T(x) = 1.$$

In plain words, T halts on every input x and accepts x if and only if $x \in L$. This means that there is a finite procedure that decides whether $f(x) = 0$ or $f(x) = 1$ reaching the accept or the reject state after finitely many computation steps. While there are many problems that can be shown to be decidable by giving an explicit description of such a computation procedure, there are many more problems which are *undecidable*.[3] One such problem is the halting problem [123].

3: This follows from the fact that there are uncountably many functions

$$f : \mathbb{N} \to \{0, 1\},$$

while there are only countably many Turing machines which model decidable functions.

Theorem 6.1.1

The Halting problem HALT is undecidable.

This statement is proven via contradiction: If there exists a procedure that decides HALT, this implies a logical contradiction. For a comprehensive proof, we refer to [2, Section 1.4].

Analogous to decision problems (i.e., Boolean functions yielding two outcomes), there exists a concept of computability for functions

$$g : \Sigma^* \to \Sigma^*.$$

We say that g is *computable* if there exists a Turing machine T which halts on every input x and the outcome of the function

$$y = f(x)$$

is finally written on the tape. This notion will be an important ingredient for reductions in Definition 6.1.4.

6.1.3 Computational complexity classes

Thus far, we have observed that decision problems fall into the categories of decidable or undecidable. However, the decidable nature of a problem does not guarantee practical solvability — that is, the ability to solve the problem efficiently or within a reasonable timeframe, in practice. It is plausible for a problem to demand an exorbitantly large number of computation steps, even for relatively small inputs. For this purpose, there exist further complexity classes that capture efficiently solvable problems, namely the class of polynomial-time problems, as well as efficiently verifiable problems (so-called NP-problems).

Polynomial-time problems

We say that a Turing machine T is polynomial-time if there exists a polynomial $p : \mathbb{N} \to \mathbb{N}$ such that for every input $x \in \Sigma^*$, T halts within $p(|x|)$ steps, where $|x|$ denotes the length of the string.

This definition is meant to reflect that the computation time is reasonably small for every input. However, note that p can be arbitrary in this definition, including also Turing machines whose runtime scales for example with $|x|^{1000}$. We refer to [2, Section 1.5] for a detailed discussion on the philosophical importance and criticism on this definition.

Definition 6.1.1 (Polynomial-time decision problems)

Let L be a language. We say that L is polynomial-time decidable (in short $L \in \mathsf{P}$) if there exists a polynomial-time Turing machine T such that

$$x \in L \iff T(x) = 1.$$

Many efficiently solvable problems are in P, for example, multiplying matrices or finding the shortest path between two vertices on a graph. Most of these problems have in common that the exponent in the polynomial p is reasonably small, which makes them also efficiently solvable in practice.

Non-deterministic polynomial-time problems

When solving a puzzle, it makes a huge difference in solving this puzzle from scratch versus verifying if a given solution is correct. In physics and mathematics, many problems share a similar behavior. We now review the complexity class NP (and subsequently coNP), which precisely captures this.

Definition 6.1.2
Let L be a language. We say the L is non-deterministic polynomial-time (in short $L \in \mathsf{NP}$) if there exists a polynomial-time Turing machine T and a polynomial p such that

$$x \in L \quad \Longleftrightarrow \quad \exists y \in \Sigma^* \colon |y| \leqslant p(|x|) \colon T(\langle x, y \rangle) = 1$$

Here, $\langle x, y \rangle$ means that the strings x and y are merged with each other, separated through a colon.

In simple terms, if a problem L is in the complexity class NP, it means that for every instance where the answer is yes, there exists a short, easily checkable proof, also called a certificate. This certificate provides evidence that the instance indeed belongs to the set of yes-instances. Think of it like having the solution to a puzzle — if you have the solution, it is quick to check that it is correct. However, if the instance is a no-instance, there is no straightforward way to certify it. In other words, there is no quick, easily verifiable proof that the instance does not belong to the set of yes-instances. This mirrors the situation where verifying that a puzzle has no valid solution is challenging.

When the complement of a language L, defined as $L^c := \Sigma^* \setminus L$, belongs to the class NP, we say that $L \in \mathsf{coNP}$. This complexity class operates similarly to NP, but it focuses on verifying no-instances, contrasting with NP which verifies yes-instances. Namely, $L \in \mathsf{coNP}$ if there exists a polynomial-time Turing machine T and a polynomial p such that

$$x \in L \quad \Longleftrightarrow \quad \forall y \in \Sigma^* \colon |y| \leqslant p(|x|) \colon T(\langle x, y \rangle) = 1. \tag{6.2}$$

Since yes- and no-instances are asymmetric in the definition of NP, the complexity class coNP might be very different from NP.

There are many problems in NP that are unknown to be in P. Examples include the 3-satisfiability problem SAT, graph problems like MaxCut, and the 3-coloring problem (see Figure 6.3). We refer to [54] for details and many more examples of such problems. Despite extensive investigation, the existence of efficient algorithms for these problems remains uncertain. In fact, the conjecture $\mathsf{P} \neq \mathsf{NP}$ is one of the most significant unsolved problems in computer science.

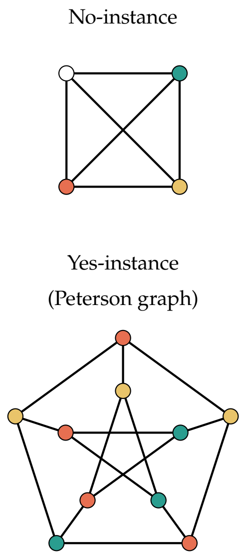

Figure 6.3: The 3-coloring problem. Can the vertices of a graph be colored with three colors such that its adjacent vertices have a different color? These figures show a yes-instance and a no-instance. The fully connected graph with four vertices cannot be 3-colored since the upper left vertex does not admit any color that is not used already for an adjacent vertex. The 3-coloring problem is an NP-complete decision problem.

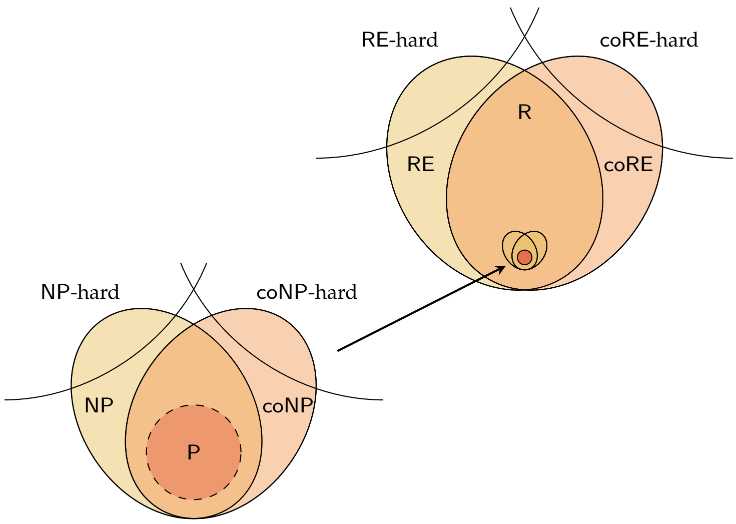

Figure 6.4: The complexity classes introduced in this chapter. R is the set of decidable languages that corresponds to the intersection of RE and coRE. RE-hard and coRE-hard problems are harder than all RE and coRE problems. A subset of decidable problems are P, NP, and coNP problems. In contrast to R, the set P might be a strict subset of NP $\cap$ coNP.

An example of an NP-problem is the non-deterministic bounded Halting problem.

Example 6.1.2

The non-deterministic bounded halting problem BNHalt is a decision problem on $\mathcal{T}_N \times \mathbb{N}$ defined as

$$\langle T, n \rangle \in \text{BNHalt} \quad \Longleftrightarrow \quad T \text{ halts on the empty input in } n \text{ steps.}$$

Here, $\mathcal{T}_N$ is the set of all non-deterministic Turing machines.

BNHalt is in NP because it is easy to verify if T halts within n steps by giving the computational path as a certificate. However, it is hard to verify that T does not halt within n steps since one has to check that it does not halt on any of the exponentially many computational paths (see Figure 6.2).

Semi-decidable problems

We now present an analog notion to NP at the level of decidable problems, namely the set of *recursively enumerable* languages.

Definition 6.1.3 (Recursively enumerable)

A language $L \subseteq \Sigma^*$ is called *recursively enumerable* (in short $L \in$ RE) if there exists a Turing machine T such that

$$x \in L \quad \Longleftrightarrow \quad \exists y \in \Sigma^* \colon T(\langle x, y \rangle) = 1.$$

Moreover, L is called *co-recursively enumerable* (in short $L \in$ coRE) if $L^c \in$ RE, i.e. there exists a Turing machine T' such that

$$x \in L \quad \Longleftrightarrow \quad \forall y \in \Sigma^* \colon T'(\langle x, y \rangle) = 1.$$

In simple terms, for a recursively enumerable language it is possible to verify that x is a yes-instance by checking whether there exists a finite certificate y that verifies x via the Turing machine T. However, certifying that $x \notin L$ might not be possible in finite time, since one must check whether none of the (infinitely many) certificates verifies x. For this purpose, RE-problems are called *semi-decidable*, since they can only verify one possibility (namely $x \in L$) in finite time.

The halting problem HALT is an example that is semi-decidable, but not decidable. If $\langle T, x \rangle$ is a yes-instance of HALT, i.e. T halts on x, then there exists a finite number $n \in \mathbb{N}$ such that T halts on x within n computation steps. Note that the number n can be arbitrarily large (independent of $|x|$). Using the halting time n as a certificate shows that HALT $\in$ RE since checking that T halts on x within n steps can be done in finite time via using a universal Turing machine that simulates T.[4]

4: We refer to [2, Section 1.3] for an elaborate discussion on the notion of a universal Turing machine.

The intersection of recursively enumerable (RE) and co-recursively enumerable (coRE) languages coincides with the set of decidable languages, denoted as R, i.e.

$$\mathsf{R} = \mathsf{RE} \cap \mathsf{coRE}.$$

For a language $L \in \mathsf{RE}$, there is a Turing machine T_1 such that

$$x \in L \quad \Longleftrightarrow \quad \exists y \in \Sigma^* : T_1 (\langle x, y \rangle) = 1.$$

Similarly, since $L \in \mathsf{coRE}$, there is a Turing machine T_2 such that

$$x \notin L \quad \Longleftrightarrow \quad \exists y \in \Sigma^* : T_2 (\langle x, y \rangle) = 1.$$

Enumerating among all strings $y \in \Sigma^*$ and letting T_1 and T_2 run in parallel leads to an algorithm that halts for every input in finite time. If $x \in L$, then T_1 will accept in after finite iterations, if $x \notin L$, then T_2 accepts in finite time.

In Section 7.2.2, we use this observation to establish the decidability of the moment membership problem. Specifically, we present an algorithm to verify yes-instances in finite time and a method to verify no-instances in finite time to construct an algorithm for the problem.

It is worth noting that a similar statement for NP and coNP is not true. While we have

$$\mathsf{P} \subseteq \mathsf{NP} \cap \mathsf{coNP},$$

the inclusion is believed to be strict.

Complexity lower bounds

Mathematics, physics, and computer science are full of problems that do not seem to have a simple solution; it is even impossible to construct an algorithm to solve them. For this purpose, it is relevant to classify the *hardness* of a decision problem. Computational complexity relies on a many conjectures, like the famous $\mathsf{P} \overset{?}{=} \mathsf{NP}$. This exemplifies the difficulty of proving that an NP-language L is not in P.[5]

5: If one finds a single example where this is the case, this implies $\mathsf{P} \neq \mathsf{NP}$.

Although there is no immediate hope to solve the above conjecture, there are techniques to classify problems that are most probably not in P, so-called NP-hard problems. If $\mathsf{P} \neq \mathsf{NP}$, then the NP-hard problems are automatically not in P then these are the hardest problems among all NP-problems.

We say that Problem B is harder than Problem A, if every algorithm for Problem B automatically gives rise to an algorithm for Problem A. In other words, we can embed the instances of the easier Problem A into instances of Problem B. This is formalized by the notion of a *reduction*.

Definition 6.1.4

Let $L_1, L_2 \subseteq \Sigma^*$ be two languages. A reduction $\mathcal{R} : L_1 \to L_2$ is a computable function

$$\mathcal{R} : \Sigma^* \to \Sigma^*$$

that satisfies

$$x \in L_1 \iff \mathcal{R}(x) \in L_2.$$

If $\mathcal{R}$ is in addition computable in polynomial-time, then $\mathcal{R}$ is a *polynomial-time reduction*.

Note that reductions are transitive, i.e. if there is a reduction $\mathcal{R} : L_1 \to L_2$ and a reduction $\mathcal{Q} : L_2 \to L_3$, then $\mathcal{Q} \circ \mathcal{R} : L_1 \to L_3$ defines a reduction from L_1 to L_3. For this reason, if there is a reduction $L_1 \to L_2$, we will denote this by $L_1 \leqslant L_2$. If the reduction $L_1 \to L_2$ is in addition poly-time, we denote this by $L_1 \leqslant_{\mathsf{poly}} L_2$.

This gives rise to the notion of NP-hard and RE-hard problems.

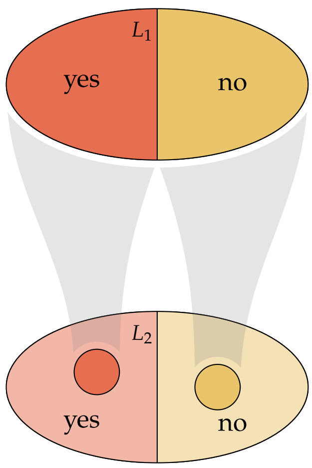

Figure 6.5: Illustration of a reduction $\mathcal{R} : L_1 \to L_2$. The yes-instances in of the first problem (i.e. elements of L_1) are mapped to the yes-instance of the second problem (i.e. elements of L_2), and no-instance of the first problem are mapped to no-instance of the second problem. Therefore, the language L_1 can be decided with an algorithm for L_2 via the reduction.

Definition 6.1.5

We call a problem L

- NP-hard (coNP-hard), if $L' \leqslant_{\mathsf{poly}} L$ for every language $L' \in$ NP ($L' \in$ coNP).
- RE-hard (coRE-hard), if $L' \leqslant L$ for every language $L' \in$ RE ($L' \in$ coRE).

Note that NP-hardness and coNP-hardness require polynomial-time reductions. This is essential because only polynomial-time computations are negligible when considering problems in these complexity classes. If a problem L is NP-hard and in addition in NP, the problem is NP-complete. We use a similar convention for all other complexity classes.

Many graph problems are NP-complete, for example the MaxCut problem or the 3-coloring problem of graphs (see Figure 6.3). Also, the non-deterministic halting problem BNHalt is NP-complete; we refer to Section 8.2 for a proof of this statement. The halting problem Halt is an example of an RE-complete problem:

Proposition 6.1.2

The Halting problem is RE-complete.

Proof. We start by showing that Halt $\in$ RE-hard. Let T be a Turing machine that decides the RE-language L. We construct a Turing machine T' from T such that

- If T accepts x, then T' halts.

- If T rejects x, then T' loops.

This construction can be performed by adding a finite number of additional states in T', and shows that

$$x \in L \iff \langle T, x \rangle \in \textsc{Halt}.$$

$\textsc{Halt} \in \mathsf{RE}$ is clear by definition. □

Note that not every undecidable (i.e. $L \notin \mathsf{R}$) language is RE-hard or coRE-hard; however, Proposition 6.1.2 implies that every RE-hard (and every coRE-hard) problem is also undecidable.

6.2 Computational aspects in semi-algebraic geometry

Many decision problems in physics and mathematics reduce to verifying whether a specific set of polynomial equations or inequalities is true. For instance, determining whether a matrix $A \in \mathrm{Mat}_s(\mathbb{Q})$ is psd involves checking the infinitely many polynomial inequalities of the form

$$\langle v | A | v \rangle \geqslant 0$$

for every vector $|v\rangle \in \mathbb{R}^s$. However, taking this definition literally as an algorithm is impossible, as it would entail verifying uncountably many inequalities — a task that cannot be accomplished in finite time.

A similar problem appears when classifying separable matrices. A matrix $A \in \mathrm{Mat}_s(\mathbb{C}) \otimes \mathrm{Mat}_s(\mathbb{C})$ is separable if there exists $r \in \mathbb{N}$ and psd matrices $A_\alpha^{[i]} \in \mathrm{Psd}_s(\mathbb{C})$ such that

$$A = \sum_{\alpha=1}^{r} A_\alpha^{[1]} \otimes A_\alpha^{[2]}.$$

Once again, deciding whether A is separable using this definition is unfeasible as an algorithm due to the infinite range of quantified variables involved.

In this chapter, we present two key results from (semi-)algebraic geometry that can be leveraged to construct algorithms solving such problems.

- The *Tarski–Seidenberg theorem* offers a method for handling statements involving polynomials and quantifiers over real numbers. This theorem enables to develop algorithms for addressing a wide range of problems in quantum information and beyond.
- *Hilbert's basis theorem* shows that every set described by infinitely many polynomial equations can be recovered by a finite subset of these polynomials.

6.2.1 The Tarski–Seidenberg theorem

In the following, we present the *Tarski–Seidenberg theorem*, which provides insights into the structure of sets X of the following form:

$$x \in X \subseteq \mathbb{R}^n \quad \Longleftrightarrow \quad \exists y \in \mathbb{R}^m \colon p(x, y) \geqslant 0$$

where $p\colon \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}$ is a polynomial with integer coefficients. Essentially, these sets are projections of sets arising via a polynomial inequality. Naively, checking membership (i.e. whether $x \in X$ or $x \notin X$) is not doable in finite time, since one has to explore the space of parameters $y \in \mathbb{R}^m$ which can take uncountably many values.

However, the Tarski–Seidenberg theorem provides an algorithm to decide membership in X. Intuitively, the Tarski–Seidenberg theorem asserts that sets like X are *semi-algebraic*, meaning they can be represented by a finite number of polynomial inequalities without any quantifiers involved. Furthermore, these polynomials can be derived in a computable way from the original description of X. Verifying membership in X via these finitely many polynomial inequalities can be achieved in finite time. We first introduce the notion of a semi-algebraic set and then present the statement of the Tarski–Seidenberg theorem with its implications.

Definition 6.2.1
A set $S \subseteq \mathbb{R}^n$ is called *semi-algebraic* if there exist polynomials

$$p_1, \ldots, p_k, q_{ij} \colon \mathbb{R}^n \to \mathbb{R}$$

for every $i, j \in \{1, \ldots k\}$ such that

$$S = \bigcup_{i=1}^{k} \{a \in \mathbb{R}^n \colon p_i(a) = 0,\ q_{i1}(a) > 0,\ \ldots,\ q_{ik}(a) > 0\}$$

An example of a semi-algebraic set is for instance

$$\left\{(x_1, x_2) \in \mathbb{R}^2 \colon 1 < x_1^2 + x_2^2 \leqslant 4, x_2 \geqslant x_1^2 \text{ or } x_1^2 + x_2^2 \leqslant \frac{1}{2}\right\}$$

which is illustrated in Figure 6.6.

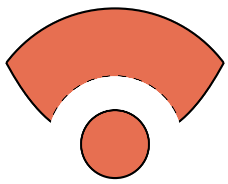

Figure 6.6: Example of a semi-algebraic set. Note that semi-algebraic sets do not have to be open or closed. They also do not have to be connected. Only the boundary of these sets has to be described via polynomials.

A semi-algebraic set can be expressed via polynomial inequalities of the form:

$$p(a) \geqslant 0,\ p(a) > 0,\ p(a) = 0,\ p(a) < 0,\ p(a) \leqslant 0$$

as well as Boolean combinations thereof. Each of these conditions can be transformed into the standard form of Definition 6.2.1.

Verifying membership in a semi-algebraic set is a straightforward task. Let

$$S := \{a \in \mathbb{R}^n : q_1(a) > 0, \ldots, q_k(a) > 0, p(a) = 0\}.$$

To check whether a belongs to S, it suffices to check

$$p_1(a) > 0, \ldots, p_k(a) > 0, q_1(a) = 0, \ldots, q_s(a) = 0,$$

which can be done in finite time.

We now consider the more complex membership problem involving semi-algebraic sets: Checking membership in projections of semi-algebraic sets.

Problem 6.2.1 (Membership in projections of semi-algebraic sets)
Let $S \subseteq \mathbb{R}^n \times \mathbb{R}^m$ be a semi-algebraic set (given by polynomials $p_1, \ldots, p_k, q_{i1}, \ldots, q_{ik}$). Moreover, let $x := (x_1, \ldots, x_n) \in \mathbb{Q}^n$. Decide whether

$$x \in \pi_1(S) := \{a \in \mathbb{R}^n : \quad \exists y \in \mathbb{R}^m : (a, y) \in S\}.$$

Here π_1 denotes the projection map on the first component of the space $\mathbb{R}^n \times \mathbb{R}^m$. This problem is naively harder to decide since it involves a quantifier. While sets like $\pi_1(S)$ may appear more general than semi-algebraic sets, the Tarski–Seidenberg theorem reveals that $\pi_1(S)$ is also semi-algebraic. Furthermore, this theorem yields an algorithm to decide Problem 6.2.1.

Theorem 6.2.2 (Tarski–Seidenberg theorem)
Let $S \subseteq \mathbb{R}^n \times \mathbb{R}^m$ be semi-algebraic and

$$\pi_1 : \mathbb{R}^n \times \mathbb{R}^m \to \mathbb{R}^n : (x, y) \mapsto x$$

be the projection map. Then, the set $\pi_1(S) \subseteq \mathbb{R}^n$ is again semi-algebraic. Moreover, the defining polynomials of $\pi_1(S)$ can be constructed explicitly.

We refer to [7, Section 2] or to [19] for a proof of this statement. Theorem 6.2.2 immediately gives rise to the following corollary.

Corollary 6.2.3
Problem 6.2.1 is decidable.

It is worth noting that the polynomials describing $\pi_1(S)$ are generally more complex than those describing S. Consequently, the computational complexity of deciding membership in $\pi_1(S)$ surpasses that of NP.

Further note that a similar decision procedure applies to membership problems involving the all-quantifier instead of the existential quantifier. This follows from the fact that complements of semi-algebraic sets are again semi-algebraic. Specifically, Theorem 6.2.2 implies that the statement, given $x \in \mathbb{Q}^n$, decide whether

$$\forall y \in \mathbb{Q}^m : (x, y) \in S$$

is decidable.

We now present some examples where Theorem 6.2.2 can be applied.

Example 6.2.1 (The set of psd matrices is semi-algebraic)
A matrix $A \in \mathrm{Mat}_s(\mathbb{R})$ is psd, if it satisfies

$$\forall \left|v\right\rangle \in \mathbb{R}^s : p(v, A) := \left\langle v\right| A \left|v\right\rangle \geqslant 0,$$

where p is a polynomial expression in the entries of $|v\rangle$ and A.
By Theorem 6.2.2, it follows that the set of psd matrices $\mathrm{Psd}_s(\mathbb{R})$ is semi-algebraic.
There are also explicit polynomial descriptions of $\mathrm{Psd}_s(\mathbb{R})$ known. For example, the set of psd matrices can be characterized as follows [65, Theorem 7.2.5]:

$$A \in \mathrm{Psd}_s(\mathbb{R}) \Leftrightarrow \det(M) \geqslant 0 \text{ for the } 2^s - 1 \text{ principal minors } M \text{ of } A.$$

A principal minor M of a matrix $A \in \mathrm{Mat}_s(\mathbb{R})$ is, given a sequence

$$1 \leqslant i_1 < i_2 < \ldots < i_k \leqslant s,$$

the matrix

$$M = \left(A_{i,j}\right)_{i,j=i_1,\ldots,i_k}.$$

There are $2^s - 1$ principal minors of an $s \times s$ matrix.

