

Quantum Information Methods for Many-Body Physics

Xhek Turkeshi

Markus Heinrich

Institute for Theoretical Physics, University of Cologne

CONTENTS

1	Introduction	2
1.1	Overview	2
1.2	Permutations and their combinatorics	3
1.2.1	Permutations and Cycles	3
1.2.2	The Action of Permutations on Quantum States	8
2	Quantum Randomness I	13
2.1	What is a Haar measure?	13
2.1.1	The unitary commutant	14
2.1.2	Weingarten calculus	15
2.2	Approximating the Haar measure: Unitary designs	17
2.2.1	Unitary designs	18
2.2.2	The Clifford group	18
2.3	Further reading	20
3	Measuring Properties of Many-Body States using Classical Shadows	21
3.1	Reconstructing quantum states from measurements	21
3.2	Shadow estimation with randomized measurements	23
3.2.1	Clifford measurements	24
3.2.2	Pauli measurements	28
3.2.3	Beyond linear observables: the purity	30
3.3	Further reading	31
4	Quantum Randomness II	32
4.1	Random quantum circuits	32
4.2	More on spectral gaps*	36
4.3	Further reading	36
	Bibliography	37
A	Some linear algebra	38
A.1	States, operators, superoperators	38
A.2	Non-orthonormal bases	40

CHAPTER 1

INTRODUCTION

Quantum information science is the field that studies how information is stored, processed, and transmitted when it is governed by the laws of quantum mechanics. It includes areas such as *quantum computing*, *quantum communication*, *quantum cryptography*, *quantum sensing*, and *quantum error correction*. Although still relatively young, quantum information science has already had a profound impact on many areas of physics, especially on the study of quantum systems composed of many interacting particles, commonly referred to as *many-body quantum systems*.

Within this context, *quantum circuits* and *tensor networks* have emerged as essential tools to tackle fundamental questions in quantum dynamics—from the mechanisms of *thermalization* and the emergence of *statistical mechanics* in isolated systems, to the onset of *quantum chaos* and its deep connections with *black hole physics* and *holography* in quantum gravity. Studying these phenomena in a concrete setting is notoriously difficult, due to the exponential growth of the Hilbert space and the inherently non-linear structure of quantum correlations. *Random quantum circuits* and *tensor networks* offer a powerful way to overcome these challenges: they enable analytical and numerical progress through disorder averaging, while capturing the typical behavior of generic quantum systems thanks to *quantum typicality arguments*.

Crucially, the interplay between quantum information and many-body physics has not only refined our understanding of traditional problems, but has also uncovered entirely new dynamical phases of matter – phases that arise uniquely in *programmable quantum devices*. This so-called *synthetic quantum matter* cannot be characterized by conventional local order parameters such as magnetization or current. Instead, its defining features are quantum informational, such as the structure of entanglement or nonstabilizer (magic state) resources, or the system’s ability to preserve quantum information against noise and local errors. Understanding and classifying such phases requires a shift in perspective –from symmetry and energetics to *information content* and *computational complexity*.

At the same time, random unitary dynamics, especially in the form of random quantum circuits, have become indispensable tools in the *NISQ (noisy intermediate-scale quantum)* era. They provide efficient and versatile frameworks for a wide range of applications, including the benchmarking and verification of quantum computations, the characterization of noise in experimental platforms, and the estimation of observables via shadow tomography. Far from being purely theoretical constructs, these methods are implemented across various platforms – from superconducting qubits to cold atoms – and are central to the ongoing development of near-term quantum technologies.

This course provides a pedagogical introduction to random unitaries and to several key methods from quantum information theory, with a focus on their application to many-body physics. A substantial part of the course covers research-level topics introduced only in the past few years, offering a unique opportunity to engage with current questions at the interface between two rapidly evolving fields. It is also intended to serve as a solid preparation to pursue a Master’s thesis, a doctorate or work in the private for these areas.

1.1 Overview

The structure of the course is as follows. We begin with a chapter on permutations and a graphical calculus, which will provide the foundation for the treatment of randomization methods throughout the course. We then introduce the core randomization concepts in two parts, Quantum Randomness I and II—each followed by a chapter that connects the methods to applications in many-body systems.

1. In Quantum Randomness I (Ch. 2), we introduce *Weingarten calculus*, the central toolbox for computing averages over the unitary group, which naturally arise when considering statistical

properties of random evolutions. We also *study unitary designs*, which provide efficient approximations of Haar randomness and play an important role in practical implementations.

2. In Ch. 3, we present the *framework of classical shadows*. Introduced around 2020, classical shadows offer an efficient way to extract information about quantum states using only a small number of randomized measurements. This technique has already found widespread use in experimental platforms and continues to inspire a growing body of research.
3. In Quantum Randomness II (Ch. 4), we explore random quantum circuits in detail. These models provide an efficient and physically motivated approach to generate randomness in quantum many-body systems, and they serve as minimal models for chaotic quantum dynamics.
4. In Ch. ??, we show how random circuits can be used to model scrambling, thermalization, and information spreading in interacting quantum systems. These models also offer connections to quantum chaos, complexity growth, and typicality in many-body physics.

Each chapter concludes with a short guide to the research literature, primarily in the form of original articles, allowing students to explore further and connect the course material to current work in the field.

1.2 Permutations and their combinatorics

Why permutations? A central theme of this course is the study of random unitaries and their applications in quantum many-body systems. To understand their behavior, we need to analyze the statistics of random unitaries—specifically, we are interested in computing averages, variances, and higher moments of functions involving random unitary matrices.

At first glance, this might seem daunting: computing integrals over the unitary group is, in general, a highly nontrivial task. However, a powerful insight from representation theory, known as *Schur-Weyl duality*, provides a way forward. This duality reveals a deep connection *between the action of the unitary group and the action of the permutation group*, which allows us to reformulate complicated integrals in terms of combinatorics of permutations.

This leads to the framework known as *Weingarten calculus*, which enables the exact evaluation of many relevant averages over the unitary group. As a result, permutations will play a key role throughout this course – not for abstract mathematical reasons, but because they offer a concrete and computable handle on random quantum processes.

In this section, we will introduce the essential properties of permutations needed for our purposes. While there is a rich and beautiful mathematical structure behind these ideas, we will focus only on the aspects that are directly relevant for our discussion and applications. The interested reader can consult the bibliography.