Example 6.2.2 (The set of separable states is decidable)
A matrix $A \in \mathrm{Mat}_s(\mathbb{C}) \otimes \mathrm{Mat}_s(\mathbb{C})$ is separable if there exists $r \in \mathbb{N}$ and psd matrices $A_\alpha^{[i]} \in \mathrm{Psd}_s(\mathbb{C})$ such that

$$A = \sum_{\alpha=1}^{r} A_\alpha^{[1]} \otimes A_\alpha^{[2]}. \tag{6.3}$$

Note that r can be upper bounded by s^4 by the Carathéodory theorem since the dimension of the matrix space is s^4 and the set of separable matrices is a convex cone generated by elementary tensors consisting of psd matrices.
But this shows that $\mathrm{Sep}_{s,s}(\mathbb{C}) \subseteq \mathrm{Mat}_s(\mathbb{C})^{\otimes 2}$ is semi-algebraic by Theorem 6.2.2 since it is a projection of a set generated by polynomial equations (namely Equation (6.3)).[6]
Note that no simple explicit polynomial description for separable states is known. Consequently, applying the construction of the Tarski–Seidenberg theorem to the set of separable states becomes necessary. However, this approach is not efficient in practice.
Alternatively, one can employ hierarchies of semidefinite programs [44] to determine membership in $\mathrm{Sep}_{s,s}(\mathbb{C})$.

6: The definition of semi-algebraic sets over $\mathbb{C}$ can be understood by taking real and imaginary part separately, using $\mathbb{C} \cong \mathbb{R}^2$.

Example 6.2.3 (Membership in the set of nonnegative polynomials is decidable)
We now consider the set $\mathcal{N} \subseteq \mathbb{R}[x_1, \ldots, x_n]_d$ defined as

$$\mathcal{N} := \left\{ p \in \mathbb{R}[x_1, \ldots, x_n]_d \colon \forall a \in \mathbb{R}^n : p(a) \geqslant 0 \right\}.$$

where $\mathbb{R}[x_1, \ldots, x_n]_d$ is the space of polynomials with degree at most d.
In the following, we associate $\mathbb{R}[x_1, \ldots, x_n]_d$ with the coordinate space $\mathbb{R}^k$, i.e. every entry corresponds to a coefficient of a different monomial.
In this sense, $\mathcal{N}$ is a semi-algebraic, as it can be written as a quantified formula with a polynomial inequality.
This implies that deciding whether p is nonnegative on $\mathbb{R}^n$ is decidable.

Deciding statements in first-order logic

Theorem 6.2.2 allows us to even decide more general statements, more precisely, *statements in first-order logic*. A statement φ in first-order logic is defined by a set of polynomials $p_1, \ldots, p_n : \mathbb{R}^n \to \mathbb{R}$ in n variables together with quantifiers on all variables, i.e.

$$\varphi := Q_1 x_1 \colon Q_2 x_2 \colon \cdots Q_n x_n \colon B(x_1, \dots, x_n)$$

where $Q_i \in \{\exists, \forall\}$ is a quantifier and where $B(x_1, \dots, x_n)$ is a Boolean expression of polynomial inequalities, consisting of conjunctions and disjunctions of statements:

$$p(a) \geqslant 0, p(a) > 0, p(a) = 0, p(a) < 0, p(a) \leqslant 0.$$

Corollary 6.2.4 (Tarski–Seidenberg quantifier elimination)
Statements in *first-order logic* are decidable. Moreover, for every formula in first-order logic, there exists a quantifier-free formula ψ in first-order logic such that $\varphi = \psi$.

Proof. First note that

$$\{(x_1, \dots, x_n) \in \mathbb{R}^n \colon B(x_1, \dots, x_n) \text{ is true}\}$$

is semi-algebraic. Iteratively applying Theorem 6.2.2 to all variables shows the statement. □

6.2.2 Hilbert's basis theorem

In the following, we present *Hilbert's basis theorem* and its computational consequences. We start by introducing the notion of algebraic varieties.

Definition 6.2.2
A set $V \subseteq \mathbb{R}^n$ is called an *algebraic variety* if there exists a subset[7] $P \subseteq \mathbb{R}[x_1, \dots, x_n]$ such that

$$V = \{a \in \mathbb{R}^n \colon f(a) = 0 \text{ for all } f \in P\}.$$

7: This set does not have to be finite.

Let $(f_i)_{i \in \mathbb{N}}$ be a recursively enumerable sequence[8] of polynomials in n variables generating the algebraic variety

$$V\Big((f_i)_{i \in \mathbb{N}}\Big) = \{a \in \mathbb{R}^n \colon f_i(a_1, \dots, a_n) = 0 \text{ for all } i \in \mathbb{N}\}$$

8: A sequence is called recursively enumerable if there exists a Turing machine that computes the first n sequence elements in finite time for every n.

We consider now the following decision problem:

Problem 6.2.5
Given a recursively enumerable sequence of polynomials $(f_i)_{i \in \mathbb{N}}$, decide the following statement:

$$\forall x \in V\Big((f_i)_{i \in \mathbb{N}}\Big) \colon p(x) \geqslant 0. \tag{6.4}$$

Note that there is no obvious way to verify both yes- and no-instances. We will now present *Hilbert's basis theorem* that will show that Problem 6.2.5 is in RE, i.e. the problem is semi-decidable.

Theorem 6.2.6 (Hilbert's basis theorem)
Let V be an algebraic variety generated by $S \subseteq \mathbb{R}[x_1, \dots x_n]$.
Then there exist finitely many polynomials $f_1, \dots, f_k \in S$ such that

$$V = \{x \in \mathbb{R}^n : f_1(x) = \dots = f_k(x) = 0\}.$$

For a proof, we refer to [105].[9] Note that in contrast to the Tarski–Seidenberg Theorem (Theorem 6.2.2) there is no constructive way to obtain the polynomials $f_1, \dots, f_k$ or the upper bound k. However, if $(f_i)_{i \in \mathbb{N}}$ is a sequence of polynomials generating V, then Theorem 6.2.6 shows that there exists $N \in \mathbb{N}$ such that

$$V = \{x \in \mathbb{R}^n : f_1(x) = f_2(x) = \dots = f_N(x) = 0\}.$$

9: Hilbert's basis theorem is usually stated as follows: every ideal $I \subseteq \mathbb{R}[x_1, \dots, x_n]$ is finitely generated. The ideal in our formulation is the set of polynomials that generate the algebraic variety and the finitely many generators of V correspond to the generators of the ideal.

This allows us to verify yes-instances of Equation (6.4) in finite time via the following algorithm:

(i) Consider the algebraic variety

$$V_N = \{x \in \mathbb{R}^n : f_1(x) = \dots = f_N(x) = 0\}.$$

(ii) Decide the following statement in first-order logic:

$$\forall x \in V_N : p(x) \geqslant 0. \tag{6.5}$$

If Equation (6.5) is true, halt and accept the input.
If Equation (6.5) if false, increment N to $N + 1$.

Due to Theorem 6.2.6, for every yes-instance, there exists $N \in \mathbb{N}$ such that Equation (6.5) holds true. Therefore, the algorithm eventually halts for yes-instances.

In simpler terms, this algorithm leverages the decidability of statements in first-order logic to assess the truth of the statement:

$$\forall x \in V(f_1, \dots, f_N) : p(x) \geqslant 0$$

for fixed N. If the statement holds true, then we can infer that

$$\forall x \in V\Big((f_i)_{i \in \mathbb{N}}\Big) : p(x) \geqslant 0$$

since $V\Big((f_i)_{i \in \mathbb{N}}\Big) \subseteq V(f_1, \dots, f_N)$. Conversely, if the statement is false, then we increment N by one and repeat the procedure. According to Theorem 6.2.6, there exists a value k such that

$$V(f_1, \dots, f_k) = V\Big((f_i)_{i \in \mathbb{N}}\Big).$$

Therefore, if the sequence is a yes-instance, then the algorithm also halts at $N = k$.

Note that this procedure cannot be used to verify no-instances because it is unclear what the number k is. It might be the case that determining k is undecidable.

Positivity of matrix moments 7

This chapter is based on [42] and Section 6 of [39].

Some problems may look very innocent yet be formally very difficult—perhaps uncomputable—or even worse, their computability may be unknown. *Skolem's problem* exemplifies this uncertainty, focusing on the behavior of *linear recurrence sequence (LRS)*, where each term in the sequence is generated linearly from its predecessors. Examples of LRS include well-known sequences like the Fibonacci sequence or those derived from discretizing differential equations. Despite their simplicity, LRS are fundamental in various mathematical and computer science domains, notably in generating pseudo-random numbers [120], describing the dynamics of cellular automata [82], and many other applications [47].

More specifically, an LRS of order s is given by

$$u_n = a_1 u_{n-1} + a_2 u_{n-2} + \cdots + a_s u_{n-s}$$

where $a_1, \ldots, a_s \in \mathcal{R}$ are fixed elements in a ring $\mathcal{R}$, usually commutative. Together with initial values $u_1, \ldots, u_s \in \mathcal{R}$, this gives rise to a full sequence $(u_n)_{n\in\mathbb{N}}$ in $\mathcal{R}$. While several important examples of LRS are over the ring $\mathcal{R} = \mathbb{Z}$, many interesting examples are also defined over other rings. For example, the Chebyshev polynomials are defined via the LRS

$$T_n(x) = 2x \cdot T_{n-1}(x) - T_{n-2}(x) \quad \text{with } T_1(x) := x \text{ and } T_0(x) := 1$$

over the commutative ring $\mathbb{Z}[x]$ of univariate polynomials.

Skolem's problem is a long-standing open question concerning LRS over $\mathbb{Z}$ [93]. It asks whether an algorithm exists that decides if an LRS attains the value 0 for some $n \in \mathbb{N}$. While partial solutions to Skolem's problem are known, implying decidability for order $s \leqslant 4$ [122, 125], they do not apply to recurrences of order five or more. A modification of Skolem's problem is the *positivity problem* for LRS. Instead of asking whether the LRS is non-zero, it asks whether an LRS stays non-negative. In this case it is also unclear whether an algorithm exists that decides the positivity problem, as decidability is proven only for $s \leqslant 5$ [92, 91].

Examples for LRS are *moment sequences*, in which we have

$$u_n = \mathrm{tr}(A^n),$$

or *generalized moment sequences*, in which

$$u_n = \varphi(A^n)$$

for a given matrix $A \in \mathrm{Mat}_s(\mathcal{R})$ and a linear functional φ on $\mathrm{Mat}_s(\mathcal{R})$. Over a commutative ring $\mathcal{R}$, such generalized moment sequences are as expressive as LRS, i.e. every LRS can be expressed as a moment sequence and vice versa. For this reason, decidability results for generalized moment sequences translate to decidability results for Skolem's problem and the positivity problem.

Table 7.1: For which instance sets is the (generalized) moment membership problem decidable or undecidable? This table summarizes the results of this paper.

Decidable cases	**Undecidable cases**
Unitary and Orthogonal matrices (Section 7.2.2)	Comm. polynomials $\mathbb{Z}[x_1, \dots, x_d]$ (Section 7.3.1)
Dominant or real eigenvalue matrices (Section 7.2.3)	Non-comm. polynomials $\mathbb{Z}\langle z_1, \dots, z_d \rangle$ (Section 7.3.2)

In this paper, we study the decidability of the moment membership problem. That is, we consider the problem: For an $s \times s$ matrix A, decide whether

$$\operatorname{tr}(A^n) \in \mathcal{P} \quad \forall n \in \mathbb{N}$$

where $\mathcal{P}$ is a fixed set. This set usually contains elements that are positive in some sense, so we call the problem also the moment positivity problem. Most of our results also hold for generalized moments of the form $\varphi(A^n)$ as above.

One decisive factor in the complexity of the problem is the instance set $\mathcal{D}$ of the matrices, which allows us to distinguish between our two main results:

- We restrict the instance set $\mathcal{D} \subseteq \operatorname{Mat}_s(\mathbb{Z})$ and prove decidability of the problem for a large subclass of integer matrices.
- We enlarge the instance set $\operatorname{Mat}_s(\mathbb{Z}) \subseteq \mathcal{D}$ and prove that the problem is undecidable for matrices whose entries are elements of certain unital rings $\mathcal{R}$, for certain $\mathcal{P} \subseteq \mathcal{R}$.

Contributions: Specifically, we determine the complexity of the moment membership problem in the following cases (see Table 7.1):

- **Decidability:** The moment positivity problem is decidable for orthogonal matrices (Theorem 7.2.3), unitary matrices (Corollary 7.2.5), and matrices with a unique dominant eigenvalue or only real eigenvalues (Theorem 7.2.7). It follows that the positivity problem is decidable for simple unitary LRS, i.e. LRS whose characteristic polynomial only has simple roots of modulus 1, as well as for LRS whose characteristic polynomial has a unique dominant root, or only real roots.
- **Undecidability:** The generalized problem is undecidable for the ring of multivariate commutative polynomials (Theorem 7.3.2) as well as for non-commutative polynomials, where $\mathcal{P}$ is the set of polynomials with nonnegative coefficients (Theorem 7.3.6). This implies that the corresponding positivity problem for LRS over commutative polynomials is undecidable.
- **Free Pólya's Theorem:** As a side result, we prove a free version of Pólya's theorem (Theorem 7.3.5). We show that a non-commutative polynomial has nonnegative coefficients if and only if it is entrywise nonnegative on the set of entrywise nonnegative matrices.

This paper is structured as follows. In Section 7.1 we introduce the problem statement and show the relation of moment problems to LRS. In Section 7.2 we present cases in which the moment problem is decidable. This includes a review of known results (Section 7.2.1), the decidability for orthogonal and unitary matrices (Section 7.2.2), and the decidability for matrices with unique largest eigenvalues or only real roots (Section 7.2.3).

In Section 7.3, we present examples of commutative and non-commutative rings where the moment problem is undecidable, as well as a non-commutative version of Pólya's Theorem. Moreover, we present a related undecidable problem on commutative polynomials, from [39].

7.1 Problem statement

Let $\mathcal{R}$ be a unital ring, and let $A \in \mathrm{Mat}_s(\mathcal{R})$ be an $s \times s$ square matrix with entries from $\mathcal{R}$. For $n \geqslant 0$ the n^{th} *moment of* A is defined as

$$\mu_n(A) := \mathrm{tr}\,(A^n)$$

where tr denotes the usual trace of a matrix, i.e. the sum of its diagonal entries. The moments of A are clearly elements from $\mathcal{R}$, as for $A = (a_{ij})_{i,j=1,\dots,s}$ we have

$$\mu_0(A) = \underbrace{1_{\mathcal{R}} + \cdots + 1_{\mathcal{R}}}_{s} \quad \text{and} \quad \mu_1(A) = \sum_{i=1}^{s} a_{ii}.$$

and for $n \geqslant 2$

$$\mu_n(A) = \sum_{i_1,\dots,i_n=1}^{s} a_{i_1 i_2} \cdot a_{i_2 i_3} \cdot \cdots \cdot a_{i_{n-1} i_n} \cdot a_{i_n i_1}.$$

Depending on the ring $\mathcal{R}$, the moments are studied in different contexts, as the following example shows.

A *unital ring* $\mathcal{R}$ is an algebraic structure that generalizes the notion of a field. Specifically, multiplication needs not to be commutative and inverses do not have to exist. It consists of the binary operations $+$ and $\cdot$, that satisfy the following:

- $(\mathcal{R}, +)$ forms an abelian group
- $(\mathcal{R}, \cdot)$ is a monoid, i.e. it is associative and it contains a multiplicative identity.

We denote the multiplicative identity by $1_{\mathcal{R}}$.

Example 7.1.1

Let $\mathcal{V}$ be a $\mathbb{C}$-vector space. Consider the tensor algebra

$$\mathcal{R} := T(\mathcal{V}) := \bigoplus_{m \geqslant 0} \mathcal{V}^{\otimes m} = \mathbb{C} \oplus \mathcal{V} \oplus (\mathcal{V} \otimes \mathcal{V}) \oplus \cdots$$

$\mathcal{R}$ forms a unital ring with tensor product as multiplication. Actually, $\mathcal{R}$ is an $\mathbb{N}$-graded unital $\mathbb{C}$-algebra.
For $A = (|a_{ij}\rangle)_{i,j} \in \mathrm{Mat}_s(\mathcal{V})$ (which embeds into $\mathrm{Mat}_s(\mathcal{R})$), we obtain for the moments

$$\mu_n(A) = \sum_{\alpha_1,\dots,\alpha_n=1}^{s} |a_{\alpha_1,\alpha_2}\rangle \otimes |a_{\alpha_2,\alpha_3}\rangle \otimes \cdots \otimes |a_{\alpha_{n-1},\alpha_n}\rangle \otimes |a_{\alpha_n,\alpha_1}\rangle\,, \quad (7.1)$$

where the n^{th} moment is homogeneous of degree n in $\mathcal{R}$. The expression in Equation (7.1) corresponds to the *translational invariant matrix product state* introduced in Section 2.3.1, where s corresponds to the (Θ_n, C_n)-rank of $\mu_n(A)$.

Now assume that $\mathcal{R}$ is also equipped with a subset $\mathcal{P} \subseteq \mathcal{R}$. In our results and applications, this will always be a set of elements that are *positive* in some sense. Further $\mathcal{D} \subseteq \mathrm{Mat}_s(\mathcal{R})$ will be the set containing all instances of our decision problem. The general decision problem that we will address in this paper is the following:

Problem 7.1.1 (Moment Positivity Problem)
Let $s, \mathcal{P}, \mathcal{D}$ be fixed as above. For $A \in \mathcal{D}$ decide whether *all* moments $\mu_n(A)$ belong to $\mathcal{P}$.

Note that $\mathcal{D}, \mathcal{P}, s$ are fixed in our formulation of the decision problem. We are thus looking for an algorithm (tailored to $\mathcal{R}, \mathcal{P}, s$ and $\mathcal{D}$) that upon an input of any instance $A \in \mathcal{D}$ stops after a finite time, and returns yes if all moments of A belong to $\mathcal{P}$, and no if at least one moment of A does not belong to $\mathcal{P}$. If such an algorithm exist, we call the moments membership problem *decidable*, otherwise we call it *undecidable*.

Note that if the ring operations are computable and membership of single elements in $\mathcal{P}$ is decidable, the moments membership problem is clearly semi-decidable in the following sense. Given $A \in \mathrm{Mat}_s(\mathcal{R})$, we simply compute higher and higher moments of A, and check membership in $\mathcal{P}$. If some moment does *not* belong to $\mathcal{P}$, we will know after a finite time. However, this algorithm runs forever in case that all moments do belong to $\mathcal{P}$. So the hard part of the problem is certifying membership of all moments in $\mathcal{P}$. We will make use of the semi-decidability in Theorem 7.2.3.

7.1.1 Relation to the membership problem for linear recurrence sequences

In the following, we review the relation of the moment problem with the positivity problem for linear recurrence sequences. An *LRS* $(u_n)_{n\in\mathbb{N}} \in \mathcal{R}^{\mathbb{N}}$ is a sequence whose elements are related to each other linearly, i.e.

$$u_n = a_1 \cdot u_{n-1} + a_2 \cdot u_{n-2} + \cdots + a_s \cdot u_{n-s} \tag{7.2}$$

for all $n > s$. We call s the *order* of the recurrence relation. The positivity problem for LRS is the following:

Problem 7.1.2 (Positivitiy for LRS)
Given an LRS as in Equation (7.2) with parameters $a_1, \dots, a_s \in \mathcal{R}$ and initial values $u_1, \dots, u_s \in \mathcal{R}$, decide whether $u_n \in \mathcal{P}$ for all $n \in \mathbb{N}$.

We start with the (well-known) observation that every generalized moment sequence is an LRS, if $\mathcal{R}$ is commutative.

Lemma 7.1.3 (Moment sequences are LRS)
Let $\mathcal{R}$ be a commutative unital ring, and let $A \in \mathrm{Mat}_s(\mathcal{R})$. Then $(\varphi(A^n))_{n\in\mathbb{N}}$ is an LRS of order s, for every $\mathcal{R}$-linear map $\varphi\colon \mathrm{Mat}_s(\mathcal{R}) \to \mathcal{R}$.

Proof. Let $p(x) = x^s - a_1x^{s-1} - \cdots - a_s$ be the characteristic polynomial of the matrix A. By the Cayley–Hamilton theorem for commutative rings (see for example [79, Chapter XIV.3]), we have that

$$A^s = a_1A^{s-1} + a_2A^{s-2} + \cdots + a_sI \tag{7.3}$$

and therefore

$$A^n = a_1 A^{n-1} + a_2 A^{n-2} + \cdots + a_s A^{n-s}$$

for all $n \geqslant s$. Applying φ proves the statement. □

It is unclear whether a similar statement to Lemma 7.1.3 is true for non-commutative rings. While there exist versions of the Cayley–Hamilton theorem for non-commutative rings (see for example [61, 119]), they cannot be applied to obtain an equation similar to Equation (7.3).

The next observation states that LRS are equivalent to generalized moment sequences as introduced above. It can also be found in [92]:

Lemma 7.1.4 (LRS are moment sequences)
Let $(u_n)_{n\in\mathbb{N}}$ be a sequence in a commutative unital ring $\mathcal{R}$. Then the following are equivalent:

(i) $(u_n)_{n\in\mathbb{N}}$ is an LRS of order s.
(ii) There is a matrix $A \in \mathrm{Mat}_s(\mathcal{R})$ and two vectors $|v\rangle, |w\rangle \in \mathcal{R}^s$ such that $u_n = \langle v| A^{n-s} |w\rangle$ for all $n > s$.

Proof. For (i) ⇒ (ii) assume that the recurrence is given by

$$u_n = a_1 u_{n-1} + a_2 u_{n-2} + \cdots + a_s u_{n-s}.$$

Using the companion matrix

$$A = \begin{pmatrix} a_1 & 1 & & & \\ a_2 & & 1 & & \\ \vdots & & & \ddots & \\ a_{s-1} & & & & 1 \\ a_s & & & & \end{pmatrix}$$

we have that $u_n = \langle v| A^{n-s} |w\rangle$ where $|v\rangle = (u_s, u_{s-1}, \ldots, u_1)^t$ and $|w\rangle = (1, 0, \ldots, 0)^t$.

The proof of (ii) ⇒ (i) is analogous to Lemma 7.1.3, by replacing tr by the function $A \mapsto \langle v| A |w\rangle$. Note that the recurrence starts to hold only for $n > 2s$, but for our purposes this is irrelevant. □

7.2 Decidable cases

In the following we present cases in which the moment membership problem is decidable. This includes known results for small s (Section 7.2.1), the moment positivity problem for unitary and orthogonal matrices (Section 7.2.2), and for matrices with a unique largest eigenvalue or only real eigenvalues (Section 7.2.3). Throughout this section we will always choose $\mathcal{P} = \mathbb{R}_{\geqslant 0}$.

7.2.1 Known results: small order

We first review known results on the decidability of the moment positivity problem. The known results are all about LRS, in view of Lemma 7.1.3 they immediately transfer to moments.

> **Theorem 7.2.1**
> The moment positivity problem is decidable in the following cases:
>
> (i) $s \leqslant 5, \mathcal{D} = \mathrm{Mat}_s(\mathbb{Q})$.
> (ii) $s \leqslant 9, \mathcal{D} \subseteq \mathrm{Mat}_s(\mathbb{Q})$ is the set of matrices with simple eigenvalues.

The proof of (i) is contained in [92], the proof of (ii) goes back to [91]. Decidability for other values of s is unknown.

The positivity problem of LRS is closely related to *Skolem's Problem* which asks if some sequence element equals 0. The best result in this direction is that Skolem's Problem is NP-hard [16]. The decidability of the positivity problem implies decidability of Skolem's Problem. This follows for an integer LRS because $u_n \neq 0$ if and only if $u_n^2 - 1 \geqslant 0$. If $(u_n)_{n\in\mathbb{N}}$ is an LRS of order s, then $u_n^2 - 1$ is an LRS of order s^2. Moreover, since Skolem's Problem is NP-hard, the positivity problem is NP-hard as well.

7.2.2 Orthogonal and unitary matrices

We now show that the moment positivity problem for orthogonal (Theorem 7.2.3) and unitary matrices (Corollary 7.2.5) is decidable. The proof strategy is very similar to [15].

We say that a set $X \subseteq \mathbb{R}^m$ is algebraic if there are polynomials

$$p_1, \ldots, p_n \colon \mathbb{R}^m \to \mathbb{R}$$

such that

$$X = \{x \in \mathbb{R}^m \colon p_1(x) = \cdots = p_n(x) = 0\}.$$

In this case, we call X the algebraic variety defined by $p_1, \ldots, p_n$, and write $X = \mathcal{V}(p_1, \ldots, p_n)$. Even if the set of defining polynomials is infinite, there always exists a finite choice of polynomials defining the same algebraic variety, by Hilbert's basis theorem. Since we work over $\mathbb{R}$, we can even reduce to a single polynomial, by taking the sum of squares of the defining polynomials.

For matrices $A_1, \ldots, A_d \in \mathrm{Mat}_s(\mathbb{R})$, let

$$\langle A_1, A_2, \ldots, A_d \rangle := \{A_{k_1} \cdots A_{k_\ell} \colon \ell \in \mathbb{N},\ k_1, \ldots, k_\ell = 1, \ldots, d\}$$

be the semigroup generated by $A_1, \ldots, A_d$. We denote by $\overline{\langle A_1, \ldots, A_d \rangle}$ the topological closure inside $\mathrm{Mat}_s(\mathbb{R})$ with respect to the Euclidean topology.

> **Lemma 7.2.2**
> Let $A_1, \ldots, A_d \in \mathrm{O}_s(\mathbb{Q})$ be orthogonal $s \times s$ matrices with ratio-

nal entries. Then $\mathcal{G} := \overline{\langle A_1, \ldots, A_d \rangle}$ is a compact algebraic group. Moreover there is a recursively enumerable sequence of rational polynomials $(p_k)_{k\in\mathbb{N}}$ defining $\mathcal{G}$ inside $\mathrm{Mat}_s(\mathbb{R})$.

Proof. Compactness of $\mathcal{G}$ is obvious. To prove that $\mathcal{G}$ is a group we only have to show that $A^{-1} \in \mathcal{G}$ for every $A \in \mathcal{G}$. Consider the sequence $(A^k)_{k\in\mathbb{N}}$. By compactness, there exists a converging subsequence. In other words, for every $\varepsilon > 0$, there exists $n_2 > n_1 + 1$ such that

$$\|A^{n_1} - A^{n_2}\| < \varepsilon$$

where $\|\cdot\|$ is the operator norm. Since $\|A \cdot B\| = \|B\|$ for every matrix B, we obtain

$$\|A^{-1} - A^{n_2 - n_1 - 1}\| < \varepsilon.$$

This shows that $A^{-1} \in \mathcal{G}$.

Now note that every compact group $\mathcal{G} \subseteq \mathrm{Mat}_s(\mathbb{R})$ is algebraic (see for example [88, Chapter 3, Section 4.4]). In particular, it is shown there that

$$\begin{aligned} \mathcal{G} =& \mathcal{V}\left(\mathbb{R}[X]^{\mathcal{G}}\right) \\ :=& \mathcal{V}\Big(p \in \mathbb{R}[X] \colon p(I_s) = 0,\ p(gX) = p(X) \text{ for all } g \in \mathcal{G}\Big), \end{aligned}$$

where I_s is the identity matrix of size s.

Now note that if $\mathcal{G}$ is generated by $A_1, \ldots, A_d$, then the invariance only needs to be checked w.r.t. the generators, i.e.

$$\mathcal{G} = \mathcal{V}\Big(p \in \mathbb{R}[X] \colon p(I_s) = 0,\ p(A_i X) = p(X) \ \text{ for } i = 1, \ldots, d\Big).$$

Since the conditions $p(I_s) = 0$ and $p(A_i X) = p(X)$ are linear in the coefficients of p, there exists a basis $(p_k)_{k\in\mathbb{N}}$ of the space of solutions of these conditions. Moreover, the coefficients of the basis vectors p_k can be chosen from $\mathbb{Q}$, since all conditions are rational. We now clearly have

$$\mathcal{G} = \mathcal{V}(p_k \colon k \in \mathbb{N}).$$

The polynomials p_k can be computed recursively by solving the system of linear equations over the space of polynomials with degree d, and by increasing d iteratively. □

Note that the statement is not true anymore when replacing $\mathbb{R}$ by $\mathbb{C}$. For example the group

$$\mathcal{G} := \left\{e^{i\theta} \colon \theta \in [0, 2\pi)\right\},$$

seen as a subset of 1×1 matrices, is not algebraic. However, we show that the moment problem also generalizes to unitary matrices (see Corollary 7.2.5).

Since $\mathbb{R}[X]$ is a Noetherian ring, there exists $n \in \mathbb{N}$ such that

$$\mathcal{G} = \mathcal{V}(p_1, \ldots, p_n).$$

This will be an important ingredient to show the decidability of the moment problem. Note however, that n can be arbitrarily large and it is unclear whether n is computable or not.

Theorem 7.2.3

The moment positivity problem for $\mathcal{D} = \mathrm{O}_s(\mathbb{Q})$ is decidable.

Proof. We will present two procedures, each certifying either yes- or no-instances in finite time. Letting these algorithms run in parallel will result in a decision algorithm for the problem.

Certifying no-instances for $A \in \mathrm{O}_s(\mathbb{Q})$ is achieved by iteratively checking whether $\mathrm{tr}(A^n) \geqslant 0$ holds, for every n. If A is a no-instance, then this algorithm will halt after detecting $\mathrm{tr}(A^n) < 0$ for the first time.

We now present an algorithm to certify yes-instances in finite time. For a given $A \in \mathrm{O}_s(\mathbb{Q})$, the moment membership problem can be rephrased as

$$\forall B \in \langle A \rangle : \mathrm{tr}(B) \geqslant 0.$$

By the continuity of the trace, this is equivalent to

$$\forall B \in \overline{\langle A \rangle} : \mathrm{tr}(B) \geqslant 0. \tag{7.4}$$

By Lemma 7.2.2 there exists a recursively enumerable sequence of polynomials $(p_k)_{k\in\mathbb{N}}$ and some $n \in \mathbb{N}$ such that

$$\overline{\langle A \rangle} = \mathcal{V}(p_1, \ldots, p_n).$$

Now step k of the algorithm verifies the statement

$$\forall B \in \mathcal{V}(p_1, \ldots, p_k)\colon\ \mathrm{tr}(B) \geqslant 0 \tag{7.5}$$

which is decidable by the Tarski–Seidenberg Theorem, since it is a statement in first order logic. As soon as Equation 7.5 is true for the first time, the algorithm halts and outputs a correct yes-answer. This will indeed be the case after at most n steps, if A is a yes-instance. □

Remark 7.2.1

This statement can be generalized in two directions:

(i) By the same argument, the following problem is also decidable: Given $A_1, \ldots A_d \in \mathrm{O}_s(\mathbb{Q})$ for a fixed matrix size s, decide if:

$$\forall \ell \in \mathbb{N}\ \forall k_1, \ldots, k_\ell \in \{1, \ldots, d\}\colon\ \mathrm{tr}(A_{k_1} \cdots A_{k_\ell}) \geqslant 0.$$

Note that generalizing this decision problem to arbitrary matrices makes it undecidable [36].

(ii) The proof remains true if tr is replaced by any other continuous function. This in particular implies that the generalized problem

$$\forall n \in \mathbb{N}\colon\ \varphi(A^n) \geqslant 0$$

is decidable.

We now generalize the result to unitary matrices, by embedding them into orthogonal matrices of larger size. We denote by $\mathbb{Q}[i]$ the field of complex numbers with rational real and imaginary parts, and we denote the set of $s \times s$ unitary matrices with entries in $\mathbb{Q}[i]$ by $\mathrm{U}_s(\mathbb{Q}[i])$.

Lemma 7.2.4
The map

$$\begin{aligned} \Psi\colon \mathrm{U}_s(\mathbb{Q}[i]) &\to \mathrm{O}_{2s}(\mathbb{Q}) \\ U = A + iB &\mapsto \begin{pmatrix} A & -B \\ B & A \end{pmatrix} \end{aligned}$$

is a group homomorphism. Moreover we have

$$\operatorname{tr}(U) = \frac{1}{2}\operatorname{tr}\left(\Psi(U) \cdot \begin{pmatrix} I_s & iI_s \\ -iI_s & I_s \end{pmatrix}\right),$$

where I_s is the identity matrix of size s.

Proof. The map is well defined since $\Psi(U)$ is orthogonal if and only if U is unitary. The rest is immediate. □

The main results of this section are summarized in the following two corollaries.

Corollary 7.2.5
For each $s \geqslant 1$, the moment positivity problem for matrices from $\mathrm{U}_s(\mathbb{Q}[i])$ is decidable.

Proof. It follows immediately from Lemma 7.2.4, Theorem 7.2.3 and Remark 7.2.1 (ii). □

Corollary 7.2.6
The positivity problem is decidable for simple unitary LRS, i.e.

$$u_n = a_1 u_{n-1} + \cdots + a_s u_{n-s}$$

with $a_1, \ldots, a_s \in \mathbb{Q}[i]$, where the roots of

$$p(x) = x^s - a_1 x^{s-1} - a_2 x^{s-2} - \ldots - a_s$$

are all simple and of modulus 1.

Proof. We choose a unitary matrix $A \in \mathrm{U}_s(\mathbb{C})$ whose eigenvalues are the roots of p, and whose entries are computable numbers. For example, one can take a diagonal matrix with the specified roots on the diagonal. We obtain the reccurence

$$A^n = a_1 A^{n-1} + \cdots + a_s A^{n-s}$$

for all $n \geqslant s$, and since the roots are all simple, p is actually the minimal polynomial of A. So $I_s, A, A^2, \ldots, A^{s-1}$ are linearly independent, and we can thus find a linear functional φ on $\mathrm{Mat}_s(\mathbb{C})$ with $\varphi(A^i) = u_i$ for $i = 0, \ldots, s-1$. Now as stated in Remark 7.2.1 and Lemma 7.2.4 above, it is decidable whether $\varphi(A^i) \geqslant 0$ holds for all i, and since this sequence fulfills the same recurrence and initial conditions as $(u_i)_{i \geqslant 1}$, the two sequences coincide. □

7.2.3 Matrices with a unique dominant eigenvalue or real eigenvalues

In the following, we show that for matrices with a unique dominant eigenvalue, and for matrices with only real eigenvalues, the moment problem is decidable. Note that the idea for the case of a unique dominating eigenvalue is already present in [92], but restricted to multiplicity 1 and matrices of size at most $s = 5$.

Theorem 7.2.7

The moment positivity problem is decidable in the following cases:

(i) $\mathcal{R} = \mathbb{Q}$, s arbitrary, and the set of instances restricted to matrices with a unique dominant eigenvalue.
(ii) $\mathcal{R} = \mathbb{Q}$, s arbitrary, and the set of instances restricted to matrices with only real eigenvalues.

Proof. We provide algorithms that decide the moments positivity problem for the stated instance sets. Note that we can assume without loss of generality that $A \in \mathrm{Mat}_s(\mathbb{Z})$, by possibly multiplying the matrix with the largest denominator of its entries.

For (i) let $A \in \mathrm{Mat}_s(\mathbb{Z})$ have a unique dominant eigenvalue. Since A has real entries, the non-real eigenvalues of A come in conjugate pairs. Since there is exactly one eigenvalue λ_1 of largest absolute value, it must therefore be real. We let k denote its multiplicity and obtain

$$|\mu_n(A) - k \cdot \lambda_1^n| \leqslant (s-k)|\lambda_2|^n,$$

where λ_2 denotes the second largest eigenvalue in absolute value. Thus it suffices to check $\mu_n(A) \geqslant 0$ for n up to

$$\frac{\log(s/k - 1)}{\log(|\lambda_1|) - \log(|\lambda_2|)}.$$

(ii): In this case only odd moments matter, since the even moments are always nonnegative. If the dominant eigenvalues all have the same sign, then we can apply (i). Otherwise, since odd powers of eigenvalues with the same absolute values but different signs cancel out, we can reduce the problem to a smaller matrix, where the dominant eigenvalues do have the same sign. □

7.2.4 Further generalizations

In the following, we present a generalization of the statements in Section 7.2.2 and Section 7.2.3. For a matrix $A \in \mathrm{Mat}_s(\mathbb{R})$, we denote by

$\mathsf{spec}(A)$ the multi-set of all eigenvalues of A (where multiple eigenvalues are represented by multiple elements of $\mathsf{spec}(A)$). Express

$$\mathsf{spec}(A) = \mathsf{per}_1(A) \cup \mathsf{per}_2(A) \cup \cdots \cup \mathsf{per}_s(A)$$

as a partition of peripheral spectra, i.e. eigenvalues of the same absolute value, in decreasing order (i.e. $\mathsf{per}_1(A)$ contains the dominant eigenvalues, $\mathsf{per}_2(A)$ the eigenvalues of second largest absolute value...). Note that $\mathsf{per}_i(A)$ can be empty if A has multiple eigenvalues of same absolute value. Moreover, let

$$\mu_n^{(i)}(A) := \sum_{\lambda \in \mathsf{per}_i(A)} \left(\frac{\lambda}{|\lambda|}\right)^n.$$

We define

$$\eta_i(A) = \begin{cases} \inf_{n\in\mathbb{N}} \mu_n^{(i)}(A) & : \text{if } \mathsf{per}_i(A) \neq \emptyset \\ \infty & : \text{if } \mathsf{per}_i(A) = \emptyset \end{cases}$$

and

$$\gamma_i(A) = \begin{cases} \sup_{n\in\mathbb{N}} \mu_{pn+q}^{(i)}(A) & : \text{if } \mathsf{per}_i(A) \neq \emptyset \\ -\infty & : \text{if } \mathsf{per}_i(A) = \emptyset. \end{cases},$$

where $p, q \geqslant 1$ are arbitrary but fixed integers. So we compute the supremum along an arithmetic progression.

Lemma 7.2.8 (Computability of η_i and γ_i)
The following two problems are decidable:

(i) Given[1] $A \in \mathrm{Mat}_s(\mathbb{R})$, $c \in \mathbb{R}$, decide whether $\eta_i(A) \geqslant c$.
(ii) Given $A \in \mathrm{Mat}_s(\mathbb{R})$, $c \in \mathbb{R}$, decide whether $\gamma_i(A) \leqslant c$.

1: To assume that the inputs attain a finite description, we restrict to algebraic numbers, i.e. numbers that can be represented as roots of an integer polynomial. This is enough for applying Lemma 7.2.8 in the proof of Theorem 7.2.9.

Proof. The decision algorithms are very similar to one from the proof of Theorem 7.2.3. To construct an algorithm for (i), let the following two procedures run in parallel:

(a) Evaluate $\mu_n^{(i)}(A)$ for increasing $n \in \mathbb{N}$. Halt if $\mu_n^{(i)}(A) < c$.
(b) Check the statement

$$\forall B \in \mathcal{V}(p_1, \ldots, p_k): \ \mathrm{tr}(B) \geqslant c$$

for increasing $k \in \mathbb{N}$, where $(p_\ell)_{\ell\in\mathbb{N}}$ define the variety $\overline{\langle U \rangle}$, where U is the diagonal matrix with eigenvalues $\lambda/|\lambda|$ for $\lambda \in \mathsf{per}_i(A)$. Halt if the statement is true.

If A, c is a no-instance of (i), then (ii) will eventually halt; if A, c is a yes-instance, (b) will eventually halt, for the same reason as in the proof of Theorem 7.2.3.

The algorithm for (ii) is very similar. Let the following two procedures run in parallel:

(a) Evaluate $\mu_{pn+q}^{(i)}(A)$ for increasing $n \in \mathbb{N}$. Halt, if $\mu_{pn+q}^{(i)}(A) > c$.
(b) Check the statement

$$\forall B \in \mathcal{V}(p_1, \ldots, p_k): \ \mathrm{tr}(U^q B) \leqslant c$$

for increasing $k \in \mathbb{N}$, where $(p_\ell)_{\ell\in\mathbb{N}}$ define the group $\overline{\langle U^p \rangle}$, where U is the diagonal matrix with eigenvalues $\lambda/|\lambda|$ for $\lambda \in \mathsf{per}_i(A)$. Halt, if the statement is true.

In we evaluate a generalized moment, so recall Remark 7.2.1 (ii). □

It is unclear whether $\eta_i(A) > c$ or $\eta_i(A) = c$ is decidable. This is due to the fact that we do not know whether $\mu_n^{(i)}(A)$ attains the infimum/-supremum for finite n.

Theorem 7.2.9

For a fixed parameter $\varepsilon > 0$, the moment positivity problem is decidable for all non-zero matrices A satisfying one of the following conditions:

(i) $\exists k \in \mathbb{N}\colon \eta_1(A), \dots, \eta_k(A) \geqslant 0, \eta_{k+1}(A) \geqslant \varepsilon$.
(ii) $\exists k \in \mathbb{N}\colon \gamma_1(A), \dots, \gamma_k(A) \leqslant 0, \gamma_{k+1}(A) \leqslant -\varepsilon$.
(iii) $\eta_1(A) < 0$.

If (ii) or (iii) are satisfied, then A is automatically a no-instance. If (i) is satisfied, then A can be a yes or a no-instance. Moreover, each of the above criteria is decidable.

Proof. First, checking whether A satisfies (i), (ii) or (iii) is decidable by Lemma 7.2.8, and since there are only finitely many of these statements to check.

To prove (i), assume that $\eta_{k+1}(A) \neq \infty$ (the other case is trivial). Let $\lambda_i \in \mathsf{per}_i(A)$. We have that

$$\mu_n(A) = \sum_{i=1}^{d} |\lambda_i|^n \mu_n^{(i)}(A) \geqslant |\lambda_{k+1}|^n \left(\varepsilon - s \sum_{i=k+2}^{d} \left(\frac{|\lambda_i|}{|\lambda_{k+1}|} \right)^n \right)$$

which is positive for

$$n \geqslant \frac{\log(\varepsilon) - \log(sd)}{\log(|\lambda_{k+2}|) - \log(|\lambda_{k+1}|)}.$$

So we only need to check finitely many instances of the problem.

For (ii) we have that

$$\mu_m(A) = \sum_{i=1}^{d} |\lambda_i|^m \mu_m^{(i)}(A) < |\lambda_{k+1}|^m \left(-\varepsilon + s \sum_{i=k+2}^{d} \left(\frac{|\lambda_i|}{|\lambda_{k+1}|} \right)^m \right).$$

Now there clearly exists some m of the form $pn + q$ such that the right hand side is negative.

For (iii) note that $\eta_1(A) < 0$ is decidable, since $\eta_1(A) \geqslant 0$ is decidable by Lemma 7.2.8. Let $0 < \delta < -\eta_1(A)$. Then there exists an increasing sequence $(n_\ell)_{\ell\in\mathbb{N}}$ such that $\mu_{n_\ell}^{(1)}(A) < \eta_1(A) + \delta < 0$ for all ℓ. This follows from the fact that for a unitary matrix U, the group $\overline{\{U^n \colon n \in \mathbb{N}\}}$ is either finite or contains no isolated points. This follows from the fact that if the set contains an isolated point, then all elements are isolated.

But a compact set which contains only isolated points is finite. Hence there exists an increasing sequence $(n_\ell)_\ell$ such that

$$\eta_i(U) \leqslant \operatorname{tr}(U^{n_\ell}) \leqslant \eta_i(U) + \frac{1}{\ell}.$$

Therefore we have

$$\mu_{n_\ell}(A) = \sum_{i=1}^{d} |\lambda_i|^{n_\ell} \mu_{n_\ell}^{(i)}(A) < |\lambda_1|^{n_\ell} \left(\eta_1(A) + \delta + s \sum_{i=2}^{d} \left(\frac{|\lambda_i|}{|\lambda_1|} \right)^{n_\ell} \right).$$

Again there exists ℓ_0 such that $\mu_{n_{\ell_0}}(A) < 0$. □

7.3 Undecidable cases

We now present two finitely generated rings, for which the moment membership problem is undecidable. Specifically, in Section 7.3.1 we prove that the moment membership problem is undecidable for the ring of commutative polynomials $\mathcal{R} = \mathbb{Z}[x_1, \ldots, x_d]$ if n is sufficiently large. In Section 7.3.2, we show that the moment membership problem is also undecidable for the space of non-commutative polynomials

$$\mathcal{R} = \mathbb{Z}\langle z_1, \ldots, z_d \rangle.$$

7.3.1 Commutative polynomial rings

In the following, we show that the generalized moment membership problem for $\mathcal{R} = \mathbb{Z}[x_1, \ldots, x_d]$ and the cone

$$\mathcal{P}_{\text{coeff}} = \{ p \in \mathbb{Z}[x_1, \ldots, x_d] : \text{all coefficients of } p \text{ are nonnegative} \}$$

is undecidable. In particular, we consider the following problem

Problem 7.3.1
Let $M \in \operatorname{Mat}_s(\mathcal{R})$ be a fixed matrix. For an input $A \in \operatorname{Mat}_s(\mathcal{R})$, decide whether

$$\operatorname{tr}(A^n \cdot M) \in \mathcal{P}$$

holds for all $n \geqslant 1$.

For generalized moments of the form $A \mapsto \operatorname{tr}(A^n \cdot M)$, we obtain the following result:

Theorem 7.3.2
If $s, d \in \mathbb{N}$ are large enough, and M is chosen suitably, then Problem 7.3.1 is undecidable for $\mathcal{R} = \mathbb{Z}[x_1, \ldots, x_d]$ and $\mathcal{P}_{\text{coeff}}$.

In order to prove this theorem, we present a chain of two reductions. We start with a known undecidable problem, a version of the matrix mortality problem:

Proposition 7.3.3

If s and d are integers that are large enough, then the following problem is is undecidable:
Let $A_1, \ldots, A_d \in \mathrm{Mat}_s(\mathbb{Z})$. Does there exist a choice of $n_1, \ldots, n_d \in \mathbb{N}$ such that

$$A_1^{n_1} \cdot A_2^{n_2} \cdots A_d^{n_d} = 0.$$

For a proof of Proposition 7.3.3 we refer to [10]. We now present the first reduction, that shows that a positivity problem for traces is undecidable.

Lemma 7.3.4

For large enough values of s and d, and a suitable matrix $N \in \mathrm{Mat}_s(\mathbb{Z})$, the following problem is undecidable: Given $A_1, \ldots, A_d \in \mathrm{Mat}_s(\mathbb{Z})$, do there exist $n_1, \ldots, n_d \in \mathbb{N}$ with

$$\mathrm{tr}(A_1^{n_1} \cdot A_2^{n_2} \cdots A_d^{n_d} \cdot N) < 0?$$

Proof. We prove the statement by a reduction from Proposition 7.3.3. First, fix the matrix

$$N = \begin{pmatrix} \mathbf{0} & 0 \\ 0 & 1 \end{pmatrix} + \sum_{i,j=1}^{s} \begin{pmatrix} E_{ij} \otimes E_{ij} & 0 \\ 0 & 0 \end{pmatrix} \in \mathrm{Mat}_s(\mathbb{Z})^{\otimes 2} \oplus \mathbb{Z} \subseteq \mathrm{Mat}_{s^2+1}(\mathbb{Z})$$

where $E_{ij} = |i\rangle\langle j|$ with $|k\rangle$ being the k^{th} standard vector. For every matrix in $\mathrm{Mat}_{s^2+1}(\mathbb{Z})$ of the form

$$Y = \begin{pmatrix} X \otimes X & 0 \\ 0 & a \end{pmatrix}$$

we have

$$\mathrm{tr}(YN) = a + \sum_{i,j=1}^{s} X_{ij}^2.$$

For an instance $A_1, \ldots, A_d \in \mathrm{Mat}_s(\mathbb{Z})$ of Proposition 7.3.3, define the following $d+1$ matrices:

$$B_i = \begin{pmatrix} A_i \otimes A_i & 0 \\ 0 & 1 \end{pmatrix} \quad \text{for } i = 1, \ldots, d$$

and

$$B_{d+1} = \begin{pmatrix} I_s \otimes I_s & 0 \\ 0 & -1 \end{pmatrix}$$

where I_s is the identity matrix of size s.