1.2.1 Permutations and Cycles

A *permutation* is a reordering of a finite set of elements. In this course, we consider permutations of the set $\{1, 2, \dots, k\}$. We will denote permutations by Greek letters such as π, σ, τ , and so on. The set of all permutations of k elements forms a group under composition, called the *symmetric group*, and denoted by S_k .

Given a permutation $\sigma \in S_k$ and an element $x \in \{1, \dots, k\}$, we write $\sigma(x)$ to indicate the image of x under σ . A common and explicit way to write a permutation is the *two-line notation*, where the first row lists the original elements and the second row gives their images under the permutation:

$$\sigma = \begin{pmatrix} 1 & 2 & \dots & k \\ \sigma(1) & \sigma(2) & \dots & \sigma(k) \end{pmatrix}. \quad (1.1)$$

Example 1.1: Explicit Notation of Permutations

An example of a permutation of four elements is:

$$\sigma = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 3 & 1 & 4 & 2 \end{pmatrix}. \quad (1.2)$$

This means that the permutation acts as:

$$\sigma(1) = 3, \quad \sigma(2) = 1, \quad \sigma(3) = 4, \quad \sigma(4) = 2. \quad (1.3)$$

The group operation in S_k is the composition of permutations, denoted by $\sigma \cdot \tau$ for $\sigma, \tau \in S_k$, and defined by

$$(\sigma \cdot \tau)(x) = \sigma(\tau(x)) \quad \text{for all } x \in \{1, \dots, k\}. \quad (1.4)$$

Note that permutation composition is applied from right to left: τ acts first, followed by σ .

Example 1.2: Product of Permutations

Consider the following two permutations of $k = 4$ elements:

$$\sigma = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 3 & 1 & 4 & 2 \end{pmatrix}, \quad \tau = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 4 & 2 & 3 & 1 \end{pmatrix}. \quad (1.5)$$

To compute the composition $(\tau \cdot \sigma)(x) = \tau(\sigma(x))$, we apply σ first and then τ :

$$\begin{aligned} \sigma(1) &= 3, & \tau(\sigma(1)) &= \tau(3) = 3, \\ \sigma(2) &= 1, & \tau(\sigma(2)) &= \tau(1) = 4, \\ \sigma(3) &= 4, & \tau(\sigma(3)) &= \tau(4) = 1, \\ \sigma(4) &= 2, & \tau(\sigma(4)) &= \tau(2) = 2. \end{aligned} \quad (1.6)$$

Putting everything together, we find:

$$\tau \cdot \sigma = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 3 & 4 & 1 & 2 \end{pmatrix}. \quad (1.7)$$

It is important to emphasize that for $k \geq 3$, the symmetric group S_k is *non-abelian*, meaning that the order of composition matters – in general, $\sigma \cdot \tau \neq \tau \cdot \sigma$.

Exercise 1.1. Consider the permutations from Example 1.2. Compute the product $\sigma \cdot \tau$, and verify that it differs from $\tau \cdot \sigma$.

The symmetric group contains a special element called the *identity permutation*, denoted by ι , which leaves all elements unchanged:

$$\iota = \begin{pmatrix} 1 & 2 & \dots & k \\ 1 & 2 & \dots & k \end{pmatrix}. \quad (1.8)$$

By definition, composition with the identity does not change the permutation:

$$\iota \cdot \sigma = \sigma = \sigma \cdot \iota. \quad (1.9)$$

Moreover, every permutation $\sigma \in S_k$ has an *inverse* σ^{-1} such that:

$$\sigma \cdot \sigma^{-1} = \iota = \sigma^{-1} \cdot \sigma. \quad (1.10)$$

To compute the inverse of a permutation π , for each index $i \in \{1, 2, \dots, k\}$ we set $\pi^{-1}(\pi(i)) = \iota(i) = i$. The resulting list π^{-1} is the inverse permutation.

Example 1.3: Inverse of a Permutation

Consider the following two permutations of $k = 5$ elements:

$$\sigma = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 3 & 5 & 1 & 2 & 4 \end{pmatrix}. \quad (1.11)$$

To compute the inverse σ^{-1} we use the rule $\sigma^{-1}(\sigma(i)) \equiv i$

$$\begin{aligned} \sigma(1) = 3 &\Rightarrow \sigma^{-1}(3) = 1 \\ \sigma(2) = 5 &\Rightarrow \sigma^{-1}(5) = 2 \\ \sigma(3) = 1 &\Rightarrow \sigma^{-1}(1) = 3 \\ \sigma(4) = 2 &\Rightarrow \sigma^{-1}(2) = 4 \\ \sigma(5) = 4 &\Rightarrow \sigma^{-1}(4) = 5 \end{aligned} \quad (1.12)$$

Reordering the list in terms of the argument, we find

$$\sigma^{-1} = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 3 & 4 & 1 & 5 & 2 \end{pmatrix}. \quad (1.13)$$

A *cyclic permutation*, or simply a *cycle*, is a specific type of permutation in which a subset of elements is permuted in a closed loop, while all remaining elements remain fixed. Formally, an l -cycle is a permutation that permutes r elements cyclically and leaves the other $k - l$ elements unchanged. The number l is called the *length of the cycle*.

Concretely, a cycle of length l means that there exists a subset $\{i_1, i_2, \dots, i_l\} \subset \{1, 2, \dots, k\}$ such that

$$\sigma(i_1) = i_2, \quad \sigma(i_2) = i_3, \quad \dots, \quad \sigma(i_{l-1}) = i_l, \quad \sigma(i_l) = i_1, \quad (1.14)$$

and for all other elements $x \notin \{i_1, \dots, i_l\}$, we have $\sigma(x) = x$.

Exercise 1.2 (Cyclic permutations). *The following are examples of cyclic permutations:*

$$\tau = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 2 & 3 & 4 & 1 \end{pmatrix}, \quad \sigma = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 4 & 2 & 1 & 3 & 5 \end{pmatrix}. \quad (1.15)$$

Can you identify the subset $\{a_1, \dots, a_l\}$ that is cyclically permuted in each case? What is the length l of each cycle? [Answer: For τ , $r = l$; for σ , $r = l$.]

One fundamental property of cycles is that any permutation can be decomposed into a product of disjoint cycles. That is, for any $\sigma \in S_k$, there exists a unique set of cycles that act on mutually disjoint subsets of $\{1, 2, \dots, k\}$. This decomposition is unique up to the order in which the cycles are written. This motivates the *cycle notation* of permutations:

$$\sigma = (a_1 a_2 \dots a_{j_1})(a_{j_1+1} a_{j_1+2} \dots a_{j_2}) \dots (a_{j_{r-1}+1} a_{j_{r-1}+2} \dots a_{j_r}). \quad (1.16)$$

Here, each tuple represents a cycle, and all elements a_m are drawn from $\{1, 2, \dots, k\}$ without repetition. The integer $r = \#(\sigma)$ denotes the *number of disjoint cycles* in the decomposition of σ .

Let us now describe an explicit algorithm to obtain the cycle decomposition of a permutation, as in Eq. (1.16). The idea is simple: we iteratively follow the action of the permutation until we return to the starting point, keeping track of all visited elements.

- (i) **Start from the smallest unvisited element.** Initially, this is $x = 1$. If 1 has already been included in a previous cycle, move to the next smallest unvisited element.

- (ii) **Construct a cycle by iterating the permutation.** Begin by writing down x . Then repeatedly apply the permutation π to generate the sequence

$$x, \pi(x), \pi(\pi(x)), \pi(\pi(\pi(x))), \dots$$

Continue this process until you return to the starting point x . The list of elements

$$(x \ \pi(x) \ \pi^2(x) \ \dots \ \pi^{j-1}(x)) \quad (1.17)$$

forms a cycle of length j . Mark all of these elements as visited.

- (iii) **Repeat the process.** Find the next smallest unvisited element and return to step (ii). Continue until all elements have been visited. The full cycle decomposition of π is then obtained by combining the individual cycles found in each iteration.

Example 1.4: Cycle decomposition

Consider the permutation

$$\tau = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 \\ 4 & 5 & 3 & 1 & 2 \end{pmatrix}. \quad (1.18)$$

We apply the cycle decomposition algorithm step by step:

- Start with 1:

$$1 \xrightarrow{\tau} 4 \xrightarrow{\tau} 1.$$

This gives the first cycle: $(1 \ 4)$. Mark 1 and 4 as visited.

- Next smallest unvisited element is 2:

$$2 \xrightarrow{\tau} 5 \xrightarrow{\tau} 2.$$

This gives the second cycle: $(2 \ 5)$. Mark 2 and 5 as visited.

- The last unvisited element is 3, and since

$$3 \xrightarrow{\tau} 3,$$

this is a fixed point (a 1-cycle), written as (3) .

Combining the above, the full cycle decomposition is:

$$\tau = (1 \ 4)(2 \ 5)(3). \quad (1.19)$$

Note that the *order* in which disjoint cycles are written is irrelevant. For instance, in Example 1.4, all of the following represent the same permutation:

$$\tau = (1 \ 4)(2 \ 5)(3) = (2 \ 5)(1 \ 4)(3) = (3)(1 \ 4)(2 \ 5).$$

When the total number of elements k is clear from the context (e.g., $k = 5$ in this case), it is common to omit one-cycles, also denoted *fixed points*, because these elements are understood to remain unchanged under the permutation. Using this convention, the permutation in Example 1.4 is simply written as:

$$\tau = (1 \ 4)(2 \ 5). \quad (1.20)$$

With this notation, the identity permutation is denoted by $\iota = ()$, which is shorthand for $\iota = (1)(2) \cdots (k)$ – that is, all elements are fixed.

The *cycle structure* of a permutation, denoted by $\lambda(\pi)$, is the list of the lengths of its disjoint cycles. For example, the permutation $\tau = (1\ 4)(2\ 5)$ has cycle structure $\lambda(\tau) = (2, 2, 1)$. The number of disjoint cycles is then given by the length of this list:

$$\#(\pi) = |\lambda(\pi)|, \quad \text{for any } \pi \in S_k. \quad (1.21)$$

Example 1.5: Cycle Structure

Consider the following permutations of 6 elements

$$\tau = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 5 & 4 & 1 & 6 & 3 & 2 \end{pmatrix}, \quad \sigma = \begin{pmatrix} 1 & 2 & 3 & 4 & 5 & 6 \\ 5 & 2 & 6 & 4 & 1 & 3 \end{pmatrix}. \quad (1.22)$$

It is a simple exercise to show that their cycle decomposition is $\tau = (1, 5, 3)(2, 4, 6)$ and $\sigma = (1, 5)(3, 6)$ [which is a shorthand notation for $\sigma = (1, 5)(3, 6)(2)(4)$]. The permutation τ has two cycles of length 3, hence the cycle structure is $\lambda(\tau) = (3, 3)$. Instead, σ is composed of two 2-cycles and two 1-cycles, so the cycle structure is $\lambda(\sigma) = (2, 2, 1, 1)$.

Exercise 1.3. Show that conjugation preserves the cycle structure of a permutation. That is, for any $\sigma, \pi \in S_k$, prove that

$$\lambda(\pi\sigma\pi^{-1}) = \lambda(\sigma).$$

Transpositions, also known as *swaps*, are a special class of permutations that exchange exactly two elements and leave all others unchanged. A transposition has the form:

$$\sigma = (i\ j) = \begin{pmatrix} 1 & \dots & i & \dots & j & \dots & k \\ 1 & \dots & j & \dots & i & \dots & k \end{pmatrix}. \quad (1.23)$$

It is straightforward to verify that *any permutation* can be written as a product of transpositions. However, unlike the decomposition into disjoint cycles, this representation is *not unique* – a given permutation can be written in many different ways as a product of transpositions. While this makes the transposition decomposition less suitable for labeling permutations, it is extremely useful in algebraic manipulations. For instance, it turns out that the number of transpositions in any representation is always odd or always even, which justifies the definition of the *sign of a permutation*:

$$\text{sgn}(\sigma) = (-1)^{\#\text{transpositions in } \sigma}. \quad (1.24)$$

The sign function is important for the construction of representations of S_k and plays an important role in multilinear algebra, in particular in the definition of the *determinant of a matrix*.

We conclude this section by reviewing some structural aspects of the symmetric groups S_k for varying values of k . A key property is that the group S_k naturally embeds into S_{k+1} : that is, every permutation of k elements can be viewed as a permutation of $k+1$ elements that leaves the $(k+1)$ -th element fixed. More formally, we can write

$$S_{k+1} = S_k \sqcup \{(j\ k+1) \cdot \sigma : \sigma \in S_k, j = 1, 2, \dots, k\}, \quad (1.25)$$

where the union is disjoint and the second term represents all permutations obtained by composing an element of S_k with a transposition that swaps $k+1$ with one of the first k elements.