Let $n_1, \ldots, n_d \in \mathbb{N}$ such that

$$A_1^{n_1} \cdots A_d^{n_d} = 0.$$

Choosing $n_{d+1} = 1$ we obtain

$$\mathrm{tr}\left(B_1^{n_1} \cdots B_{d+1}^{n_{d+1}} \cdot N\right) = -1 + \sum_{i,j=1}^{s} (A_1^{n_1} \cdots A_d^{n_d})_{ij}^2 = -1 < 0.$$

Conversely, let $n_1, \ldots, n_{d+1} \in \mathbb{N}$ such that $\operatorname{tr}(B_1^{n_1} \cdots B_{d+1}^{n_{d+1}} \cdot N) < 0$. This is clearly only possible for n_{d+1} odd and

$$\sum_{i,j=1}^{s} \left(A_1^{n_1} \cdots A_d^{n_d}\right)_{ij}^2 = 0,$$

which implies $A_1^{n_1} \cdots A_d^{n_d} = 0$. □

We are now ready to prove the main result of this section.

Proof of Theorem 7.3.2. Given $A_1, \ldots, A_d \in \mathrm{Mat}_s(\mathbb{Z})$, set

$$A = \sum_{i=1}^{d} \left(\sum_{1 \leqslant j \leqslant i} |j\rangle \langle i| \right) \otimes A_i \cdot x_i \in \mathrm{Mat}_{ds}(\mathcal{R}).$$

Moreover, define

$$M = |\phi\rangle \langle \phi| \otimes N$$

with

$$|\phi\rangle = |1\rangle + |2\rangle + \ldots + |s\rangle$$

and N as in Lemma 7.3.4. We have that

$$\begin{aligned} \operatorname{tr}(A^n M) &= \sum_{1 \leqslant i_1 \leqslant \cdots \leqslant i_n \leqslant d} i_1 \cdot \operatorname{tr}(A_{i_1} \cdots A_{i_n} \cdot N) \cdot x_{i_1} \cdots x_{i_n} \\ &= \sum_{n_1 + \cdots + n_d = n} \underbrace{c_{n_1,\ldots,n_d}}_{\geqslant 1} \cdot \operatorname{tr}(A_1^{n_1} \cdots A_d^{n_d} \cdot N) \cdot x_1^{n_1} \cdots x_d^{n_d} \end{aligned}$$

where $c_{n_1,\ldots,n_d} = \min\{i \colon n_i \neq 0\}$. Thus Problem 7.3.1 reduces to the undecidable problem from Lemma 7.3.4. □

Note that since the sequence $\operatorname{tr}(A^n M)$ is clearly an LRS (see Lemma 7.1.3), the last result shows that positivity of LRS over $\mathcal{R} = \mathbb{Z}[x_1, \ldots, x_d]$ is undecidable in general.

7.3.2 Non-commutative polynomial rings

We now consider the ring $\mathcal{R} = \mathbb{Z}\langle z_1, \ldots, z_d\rangle$ of non-commutative polynomials, and show that its moment membership problem is undecidable for the cone of polynomials with positive coefficients. As a $\mathbb{Z}$-module, a basis of $\mathcal{R}$ consists of all words in the letters $z_1, \ldots, z_d$, where the order of letters *does* matter. Concatenation of words extends to a multiplication making $\mathcal{R}$ a unital ring, where 1 corresponds to the empty word. There is a slightly different way to define this object, namely just as the tensor algebra

$$\mathbb{Z}\langle z_1, \ldots, z_d\rangle = T(\mathbb{Z}^d).$$

The equivalence of definitions is apparent when identifying a word $z_{k_1} \cdots z_{k_m}$ with the element $|k_1, \ldots, k_m\rangle \in \left(\mathbb{Z}^d\right)^{\otimes m}$, where $|r\rangle$ denotes the r-th standard basis vector in $\mathbb{Z}^d$.

We equip $\mathcal{R}$ with two (a priori) different sets of positive elements:

$$\begin{aligned}\mathcal{P}_{\text{coeff}} &:= \mathbb{Z}_{\geqslant 0}\langle z_1, \dots, z_d\rangle \\ &= \{p \in \mathbb{Z}\langle z_1, \dots, z_d\rangle : \text{all coefficients of } p \text{ are nonnegative}\}\end{aligned}$$

$$\mathcal{P}_{\text{eval}} := \left\{ p \in \mathbb{Z}\langle z_1, \dots, z_d\rangle : \begin{array}{l} \forall \ell, A_1, \dots, A_d \in \operatorname{Mat}_\ell(\mathbb{Z}_{\geqslant 0}) : \\ p(A_1, \dots, A_d) \in \operatorname{Mat}_\ell(\mathbb{Z}_{\geqslant 0}) \end{array} \right\}.$$

We first show that both cones coincide, which is a free version of Pólya's Theorem.

Theorem 7.3.5 (Free Pólya's Theorem)
Let $p \in \mathbb{C}\langle z_1, \dots, z_d\rangle$ with $m := \deg(p)$. Then the following are equivalent:

(i) All coefficients of p are nonnegative reals.
(ii) For all $A_1, \dots, A_d \in \operatorname{Mat}_{m+1}(\mathbb{Z}_{\geqslant 0})$ we have

$$p(A_1, \dots, A_d) \in \operatorname{Mat}_{m+1}(\mathbb{R}_{\geqslant 0}).$$

In particular, $\mathcal{P}_{\text{coeff}} = \mathcal{P}_{\text{eval}}$, and in the definition of $\mathcal{P}_{\text{eval}}$ one can restrict ℓ to $\deg(p) + 1$.

Proof. (i) $\Rightarrow$ (ii) is obvious (even without the restriction on the matrix size $m + 1$). For (ii) $\Rightarrow$ (i) we construct matrices $A_1, \dots, A_d$ that allow us to isolate a single coefficient of p.

Let $z_{k_1} \cdots z_{k_\ell}$ be a word in the letters $z_1, \dots, z_d$. For $j = 1, \dots, d$ define

$$A_j := \sum_{i=1,\dots,\ell;\ k_i = j} E_{i,i+1} \in \operatorname{Mat}_{\ell+1}(\mathbb{Z}_{\geqslant 0}),$$

where $E_{i,j}$ denotes the matrix (of size $\ell + 1$) with a 1 in position (i, j) and zeros elsewhere. For $t_1, \dots, t_r \in \{1, \dots, d\}$ we have

$$A_{t_1} \cdots A_{t_r} = \sum_{\substack{i \\ k_i = t_1 \\ k_{i+1} = t_2 \\ \vdots \\ k_{i+r-1} = t_r}} E_{i,i+r} \in \operatorname{Mat}_{\ell+1}(\mathbb{Z}_{\geqslant 0}).$$

In particular, the $(1, \ell + 1)$-entry of a product $A_{t_1} \cdots A_{t_r}$ is 1 if and only if $r = \ell$ and $(k_1, \dots, k_\ell) = (t_1, \dots, t_\ell)$; in all other cases it is zero. So $p(A_1, \dots, A_d)$ contains in its upper right entry precisely the coefficient of p at the word $z_{k_1} \cdots z_{k_\ell}$.

Since all words appearing in p are of length at most $\deg(p) = m$, we can do this procedure with matrices A_j of size at most $m + 1$, and thus clearly with matrices of size exactly $m + 1$. □

Remark 7.3.1 (Pólya's theorem for commutative polynomials)
Pólya's theorem [98, 64] states that for every homogeneous polyno-

mial $p \in \mathbb{R}[x_1, \ldots, x_d]$ that is strictly positive on the d-simplex

$$\Delta_d := \left\{ (a_1, \ldots, a_d) \in \mathbb{R}^d : a_i \geqslant 0, \sum_{i=1}^{d} a_i = 1 \right\},$$

the polynomial

$$(x_1 + \cdots + x_d)^n \cdot p(x_1, \ldots, x_d)$$

has positive coefficients, for sufficiently large $n \in \mathbb{N}$. In Theorem 7.3.5, the space of nonnegative matrices takes the role of the d-simplex. While in the commutative case we have to multiply p with an additional polynomial, this is not the case in the free version.

We now show that for these cones, the moment membership problem is undecidable.

Theorem 7.3.6

Let $d, s \geqslant 7$. Then the moment membership problem for $\mathcal{R} = \mathbb{Z}\langle z_1, \ldots, z_d \rangle$, $\mathcal{P}_{\text{coeff}} = \mathcal{P}_{\text{eval}}$ and s is undecidable. This remains true if we restrict the instances to linear matrix polynomials, i.e. $A \in \text{Mat}_s(\mathbb{Z}\langle z_1, \ldots, z_d \rangle)$ whose entries are linear forms in $z_1, \ldots, z_d$.

Proof. For $A = \sum_{k=1}^{d} z_k A_k$ with $A_k \in \text{Mat}_s(\mathbb{Z})$ we have

$$\mu_n(A) = \sum_{k_1, \ldots, k_n = 1}^{d} \text{tr}(A_{k_1} \cdots A_{k_n}) \cdot z_{k_1} \cdots z_{k_n}.$$

So $\mu_n(A) \in \mathcal{P}_{\text{coeff}}$ means that $\text{tr}(A_{k_1} \cdots A_{k_n}) \geqslant 0$ for all $k_1, \ldots, k_n = 1, \ldots, d$. Undecidability of this problem was proven in [36, Lemma 3]. □

7.3.3 Commutative polynomials with an unbounded number of variables

It is an open question whether Problem 7.3.1 remains undecidable for commutative polynomials if the set $\mathcal{P}$ is specified to be sos polynomials or to be nonnegative polynomials. In this part, we show that a certain generalization of the problem becomes undecidable, even for sos and nonnegative polynomials.

More specifically, we show that the invariant unconstraint (Θ_n, C_n)-decomposition for polynomials[2] has no local and computable certificate of positivity. We will reach this conclusion by proving that Problem 7.3.7 is undecidable.

2: For the definition of (Θ_n, C_n)-decompositions of polynomials, we refer to Section 2.3.1.

Given a collection of s^2 polynomials in $\mathbb{Z}[\mathbf{x}]$, denoted $(p_{\alpha,\beta})_{\alpha,\beta=1}^{s}$, define

$$p_n := \sum_{\alpha_1, \ldots, \alpha_n = 1}^{s} p_{\alpha_1,\alpha_2}(\mathbf{x}^{[1]}) \cdot p_{\alpha_2,\alpha_3}(\mathbf{x}^{[2]}) \cdots p_{\alpha_n,\alpha_1}(\mathbf{x}^{[n]}). \tag{7.6}$$

We have that $p_n \in \mathbb{R}[\mathbf{x}^{[1]}, \ldots, \mathbf{x}^{[n]}]$. Note that the summation indices are arranged in a circle Θ_n, and that the local polynomials do not depend on the site, so p_n is invariant under the cyclic group C_n. The previous expression is thus a (Θ_n, C_n)-decomposition of p_n. Moreover, p_n generalizes the moment problem in the following sense: If

$$\mathbf{x} := \mathbf{x}^{[1]} = \mathbf{x}^{[2]} = \ldots = \mathbf{x}^{[n]},$$

then we obtain a moment sequence

$$p_n(\mathbf{x}) = \mathrm{tr}(A^n)$$

with

$$A = \begin{pmatrix} p_{1,1}(\mathbf{x}) & p_{1,2}(\mathbf{x}) & \cdots & p_{1,s}(\mathbf{x}) \\ p_{2,1}(\mathbf{x}) & p_{2,2}(\mathbf{x}) & & \\ \vdots & & \ddots & \\ p_{s,1}(\mathbf{x}) & & & p_{s,s}(\mathbf{x}) \end{pmatrix}.$$

Problem 7.3.7 (Positivity of (Θ_n, C_n)-decompositions)
Given positive integers m and s and a collection of polynomials $(p_{\alpha,\beta})_{\alpha,\beta=1}^{s} \in \mathbb{Z}[\mathbf{x}]$ (where $\mathbf{x}$ denotes a vector of m variables $(x_1, \ldots, x_m)$),

(a) Is p_n sos for all $n \in \mathbb{N}$?
(b) Is p_n nonnegative for all $n \in \mathbb{N}$?

Theorem 7.3.8 (Undecidability of Problem 7.3.7)
Problem 7.3.7 (a) and Problem 7.3.7 (b) are undecidable. This is true even if $m, D \geqslant 7$ and if the polynomials are of the form

$$p_{\alpha,\beta}(\mathbf{x}) = \sum_{j=1}^{m} p_{\alpha,\beta,j} \cdot x_j^2$$

with $p_{\alpha,\beta,j} \in \mathbb{Z}$ for all $\alpha, \beta \in \{1, \ldots, D\}$.

So there does not exist an algorithm that can decide in finite time whether p_n is sos or nonnegative for all n, given the local polynomials as input. We will prove Theorem 7.3.8 by a reduction from the following undecidable problem:

Theorem 7.3.9 (Undecidability of positivity for all system sizes [36])
Let $|T_{\alpha,\beta}\rangle \in \mathbb{Z}^m$ for $\alpha, \beta \in \{1, \ldots, D\}$ be a collection of vectors. For $n \geqslant 0$ define

$$|T_n\rangle := \sum_{\alpha_1,\ldots,\alpha_n=1}^{D} |T_{\alpha_1,\alpha_2}\rangle \otimes |T_{\alpha_2,\alpha_3}\rangle \otimes \cdots \otimes |T_{\alpha_n,\alpha_1}\rangle.$$

For $m, D \geqslant 7$, the following problem is undecidable:

$$\text{Is } |T_n\rangle \text{ nonnegative for all } n \in \mathbb{N}?$$

Proof of Theorem 7.3.8. Let $|T_{\alpha,\beta}\rangle \in \mathbb{Z}^m$ be a collection of vectors for $\alpha, \beta \in \{1, \ldots, D\}$. We apply the construction from Section 5.2.3 to obtain the collection of polynomials

$$p_{\alpha,\beta} = \sum_{j=1}^{m} \langle j \mid T_{\alpha,\beta}\rangle \, x_j^2$$

and generate the polynomials $p_n \in \mathbb{Z}[\mathbf{x}^{[1]}, \ldots, \mathbf{x}^{[n]}]$. It is obvious that $p_{|T_n\rangle} = p_n$ for all n, and from Lemma 5.2.4 we thus know that $|T_n\rangle$ is nonnegative if and only if p_n is a sum of squares/nonnegative. So decidability of Problem 7.3.7 (a) or (b) contradicts Theorem 7.3.9. $\square$

We remark that Problem 7.3.7 remains undecidable if the input polynomials are in $\mathbb{Q}[\mathbf{x}]$, since multiplying all polynomials by a positive constant does not change the positivity/sos property.

It can also be shown that a *bounded* version of the questions of Problem 7.3.7 — i.e. where n is fixed — result in an NP-hard problem (see Chapter 8).

7.4 Conclusion

We have studied the moment membership problem (Problem 7.1.1) for matrices over a ring. We have shown that there is a relation to LRS for commutative rings (Lemma 7.1.3 and Lemma 7.1.4) and that the moments positivity problem is decidable in many cases, including unitary and orthogonal matrices (Theorem 7.2.3 and Corollary 7.2.5) as well as matrices with a unique dominating eigenvalue or only real eigenvalues (Theorem 7.2.7). Finally, we have shown that the generalized moment membership problem is undecidable over the ring of commutative and non-commutative polynomials, where the positivity cone is given by the set of polynomials with non-negative coefficients (Theorem 7.3.2 and Theorem 7.3.6).

The central open question is still whether the moment membership problem is decidable or undecidable for $\mathcal{R} = \mathbb{Q}$ and $\mathcal{P} = [0, \infty)$. In the context of rings it would be interesting, whether it is also undecidable for commutative polynomials for the cone of sum-of-square polynomials or the non-negative polynomials. This might be the case since these cones have a richer structure than that of polynomials with nonnegative coefficients.

Bounded versions of undecidable problems | 8

Many problems in quantum information and quantum many-body physics are undecidable. This includes the spectral gap of physical systems [34, 8], membership problems for quantum correlations [116, 117, 70, 53, 86], properties of tensor networks [36, 72, 108], measurement occurrence and reachability problems [46, 129], and many more [39, 48, 15, 107, 51]. In addition, other problems are believed to be undecidable, such as detecting quantum capacity [33], distillability of entanglement [129], or tensor-stable positivity [48].

This chapter is based on [73].

All these problems have a common theme: They ask for a property that includes an unbounded parameter. For example, in a quantum correlation scenario, the dimension of the shared quantum state between the two parties may be unbounded. Similarly, properties characterizing many-body systems, such as the spectral gap, inherently involve assertions across arbitrarily large system sizes.

On the other hand, many problems in science, engineering, and mathematics fall under the umbrella of NP-hard problems [128]. Some examples relevant for physics are finding the ground state energy of an Ising model [4], the training of variational quantum algorithms [14], or the quantum separability problem [62, 56], and many more. These problems typically concern properties where all size parameters are bounded or even fixed. For example, the ground state energy problem concerns the minimal energy of Hamiltonians with fixed system size.

This highlights an analogy between certain classes of problems: an *unbounded problem* tests a property for an unbounded number of occurrences (which can be generated recursively), whereas the corresponding *bounded version* tests the same property for a bounded number of situations. This includes, for example, testing a certain property of a translational invariant spin system for all system sizes, or up to a certain size. A common observation in this context is that bounded versions of undecidable problems tend to be NP-hard. This insight has been noted in various examples, as documented in [72, 17, 108], as well as discussed in [128, Chapter 3].

Despite this analogy, the techniques used to prove NP-hardness and undecidability often differ. While proofs of undecidability predominantly hinge on reductions from the halting problem, the Post correspondence problem or the Wang tiling problem, NP-hardness proofs mainly rely on reductions from the satisfiability problem SAT, or from NP-complete graph problems like the 3-coloring problem or MaxCut.[1]

1: We refer to [2, Chapter 2] for the details on these problems.

In this work, we establish a relation between undecidable problems and certain NP-hard problems. Specifically, we define the notion of a bounded version of a problem and a method to leverage the reduction from unbounded problems to their corresponding bounded problems (see Figure 8.1). Subsequently, we present two versions of the halting problem whose bounded versions are NP-hard, and use these, together with our method, to provide simple and unified proofs of the NP-hardness of the bounded version of the Post correspondence problem,

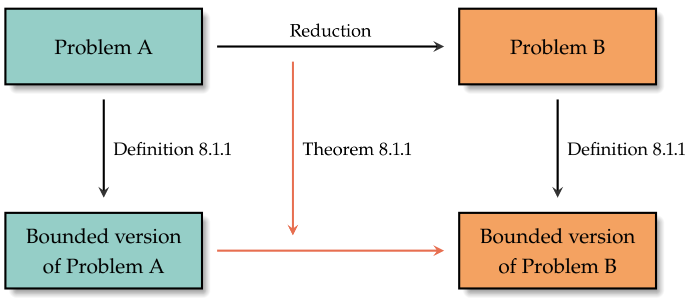

Figure 8.1: If Problem B is at least as hard as Problem A (i.e. there is a reduction from A to B), is the bounded version of Problem B at least as hard as the bounded version of Problem A? Theorem 8.1.1 gives a sufficient condition when this is the case by reusing the reduction between their unbounded versions.

the matrix mortality problem, the positivity of matrix product operators, the reachability problem, the tiling problem, and the ground state energy problem.

This work sheds light on the various intractability levels of problems used in theoretical physics by highlighting the computational consequences of bounding a parameter. More generally, this work is part of a tradition of studying problems from a computational perspective, which has proven extremely successful in mathematics and beyond [128]. For example, the hardness results of the ground state energy problem rule out a tractable solution of the ground state for a given Hamiltonian, both for unbounded system sizes as well as a fixed system size.

8.1 Bounding

In this section, we present a definition of a bounded version of a language (Section 8.1.1), and a method to leverage the reduction from unbounded problems to their corresponding bounded versions (Section 8.1.2).

8.1.1 Definition of bounding

Let Σ be a finite alphabet and Σ^* the set of all words generated from Σ. A language $L \subseteq \Sigma^*$ encodes all the yes-instances of a given problem, i.e. $x \in L$ if x is a yes-instance and $x \notin L$ if x is a no-instance.

We now define a bounded version L_B of L. For this purpose, we add a second parameter $n \in \mathbb{N}$ to every yes-instance in L. This parameter acts as an acceptance threshold for every yes-instance $x \in L$ and is encoded in unary, i.e. for $1 \in \Sigma$, every element of L_B is of the form $\langle x, 1^n \rangle$, where 1^n represents the n-fold concatenation of 1.

Definition 8.1.1 (Bounded version)
Let $L \subseteq \Sigma^*$ be a language. A language

$$L_B \subseteq \left\{ \langle x, 1^n \rangle \mid x \in \Sigma^*, n \in \mathbb{N} \right\}$$

is called a *bounded version of* L if

(i) $x \in L \iff \exists n \in \mathbb{N} : \langle x, 1^n \rangle \in L_B$.
(ii) $\langle x, 1^n \rangle \in L_B \implies \langle x, 1^{n+1} \rangle \in L_B$.

We shall often refer to L as the *unbounded* language of L_B.

First, note that the definition of bounded versions relies only on the existence of a parameter n in the problem that acts accordingly. While most problems we consider in this paper are RE-complete, Definition 8.1.1 applies to languages of arbitrary complexity. Moreover, note that the bounding parameter can also be encoded differently. For example, if the parameter is encoded in binary, most of the bounded version would be NEXP-hard instead of NP-hard. Finally, we remark that the process of bounding a language can be reversed. Given a language L_B with instances of the form $\langle x, 1^n \rangle$ satisfying only Condition (ii), there is a unique language L, defined via (i), which is the unbounded language of L_B.

Many problems mentioned in the introduction contain a parameter that gives rise to a bounded version according to Definition 8.1.1. This parameter can be the system size for tensor network and spectral gap problems, or the dimension of the entangled state for quantum correlation scenarios; we will present many such examples in Section 8.3.

As an example, let us consider the halting problem Halt with its known bounded version BHalt. The former takes instances $\langle T, x_0 \rangle$ with a description T of a Turing machine and an input x_0. An instance $\langle T, x_0 \rangle$ is accepted if and only if the Turing machine T halts on x_0. The bounded halting problem takes instances $\langle T, x_0, 1^n \rangle$, which are accepted if and only if the Turing machine halts on x_0 within n computational steps. BHalt is indeed a bounded version according to Definition 8.1.1 since halting of a Turing machine is equivalent to the existence of a finite halting time, and halting within n steps implies halting within $n + 1$ steps.

We remark that in Definition 8.1.1 there is some freedom in the choice of the bounding parameter. For example, for every non-decreasing, unbounded function $f : \mathbb{N} \to \mathbb{N}$, the language

$$\text{BHALT}_f := \left\{ \langle T, x_0, 1^n \rangle \mid T \text{ halts on } x_0 \text{ in } f(n) \text{ steps} \right\}$$

is also a bounded version of Halt. In this paper, we will focus on the simplest versions setting $f = \text{id}$ in all examples.

8.1.2 Leveraging reductions to the bounded case

Given the hardness of the unbounded languages, what can we say about the bounded ones? We will now give a condition to leverage a reduction of unbounded problems to a reduction between the corresponding bounded problems. This results in a method to prove hardness results of many bounded versions of undecidable problems, as we will see in Section 8.3.

Let L_B be a bounded version of $L \subseteq \Sigma^*$. For $x \in \Sigma^*$, we define the threshold parameter

$$n_{\min,L}[x] := \inf\{n \in \mathbb{N} : \langle x, 1^n \rangle \in L_B\}$$

where we set $\inf \emptyset = \infty$. In other words, $n_{\min}[x]$ denotes the minimum value of n leading to an accepting instance of L_B. Note that

$$n_{\min}[x] < \infty$$

for every $x \in L$ due to (i) of Definition 8.1.1 and

$$n_{\min}[x] = \infty$$

if $x \notin L$. Moreover, $\langle x, 1^n \rangle \in L_B$ if and only if $n \geqslant n_{\min}[x]$ due to (ii) of Definition 8.1.1.

Theorem 8.1.1 (Hardness of bounded versions)
Let $L_1, L_2 \subseteq \Sigma^*$ be two languages and $\mathcal{R} : L_1 \to L_2$ a polynomial-time reduction[2] from L_1 to L_2, i.e. $L_1 \leqslant_{\text{poly}} L_2$. Furthermore, let L_{B1} and L_{B2} be bounded versions of L_1 and L_2, respectively.
If there is a strictly increasing polynomial $p : \mathbb{N} \to \mathbb{N}$ such that

$$n_{\min,L_2}[\mathcal{R}(x)] \leqslant p\big(n_{\min,L_1}[x]\big) \tag{8.1}$$

for every $x \in L$, then

$$\langle x, 1^n \rangle \mapsto \langle \mathcal{R}(x), 1^{p(n)} \rangle \tag{8.2}$$

is a polynomial-time reduction from L_{B1} to L_{B2}, hence $L_{B1} \leqslant_{\text{poly}} L_{B2}$.