This recursive structure is useful for establishing many properties of permutations by induction on k . A simple but important example is the total number of elements in the symmetric group.

Theorem 1.1 (Counting of Permutations). Given $k \geq 1$, the total number of permutations in S_k is

$$|S_k| = k!. \quad (1.26)$$

Proof. The total number of permutations is given by the factorial $k! = k(k-1) \cdots 2 \cdot 1$, with the convention that $0! = 1$. To prove this, we use induction.

For $k = 1$, the symmetric group $S_1 = \{\iota\}$ consists only of the identity permutation, so $|S_1| = 1 = 1!$.

Assume now that $|S_k| = k!$ for some $k \geq 1$. From Eq. (1.25), the next symmetric group can be written as

$$S_{k+1} = S_k \sqcup \{(j \ k+1) \cdot \sigma : \sigma \in S_k, j = 1, \dots, k\}. \quad (1.27)$$

There are k possible values of j , and for each j , σ runs over all $k!$ permutations in S_k . Hence,

$$|S_{k+1}| = |S_k| + k \cdot |S_k| = k! + k \cdot k! = (k+1) \cdot k! = (k+1)!. \quad (1.28)$$

This completes the proof by induction. \square

Consider now the following permutations of six elements:

$$\sigma_1 = (1\ 2\ 3)(5\ 6), \quad \sigma_2 = (2\ 4\ 6\ 1), \quad \sigma_3 = (1\ 2)(3\ 4)(5\ 6). \quad (1.29)$$

While their cycle structures differ, all three permutations have exactly three disjoint cycles. This illustrates that simply counting the number of cycles is a coarser classification than specifying the full *cycle structure*.

In many computations throughout this course, we will be interested in the number of permutations in S_k with a fixed number of cycles $r = \#(\sigma)$. This quantity is given by the *unsigned Stirling numbers of the first kind*, denoted by $c(k, r)$. These numbers satisfy the recursive relation:

$$c(k+1, r) = k \cdot c(k, r) + c(k, r-1), \quad (1.30)$$

which allows them to be computed inductively, without explicitly listing all permutations. These numbers form a triangle similar to Pascal's triangle and are tabulated up to $k = 10$ in the Appendix.

Starting with the base case $k = 1$, where $S_1 = \{()\}$, we find:

$$c(1, 0) = 0, \quad c(1, 1) = 1. \quad (1.31)$$

Using the recurrence, the next values for $k = 2$ are:

$$c(2, 0) = 0, \quad c(2, 1) = 1, \quad c(2, 2) = 1. \quad (1.32)$$

Exercise 1.4. Compute the values $c(k, r)$ for $1 \leq r \leq k$ when $k = 3$ and $k = 4$. Verify that:

$$\sum_{r=1}^k c(k, r) = k!, \quad \sum_{r=1}^k c(k, r) x^r = x(x+1) \cdots (x+k-1). \quad (1.33)$$

The first identity reflects the fact that summing over all permutations with a fixed number of cycles r recovers the total number of permutations in S_k , while the second gives another combinatorial interpretation of $c(k, r)$ as the coefficients in the power series of the 'rising factorial'.

1.2.2 The Action of Permutations on Quantum States

Throughout this course, we work with a d -dimensional Hilbert space $\mathcal{H} = \mathbb{C}^d$, equipped with the standard orthonormal basis $\{|x\rangle\}_{x=0}^{d-1}$. Our main object of interest is the k -fold tensor product space $\mathcal{H}^{\otimes k} = (\mathbb{C}^d)^{\otimes k}$, often referred to as the *replica space* in the many-body literature.

Elements of the symmetric group S_k act naturally on this space by permuting the tensor factors. Concretely, given a permutation $\sigma \in S_k$, its action on a product basis state is defined as:

$$R_\sigma |x_1, x_2, \dots, x_k\rangle = |x_{\sigma(1)}, x_{\sigma(2)}, \dots, x_{\sigma(k)}\rangle. \quad (1.34)$$

Consider an example with $k = 4$ and $\sigma = (1\ 4\ 3)(2) \in S_4$. We have

$$R_\sigma |x_1, x_2, x_3, x_4\rangle = |x_4, x_2, x_1, x_3\rangle. \quad (1.35)$$

Equation (1.34) defines a linear (even unitary) operator R_σ on $\mathcal{H}^{\otimes k}$. The map

$$R: S_k \longrightarrow U(\mathcal{H}^{\otimes k}), \quad \sigma \mapsto R_\sigma \quad (1.36)$$

is a so-called *group representation* of S_k . This means it respects the group structure of S_k , more specifically:

$$R_{\sigma\pi} = R_\pi R_\sigma, \quad R_I = \mathbb{1}, \quad R_{\sigma^{-1}} = R_\sigma^{-1}, \quad R_\sigma^\dagger = R_\sigma^{-1}. \quad (1.37)$$

These properties follow directly from the definition in Eq. (1.34). As an explicit verification of the composition law, let us define $\tilde{x}_i := x_{\sigma(i)}$ and compute:

$$R_\pi R_\sigma |x_1, x_2, \dots, x_k\rangle = R_\pi |x_{\sigma(1)}, \dots, x_{\sigma(k)}\rangle = |\tilde{x}_{\pi(1)}, \dots, \tilde{x}_{\pi(k)}\rangle \quad (1.38)$$

$$= |x_{(\sigma\cdot\pi)(1)}, \dots, x_{(\sigma\cdot\pi)(k)}\rangle = R_{\sigma\cdot\pi} |x_1, x_2, \dots, x_k\rangle. \quad (1.39)$$

We leave the verification of the other axioms – such as unitarity and inverse consistency – as an exercise.

Exercise 1.5. Verify that R defines a unitary representation of the symmetric group S_k , i.e., check Eq. (1.37).

Exercise 1.6. Show that permutations act on tensor products of operators as follows:

$$R_\sigma(A_1 \otimes A_2 \otimes \dots \otimes A_k)R_\sigma^\dagger = A_{\sigma(1)} \otimes A_{\sigma(2)} \otimes \dots \otimes A_{\sigma(k)}. \quad (1.40)$$

Hint: Apply both sides to a product basis state.

Composite Systems In practice, we often work with *multi-qudit systems* described by a Hilbert space of the form $\mathcal{H} = (\mathbb{C}^q)^{\otimes n}$, corresponding to n qudits of local dimension q . In this setting, the k -fold copy of the system is given by the tensor product

$$((\mathbb{C}^q)^{\otimes n})^{\otimes k}, \quad (1.41)$$

which can be naturally visualized as a $k \times n$ grid of qudits (see Fig. 1.1). Each row represents one replica of the full system, and each column represents the k copies of a single local qudit.

Quantum operations such as global unitaries typically act *row-wise*, that is, identically and independently on each copy. Such operations are of the form

$$U^{\otimes k}, \quad \text{where } U \in U((\mathbb{C}^q)^{\otimes n}), \quad (1.42)$$

meaning that the unitary acts in parallel across the k rows.

In contrast, permutations act by permuting the *rows*, i.e., the k copies of each local qudit. This operation is performed *column-wise*, and can be implemented by applying the same permutation operator to each column in parallel. This leads to a convenient factorized structure: if we reinterpret the total space via the isomorphism

$$((\mathbb{C}^q)^{\otimes n})^{\otimes k} \simeq ((\mathbb{C}^q)^{\otimes k})^{\otimes n}, \quad (1.43)$$

then the permutation operator R_π acting on the full system decomposes as

$$R_\pi = r_\pi^{\otimes n}, \quad (1.44)$$

where r_π acts on the k -dimensional replica space associated with each local qudit. This “horizontal” factorization is exactly what is depicted in Fig. 1.1 and will be essential in constructing efficient representations of randomized operations throughout the course.




Figure 1.1: The Hilbert space $((\mathbb{C}^q)^{\otimes n})^{\otimes k}$ depicted as a $k \times n$ grid where every point corresponds to a copy of \mathbb{C}^q . Unitaries $U \in U(q^n)$ act row-wise on the grid, while permutations $\pi \in S_k$ act column-wise.

Traces Traces will play a fundamental role for explicit computations. Here, we derive a formula for the trace of permutation operators and for the trace of a product operator multiplied with a permutation. Let us first consider the cyclic permutation $\gamma = (1, 2, \dots, k)$. Then, we compute

$$\begin{aligned} \text{tr}(R_\gamma) &= \sum_{x_1, \dots, x_k=0}^{d-1} \langle x_1, \dots, x_k | R_\gamma | x_1, \dots, x_k \rangle \\ &= \sum_{x_1, \dots, x_k} \langle x_1, \dots, x_k | x_2, \dots, x_k, x_1 \rangle = \sum_{x_1, \dots, x_k} \delta_{x_1, x_2} \delta_{x_2, x_3} \dots \delta_{x_{k-1}, x_k} \delta_{x_k, x_1} = d. \end{aligned} \quad (1.45)$$

Next, consider an arbitrary permutation σ . Then, we can find a permutation π such that $\pi\sigma\pi^{-1} = (1, \dots, b_1)(b_1 + 1, b_1 + 2, \dots, b_2) \dots (b_{r-1} + 1, b_{r-1} + 2, \dots, b_r)$ where $r = \#\sigma$ is the number of cycles in σ (recall that the cycle structures of σ and $\pi\sigma\pi^{-1}$ have to necessarily match, cf. Ex. 1.3). But $R_{\pi\sigma\pi^{-1}}$ is simply a tensor product of cyclic permutations on $\mathcal{H}^{\otimes b_1}, \dots, \mathcal{H}^{\otimes (b_2-b_1)}, \dots, \mathcal{H}^{\otimes (b_r-b_{r-1})}$ and thus

$$\text{tr}(R_\sigma) = \text{tr}(R_{\pi\sigma\pi^{-1}}) = \prod_{i=1}^r d = d^r = d^{\#\sigma}. \quad (1.46)$$

Beyond this simple situation, we often need to compute the trace of product operators multiplied with a permutation operator, i.e. an expression of the form $\text{tr}(A_1 \otimes A_2 \otimes \dots \otimes A_k R_\sigma)$. Here, we can follow the same arguments as above: First, if $\sigma = \gamma = (1, \dots, k)$ is the cyclic permutation, then

$$\text{tr}(A_1 \otimes A_2 \otimes \dots \otimes A_k R_\gamma) = \sum_{x_1, \dots, x_k} \langle x_1, \dots, x_k | A_1 \otimes A_2 \otimes \dots \otimes A_k | x_2, x_3, \dots, x_k, x_1 \rangle \quad (1.47)$$

$$= \sum_{x_1, \dots, x_k} (A_1)_{x_1, x_2} (A_2)_{x_2, x_3} \dots (A_k)_{x_k, x_1} = \text{tr}(A_1 A_2 \dots A_k). \quad (1.48)$$

Next, for an arbitrary $\sigma = (a_1, \dots, a_{j_1})(a_{j_1+1}, \dots, a_{j_2}) \dots (a_{j_{r-1}+1}, \dots, a_k)$, find again a permutation π that ‘orders the cycles’ as $\pi\sigma\pi^{-1} = (1, \dots, b_1)(b_1 + 1, b_1 + 2, \dots, b_2) \dots (b_{r-1} + 1, b_{r-1} + 2, \dots, b_r)$ and then use Ex. 1.6 to conclude that

$$\text{tr}(A_1 \otimes A_2 \otimes \dots \otimes A_k R_\sigma) = \text{tr}(A_{\pi(1)} \otimes \dots \otimes A_{\pi(k)} R_{\pi\sigma\pi^{-1}}) \quad (1.49)$$

$$= \text{tr}(A_{\pi(1)} \dots A_{\pi(b_1)}) \dots \text{tr}(A_{\pi(b_{r-1}+1)} \dots A_{\pi(b_r)}) \quad (1.50)$$

$$= \text{tr}(A_{a_1} \dots A_{a_{j_1}}) \dots \text{tr}(A_{a_{j_{r-1}+1}+1} \dots A_{a_k}). \quad (1.51)$$

For the important case when $A_1 = A_2 = \dots = A_k$, the final result is simplified to

$$\text{tr}(A^{\otimes k} R_\sigma) = \prod_{c \in \lambda(\sigma)} \text{tr}(A^c), \quad (1.52)$$

where $\lambda(\sigma)$ is the cycle structure of the permutation σ .