2: We refer to Definition 6.1.4 for the notion of polynomial-time reductions.

Theorem 8.1.1 also generalizes to other types of reductions. For example, we obtain an exponential-time reduction between the bounded versions when $\mathcal{R}$ is considered a exponential-time reduction and p being a strictly increasing function that can be computed in exponential time.

Proof. Since $\mathcal{R}$ and p are polynomial-time maps, the map in Equation (8.2) is also polynomial-time. It remains to show that yes/no-instances are preserved via this map. We have that $\langle x, 1^n \rangle \in L_{B1}$ if and only if $n \geqslant n_{\min,L_1}[x]$. This is equivalent to

$$p(n) \geqslant p\big(n_{\min,L_1}[x]\big) \geqslant n_{\min,L_2}[\mathcal{R}(x)]$$

since p is a strictly increasing function. But this is again equivalent to $\langle \mathcal{R}(x), 1^{p(n)} \rangle \in L_{B2}$. □

We require that p is strictly increasing instead of mere non-decreasing as we need the equivalence of the statements $n \geqslant m$ and $p(n) \geqslant p(m)$ in the proof.

In words, Condition (8.1) demands that there is a polynomial that relates thresholds of x and $\mathcal{R}(x)$ for all x.

Many known reductions of undecidable problems implicitly contain such a polynomial p in their construction. This gives an almost-for-free proof of the NP-hardness of their bounded problems. However, most of these works do not make this polynomial explicit and therefore do not obtain the NP-hardness results. While the theorem only assumes that $p(n_{\min,L_2}[x])$ upper bounds $n_{\min,L_1}[\mathcal{R}(x)]$, in all examples, we have an equality between these expressions. In Section 8.3, we will present many examples of this behavior.

8.2 Halting problems as root problems

The result of Theorem 8.1.1 gives only relative statements about hardness. Specifically, it allows to construct a reduction between bounded versions given a reduction between their original problems. To prove NP/coNP-hardness of bounded problems, we need root problems with bounded versions whose complexities are already known. In this section, we review two fundamental undecidable problems and their bounded versions, namely two variants of the halting problem.

While Halt and BHalt are the most basic versions of halting problems, we need variations of the halting problem that take non-deterministic Turing machines as inputs. This is due to the fact that, while Halt is undecidable, BHalt is in P.[3] Since we want to prove NP/coNP-hardness of bounded problems, we need root problems with a NP/coNP-hard bounded version to start the reduction from. Therefore, we introduce two non-deterministic versions of Halt, called NHalt and NHaltAll, with an NP-hard and a coNP-hard bounded version, respectively.

3: An efficient algorithm to decide BHalt is simply letting the the Turing machine with description T run on a universal Turing machine. Since the simulation only needs a polynomial-time overhead, this procedure checks whether T halts within n steps after polynomially many steps in the size of $\langle T, x_0, 1^n \rangle$.

- The problem NHalt checks the halting of a non-deterministic Turing machine on the empty tape. An instance is given by a description of a non-deterministic Turing machine T, which is accepted if and only if T halts on the empty tape[4]. Its bounded version BNHalt takes instances $\langle T, 1^n \rangle$ and accepts if and only if T halts on the empty tape in at most n steps. The unbounded problem is RE-hard since it contains the (deterministic) halting problem on the empty tape, which is itself RE-hard. Its bounded version BNHalt is NP-hard.
- The problem NHaltAll takes a description of a non-deterministic Turing machines T as an instance, which is accepted if and only if T halts on the empty tape along *all* computation paths. Its bounded version BNHaltAll is given by instances $\langle T, 1^n \rangle$ which are accepted if and only if T halts on the empty tape within n computational steps along *all* computation paths. The unbounded problem is RE-hard, and the bounded version is coNP-hard.

4: In other words, it accepts if and only if there is a computation path such that T halts along this path.

NHalt will be the root problem to prove the hardness of the bounded Post correspondence problem (Section 8.3.1) and the bounded matrix mortality problem (Section 8.3.2). NHaltAll will be the root problem to prove the hardness of the bounded Tiling problem (Section 8.3.7).

Let us now provide a detailed analysis of the two halting problems NHalt and NHaltAll together with their bounded versions which act as root problems. We start with the unbounded problems showing their undecidability, and continue with their bounded version's complexity.

Note that the inputs of NHalt and NHaltAll are just a Turing machines T, as we ask whether T halts on the empty tape.

Definition 8.2.1 (Non-deterministic Halting problems)
Let T be a description of a non-deterministic Turing machine.

$$T \in \text{NHalt} \quad :\Longleftrightarrow \quad T \text{ halts on the empty tape.}$$

$$T \in \text{NHaltAll} \quad :\Longleftrightarrow \quad T \text{ halts on the empty tape along all paths.}$$

Both problems are undecidable, as the following reduction from the halting problem Halt shows.

Theorem 8.2.1
NHalt and NHaltAll are RE-complete.

Proof. We prove RE-hardness only for NHalt, as the same argument applies to NHaltAll. To this end, we provide a reduction from Halt. Recall that Halt takes $\langle T, x_0 \rangle$ as input (where T is a description of a deterministic Turing machine T, and x_0 is an input) and accepts if and only if T halts on x_0.[5] The reduction transforms instance $\langle T, x_0 \rangle$ to a Turing machine $T' = \mathcal{R}(\langle T, x_0 \rangle)$ which first writes x_0 on the tape, and then does the same computation as T on the given input. By construction, $\langle T, x \rangle \in$ Halt if and only if $T' \in$ NHalt, i.e. $\mathcal{R}$ is a valid reduction.

5: We refer to Section 6.1.2 for the definition of the (deterministic) halting problem Halt.

That NHalt $\in$ RE follows by taking the halting computation path as a certificate, and a verifier that verifies the computation along the path. That NHaltAll $\in$ RE follows by taking the halting time as a certificate, and a verifier that verifies that the computation halts along all paths within this halting time. □

Let us now consider the bounded versions BNHalt and BNHaltAll. Since these problems have different complexities, we will treat them separately.

Definition 8.2.2 (Bounded non-deterministic halting problem I)
Let T be a description of a non-deterministic Turing machine, and $n \in \mathbb{N}$.

$$\langle T, 1^n \rangle \in \text{BNHalt} \quad :\Longleftrightarrow \quad \begin{array}{l} T \text{ halts on the empty tape} \\ \quad \text{in } n \text{ steps.} \end{array}$$

Theorem 8.2.2
BNHalt is NP-complete.

Proof. To show that BNHalt is NP-hard, we prove that every NP-language L has a polynomial-time reduction to BNHalt. Since L is in NP, there exists a non-deterministic polynomial-time Turing machine M which accepts x within time $p(|x|)$ if and only if $x \in L$. We construct a non-deterministic Turing machine $P_{M,x}$ that (i) writes x on the tape, (ii) does the same computation as M on the tape with input x, and (iii) if M accepts x along a path, $P_{M,x}$ halts along this path, and if M rejects x along a path, $P_{M,x}$ loops along this path. Since step (i) needs a polynomial number $q(|x|)$ steps, and step (iii) needs a constant number k of steps, we have that $x \in L$ if and only if

$$\langle P_{M,x}, 1^{q(|x|)+k+p(|x|)} \rangle \in \text{BNHalt}.$$

Completeness follows from Definition 6.1.2 by choosing the halting computation path as a certificate, and a polynomial-time verifier which verifies the computation along this path. □

Similarly, we define the problem BNHaltAll as the language accepting the instance $\langle T, 1^n \rangle$ if and only if T halts on the empty tape along *all* computation paths in at most n steps.

Definition 8.2.3 (Bounded non-deterministic halting problem II)
Let T be a description of a non-deterministic Turing machine T, and $n \in \mathbb{N}$.

$$\langle T, 1^n \rangle \in \text{BNHaltAll} \quad :\Longleftrightarrow \quad \begin{array}{c} T \text{ halts on the empty tape} \\ \text{along all paths in } n \text{ steps.} \end{array}$$

While NHalt and NHaltAll are in the same complexity class, their bounded versions are in different ones.

Theorem 8.2.3
BNHaltAll is coNP-complete.

Proof. The hardness proof is very similar to Theorem 8.2.2. Namely, we prove that every coNP-language L has a polynomial-time reduction to BNHaltAll. Since L is in coNP, there exists a non-deterministic polynomial-time Turing machine M which accepts x along every computation path of length at most $p(|x|)$ if and only if $x \in L$. We construct the non-deterministic Turing machine $P_{M,x}$ which (i) writes x on the tape, (ii) does the same computation as M on the tape with input x, and (iii) if M accepts x along a path, $P_{M,x}$ halts along this path. If M rejects x along a path, $P_{M,x}$ loops along this path. Since (i) needs a polynomial number $q(|x|)$ steps and (iii) needs a constant number k of steps, we have that $x \in L$ if and only if

$$\langle P_{M,x}, 1^{q(|x|)+k+p(|x|)} \rangle \in \text{BNHaltAll}.$$

Completeness again follows from Equation (6.2) by choosing computation paths as a certificate, and a polynomial-time verifier that verifies the computation along the given path. □

While reductions for undecidable problems usually stem from the deterministic halting problem Halt, here we need non-deterministic halting problems in order to prove NP-hardness of the bounded versions. Canonical extensions of the reductions from Halt to a non-deterministic halting problem lead to different choices of root problems. For example, the Post correspondence problem has a similar structure as NHalt, while the structure of the tiling problem relates to NHaltAll. We will elaborate on these structures in the corresponding sections.

We expect that other variants of the halting problem serve as root problems for other complexity results; see Section 8.4 for further discussion.

8.3 A tree of undecidable problems and their bounded versions

In this section, we apply Theorem 8.1.1 to several undecidable problems in order to prove the NP-hardness of the bounded versions. The problems studied in this paper are summarized in Figure 8.2, where every edge corresponds to one application of the theorem.

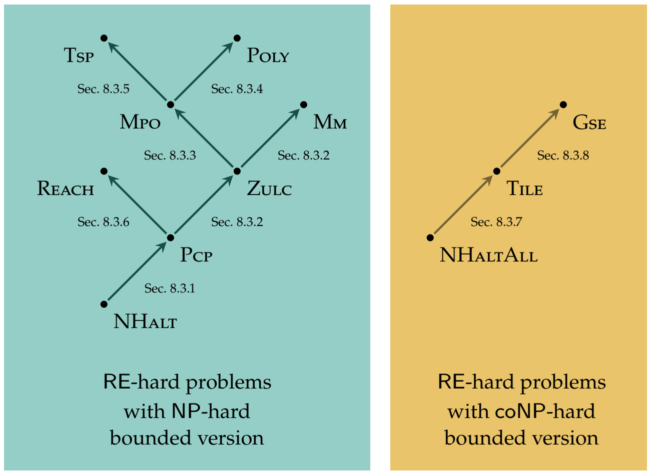

Figure 8.2: The problems and reductions considered in this work. NHalt is the non-deterministic halting problem, Pcp is the Post correspondence problem, Reach is the reachability problem for resource theories, Zulc is the zero in the upper left corner problem, Mm is the matrix mortality problem, Mpo is the positivity of Matrix product operators problem, Tsp is the stability of positive maps problem and Poly is the polynomial positivity problem. NHaltAll is the non-deterministic halting problem on all computational paths, Tile is the Wang tiling problem, and Gse is the ground state energy problem. NHalt and NHaltAll are the root problems, and every arrow corresponds to a reduction, explained in the referenced subsection.

8.3.1 The Post correspondence problem

The Post correspondence problem (Pcp) [99] is an undecidable problem with a particularly simple and intuitive formulation. For this reason, it is often used to prove undecidable results in quantum information theory [129], including a version of the matrix product operator positivity problem [72], threshold-problems for probabilistic and quantum finite automata [15], or reachability problems in resource theories [107]. It is stated as follows:

Problem 8.3.1 (The Post correspondence problem)
Given two finite sets of words, $\{a_1, \dots, a_k\}$ and $\{b_1, \dots, b_k\} \subseteq \Sigma^*$, is there a finite sequence of indices $i_1, \dots, i_\ell$ such that

$$a_{i_1} a_{i_2} \dots a_{i_\ell} = b_{i_1} b_{i_2} \dots b_{i_\ell} \text{ ?}$$

This decision problem is known to be RE-complete via a reduction from the halting problem. Since a_i and b_i only appear in fixed pairs, this problem has an equivalent formulation in terms of dominoes

$$d_i = \left[\frac{a_i}{b_i}\right].$$

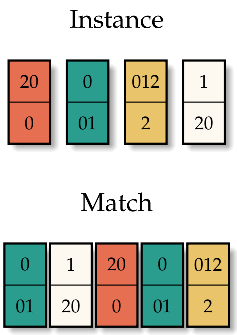

Figure 8.3: An instance of Pcp is a set of dominoes (top). This is a yes-instance if they form a match (bottom), i.e. the words on the top and the bottom coincide.

The question is whether there exists a finite arrangement of dominoes that form a match, i.e. where the upper and lower parts coincide when the words are read across the dominoes (see Figure 8.3).

We define a bounded version of Pcp that checks for sequences of length at most n:

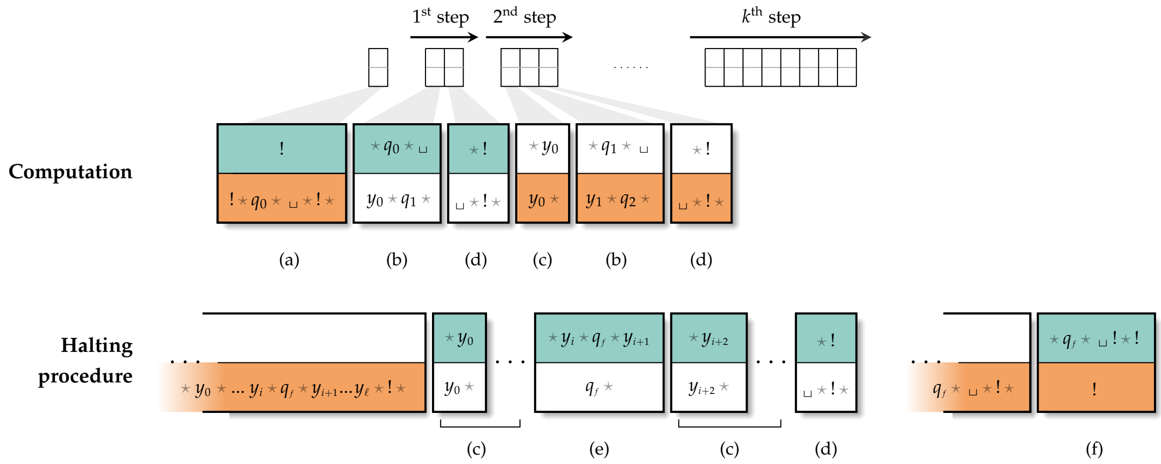

Figure 8.4: (Top) In the reduction NHALT → PCP, domino (a) contains the initial configuration of the TM, i.e. an empty tape with head at position zero. Each computation step is simulated by copying the lower string to the upper part in green. This is done by applying a transition domino (b), reproducing the tape (c), and adding a new empty tape slot (d). This generates a new string on the bottom, showing the new instantaneous description (white). Repeating the procedure simulates the computation. (Bottom) The halting of the Turing machine is mapped to the following match of tiles. When the Turing machine reaches the final state q_f, the instantaneous description is successively removed by dominoes (e). Adding a final domino (f) guarantees the match.

Problem 8.3.2 (The bounded Post correspondence problem)
Given a finite set of dominoes $\{d_1, \ldots, d_k\}$ and a number $n \in \mathbb{N}$ in unary, is there a matching arrangement of dominoes $d_{i_1}, \ldots, d_{i_\ell}$ with $\ell \leqslant n$?

This problem, denoted BPCP, is a bounded version of PCP according to Definition 8.1.1. It is known to be NP-complete (see [54, 67, 72] for the ideas of the reductions). The basic idea of the reduction is analogous to Theorem 8.1.1, i.e. using the reduction of the (unbounded) undecidable problems to relate the bounding parameters via a polynomial-time map. Yet, the usual reductions do not directly give rise to a polynomial relation between the bounding parameters, contrary to what is claimed in [72]. We will now provide a reduction NHALT → PCP leading to such a relation. Our approach is similar to that of [115].

We define a map $\mathcal{R}$ mapping a description of a Turing machine to a set of dominoes, $\mathcal{R}(T) := \langle d_1, \ldots, d_k \rangle$. This map mimics the description of T (see Figure 8.4). For example, d_1 is a domino whose lower string is given by

$$! \star q_0 \star \sqcup \star ! \star$$

where ! and $\star$ are separator symbols, and q_0 and $\sqcup$ indicate that the Turing machine head is initially in state q_0 acting on an empty tape.

Let us now provide the reduction NHALT → PCP in greater detail. The following reduction modifies that of Ref. [115], so that the bounding parameters of both problems are polynomially related.

We consider a Turing machine given by a tape alphabet Σ with blank symbol $\sqcup \in \Sigma$, a state set Q with an initial state q_0, final states $F \subseteq Q$, and a transition function

$$\delta : \Sigma \times (Q \setminus F) \to \Sigma \times Q \times \{L, R\}.$$

Without loss of generality, we consider here only semi-infinite tape Turing machines, i.e. having a tape with a left end but no right end. This is no restriction for the complexity since semi-infinite tape Turing machines are equivalent to standard Turing machines [2, Claim 1.4]. The set of dominoes $\mathcal{D}$ is defined in Figure 8.5.

(i) An initial domino

$$\left[\frac{!}{! \star q_0 \star \sqcup \star ! \star}\right]$$

(ii) For every $x \in \Sigma$, a copy domino

$$\left[\frac{\star x}{x \star}\right]$$

(iii) Transitions $(q, x) \mapsto (\hat{q}, y, L)$

$$\left[\frac{\star x \star q}{\hat{q} \star y \star}\right]$$

(iv) Transitions $(q, x) \mapsto (\hat{q}, y, R)$

$$\left[\frac{\star q \star x}{y \star \hat{q} \star}\right]$$

(v) A tape expander

$$\left[\frac{\star !}{\sqcup \star ! \star}\right]$$

(vi) For every $q_f \in F, y_1, y_2 \in \Sigma$

$$\left[\frac{\star y_1 \star q_f \star y_2}{q_f \star}\right]$$

(vii) For every $q_f \in F, y_1, y_2 \in \Sigma$

$$\left[\frac{\star q_f \star y_1 \star y_2}{q_f \star}\right]$$

(viii) A final domino

$$\left[\frac{\star q_f \star \sqcup \star ! \star !}{!}\right]$$

Figure 8.5: The necessary dominoes for the reduction NHALT → PCP as well as BNHALT → BPCP.

Note that the domino set $\mathcal{D}$ can be constructed in polynomial time from T, and that $|\mathcal{D}|$ is polynomial in $|Q|$ and $|\Sigma|$.

Let us now apply this reduction to a non-deterministic Turing machine, as the bounded version needs the latter. First note that the exclamation marks serve as a separator between the instantaneous descriptions of different computation steps, while the grey star separates every symbol in the string. The lower part of the initial domino (i) represents the initial tape configuration of the Turing machine together with its current head state and position. Since the initial domino (i) is the only domino whose first upper and lower symbols coincide, every match has to start with the initial domino. A computation step along some computation path is simulated by applying copy-dominoes (ii), transition dominoes (iii), (iv), and tape expanders (v), according to Figure 8.4. If a computation reaches a final state q_f, the final instantaneous description is successively removed by applying dominoes (ii), (vi), (vii), and (v) according to Figure 8.4. Finally, a match is obtained by adding (viii).

This implies that T halts on the empty tape along a computation path if and only if $\mathcal{D}$ forms a match. Hence, $\mathcal{R}$: NHALT → PCP is a reduction. It follows that PCP is RE-hard.

Note that simulating the k^{th} computation step by a domino arrangement requires precisely $k+1$ dominoes. When T reaches the final state after n computation steps, the post-simulation procedure requires another $n+1$ repetitions, where each procedure needs precisely $m = n+1$ arrangements with length starting with m and decreasing by 1. So T halts after n computation steps on the empty tape if and only if the corresponding domino set forms a match in at most

$$q(n) := 1 + \sum_{k=1}^{n}(k+1) + \sum_{k=1}^{n+1} k = (n+1)\cdot(n+2)$$

steps, where the first sum represents the computation procedure and the second sum the post-simulation procedure. Since $\mathcal{R}$ is a polynomial-time reduction, using Theorem 8.1.1, this implies that

$$\langle T, 1^n \rangle \mapsto \langle \mathcal{R}(T), 1^{(n+1)\cdot(n+2)} \rangle$$

is a polynomial-time reduction from BNHALT to BPCP, which shows that BPCP is NP-hard.

The map $\mathcal{R}$ is a polynomial-time map; in particular, the number of dominoes k is polynomial in the description size of T. From the construction of $\mathcal{R}$ it follows that T halts on the empty tape if and only if there exists a match of dominoes $d_1, \dots, d_k$. This implies that $\mathcal{R}$ is a polynomial-time reduction from the non-deterministic halting problem, which implies that PCP is RE-hard.

Refining this argument and using Theorem 8.1.1, we obtain that $\mathcal{R}$ can be used as a reduction from BNHalt to BPcp. Each computation step of T on the empty tape is simulated by attaching dominoes, as shown in Figure 8.4. This procedure guarantees that T halts within n steps if and only if $d_1, \ldots, d_k$ form a match within

$$p(n) := (n+1) \cdot (n+2)$$

steps. Hence, the halting time of T is polynomially related to the length of a minimal match of $\mathcal{R}(T)$. This proves that BPcp is NP-hard by Theorem 8.1.1.

Moreover, Pcp is RE-complete and BPcp is NP-complete, by taking matching domino arrangements as certificates, and a polynomial-time verifier that checks arrangements.

8.3.2 The zero in the upper left corner and the matrix mortality problem

We now present the matrix mortality problem (short Mm) and the zero in the upper left corner problem (short Zulc) with their bounded versions. Both problems are undecidable and have been applied to prove the undecidability of quantum information problems such as the positivity of Matrix product operators [36] (see Section 8.3.3), the reachability problem [129] (see Section 8.3.6), or the measurement occurrence problem [46].

Problem 8.3.3 (The matrix mortality problem)
Given $A_1, \ldots, A_k \in \mathrm{Mat}_d(\mathbb{Q})$, is there a finite sequence $i_1, \ldots, i_\ell \in \{1, \ldots, k\}$ such that

$$A_{i_1} \cdot A_{i_2} \cdots A_{i_\ell} = \mathbf{0}\ ?$$

Here, $\mathbf{0}$ denotes the zero matrix, and $\mathrm{Mat}_d(\mathbb{Q})$ the set of $d \times d$ matrices over $\mathbb{Q}$. Zulc is almost identical to Mm, the only difference is that only the upper left corner of the product $A_{i_1} \cdot A_{i_2} \cdots A_{i_\ell}$ is asked to be zero. We define the bounded matrix mortality problem (BMm) and the bounded zero in the upper left corner problem (BZulc) by adding a parameter $n \in \mathbb{N}$ to every instance, and asking whether the desired zeros can be realized within n matrix multiplications.

The undecidability of Mm was first proven by Paterson [95]. Since then, many tighter bounds on the number and size of matrices for both problems have been found (see [24] and references therein). It is also known that BMm is NP-hard [17]. However, the proof relies on a reduction from the NP-complete problem SAT and is therefore independent of the original reduction proving undecidability. To the best of our knowledge, the following is the first proof of the NP-hardness of these bounded matrix problems using the same reductions as their unbounded versions.

We briefly sketch the reductions. Following [63], there exist polynomial-time reductions $\mathcal{R}$: Pcp $\to$ Zulc and $\mathcal{Q}$: Zulc $\to$ Mm with the following properties:

(i) The dominoes $\mathbf{d} := \langle d_1, \ldots, d_k \rangle$ form a match of length n if and only if the matrices

$$\langle N_1, \ldots, N_{k'} \rangle := \mathcal{R}(\mathbf{d})$$

multiply to a matrix with a zero in the upper left corner within n matrix multiplications.