Symmetric subspace Throughout this course, symmetries under permutations will be a fundamental role, in particular the subspace of $(\mathbb{C}^d)^{\otimes k}$ composed of vectors that are left invariant by permutations:

$$\text{Sym}_{k,d} \equiv \text{Sym}((\mathbb{C}^d)^{\otimes k}) := \{\psi \in (\mathbb{C}^d)^{\otimes k} \mid R_\sigma |\psi\rangle = |\psi\rangle \forall \sigma \in S_k\}. \quad (1.53)$$

We will now show that the projector onto $\text{Sym}_{k,d}$ is

$$P_{\text{Sym},k,d} = \frac{1}{k!} \sum_{\sigma \in S_k} R_\sigma. \quad (1.54)$$

To see this, we first check that $P_{\text{Sym},k,d}$ is an orthogonal projector:

$$P_{\text{Sym},k,d}^2 = \frac{1}{(k!)^2} \sum_{\sigma, \pi \in S_k} R_{\sigma\pi} = \frac{1}{k!} \sum_{\sigma \in S_k} \frac{1}{k!} \sum_{\tau \in S_k} R_\tau = P_{\text{Sym},k,d}, \quad (1.55)$$

$$P_{\text{Sym},k,d}^\dagger = \frac{1}{k!} \sum_{\sigma \in S_k} R_{\sigma^{-1}} = \frac{1}{k!} \sum_{\pi \in S_k} R_\pi = P_{\text{Sym},k,d}, \quad (1.56)$$

where we substituted variables as $\tau = \sigma\pi$ and $\pi = \sigma^{-1}$, respectively, and used that the sum is invariant under the change of variables. Next, note that for all $\psi \in \text{Sym}_{k,d}$:

$$P_{\text{Sym},k,d} |\psi\rangle = \frac{1}{k!} \sum_{\sigma \in S_k} R_\sigma |\psi\rangle = \frac{1}{k!} \sum_{\sigma \in S_k} |\psi\rangle = |\psi\rangle, \quad (1.57)$$

thus, $\text{Sym}_{k,d}$ is in the range of $P_{\text{Sym},k,d}$. Moreover, if $P_{\text{Sym},k,d} |\psi\rangle = |\psi\rangle$, then

$$R_\pi |\psi\rangle = R_\pi P_{\text{Sym},k,d} |\psi\rangle = \frac{1}{k!} \sum_{\sigma \in S_k} R_{\pi\sigma} |\psi\rangle = \frac{1}{k!} \sum_{\tau \in S_k} R_\tau |\psi\rangle = P_{\text{Sym},k,d} |\psi\rangle = |\psi\rangle, \quad (1.58)$$

and thus the range of $P_{\text{Sym},k,d}$ is exactly $\text{Sym}_{k,d}$. We can now compute the dimension of the symmetric subspace using Eq. (1.46) and Ex. 1.4 as

$$\dim \text{Sym}_{k,d} = \text{tr } P_{\text{Sym},k,d} = \frac{1}{k!} \sum_{\sigma \in S_k} d^{\#\sigma} = \frac{1}{k!} \sum_{r=1}^k c(k,r) d^r = \frac{d(d+1) \cdots (d+k-1)}{k!} = \binom{d+k-1}{k}. \quad (1.59)$$

Graphical representation While the above methodologies are generic and straightforward, the algebra is often cumbersome. For this reason, it is useful to introduce a graphical notation to represent permutation. Similar to the Feynman diagrammatics for perturbative expansions, this is simply a bookkeeping of the operation previously described.

We denote permutations as lines connecting the list $\{1, 2, \dots, k\}$ to the output $\{\sigma(1), \sigma(2), \dots, \sigma(k)\}$. For example, given $\tau = (12)(35)(4)$ a 5 elements permutation, we can represent it as

$$(12)(35)(4) \cong \begin{array}{c} \begin{array}{c} 1 \text{ --- } \diagup \text{ --- } \\ 2 \text{ --- } \diagdown \text{ --- } \end{array} \\ \begin{array}{c} 3 \text{ --- } \diagup \text{ --- } \\ 4 \text{ --- } \diagdown \text{ --- } \\ 5 \text{ --- } \diagup \text{ --- } \end{array} \end{array} \quad (1.60)$$

This notation is particularly useful, since it makes computing products particularly easy. For example, the product of τ with $\sigma = (123)(4)(5)$, we simply need to follow the lines after connecting them, specifically

$$\sigma \cdot \tau \cong \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \cong \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \\ 5 \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \cong (235)(1)(4) \quad (1.61)$$

Since the group structure is the same for the representations of S_k , we can use the same notation also for the operators $\{R_\sigma : \sigma \in S_k\}$. For representations, the diagrammatic notation allows also to include traces and product by operators.

Traces require to add a curve that connect initial and final endpoints. For example:

$$\text{tr}(R_\sigma) = \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} = (\bigcirc)^{\#(\sigma)} = d^{\#(\sigma)}. \quad (1.62)$$

Similarly, for operators we have

$$\text{tr}(A^{\otimes k} R_\sigma) = \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} \text{---} \\ \text{---} \\ \text{---} \\ \text{---} \end{array} \begin{array}{c} A \\ A \\ A \\ A \end{array} = \text{tr}(A^3) \text{tr}(A)^2 \quad (1.63)$$

QUANTUM RANDOMNESS I

In this section, we lay the foundation for one of the central objects in quantum information theory and random quantum systems: the *Haar twirl*. This is the average over the unitary group of a replicated quantum channel, defined as

$$M_k(A) := \int_{U(d)} U^{\otimes k} A U^{\otimes k, \dagger} dU. \quad (2.1)$$

This map, known as the k -th unitary twirl, appears repeatedly throughout these notes and plays a fundamental role in both quantum information theory and many-body physics. It arises in contexts ranging from randomized benchmarking and quantum designs to entanglement theory, thermalization, and quantum chaos.

We begin by briefly discussing the structure of the Haar measure dU on the unitary group $U(d)$, and how random unitaries can be parametrized in practice. We then turn to the evaluation of the Haar twirl integral (2.1). The key observation is that $M_k(A)$ lies in the so-called *commutant* of $U^{\otimes k}$, i.e., the space of operators that commute with all $U^{\otimes k}$ for $U \in U(d)$. A powerful result from representation theory – the *Schur-Weyl duality* – tells us that this commutant is spanned by the action of the symmetric group via permutations on $\mathcal{H}^{\otimes k}$.

This insight allows us to express the Haar twirl as a linear combination over permutations, with coefficients given by the *Weingarten calculus* – a systematic method to evaluate integrals over the unitary group. The remainder of this section is devoted to developing these tools and applying them to compute $M_k(A)$ explicitly.

2.1 What is a Haar measure?

On the real line \mathbb{R} , there is a unique measure dx satisfying two simple but fundamental properties

$$d(x + y) = dx \quad (\text{translation invariance}), \quad \int_0^1 dx = 1 \quad (\text{unit volume}). \quad (2.2)$$

This measure represents what we intuitively mean by a *uniform* distribution: translation invariance ensures that no point is preferred over any other. However, since \mathbb{R} is non-compact, this measure assigns infinite volume to the full space, and thus cannot be normalized to a probability measure. To work around this, we typically restrict to a compact interval such as $[0, 1]$, which provides a natural normalization.

Remarkably, this idea of defining an invariant, uniform measure generalizes to far more abstract settings. Whenever the underlying set carries a suitable group structure, one can often define a natural measure that is invariant under group multiplication. More precisely, if G is a compact (Hausdorff topological) group then there exists a unique measure dg on G that is both

$$\text{left/right invariant} \quad dg = d(hg) = d(gh) \quad (2.3)$$

$$\text{normalized} \quad \int_G dg = 1. \quad (2.4)$$

This unique measure is called the *Haar (or uniform) measure* on G .

In this course, our primary interest lies in the unitary group $U(d)$, which is compact and satisfied the condition above. Concretely, the Haar measure dU on $U(d)$ satisfies

$$dU = d(UV) = d(VU) \quad \int_{U(d)} dU = 1. \quad (2.5)$$

This measure provides a rigorous definition of what it means to sample a unitary matrix uniformly at random. In principle, one could perform such an integration by explicitly parametrizing the unitaries and integrating over the associated coordinates. However, this approach becomes highly impractical as the dimension d increases, due to the complicated geometry and rapidly growing number of parameters. Instead, we will introduce more powerful techniques—based on symmetry and representation theory—that allow us to compute Haar integrals in an efficient and conceptually clean way.

Example 2.1: Cycle Structure

Consider $d = 2$ and let us compute the Haar average

$$\int_{U(2)} dU UAU^\dagger \quad (2.6)$$

for a general operator

$$A = \begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}. \quad (2.7)$$

A generic unitary on $U(2)$ is given by

$$U(\theta, \phi, \psi) = \begin{pmatrix} \cos \theta & e^{i\phi} \sin \theta \\ e^{i\psi} \sin \theta & e^{i(\phi+\psi)} \cos \theta \end{pmatrix} \quad (2.8)$$

with the Haar measure $dU = (4\pi^2)^{-1} \sin(2\theta) d\theta d\phi d\psi$ with $\theta \in [0, \pi/2]$ and $\phi, \psi \in [0, 2\pi]$. (We ignore the global phase since it cancels trivially in UAU^\dagger . The integral acts on each element of the total matrix $(UAU^\dagger)_{ij} = \sum_{k,l} U_{i,k} A_{k,l} \bar{U}_{j,l}$. For concreteness, let us compute only the entry $i = j = 1$: $(U^\dagger AU)_{11} = u_{11}a_{11}\bar{u}_{11} + u_{11}a_{12}\bar{u}_{12} + u_{12}a_{21}\bar{u}_{11} + u_{22}a_{22}\bar{u}_{22}$. Using the well known formula $\int d\alpha e^{i n \alpha} = \delta_{n,0}$, we find that the terms a_{12} and a_{21} simplify. The remaining computation requires

$$\int_0^{\pi/2} d\theta \sin(2\theta) (\cos^2 \theta a_{11} + \sin^2 \theta a_{22}) = \frac{1}{2} (a_{11} + a_{22}). \quad (2.9)$$

Working out the details for the other indices, we obtain

$$\int_{U(2)} dU UAU^\dagger = \frac{1}{2} \begin{pmatrix} a_{11} + a_{22} & 0 \\ 0 & a_{11} + a_{22} \end{pmatrix} = \frac{1}{2} \text{tr}(A) \mathbb{I}. \quad (2.10)$$

2.1.1 The unitary commutant

In this section, we study the subspace of operators that commute with $U^{\otimes k}$, which we call the $(k\text{-fold})$ commutant of $U(d)$:

$$\text{Comm}_k = \{A \in L(\mathcal{H}^{\otimes k}) \mid U^{\otimes k} A = A U^{\otimes k} \forall U \in U(d)\}. \quad (2.11)$$

We care about this because of the following, elementary observation:

Lemma 2.1. $M_k(A) \in \text{Comm}_k$ for all $A \in L(\mathcal{H}^{\otimes k})$. In fact, every element in Comm_k is of the form $M_k(A)$.

Proof. Clearly, if A commutes with all $U^{\otimes k}$, then $M_k(A) = A$, i.e. every element in Comm_k is of the form $M_k(A)$. Vice versa, if $B = M_k(A)$, then, by the invariance of the Haar measure:

$$U^{\otimes k} B U^{\otimes k, \dagger} = \int_{U(d)} (UV)^{\otimes k} A (UV)^{\otimes k, \dagger} dV = \int_{U(d)} V^{\otimes k} A V^{\otimes k, \dagger} dV = B, \quad \forall U \in U(d). \quad (2.12)$$

Thus, $B \in \text{Comm}_k$. □

A central step in understanding the commutant is to note that the representation of S_k on $\mathcal{H}^{\otimes k}$ introduced in Sec. 1.2.2, this is $R_\pi |x_1, \dots, x_k\rangle = |x_{\pi(1)}, \dots, x_{\pi(k)}\rangle$, clearly commutes with $U^{\otimes k}$. In representation-theoretic terms, the representations of $U(d)$ and S_k are said to be *dual to each other*. This implies that the permutations are contained in the unitary commutant, $R_\pi \in \text{Comm}_k$ for all $\pi \in S_k$. A fundamental result in representation theory, **Schur-Weyl duality**, even states that every element in the commutant is a linear combination of permutations. We will only state this result here and refer for a proof to the literature [tbd].