(ii) The matrices $\mathbf{N} := \langle N_1, \ldots, N_\ell \rangle$ form a zero in the upper left corner using n matrix multiplications if and only if the matrices

$$\langle M_1, \ldots, M_{\ell'} \rangle := \mathcal{Q}(\mathbf{N})$$

multiply to a zero matrix within $n + 2$ matrix multiplications.

Together with Theorem 8.1.1, these observations show that

$$\langle x, 1^n \rangle \mapsto \langle \mathcal{R}(x), 1^n \rangle$$

is a polynomial-time reduction from BPcp to BZulc, and

$$\langle x, 1^n \rangle \mapsto \langle \mathcal{Q}(x), 1^{n+2} \rangle$$

is a polynomial-time reduction from BZulc to BMm. This proves that BZulc and BMm are NP-hard.

The Reduction to the Zero-in-the-upper-left-corner problem

Let us now present the reduction $\mathcal{R} : \text{Pcp} \to \text{Zulc}$ based on the ideas of [63] in greater detail. For this purpose, we consider Pcp using strings encoded in the alphabet $\Sigma = \{0, 1, 2\}$. We define the bijection $\sigma : \Sigma^* \to \mathbb{N}$ that assigns a representation in base 3 to every natural number, i.e.

$$\sigma(c_1, \ldots, c_n) := \sum_{i=1}^{n} c_i \cdot 3^{n-i}.$$

Moreover, we define a function $\gamma : \Sigma^* \times \Sigma^* \to \mathbb{N}^{3\times 3}$ via

$$\gamma(w_1, w_2) := \begin{pmatrix} 3^{|w_1|} & 0 & 0 \\ 0 & 3^{|w_2|} & 0 \\ \sigma(w_1) & \sigma(w_2) & 1 \end{pmatrix}.$$

The function γ is injective and a morphism, i.e. $\gamma(w_1u_1, w_2u_2) = \gamma(w_1, w_2) \cdot \gamma(u_1, u_2)$ where composition on Σ^* is given by concatenation of words. Let

$$d_1 = \left[\frac{a_1}{b_1}\right], \ldots, d_k = \left[\frac{a_k}{b_k}\right]$$

be an instance of Pcp where $a_i, b_i \in \Sigma^*$. For $i \in \{1, \ldots, k\}$, we define the matrices

$$A_i = X \cdot \gamma(a_i, b_i) \cdot X^{-1} \quad B_i = X \cdot \gamma(a_i, 0b_i) \cdot X^{-1}$$

with

$$X = \begin{pmatrix} 1 & 0 & 1 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

We have that

$$d_{i_1} d_{i_2} \cdots d_{i_n}$$

is a matching domino if and only if

$$(M_{i_1} \cdot M_{i_2} \cdots M_{i_n})_{11} = 0$$

where $M_{i_j} \in \{A_{i_j}, B_{i_j}\}$. We refer to [63] for details. This shows that $\mathcal{R} : \textsc{Pcp} \to \textsc{Zulc}$ with

$$\mathcal{R}\Big(\langle d_1, \ldots, d_k \rangle\Big) := \langle A_1, \ldots, A_k, B_1, \ldots, B_k \rangle$$

is a polynomial-time reduction. This implies that Zulc is RE-hard.

Since matches of length n are mapped to matrix multiplications of length n with a zero in the upper left corner, this shows that $\mathcal{R}_b : \text{BP}\textsc{cp} \to \text{BZ}\textsc{ulc}$ with

$$\mathcal{R}_b\Big(\langle d_1, \ldots, d_k, 1^n \rangle\Big) := \langle A_1, \ldots, A_k, B_1, \ldots, B_k, 1^n \rangle$$

is a polynomial-time reduction. This implies that BZulc is NP-hard.

Note that the matrices in $A_1, \ldots, A_k, B_1, \ldots, B_k$ are invertible, from which it follows that Zulc and BZulc remain RE-hard and NP-hard, respectively, when restricting the instances to invertible matrices.

The Reduction to the Matrix Mortality problem

We now construct the reduction $\mathcal{Q} : \textsc{Zulc} \to \textsc{Mm}$ following the ideas of [63]. Since Zulc remains hard when restricting the instances to invertible matrices, we construct $\mathcal{Q}$ only for invertible matrices. So let $\langle A_1, \ldots, A_k \rangle$ be an instance of invertible matrices in Zulc. We define

$$\mathcal{Q}(\langle A_1, \ldots, A_k \rangle) := \langle A_1, \ldots, A_k, B \rangle$$

with

$$B = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}.$$

We claim that $A_1, \ldots, A_k$ forms a zero in the upper left corner if and only if $A_1, \ldots, A_k, B$ multiplies to a zero matrix. This proves that Mm is RE-hard. Moreover, we show that

$$n_{\min,\textsc{Mm}}[\langle \mathbf{A}, B \rangle] = n_{\min,\textsc{Zulc}}[\langle \mathbf{A} \rangle] + 2. \tag{8.3}$$

where $\mathbf{A}$ represents the list $A_1, \ldots, A_k$.

To prove the claim, first note that if

$$(A_{i_1} \cdot A_{i_2} \cdots A_{i_n})_{11} = 0,$$

then

$$B \cdot A_{i_1} \cdot A_{i_2} \cdots A_{i_n} \cdot B = (A_{i_1} \cdot A_{i_2} \cdots A_{i_n})_{11} = 0.$$

In other words, a yes-instance of ZULC with parameter n is mapped to a yes-instance in MM with parameter $n+2$. This proves the inequality "$\leqslant$" of Equation (8.3).

Conversely, assume that there exists a sequence of n matrices in $\{A_1, \ldots A_k, B\}$ that multiplies to $\mathbf{0}$. Since $A_1, \ldots, A_k$ are invertible and B has rank 1, this sequence must contain B at least twice. The product is of the form

$$M_1 B M_2 B M_3 B \cdots B M_r = \mathbf{0}$$

where M_i is a multiplication of ℓ_i matrices in $\{A_1, \ldots, A_k\}$ for some ℓ_i.[6] Since B is idempotent, we have that

6: If it is an empty multiplication (i.e. $\ell_i = 0$), then we define M_i as the identity matrix.

$$\begin{aligned} 0 &= \left(M_1 B M_2 B M_3 B \cdots B M_r\right)_{11} \\ &= \left(B M_1 B^2 M_2 B^2 M_3 B^2 \cdots B^2 M_r B\right)_{11} \\ &= \left(M_1\right)_{11} \cdots \left(M_r\right)_{11}. \end{aligned}$$

This implies that at least one of the matrices M_i has a zero in the upper left corner, which shows that $A_1, \ldots, A_k$ form a zero in the upper left corner with a word of length n. Specifically, any minimal sequence of matrices realizing $\mathbf{0}$ must be of the form

$$B \cdot A_{i_1} \cdot A_{i_2} \cdots A_{i_n} \cdot B = \mathbf{0}.$$

Note that a shorter such product cannot exist because it would violate the proven inequality "$\leqslant$" of Equation (8.3). This representation proves the inequality "$\geqslant$" of Equation (8.3), since

$$(A_{i_1} \cdot A_{i_2} \cdots A_{i_n})_{11} = 0.$$

In summary, $\mathcal{Q}$: ZULC $\to$ MM is a reduction, which proves that MM is RE-hard. Moreover, $\mathcal{Q}_b$: BZULC $\to$ BMM with

$$\mathcal{Q}_b : \langle A_1, \ldots, A_k, 1^n \rangle \mapsto \langle A_1, \ldots, A_k, B, 1^{n+2} \rangle$$

is a polynomial-time reduction too, which proves that BMM is NP-hard.

Let us finally note that MM and ZULC are RE-complete, and their bounded versions, BMM and BZULC, are NP-complete by taking matching matrix arrangements as certificates and a polynomial-time verifier checking the statements.

8.3.3 The matrix product operator positivity problem

A Matrix Product operator (MPO) representation is a decomposition of a multipartite operator into local tensors according to a one-dimensional structure (see Section 2.3.5). A local tensor B defines a diagonal operator $\rho_n(B)$ for every system size n (see Figure 8.6). More precisely, given a family of $D \times D$ matrices (B_i) for $i \in \{1, \ldots, d\}$, the diagonal ti MPO of size n is given by

$$\rho_n(B) := \sum_{j_1,\dots,j_n=1}^{d} \operatorname{tr}\left(B_{j_1}\cdots B_{j_n}\right) |j_1,\dots,j_n\rangle \langle j_1,\dots,j_n| \,.$$

If these MPO should represent density matrices, then B should be such that $\rho_n(B)$ is psd for every n. This property cannot be decided algorithmically, not even for classical states. In other words, the following Mpo problem is undecidable:

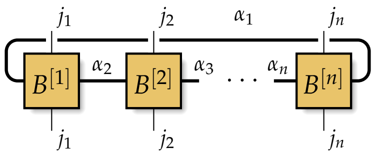

Figure 8.6: Tensor network representation of the MPO $\rho_n(B)$. The MPO problem asks: Given a tensor B, is $\rho_n(B)$ psd for all n? Note that in this setting the tensors B have only one open index in contrast to Figure 2.11.

Problem 8.3.4 (The MPO positivity problem)
Given $B_1,\dots,B_k \in \mathrm{Mat}_D(\mathbb{Q})$, is there $n \in \mathbb{N}$ such that $\rho_n(B)$ is not psd?

Note that an MPO is usually defined more generally; instead of restricting to families of diagonal (classical) matrices B_i, a general matrix product operator is defined via families of $D \times D$ matrices $(B_{i,j})$ for $i, j = 1,\dots,d$, addressing also non-diagonal entries of the matrix. However, as diagonal MPOs are contained in this definition, the undecidability of Mpo as we defined it implies that the same problem for arbitrary matrix product operators is also undecidable.

Similar to previous bounded versions, we define BMpo by bounding the system size n:

Problem 8.3.5 (The bounded MPO positivity problem)
Given $B_1,\dots,B_k$ and $n \in \mathbb{N}$, is there an $\ell \leqslant n$ such that $\rho_\ell(B)$ is not psd?

Note that Mpo is usually stated in the negated way; yet, we use this definition for consistency with the definition of bounding.

Let us now present a reduction $\mathcal{R}$: Zulc $\to$ Mpo, slightly different than [36]. The Mpo problem has as input a fixed number of $D \times D$ integer matrices $\langle B_i : i \in \{1,\dots,k\}\rangle$ and asks whether there exists a natural number $n \in \mathbb{N}$ such that

$$\rho_n(B) := \sum_{i_1,\dots,i_n=1}^{k} \operatorname{tr}\left(B_{i_1}\cdots B_{i_n}\right) |i_1\dots i_n\rangle \langle i_1\dots i_n|$$

is not psd. We define

$$\mathcal{R}(\langle A_1,\dots,A_k\rangle) = \langle B_1,\dots,B_k,B_{k+1}\rangle$$

where for $i \in \{1,\dots,k\}$

$$B_i := \begin{pmatrix} A_i \otimes A_i & 0 \\ 0 & 1 \end{pmatrix}$$

and

$$B_{k+1} := \begin{pmatrix} E_{11} & 0 \\ 0 & -1 \end{pmatrix}$$

where $E_{11} := |1\rangle\langle 1|$ with $|1\rangle = (1, 0, \dots, 0)^t$ of length D.

We now prove that the threshold parameter n in BZULC maps to the threshold parameter $n+1$ in BMPO. Let $A_{i_1}, \dots, A_{i_n}$ be the minimal sequence such that

$$\left(A_{i_1} \cdot A_{i_2} \cdots A_{i_n}\right)_{11} = 0.$$

Then,

$$\operatorname{tr}(B_{i_1} \cdots B_{i_n} \cdot B_{k+1}) = \left(A_{i_1} \cdot A_{i_2} \cdots A_{i_n}\right)_{11}^2 - 1 < 0.$$

Conversely, let $B_{i_1}, \dots, B_{i_{n+1}}$ be a minimal sequence such that

$$\operatorname{tr}(B_{i_1} \cdot B_{i_2} \cdots B_{i_{n+1}}) < 0.$$

The indices $i_1, \dots, i_{n+1}$ cannot be chosen exclusively from $\{1, \dots, k\}$, since in that case

$$\operatorname{tr}(B_{i_1} \cdot B_{i_2} \cdots B_{i_{n+1}}) = \left(\operatorname{tr}(A_{i_1} \cdots A_{i_{n+1}})\right)^2 + 1 \geqslant 0.$$

Hence, there is at least one index $i_\ell = k+1$. Assume that there is precisely one index $k+1$. Without loss of generality, we assume $i_{n+1} = k+1$ due to cyclicity of the trace. This leads to

$$0 > \operatorname{tr}(B_{i_1} \cdot B_{i_2} \cdots B_{i_{n+1}}) = \left(\left(A_{i_1} \cdot A_{i_2} \cdots A_{i_n}\right)_{11}\right)^2 - 1$$

which implies that $\left(A_{i_1} \cdot A_{i_2} \cdots A_{i_n}\right)_{11} = 0$ because the entries are integer. This shows that a threshold parameter $n+1$ in BMPO maps to a threshold parameter of a most n in BZULC. Note that having multiple indices with $k+1$ leads to a smaller threshold parameter in BZULC which contradicts the minimality assumption of $B_{i_1}, \dots, B_{i_{n+1}}$. This proves the statement.

This reduction can easily be extended to matrices with rational numbers.

In summary, $\mathcal{R}\colon$ ZULC $\to$ MPO is a reduction, which proves that MPO is RE-hard. Moreover, by Theorem 8.1.1, $\mathcal{R}_b\colon$ BZULC $\to$ BMPO with

$$\mathcal{R}_b : \langle A_1, \dots, A_k, 1^n \rangle \mapsto \langle B_1, \dots, B_k, B_{k+1}, 1^{n+1} \rangle$$

is a polynomial-time reduction too, which proves that BMPO is NP-hard.

Moreover, MPO is RE-complete and BMPO is NP-complete by defining negative diagonal entries as certificates.

While MPO precisely characterizes psd matrix product operators, in practice, algorithms distinguishing MPOs that are sufficiently positive or that violate positivity by at least an error $\varepsilon > 0$ are often acceptable. This is the idea behind weak membership problems. Along these lines, we define the approximate MPO problem MPO_ε as follows:

Problem 8.3.6 (The approximate positivity problem for MPO)
Given $C_1, \dots, C_k \in \operatorname{Mat}_D(\mathbb{Q})$ with $\operatorname{tr}(\rho_\ell(C)) \leqslant 1$ for every $\ell \in \mathbb{N}$ and a family of errors $(\varepsilon_\ell)_{\ell \in \mathbb{N}}$ with $0 < \varepsilon_\ell \leqslant 1/\exp(\ell)$. Decide the

following:

(a) Accept if $\exists n \in \mathbb{N} : \rho_n(C) + \varepsilon_n \mathbb{1}$ is not psd.
(b) Reject if $\forall n \in \mathbb{N} : \rho_n(C) - \varepsilon_n \mathbb{1}$ is psd.

$\text{M}\textsc{po}_\varepsilon$ is undecidable using the same reduction as above and the fact that $\mathrm{tr}(\rho_n(C))$ increases exponentially in n in the above reduction. Following the usual bounding process, we define $\text{BM}\textsc{po}_\varepsilon$ by bounding n:

Problem 8.3.7 (The bounded approximate positivity problem for MPO)
Given $C_1, \dots, C_k \in \mathrm{Mat}_D(\mathbb{Q})$ with $\mathrm{tr}(\rho_\ell(C)) \leqslant 1$ for every $\ell \in \mathbb{N}$, a family of errors $(\varepsilon_\ell)_{\ell \in \mathbb{N}}$ with $0 < \varepsilon_\ell \leqslant 1/\exp(\ell)$ and $n \in \mathbb{N}$. Decide the following:

(a) Accept if $\exists \ell \leqslant n : \rho_\ell(C) + \varepsilon_n \mathbb{1}$ is not psd.
(b) Reject if $\forall \ell \leqslant n : \rho_\ell(C) - \varepsilon_n \mathbb{1}$ is psd.

It follows that $\text{BM}\textsc{po}_\varepsilon$ is a bounded version of $\text{M}\textsc{po}_\varepsilon$ according to Definition 8.1.1. Moreover, Theorem 8.1.1 implies that $\text{BM}\textsc{po}_\varepsilon$ is also NP-hard.

We remark that Kliesch et al. [72] present a similar idea, by constructing a reduction from Pcp to an alternative version of Mpo and bounding both problems.

8.3.4 The polynomial positivity problem

The undecidability of Mpo leads to the undecidability of other positivity problems. One of them concerns deciding the positivity of a certain class of polynomials (see Section 7.3.3 and [39]):

Problem 8.3.8 (Polynomial positivity problem)
Given a family of polynomials $q_{\alpha,\beta}(\mathbf{x})$ for $\alpha, \beta \in \{1, \dots, D\}$ with integer coefficients, is there an $n \in \mathbb{N}$ such that the polynomial

$$p_n(\mathbf{x}^{[1]}, \dots, \mathbf{x}^{[n]}) := \sum_{\alpha_1, \dots, \alpha_n = 1}^{D} q_{\alpha_1,\alpha_2}(\mathbf{x}^{[1]}) \cdots q_{\alpha_n,\alpha_1}(\mathbf{x}^{[n]}) \tag{8.4}$$

is not nonnegative (i.e. $p_n(\mathbf{a}) < 0$ for some $\mathbf{a} \in \mathbb{R}^{d \cdot n}$)?

Here $\mathbf{x}^{[i]}$ denotes a d-tuple of variables, for every i. We define this problem as Poly and its bounded version (by restricting to checking nonnegativity of p_k for $k \leqslant n$) by BPoly.

We have that Poly is RE-hard by Theorem 7.3.8. Following the proof of Theorem 7.3.8, there exists a polynomial-time map

$$\mathcal{R}(\langle B_1, \dots, B_k \rangle) := \langle q_{\alpha,\beta} : \alpha, \beta = 1, \dots, D \rangle$$

such that

$$\rho_n(B) \geqslant 0 \text{ if and only if } p_n \text{ is nonnegative.}$$

This implies that $\langle B, 1^n \rangle \mapsto \langle \mathcal{R}(B), 1^n \rangle$ defines a reduction from BMPO to BPOLY. It follows that BPOLY is NP-hard.

Moreover, the threshold n for BMPO is mapped to the threshold n for BPOLY. It follows that BPOLY is NP-hard. Hence, BPOLY is NP-complete by taking an arrangement of the matrices leading to a negative value as a certificate, and a polynomial-time verification procedure of this statement as a verifier.

8.3.5 Stability of positive maps

Another undecidable problem related to positivity concerns tensor products of positive maps. A map

$$\mathcal{P} : \mathrm{Mat}_d(\mathbb{C}) \to \mathrm{Mat}_d(\mathbb{C})$$

is called *positive* if it maps psd matrices to psd matrices. Such a map is called *n-tensor-stable positive* if $\mathcal{P}^{\otimes n}$ is a positive map, and *tensor-stable positive* if it is n-tensor-stable positive for all $n \in \mathbb{N}$. The existence of non-trivial tensor-stable positive maps relates to the existence of NPT bound-entangled states [87].

Let us define the n-fold Matrix Multiplication tensor[7] as

7: We refer to Section 2.3.4 for its relation to the structure tensors on weighted simplicial complexes.

$$|\chi_n\rangle := \sum_{\alpha_1,\ldots,\alpha_n=1}^{s} |\alpha_1, \alpha_2\rangle \otimes |\alpha_2, \alpha_3\rangle \otimes \cdots \otimes |\alpha_n, \alpha_1\rangle$$

and denote the projection to this vector by

$$\chi_n := |\chi_n\rangle \langle \chi_n| \,. \tag{8.5}$$

The following problem is undecidable [48]:

Problem 8.3.9 (Positivity on a state problem)
Given a positive map $\mathcal{P} : \mathrm{Mat}_d(\mathbb{C}) \to \mathrm{Mat}_d(\mathbb{C})$, is $\mathcal{P}^{\otimes n}(\chi_n)$ not psd for some $n \in \mathbb{N}$?

We denote this problem by TSP. Its bounded version, BTSP takes instances $\langle \mathcal{P}, 1^n \rangle$ and asks the same question for k-fold tensor products with $k \leqslant n$.

Let us now review the reduction $\mathcal{R}$: MPO → TSP of [48], which proves that TSP is RE-hard. The same reduction also yields that BTSP is NP-hard.

We map an instance

$$\langle B_1, \ldots, B_k \rangle \in \mathrm{Mat}_{D^2}(\mathbb{Q}) \cong \mathrm{Mat}_D(\mathbb{Q}) \otimes \mathrm{Mat}_D(\mathbb{Q})$$

of MPO to a linear map

$$\begin{aligned} \mathcal{P} : \quad \mathrm{Mat}_D(\mathbb{Q}) \otimes \mathrm{Mat}_D(\mathbb{Q}) \quad &\to \quad \mathrm{Mat}_k(\mathbb{Q}) \\ X \quad &\mapsto \quad \sum_{i=1}^{k} |i\rangle \langle i| \operatorname{tr}(C_i X) \end{aligned}$$

where

$$(C_i)_{(\alpha_1,\alpha_2),(\beta_1,\beta_2)} := (B_i)_{(\alpha_1,\beta_1),(\alpha_2,\beta_2)}$$

with $\alpha_1, \alpha_2, \beta_1, \beta_2 \in \{1, \dots, D\}$. Then, we have that

$$\mathrm{tr}\left(C_{i_1} \otimes \cdots \otimes C_{i_n} \chi_n\right) = \mathrm{tr}(B_{i_1} \cdots B_{i_n})$$

where χ_n is defined in (8.5). By construction, this implies that

$$\mathcal{P}^{\otimes n}(\chi_n) = \rho_n(B).$$

In summary, $\langle B_1, \dots, B_k \rangle \in$ MPO if and only if exists $n \in \mathbb{N}$ such that $\mathcal{P}^{\otimes n}(\chi_n)$ is not psd. Furthermore, the threshold parameters in both problems coincide for this reduction. It follows that BTSP is NP-hard.

8.3.6 The reachability problem in quantum information

The reachability problem in quantum information concerns the question whether a resource state ρ (given as a density matrix) can be converted to another state σ by using only free resource operations from a fixed set $\mathcal{F} := \{\Phi_1, \dots, \Phi_k\}$. More precisely, we define REACH as follows:

Problem 8.3.10 (Reachability in resource theories)
Given density matrices $\rho, \sigma \in \mathrm{Mat}_d(\mathbb{C})$ and a set $\mathcal{F}$ of free operations $\mathrm{Mat}_d(\mathbb{C}) \to \mathrm{Mat}_d(\mathbb{C})$, is there a map

$$\Phi := \Phi_{i_n} \circ \Phi_{i_{n-1}} \circ \cdots \circ \Phi_{i_1}$$

in the free semigroup $\mathcal{F}^*$ such that $\sigma = \Phi(\rho)$?

The free semigroup $\mathcal{F}^*$ of $\mathcal{F}$ consists of all maps generated by finite compositions of maps in $\mathcal{F}$. We denote by $\mathcal{F}^n$ the set of all operations arising from at most n compositions of maps in $\mathcal{F}$, and define the bounded version BREACH by replacing $\mathcal{F}^*$ with $\mathcal{F}^n$ in the above problem statement.

REACH is undecidable via a reduction from PCP [107]. We now prove that the bounded version BREACH is NP-hard. We rely on Scandi and Surace's work [107], who provide a polynomial-time reduction $\mathcal{R}$ mapping dominoes d_i to two types of resource maps H_i^λ, G_i^λ for $\lambda \in (0, 1)$. The set of free resource operations is then specified by

$$\mathcal{F} = \left\{\mathbb{1}, H_i^\lambda, G_i^\lambda : i = 1, \dots, r \text{ and } \lambda \in (0,1)\right\}.$$

For a state $\rho \in \mathrm{Mat}_4(\mathbb{C})$, it is shown that

$$\sigma := \lambda\rho + (1-\lambda)\frac{\mathbb{1}}{4}$$

is reachable via operations in $\mathcal{F}^*$ if and only if there exists a match of the corresponding dominoes in PCP. This shows that REACH is RE-hard. More specifically, there exists a match of length n if and only if

$$\sigma = G_{i_n}^{\lambda_n} \circ \cdots \circ G_{i_1}^{\lambda_1} \circ H_{i_1}^{\lambda_1} \circ \cdots \circ H_{i_n}^{\lambda_n}(\rho)$$

for a choice $\lambda_1, \ldots, \lambda_n \in (0,1)$. In other words, a threshold parameter n in BPcp is mapped to a threshold $2n$ in BReach. This proves that BReach is NP-hard by applying Theorem 8.1.1.

8.3.7 The tiling problem

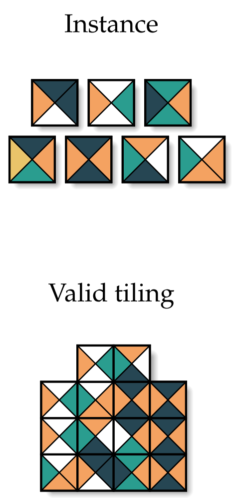

Figure 8.7: An instance of Tile is a set of tiles (top). A set of tiles is a yes-instance if there exists a valid tiling of the plane. Part of a potentially valid tiling is shown on the right. In a valid tiling, the colors of adjacent tiles must coincide and the tiles cannot be rotated.

Let us now consider the Wang tiling problem. This problem has been used to prove undecidability in many physics-related problems, like the spectral gap problem in 2D [34], 2D PEPS problems [108], or the universality of translational invariant, classical spin Hamiltonians in 2D [75].

A tile is given by a square with different colors on each side of the tile (see Figure 8.8). Given a finite set of tiles, a valid tiling is an arrangement of tiles whose adjacent edges coincide. Moreover, all tiles have a fixed orientation, i.e. they cannot rotate. We study the following variant:

> **Problem 8.3.11** (The tiling problem)
>
> Given a set of tiles $\mathcal{T} = \{t_1, \ldots, t_k\}$, is it impossible to tile the plane when t_1 is in the origin?

Note that this problem is usually stated in the negated form, but this formulation is more convenient for our purposes. The constraint on the fixed tile in the origin can also be removed [11, 104]; we stick to this version for simplicity. The corresponding bounded version is the following:

> **Problem 8.3.12**
>
> Given a set of tiles $\mathcal{T} = \{t_1, \ldots, t_k\}$ and $n \in \mathbb{N}$, is it impossible to tile $\mathbb{Z}_n^2$ when t_1 is in the origin?

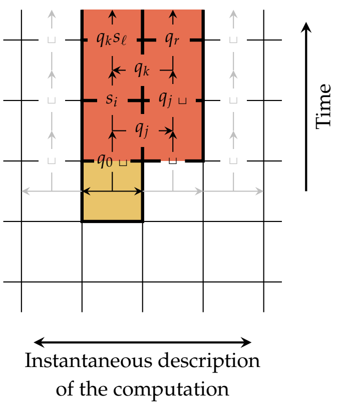

Figure 8.8: In the reduction NHaltAll → Tile, the instantaneous description of the Turing machine is mapped to a horizontal configuration of tiles, and every computational step is mapped to a valid tiling of the horizontal line above. The green tile is fixed at the origin, while the orange tiles realize the computation. The rest of the plane is filled with trivial tiles, such as the empty tiles (bottom) or tiles copying the tape information (left and right). A Turing machine halts along every path within n steps if and only if the corresponding tiling terminates after n horizontal lines.

Here we denote by $\mathbb{Z}_n^2 := \{-n, \ldots, 0, \ldots, n\}^2$ the square grid of size $(2n+1) \times (2n+1)$ around the origin.

Let us now sketch the proof that Tile is RE-hard and that BTile is coNP-hard. This will imply that the tiling problem in its usual formulation ("can the plane be tiled?") is coRE-hard and its bounded version is NP-hard.

In contrast to the previous examples, we now construct a reduction from NHaltAll instead of NHalt. While to check whether $\{d_1, \ldots, d_k\}$ is a yes-instance of BPcp, one needs to find a *single* matching arrangement, to verify whether $\{t_1, \ldots, t_k\}$ is a yes-instance of BTile one has to check (for a fixed size n) whether *all* arrangements of tiles in $\mathbb{Z}_n^2$ are invalid. This structure is similar to NHaltAll, where for a fixed computation time n, one needs to check whether a given Turing machine T halts on *all* computation steps. More precisely, there is a polynomial relation between the bounding parameters of BTile and BNHaltAll, as needed for Theorem 8.1.1.

We build a polynomial-time reduction from NHaltAll to Tile following [104]. The reduction maps a description of a Turing machine T to a set of tiles representing either a slot in the tape or a computational step. The (infinite) starting tape is mapped to the fixed origin tile representing the empty tape with head position at zero. Filling up a new line corresponds to one computational step. This reduction also applies to non-deterministic Turing machines.

The reduction is such that the tiling cannot be continued after filling up n lines if and only if T halts on all computation paths after at most n computational steps, see Figure 8.8. This proves that TILE is undecidable. By Theorem 8.1.1, we obtain that BTILE is coNP-hard, since the maximal halting time n on every computation path is mapped to the termination size $n+1$.

In addition, TILE is RE-complete by taking a system size where all tilings terminate as a certificate and an exponential-time verifier checking all tilings of this size. BTILE is coNP-complete by choosing tilings as a certificate and a polynomial-time verifier checking the validity of the tiling. This highlights that when proving completeness, *not* every construction in the unbounded case trivially translates to the bounded version.

Let us now review the reduction $\mathcal{R}$: HALT $\to$ TILE from [104]. A Turing machine, consisting of a tape alphabet Σ with blank symbol $\sqcup \in \Sigma$, a state set Q with an initial state q_0 and final states $F \subseteq Q$, and a transition function

$$\delta : \Sigma \times (Q \setminus F) \to \Sigma \times Q \times \{L, R\}$$

is mapped to the set of tiles shown in Figure 8.9.

This set of tiles captures the computation of a Turing machine on the empty tape when placing the initial tile to the origin (see Figure 8.8). The initial tile can only be extended to the left and to the right with the empty tape extension. We can also trivially tile the whole lower half of the plane by applying the empty tile.

The generated string

$$\cdots \quad \sqcup \quad \sqcup \quad \sqcup \quad q_0\sqcup \quad \sqcup \quad \sqcup \quad \sqcup \quad \cdots$$

at the top of the first line represents the instantaneous description of the Turing machine at time 0, namely an empty tape with the head at position 0 and state q_0. Simulating one step of the Turing machine corresponds to filling up the line above of the current one. Specifically, on the top of the initial tile, we need to place a transition tile $(q_0, \sqcup) \mapsto (\hat{q}, x, L/R)$. Then we need to place a state merge tile on the left/right of the transition tile. This reflects the movement of the head to the left or right. The rest of the line is filled with copy tiles.

Again, the string at the top of the second line represents the initial description after one computation step. The same procedure applies to every computation step. As soon as we apply a transition tile $(q, x) \mapsto (q_f, y, L/R)$ for some final state $q_f \in F$, there is no tile to continue the tiling procedure. In other words, every tiling procedure terminates in line n if and only if T halts on the empty tape.

The same reduction applies to non-deterministic Turing machines. In this situation, every tiling procedure terminates in n lines if and only if the Turing machine halts on the empty tape along every computation path in at most n steps. In other words, a Turing machine T halts on every path in at most n steps if and only if $\mathbb{Z}_{n+1} \times \mathbb{Z}_{n+1}$ cannot be tiled. This

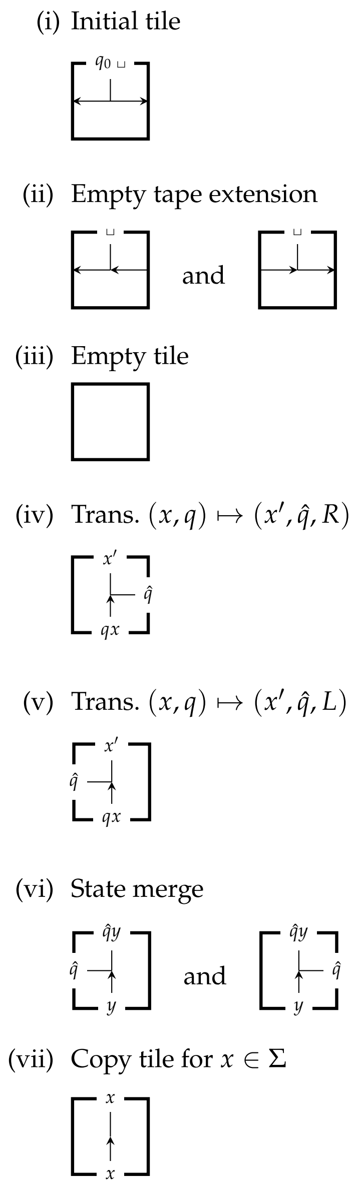

Figure 8.9: The necessary tiles for the reduction NHALTALL $\to$ TILE. State merge is defined for every $y \in \Sigma$ and $\hat{q} \in Q$, whereas transitions are defined for every such transition δ.

proves that $\mathcal{R}$: NHaltAll → Tile is a reduction. It follows that Tile is RE-hard.

Moreover, $\mathcal{R}$ is a polynomial-time map. Since the map between the threshold parameters of NHaltAll and Tile is given by $n \mapsto n+1$,

$$\langle x, 1^n \rangle \mapsto \langle \mathcal{R}(x), 1^{n+1} \rangle$$

is a reduction from BNHaltAll to BTile. This implies that BTile is coNP-hard.

8.3.8 Ground state energy problem

We now study a version of the ground state energy problem. For this purpose, we consider a spin system on a 2D grid. We assume that every spin takes values in a set $\mathcal{S}$. Given coupling functions $h^x, h^y : \mathcal{S} \times \mathcal{S} \to \mathbb{N}$ and a local field $h^{\text{loc}} : \mathcal{S} \to \mathbb{N}$, we define the Hamiltonian

$$H_n(\mathbf{s}) = h^{\text{loc}}(s_{00}) + \sum_{\langle \mathbf{a},\mathbf{b} \rangle_x} h^x(s_{\mathbf{a}}, s_{\mathbf{b}}) + \sum_{\langle \mathbf{a},\mathbf{b} \rangle_y} h^y(s_{\mathbf{a}}, s_{\mathbf{b}})$$

where $\mathbf{s} = (s_{ij})_{i,j \in \{-n,\dots,0,\dots,n\}}$ is a given spin configuration on the grid $\mathbb{Z}_n^2$ taking values in $\mathcal{S}$ and $s_{\mathbf{a}}, s_{\mathbf{b}}$ denote the elements with coordinates **a** and **b** in this array. Moreover, $\langle \mathbf{a}, \mathbf{b} \rangle_{x/y}$ denotes all neighbors in x/y-direction on $\mathbb{Z}_n^2$ where the **a** has a smaller x/y-coordinate than **b**. Hence, H_n is translational invariant except for the local field on the spin in the origin.

We start by defining the bounded version of this problem, namely the bounded ground state energy problem BGse:

Problem 8.3.13 (The bounded ground state energy problem)
Given system size $n \in \mathbb{N}$, non-negative functions h^x, h^y, h^{loc} and energy $E \in \mathbb{Q}$, is the ground state energy $E_{\min}(H_n) > E$?

A function h is non-negative if it is non-negative on its whole domain. Note that BGse is indeed a bounded version, as $E_{\min}(H_{n+1}) \geqslant E_{\min}(H_n) > E$ since all couplings are non-negative. Further note that BGse is usually formulated in the negated way, i.e. the question is if there exists a spin configuration whose energy is below the threshold E.

We now extend BGse to an unbounded ground state energy problem Gse:

Problem 8.3.14 (The (unbounded) ground state energy problem)
Given non-negative functions h^x, h^y, h^{loc} and an energy $E \in \mathbb{Q}$, is there an $n \in \mathbb{N}$ such that $E_{\min}(H_n) > E$?

Note that BGse is the bounded version of Gse according to Definition 8.1.1.

Let us show that Gse is RE-hard and BGse is coNP-hard by a reduction $\mathcal{R}$: Tile → Gse (see Figure 8.10). Given a set of tiles $\mathcal{T} = \{t_1, \dots, t_k\}$, we define the set of spin states as the set of tiles $\mathcal{S} := \mathcal{T}$. Since each tile

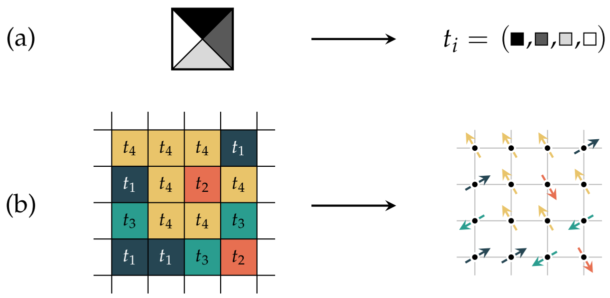

Figure 8.10: In the reduction Tile $\to$ Gse, (a) every tile t_i is mapped to a spin state s_i. (b) Every (valid and invalid) tiling maps to a spin configuration. A tiling of size n is valid iff the corresponding spin configuration is the ground state of H_n with energy 0.

is specified by four colors in a color space C, it can be represented as a 4-tuple

$$t_i = \left(t_i^N, t_i^E, t_i^S, t_i^W\right)$$

where the entries represent the colors on the top, right, bottom, and left of the tile. We define the coupling function so that a valid tiling with t_1 in the origin maps to a spin configuration of energy 0, and every inconsistent color pairing in an invalid tiling gives an additional energy penalty of 1. More precisely,

$$h^x(s,\hat{s}) := 1 - \delta(s^E, \hat{s}^W) \quad \text{and} \quad h^y(s,\hat{s}) := 1 - \delta(s^N, \hat{s}^S).$$

where $s, \hat{s} \in \mathcal{S}$. According the definition of H_n, the first component of h^x addresses the spin on the left and the second the spin on the right while the first component of h^y addresses the spin on the bottom and the second the spin on the top. Moreover, we define

$$h^{\text{loc}}(s) := 1 - \delta(s, t_1).$$

Note that H_n has a ground state of energy zero if and only if there exists a valid tiling of $\mathbb{Z}_n^2$ with tile t_1 at the origin. That is, $E_{\min}(H_n) > 0$ if and only if there is no valid tiling of size n. This guarantees that $\mathcal{R}$ is a reduction. Additionally, we obtain a reduction from BTile to BGse since the bounding parameters are identical. Similar to the tiling problem, one can show that Gse is RE-complete and BGse is coNP-complete.

Note that non-translational invariant versions of BGse are known to be coNP-hard since their negated versions are NP-hard. In particular, the ground state energy problem for 2D Ising models with fields is NP-complete [4].

8.4 Conclusions and outlook

In this work, we have shown a relation between the hardness of an (unbounded) problem and the hardness of its bounded version. In particular, we have defined a bounded version of a language (Definition 8.1.1) and given a condition under which a reduction between the unbounded problems translates to a reduction between their bounded versions (Theorem 8.1.1). We have also applied this result to two classes of examples (Section 8.3): First, we showed that RE-hard problems like Pcp, Mpo, or Reach have an NP-hard bounded version; Second, we showed that RE-hard problems like Tile and Gse have a coNP-hard bounded version.

It would be interesting to extend this work to problems in quantum physics such as the spectral gap problem [34, 8] or membership problems for quantum correlations [116, 117, 70, 53, 86]. A bounded version of the latter uses the dimension of the entangled state as the bounding parameter.

Another open question is whether the undecidability of Diophantine equations [83] and the NP-hardness of its bounded version [81] fits into our framework.[8] In this context, the unbounded problem is as follows:

8: Recall that a Diophantine equation is a polynomial over the integers whose solutions need to be integers.

Problem 8.4.1 (Solvability of Diophantine equations)
Given a Diophantine equation $p(\mathbf{x}, \mathbf{y}) = 0$ with $2k$ variables, and a k-tuple of integers $\mathbf{a} \in \mathbb{Z}^k$, does there exist $\mathbf{b} \in \mathbb{Z}^k$ such that $p(\mathbf{a}, \mathbf{b}) = 0$?

Note that here k is fixed. The bounded version would restrict to values $\mathbf{b} \in \{-n, \dots, n\}^k$, where n acts as the bounding parameter.

Are there also hard bounded versions with other types of complexity, such as QMA-hard [127] bounded versions? While we only considered the scenario of RE-hard problems with either NP-hard or coNP-hard bounded versions, there might be "root problems" whose bounded version is neither NP-hard or coNP-hard. Natural candidates for QMA-hard bounded version are the bounded/unbounded satisfiability problems of quantum circuits [21], which concerns Turing machines generating polynomial-size quantum circuits. The results of this work would imply that certain QMA-hard problems, like the ground state energy problem for k-local quantum Hamiltonians [71], relate to unbounded problems which are undecidable.

Finally, is it possible to prove the converse direction of Theorem 8.1.1? Since bounded languages give rise to a unique unbounded language, can every reduction between bounded versions be transferred to a reduction between the corresponding unbounded problems? If the bounded reduction is of the special form

$$\mathcal{R}_b : \langle x, n \rangle \mapsto \langle \mathcal{R}(x), p(n) \rangle$$

with p being a strictly increasing polynomial, then $\mathcal{R}$ is automatically a reduction between the unbounded problems. Yet, the question is open for general $\mathcal{R}_b$.

Bibliography

[1] H. Abo. *Varieties of completely decomposable forms and their secants*. J. Algebra **403** (2014), 135–153. DOI: 10.1016/j.jalgebra.2013.12.027.

[2] S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, 2009. ISBN: 1139477366.

[3] M. C. Bañuls. *Tensor Network Algorithms: A Route Map*. Annu. Rev. Condens. Matter Phys. **14** (2023), 173–191. DOI: 10.1146/annurev-conmatphys-040721-022705.

[4] F. Barahona. *On the computational complexity of Ising spin glass models*. J. Phys. A: Math. Gen. **15** (1982), 3241. DOI: 10.1088/0305-4470/15/10/028.

[5] T. Barthel, J. Lu, and G. Friesecke. *On the closedness and geometry of tensor network state sets*. Lett. Math. Phys. **112** (2022). DOI: 10.1007/s11005-022-01552-z.

[6] A. I. Barvinok. *A Course in Convexity*. American Mathematical Society, 2002. ISBN: 978-0-8218-2968-4. DOI: 10.1090/gsm/054.

[7] S. Basu, R. Pollack, and M.-F. Roy. *Algorithms in Real Algebraic Geometry*. Springer Berlin, Heidelberg, 2006. ISBN: 978-3-540-33098-1. DOI: 10.1007/3-540-33099-2.

[8] J. Bausch, T. Cubitt, A. Lucia, and D. Pérez-García. *Undecidability of the Spectral Gap in One Dimension*. Phys. Rev. X **10** (2020), 031038. DOI: 10.1103/PhysRevX.10.031038.

[9] J. S. Bell. *On the Einstein Podolsky Rosen paradox*. Phys. Phys. Fiz. **1** (1964), 195–200. ISSN: 0554-128X. DOI: 10.1103/PhysicsPhysiqueFizika.1.195.

[10] P. Bell, V. Halava, T. Harju, J. Karhumaki, and I. Potapov. *Matrix Equations And Hilbert's Tenth Problem*. Int. J. Algebra Comput. **18** (2008), 1231–1241. DOI: 10.1142/S0218196708004925.

[11] R. Berger. *The undecidability of the domino problem*. Mem. Amer. Math. Soc. **66** (1966). DOI: 10.1090/memo/0066.

[12] A. Berman and N. Shaked-Monderer. *Completely Positive Matrices*. World Scientific, 2003. DOI: 10.1142/5273.

[13] D. Bini, G. Lotti, and F. Romani. *Approximate Solutions for the Bilinear Form Computational Problem*. SIAM J. Comput. **9** (1980), 692–697. ISSN: 0097-5397. DOI: 10.1137/0209053.

[14] L. Bittel and M. Kliesch. *Training Variational Quantum Algorithms Is NP-Hard*. Phys. Rev. Lett. **127** (2021), 120502. ISSN: 10797114. DOI: 10.1103/PhysRevLett.127.120502.

[15] V. D. Blondel, E. Jeandel, P. Koiran, and N. Portier. *Decidable and undecidable problems about quantum automata*. SIAM J. Comput. **34** (2005), 1464–1473. DOI: 10.1137/S0097539703425861.

[16] V. D. Blondel and N. Portier. *The presence of a zero in an integer linear recurrent sequence is NP-hard to decide*. Linear Algebra Appl. **351** (2002), 91–98. DOI: 10.1016/S0024-3795(01)00466-9.

[17] V. D. Blondel and J. N. Tsitsiklis. *When is a pair of matrices mortal?* Inf. Process. Lett. **63** (1997), 283–286. DOI: 10.1016/S0020-0190(97)00123-3.

[18] C. Bocci, E. Carlini, and F. Rapallo. *Perturbation of matrices and nonnegative rank with a view toward statistical models*. SIAM J. Matrix Anal. Appl. **32** (2011), 1500–1512. ISSN: 08954798. DOI: 10.1137/110825455.

[19] J. Bochnak, M. Coste, and M.-F. Roy. *Real Algebraic Geometry*. Springer Berlin, Heidelberg, 1998. ISBN: 978-3-642-08429-4. DOI: 10.1007/978-3-662-03718-8.

[20] B. Bollobás. *Modern graph theory*. Springer, 1998. ISBN: 9780387984889.

[21] A. D. Bookatz. *QMA-complete problems*. Quantum Inf Comput. **14** (2012), 361–383. DOI: 10.26421/QIC14.5-6-1.

[22] J. C. Bridgeman and C. T. Chubb. *Hand-waving and interpretive dance: an introductory course on tensor networks*. J. Phys. A: Math. Theo. **50** (2017), 223001. ISSN: 1751-8113. DOI: 10.1088/1751-8121/aa6dc3.

[23] H. Buhrman, M. Christandl, and J. Zuiddam. *Nondeterministic quantum communication complexity: the cyclic equality game and iterated matrix multiplication*. LIPIcs **67** (2017), 1–24. DOI: 10.4230/LIPIcs.ITCS.2017.24.

[24] J. Cassaigne, V. Halava, T. Harju, and F. Nicolas. *Tighter Undecidability Bounds for Matrix Mortality, Zero-in-the-Corner Problems, and More*. 2014. DOI: 10.48550/arXiv.1404.0644.

[25] M. D. Choi, T. Y. Lam, and B. Reznick. *Sums of squares of real polynomials*. Proc. Sympos. Pure Math. **58** (1995), 103–126. DOI: 10.1090/pspum/058.2/1327293.

[26] M. Christandl, F. Gesmundo, and A. K. Jensen. *Border rank is not multiplicative under the tensor product*. SIAM J. Appl. Algebraic G. **3** (2019), 231–255. ISSN: 24706566. DOI: 10.1137/18M1174829.

[27] M. Christandl, F. Gesmundo, M. Michałek, and J. Zuiddam. *Border rank nonadditivity for higher order tensors*. SIAM J. Matrix Anal. Appl. **42** (2021), 503–527. ISSN: 10957162. DOI: 10.1137/20M1357366.

[28] M. Christandl, A. K. Jensen, and J. Zuiddam. *Tensor rank is not multiplicative under the tensor product*. Linear Algebra Appl. **543** (2018), 125–139. ISSN: 00243795. DOI: 10.1016/j.laa.2017.12.020.

[29] M. Christandl, A. Lucia, P. Vrana, and A. H. Werner. *Tensor network representations from the geometry of entangled states*. SciPost Phys. **9** (2020), 1–35. ISSN: 25424653. DOI: 10.21468/SCIPOSTPHYS.9.3.042.

[30] J. I. Cirac, D. Pérez-García, N. Schuch, and F. Verstraete. *Matrix Product States and Projected Entangled Pair States: Concepts, Symmetries, and Theorems*. Rev. Mod. Phys. **93** (2021), 045003. DOI: 10.1103/RevModPhys.93.045003.

[31] J. E. Cohen and U. G. Rothblum. *Nonnegative ranks, decompositions, and factorizations of nonnegative matrices*. Linear Algebra Appl. **190** (1993), 149–168. ISSN: 00243795. DOI: 10.1016/0024-3795(93)90224-C.

[32] P. Comon, G. Golub, L. H. Lim, and B. Mourrain. *Symmetric tensors and symmetric tensor rank*. SIAM J. Matrix Anal. Appl. **30** (2008), 1254–1279. ISSN: 08954798. DOI: 10.1137/060661569.

[33] T. Cubitt, D. Elkouss, W. Matthews, M. Ozols, D. Pérez-Garciá, and S. Strelchuk. *Unbounded number of channel uses may be required to detect quantum capacity*. Nat. Commun. **6** (2015), 1–11. ISSN: 20411723. DOI: 10.1038/ncomms7739.

[34] T. S. Cubitt, D. Pérez-García, and M. M. Wolf. *Undecidability of the spectral gap*. Nature **528** (2015), 207–211. ISSN: 14764687. DOI: 10.1038/nature16059.

[35] G. De las Cuevas, J. I. Cirac, N. Schuch, and D. Pérez-García. *Irreducible forms of Matrix Product States: Theory and Applications*. J. Math. Phys. **58** (2017), 121901. DOI: 10.1063/1.5000784.

[36] G. De las Cuevas, T. S. Cubitt, J. I. Cirac, M. M. Wolf, and D. Pérez-García. *Fundamental limitations in the purifications of tensor networks*. J. Math. Phys. **57** (2016), 071902. ISSN: 00222488. DOI: 10.1063/1.4954983.

[37] G. De las Cuevas, M. Hoogsteder Riera, and T. Netzer. *Tensor decompositions on simplicial complexes with invariance*. J. Symb. Comput. **124** (2024), 102299. DOI: 10.48550/arXiv.2109.06680.

[38] G. De las Cuevas, A. Klingler, and T. Netzer. *Approximate tensor decompositions: disappearance of many separations*. J. Math. Phys. **62** (2021), 093502. DOI: 10.1063/5.0033876.

[39] G. De las Cuevas, A. Klingler, and T. Netzer. *Polynomial decompositions with invariance and positivity inspired by tensors*. 2021. DOI: 10.48550/arXiv.2109.06680.

[40] G. De las Cuevas and T. Netzer. *Mixed states in one spatial dimension: decompositions and correspondence with nonnegative matrices*. J. Math. Phys. **61** (2020), 41901. DOI: 10.1063/1.5127668.

[41] G. De las Cuevas, N. Schuch, D. Pérez-García, and J. I. Cirac. *Purifications of multipartite states: Limitations and constructive methods*. New J. Phys. **15** (2013), 123021. ISSN: 13672630. DOI: 10.1088/1367-2630/15/12/123021.

[42] G. De les Coves, J. Graf, A. Klingler, and T. Netzer. *Positive Moments Forever: Undecidable and Decidable Cases*. 2024. DOI: 10.48550/arXiv.2404.15053.

[43] S. Debus and C. Riener. *Reflection groups and cones of sums of squares*. J. Symb. Comput. **119** (2023). DOI: 10.1016/j.jsc.2023.03.001.

[44] A. C. Doherty, P. A. Parrilo, and F. M. Spedalieri. *Complete family of separability criteria*. Phys. Rev. A **69** (2004), 022308. ISSN: 10941622. DOI: 10.1103/PhysRevA.69.022308.

[45] C. Eckart and G. Young. *The approximation of one matrix by another of lower rank*. Psychometrika **1** (1936), 211–218. ISSN: 00333123. DOI: 10.1007/BF02288367.

[46] J. Eisert, M. P. Mueller, and C. Gogolin. *Quantum measurement occurrence is undecidable*. Phys. Rev. Lett. **108** (2012), 26501. DOI: 10.1103/PhysRevLett.108.260501.

[47] G. Everest, A. van der Poorten, I. Shparlinski, and T. Ward. *Recurrence Sequences*. American Mathematical Soc., 2003. ISBN: 978-1-4704-1331-6.

[48] M. van der Eyden, T. Netzer, and G. De las Cuevas. *Halos and undecidability of tensor stable positive maps*. J. Phys. A: Math. Theor. **55** (2022), 264006. DOI: 10.1088/1751-8121/ac726e.

[49] H. Fawzi, J. Gouveia, P. A. Parrilo, R. Z. Robinson, and R. R. Thomas. *Positive semidefinite rank*. Math. Program. **153** (2015), 133–177. ISSN: 14364646. DOI: 10.1007/s10107-015-0922-1.

[50] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. D. Wolf. *Linear vs. semidefinite extended formulations: Exponential separation and strong lower bounds*. Proc. ACM Symp. Theory of Computing (2012), 95–106. ISSN: 07378017. DOI: 10.1145/2213977.2213988.

[51] T. Fritz. *Quantum logic is undecidable*. Arch. Math. Log. **60** (2021), 329–341. ISSN: 14320665. DOI: 10.1007/s00153-020-00749-0.

[52] R. Fröberg, G. Ottaviani, and B. Shapiro. *On the Waring problem for polynomial rings*. Proc. Natl. Acad. Sci. U.S.A. **109** (2012), 5600–5602. ISSN: 00278424. DOI: 10.1073/pnas.1120984109.

[53] H. Fu, C. A. Miller, and W. Slofstra. *The membership problem for constant-sized quantum correlations is undecidable*. 2021. DOI: 10.48550/arXiv.2101.11087.

[54] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, 1979. ISBN: 0-7167-1045-5.

[55] K. Gatermann and P. A. Parrilo. *Symmetry groups, semidefinite programs, and sums of squares*. J. Pure Appl. Algebra **192** (2004), 95–128. ISSN: 00224049. DOI: 10.1016/j.jpaa.2003.12.011.

[56] S. Gharibian. *Strong NP-Hardness of the Quantum Separability Problem*. Quantum Inf. Comput. **10** (2010), 343–360. ISSN: 1533-7146.

[57] I. Glasser, N. Pancotti, and J. I. Cirac. *From Probabilistic Graphical Models to Generalized Tensor Networks for Supervised Learning*. IEEE Access **8** (2020), 68169–68182. ISSN: 21693536. DOI: 10.1109/ACCESS.2020.2986279.

[58] I. Glasser, R. Sweke, N. Pancotti, J. Eisert, and J. I. Cirac. *Expressive power of tensor-network factorizations for probabilistic modeling, with applications from hidden Markov models to quantum machine learning*. Adv. NeurIPS **32** (2019), 1498–1510. DOI: 10.48550/arXiv.1907.03741.

[59] G. H. Golub and C. F. V. Loan. *Matrix Computations*. Johns Hopkins University Press, 1996. ISBN: 0801854148.

[60] J. Gouveia, P. A. Parrilo, and R. R. Thomas. *Lifts of convex sets and cone factorizations*. Math. Oper. Res. **38** (2013), 248–264. ISSN: 0364765X. DOI: 10.1287/moor.1120.0575.

[61] R. Grosu. *The Cayley-Hamilton Theorem for Noncommutative Semirings*. CIAA (2011), 143–153. DOI: 10.1007/978-3-642-18098-9_16.

[62] L. Gurvits. *Classical Deterministic Complexity of Edmonds' Problem and Quantum Entanglement*. Proc. Annu. ACM Symp. Theory Comput. (2003), 10–19. DOI: 10.1145/780542.780545.

[63] V. Halava and T. Harju. *Mortality in Matrix Semigroups*. Amer. Math. Monthly **108** (2001), 649–653. DOI: 10.2307/2695274.

[64] G. H. Hardy, J. E. Littlewood, and G. Pólya. *Inequalities*. Cambridge University Press, 1952. ISBN: 9780521052061.

[65] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985. ISBN: 9780521548236. DOI: 10.1017/cbo9780511810817.

[66] M. Horodecki, P. Horodecki, and R. Horodecki. *Separability of mixed states: necessary and sufficient conditions*. Phys. Lett. A **223** (1996). ISSN: 03759601. DOI: 10.1016/S0375-9601(96)00706-2.

[67] H. B. Hunt, R. L. Constable, and S. Sahni. *On the Computational Complexity of Program Scheme Equivalence*. SIAM J. Comput. **9** (1980), 396–416. DOI: 10.1137/0209031.

[68] R. Jain, Y. Shi, Z. Wei, and S. Zhang. *Efficient protocols for generating bipartite classical distributions and quantum states*. IEEE Trans. Inf. Theory **59** (2013), 5171–5178. ISSN: 00189448. DOI: 10.1109/TIT.2013.2258372.

[69] R. Jain, Z. Wei, P. Yao, and S. Zhang. *Multipartite Quantum Correlation and Communication Complexities*. Comput. Complexity **26** (2017), 199–228. ISSN: 14208954. DOI: 10.1007/s00037-016-0126-y.

[70] Z. Ji, A. Natarajan, T. Vidick, J. Wright, and H. Yuen. *MIP*= RE*. Commun. ACM **64** (2021), 131–138. ISSN: 15577317. DOI: 10.1145/3485628.

[71] J. Kempe, A. Kitaev, and O. Regev. *The Complexity of the Local Hamiltonian Problem*. SIAM J. Comput. **35** (2006), 1070–1097. DOI: 10.1137/S0097539704445226.

[72] M. Kliesch, D. Gross, and J. Eisert. *Matrix-product operators and states: NP-hardness and undecidability*. Phys. Rev. Lett. **113** (2014), 160503. ISSN: 10797114. DOI: 10.1103/PhysRevLett.113.160503.

[73] A. Klingler, M. van der Eyden, S. Stengele, T. Reinhart, and G. De las Cuevas. *Many bounded versions of undecidable problems are NP-hard*. SciPost Phys. **14** (2023), 173. DOI: 10.21468/SciPostPhys.14.6.173.

[74] A. Klingler, T. Netzer, and G. De les Coves. *Border Ranks of Positive and Invariant Tensor Decompositions: Applications to Correlations*. 2023. DOI: 10.48550/arXiv.2304.13478.

[75] T. Kohler and T. Cubitt. *Translationally Invariant Universal Classical Hamiltonians*. J. Stat. Phys. **176** (2019), 228–261. ISSN: 00224715. DOI: 10.1007/s10955-019-02295-3.

[76] D. Koller and N. Friedman. *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press, 2009. ISBN: 0262013193.

[77] J. M. Landsberg, Y. Qi, and K. Ye. *On the geometry of tensor network states*. Quantum Inf. Comput. **12** (2012), 346–354. ISSN: 15337146. DOI: 10.26421/qic12.3-4-12.

[78] J. M. Landsberg. *Tensors: Geometry and Applications*. American Mathematical Society, 2011, 439. ISBN: 9780821869079. DOI: 10.1090/gsm/128.

[79] S. Lang. *Algebra*. Springer Science & Business Media, 2012.

[80] L.-H. Lim and P. Comon. *Nonnegative approximations of nonnegative tensors*. J. Chemom. **23** (2009), 432–441. ISSN: 08869383. DOI: 10.1002/cem.1244.

[81] K. Manders and L. Adleman. *NP-complete decision problems for quadratic polynomials*. Proc. Annu. ACM Symp. Theory Comput. (1976), 23–29. DOI: 10.1145/800113.803627.

[82] O. Martin, A. M. Odlyzko, and S. Wolfram. *Algebraic properties of cellular automata*. Comm. Math. Phys. **93** (1984), 219–258. DOI: 10.1007/BF01223745.

[83] J. V. Matijasevic. *Enumerable sets are diophantine*. Soviet Math. Doklady **11** (1970), 354–358.

[84] J. Miller, G. Roeder, and T.-D. Bradley. *Probabilistic Graphical Models and Tensor Networks: A Hybrid Framework*. 2021. DOI: 10.48550/arXiv.2106.15666.

[85] L. Mirsky. *Symmetric Gauge Functions and unitarily invariant norms*. Q. J. Math. **11** (1960), 50–59. ISSN: 0033-5606. DOI: 10.1093/qmath/11.1.50.

[86] H. Mousavi, S. S. Nezhadi, and H. Yuen. *Nonlocal Games, Compression Theorems, and the Arithmetical Hierarchy*. Proc. Annu. ACM Symp. Theory Comput. (2022), 1–11. DOI: 10.1145/3519935.3519949.

[87] A. Müller-Hermes, D. Reeb, and M. M. Wolf. *Positivity of linear maps under tensor powers*. J. Math. Phys. **57** (2016), 015202. DOI: 10.1063/1.4927070.

[88] A. Onishchik and E. Vinberg. *Lie Groups and Algebraic Groups*. Springer, 1990.

[89] R. Orús. *A practical introduction to tensor networks: Matrix product states and projected entangled pair states*. Ann. Physics **349** (2014), 117–158. ISSN: 1096035X. DOI: 10.1016/j.aop.2014.06.013.

[90] R. Orús. *Tensor networks for complex quantum systems*. Nat. Rev. Phys. **1** (2019), 538–550. ISSN: 25225820. DOI: 10.1038/s42254-019-0086-7.

[91] J. Ouaknine and J. Worrell. *On the Positivity Problem for Simple Linear Recurrence Sequences*. Automata, Languages, and Programming, ICALP (2013). DOI: 10.1007/978-3-662-43951-7_27.

[92] J. Ouaknine and J. Worrell. *Positivity Problems for Low-Order Linear Recurrence Sequences*. Proc. Ann. ACM-SIAM Symp. Discr. Alg. (2013), 366–379. DOI: 10.5555/2634074.2634101.

[93] J. Ouaknine and J. Worrell. *Decision problems for linear recurrence sequences* (2012), 21–28. DOI: 10.1007/978-3-642-33512-9_3.

[94] C. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994. ISBN: 9780201530827.

[95] M. S. Paterson. *Unsolvability in 3 × 3 Matrices*. Stud. Appl. Math. **49** (1970), 105–107. ISSN: 00222526. DOI: 10.1002/sapm1970491105.

[96] A. Peres. *Separability criterion for density matrices*. Phys. Rev. Lett. **77** (1996), 1413–1415. ISSN: 10797114. DOI: 10.1103/PhysRevLett.77.1413.

[97] D. Pérez-García, F. Verstraete, M. M. Wolf, and J. I. Cirac. *Matrix product state representations*. Quantum Inf. Comput. **7** (2007), 401–430. ISSN: 15337146. DOI: 10.26421/qic7.5-6-1.

[98] G. Pólya. *Über positive Darstellung von Polynomen*. Vierteljschr. Naturforsch. Ges. Zürich **73** (1928).

[99] E. L. Post. *A variant of a recursively unsolvable problem*. Bull. Am. Math. Soc. **52** (1946), 264–268. ISSN: 0273-0979. DOI: 10.1090/S0002-9904-1946-08555-9.

[100] A. Prakash, J. Sikora, A. Varvitsiotis, and Z. Wei. *Completely positive semidefinite rank*. Math. Program. **171** (2018), 397–431. ISSN: 14364646. DOI: 10.1007/s10107-017-1198-4.

[101] Y. Qi. *A very brief introduction to nonnegative tensors from the geometric viewpoint*. Mathematics **6** (2018). ISSN: 22277390. DOI: 10.3390/math6110230.

[102] Y. Qi, P. Comon, and L.-H. Lim. *Semialgebraic Geometry of Nonnegative Tensor Rank*. SIAM J. Matrix Anal. Appl. **37** (2016), 1556–1580. ISSN: 0895-4798. DOI: 10.1137/16M1063708.

[103] E. Robeva and A. Seigal. *Duality of graphical models and tensor networks*. Inf. Inference **8** (2019), 273–288. ISSN: 20498772. DOI: 10.1093/imaiai/iay009.

[104] R. M. Robinson. *Undecidability and Nonperiodicity for Tilings of the Plane*. Invent. Math. **12** (1971), 177–209. DOI: 10.1007/BF01418780.

[105] S. Roman. *Advanced Linear Algebra*. Springer New York, 2008. ISBN: 978-0-387-72828-5. DOI: 10.1007/978-0-387-72831-5.

[106] M. Sanz, D. Pérez-García, M. M. Wolf, and J. I. Cirac. *A quantum version of Wielandt's inequality*. IEEE Trans. Inf. Theory **56** (2010). DOI: 10.1109/TIT.2010.2054552.

[107] M. Scandi and J. Surace. *Undecidability in resource theory: can you tell theories apart?* Phys. Rev. Lett. **127** (2021), 270501. DOI: 10.1103/PhysRevLett.127.270501.

[108] G. Scarpa, A. Molnár, Y. Ge, J. J. García-Ripoll, N. Schuch, D. Pérez-García, and S. Iblisdir. *Projected Entangled Pair States: Fundamental Analytical and Numerical Limitations*. Phys. Rev. Lett. **125** (2020), 210504. ISSN: 10797114. DOI: 10.1103/PhysRevLett.125.210504.

[109] A. Schönhage. *Partial and Total Matrix Multiplication*. SIAM J. Comput. **10** (1981), 434–455. ISSN: 0097-5397. DOI: 10.1137/0210032.

[110] Y.-Y. Shi, L.-M. Duan, and G. Vidal. *Classical simulation of quantum many-body systems with a tree tensor network*. Phys. Rev. A **74** (2006), 022320. ISSN: 1050-2947. DOI: 10.1103/PhysRevA.74.022320.

[111] Y. Shitov. *Counterexamples to Strassen's direct sum conjecture*. Acta Math. **222** (2019), 363–379. ISSN: 00015962. DOI: 10.4310/ACTA.2019.v222.n2.a3.

[112] Y. Shitov. *The complexity of positive semidefinite matrix factorization*. SIAM J. Optim. **27** (2017), 1898–1909. ISSN: 10526234. DOI: 10.1137/16M1080616.

[113] Y. Shitov. *The nonnegative rank of a matrix: Hard problems, easy solutions*. SIAM Rev. **59** (2017), 794–800. ISSN: 00361445. DOI: 10.1137/16M1080999.

[114] V. D. Silva and L. H. Lim. *Tensor rank and the ill-posedness of the best low-rank approximation problem*. SIAM J. Matrix Anal. Appl. **30** (2008), 1084–1127. ISSN: 08954798. DOI: 10.1137/06066518X.

[115] M. Sipser. *Introduction to the Theory of Computation*. Course Technology, 2006. DOI: 10.1145/230514.571645.

[116] W. Slofstra. *The set of quantum correlations is not closed*. Forum Math. Pi **7** (2019). ISSN: 20505086. DOI: 10.1017/fmp.2018.3.

[117] W. Slofstra. *Tsirelson's problem and an embedding theorem for groups arising from non-local games*. J. Am. Math. Soc. **33** (2019), 1–56. ISSN: 0894-0347. DOI: 10.1090/jams/929.

[118] C. J. Stark and A. W. Harrow. *Compressibility of Positive Semidefinite Factorizations and Quantum Models*. IEEE Trans. Inf. Theory **62** (2016), 2867–2880. ISSN: 00189448. DOI: 10.1109/TIT.2016.2538278.

[119] J. O. Szigeti. *Caley-Hamilton Theorem For Matrices Over an Arbitrary Ring*. Serdica Math. J **32** (2006), 269–276.

[120] R. C. Tausworthe. *Random numbers generated by linear recurrence modulo two*. Math. Comput. **19** (1965), 201–209.

[121] K. Temme and F. Verstraete. *Stochastic matrix product states*. Phys. Rev. Lett. **104** (2010). ISSN: 00319007. DOI: 10.1103/PhysRevLett.104.210502.

[122] R. Tijdeman, M. Mignotte, and T. Shorey. *The distance between terms of an algebraic recurrence sequence*. J. Reine Angew. Math. **349** (1984), 63–76.

[123] A. M. Turing. *On computable numbers, with an application to the Entscheidungsproblem*. J. of Math **58** (1936). DOI: 10.1112/plms/s2-42.1.230.

[124] S. A. Vavasis. *On the complexity of nonnegative matrix factorization*. SIAM J. Optim. **20** (2009), 1364–1377. DOI: 10.1137/070709967.

[125] N. K. Vereshchagin. *The problem of appearance of a zero in a linear recurrence sequence*. Mat. Zametki **38** (1985), 609–615. DOI: 10.1007/BF01156238.

[126] F. Verstraete, J. J. García-Ripoll, and J. I. Cirac. *Matrix product density operators: Simulation of finite-temperature and dissipative systems*. Phys. Rev. Lett. **93** (2004), 12–15. ISSN: 00319007. DOI: 10.1103/PhysRevLett.93.207204.

[127] J. Watrous. *Quantum computational complexity*. Computational Complexity: Theory, Techniques, and Applications (2012). DOI: 10.1007/978-0-387-30440-3_428.

[128] A. Wigderson. *Mathematics and computation*. Princeton University Press, 2019. ISBN: 9780691189130.

[129] M. M. Wolf, T. S. Cubitt, and D. Pérez-García. *Are problems in Quantum Information Theory (un)decidable?* 2011. DOI: 10.48550/arXiv.1111.5425.

[130] M. Zwolak and G. Vidal. *Mixed-state dynamics in one-dimensional quantum lattice systems: A time-dependent superoperator renormalization algorithm*. Phys. Rev. Lett. **93** (2004), 1–5. ISSN: 00319007. DOI: 10.1103/PhysRevLett.93.207205.

List of notations

The next list describes several symbols that are used within the body of the document.

BGse	The bounded ground state energy problem.
BHalt	The bounded halting problem.
BMm	The bounded matrix mortality problem.
BMpo	The bounded positivity problem for matrix product operators.
BNHalt	The bounded non-deterministic halting problem.
BNHaltAll	The bounded version of NHaltAll.
BPcp	The bounded Post correspondence problem.
BPoly	The bounded polynomial positivity problem.
BReach	The bounded reachability problem in resource theories.
BTsp	The bounded positivity on a state problem
BTile	The bounded tiling problem
BZulc	The bounded zero-in-the-upper-left-corner problem.
coNP	The complement of NP.
$\deg(p)$	The degree of the polynomial p.
$\deg_{\text{loc}}(p)$	The local degree of the polynomial p.
Gse	The (unbounded) ground state energy problem.
Halt	The halting problem.
$\text{Her}_d(\mathbb{C})$	The set of $d \times d$ complex hermitian matrices.
Pcp	The Post correspondence problem.
$\text{Mat}_d(K)$	The set of $d \times d$ matrices with elements from K.
$\text{Mat}_{d,k}(K)$	The set of $d \times k$ matrices with elements from K.
$\mathbb{Q}[i]$	The field of complex numbers with rational real and imaginary parts.
$\text{O}_s(K)$	The group of orthogonal matrices over the field K.
$\text{U}_s(K)$	The group of unitary matrices over the field K.
Mm	The matrix mortality problem.
Mpo	The positivity problem for matrix product operators.
NHalt	The non-deterministic halting problem.

NHaltAll	The non-deterministic halting problem on all paths.
nn-rank	The nonnegative rank.
NP	The set of non-deterministic polynomial-time decidable languages.
P	The set of polynomial-time decidable languages.
Poly	The polynomial positivity problem.
$\text{Psd}_d(\mathbb{C})$	The set of $d \times d$ positive semidefinite matrices.
psd-rank	The positive semidefinite rank.
puri-rank	The purification rank.
R	The set of recursive (decidable) languages.
rank	The (unconstrained) rank.
RE	The set of co-recursively enumerable languages.
RE	The set of recursively enumerable languages.
Reach	The reachability problem in resource theories.
sep-rank	The separable rank.
sos-rank	The sum-of-squares rank.
Tsp	The positivity on a state problem.
Tile	The tiling problem.
Zulc	The zero-in-the-upper-left-corner problem.
$A \succcurlyeq 0$	The matrix A is positive semidefinite.
A^t	The transpose of a matrix A.
$A^\dagger$	The Hermitian transpose of a matrix A.
$[n]$	The set $\{1, \ldots, n\}$.
Λ_n	The line with n vertices.
$\mathbb{N}_+$	The set of positive natural numbers $\{1, 2, 3, \ldots\}$.
Σ_n	The simplex with n vertices.
Θ_n	The cycle with n vertices.
C_n	The cyclic group with n elements.
S_n	The full permutation group on n elements.

List of abbreviations

cp	completely positive.
cpsd	completely positive semidefinite.
cpsdt	completely positive semidefinite transpose.
cptp	completely positive trace preserving.
LPDO	locally purified density operator.
LRS	linear recurrence sequence.
MaMu	matrix multiplication.
MPDO	matrix product density operator.
MPO	matrix product operator.
MPS	matrix product state.
POVM	positive operator-valued measurement.
psd	positive semidefinite.
sos	sum-of-squares.
ti	translational invariant.
WSC	weighted simplicial complex.