Theorem 2.1 (Schur-Weyl duality). *The k -fold unitary commutant Comm_k is spanned by $\{R_\sigma \mid \sigma \in S_k\}$. Vice versa, the commutant of $\{R_\sigma \mid \sigma \in S_k\}$ is spanned by $\{U^{\otimes k} \mid U \in U(d)\}$*

A natural question to ask is whether the permutations form a basis for the commutant. Intriguingly, this is the case if the dimension d is large enough:

Lemma 2.2. *The set $\{R_\sigma \mid \sigma \in S_k\}$ is linearly independent for $d \geq k$.*

Proof. We consider the standard basis of \mathbb{C}^d , which we here denote as $|1\rangle, \dots, |d\rangle$. Since $k \leq d$, we can consider the action of permutations on $|1, \dots, k\rangle \in (\mathbb{C}^d)^{\otimes k}$:

$$R(\pi)|1, \dots, k\rangle = |\pi(1), \dots, \pi(k)\rangle. \quad (2.13)$$

Now, if $R(\pi)$ and $R(\sigma)$ would be linearly dependent ($\pi \neq \sigma$), then so would be the states $|\pi(1), \dots, \pi(k)\rangle$ and $|\sigma(1), \dots, \sigma(k)\rangle$. However, these are distinct elements from a basis, thus we arrive at a contradiction. \square

Remark 2.1. *In fact, permutations become linearly dependent as soon as $k > d$. We do not need this statement in the following, thus we will not treat the proof in the lecture. We however state it here for completeness. We consider the antisymmetric subspace $\text{Alt}_{k,d} \subset (\mathbb{C}^d)^{\otimes k}$ which is the joint -1 eigenspace of all transpositions $R((ij))$ for $(ij) \in S_k$. The projector onto $\text{Alt}_{k,d}$ has the general form*

$$P_{\text{Alt},k,d} = \frac{1}{k!} \sum_{\pi \in S_k} \text{sgn}(\pi) R(\pi), \quad (2.14)$$

where the sign function $\text{sgn}(\pi)$ is given as follows: Decompose π into transpositions only, then count the number of transpositions needed. If it is even, $\text{sgn}(\pi) = 1$, else $\text{sgn}(\pi) = -1$. Now, $\dim \text{Alt}_{k,d} = \binom{d}{k} = 0$ if $k > d$, and hence $P_{\text{Alt},k,d} = 0$. This gives a non-trivial linear relation between permutations, i.e. they are linearly dependent.

2.1.2 Weingarten calculus

By the previous findings, we can always write $M_k(A)$ as a linear combination

$$M_k(A) = \sum_{\pi \in S_k} c_\pi(A) R_\pi, \quad (2.15)$$

for some coefficients $c_\pi(A)$. Note that taking the trace inner product of $M_k(A)$ with a fixed permutation R_σ^\dagger yields

$$\text{tr}(R_\sigma^\dagger M_k(A)) = \int \text{tr}(R_\sigma^\dagger U^{\otimes k} A U^{\otimes k, \dagger}) dU = \int \text{tr}(U^{\otimes k, \dagger} R_\sigma^\dagger U^{\otimes k} A) dU = \text{tr}(R_\sigma^\dagger A). \quad (2.16)$$

However, we also have

$$\text{tr}(R_\sigma^\dagger A) = \text{tr}(R_\sigma^\dagger M_k(A)) = \sum_{\pi \in S_k} c_\pi(A) \text{tr}(R_\sigma^\dagger R_\pi) =: \sum_{\pi \in S_k} c_\pi(A) G_{\sigma, \pi}, \quad (2.17)$$

where we defined the *Gram matrix*

$$G_{\pi, \sigma} := \text{tr}(R_\pi^\dagger R_\sigma) = \text{tr}(R_\pi^\dagger R_\sigma) = \text{tr}(R_{\pi^{-1}\sigma}) = d^{\#(\pi^{-1}\sigma)}. \quad (2.18)$$

Here, we used that R is a representation to combine the product of permutation operators, and the trace formula (1.46). Setting $a_\sigma := \text{tr}(R_\sigma^\dagger A)$, we can write the above equation in matrix form as $a = Gc$, which we could hope to invert to get an expression for the coefficient vector c . Note that the permutations are not orthogonal with respect to the trace inner product $(A|B) = \text{tr}(A^\dagger B)$, and hence the Gram matrix is not simply diagonal. However, we know that the permutations span the commutant and that $M_k(A)$ lies in the commutant. Hence, the equation $a = Gc$ always has a solution and it is unique if and only if the permutations form a basis, i.e. iff $d \geq k$, which is what will always assume for the remainder of this course.¹ Then, this solution is simply $c = G^{-1}a$, or put differently,

$$M_k(A) = \sum_{\pi, \sigma \in S_k} W_{\pi, \sigma} \text{tr}(R_\sigma^\dagger A) R_\pi, \quad (2.19)$$

where we defined $W := G^{-1}$, the so-called *Weingarten matrix*. Knowing the Weingarten matrix allows us to compute integrals of the form (2.1) using the Weingarten expansion (2.19).

Properties of the Gram and Weingarten matrix The Gram and Weingarten matrix have a substantial structure which directly relates to the representation theory of the symmetric group. We will not dive into these details in this course, but instead prove some concrete relations. We summarize them in the following.

Lemma 2.3. *The Gram and Weingarten matrix fulfill the following properties.*

- (a) $G_{\pi, \sigma}$ and $W_{\pi, \sigma}$ only depend on $\pi^{-1}\sigma$.
- (b) The row and column sums of G are constant:

$$\mathcal{G}_{k,d} := \sum_{\sigma} G_{\pi, \sigma} = \sum_{\pi} G_{\pi, \sigma} = \frac{(d+k-1)!}{(d-1)!} = d(d+1) \cdots (d+k-1). \quad (2.20)$$

- (c) The row and column sums of W are constant:

$$\sum_{\sigma} W_{\pi, \sigma} = \sum_{\pi} W_{\pi, \sigma} = \mathcal{G}_{k,d}^{-1} = \frac{(d-1)!}{(d+k-1)!}. \quad (2.21)$$

Proof. (a) Clearly, $G_{\pi, \sigma}$ depends only on $\pi^{-1}\sigma$ by definition (cf. Eq. (2.18)). Now note that this implies that G is invariant under simultaneous row and column permutations. Indeed, if T_τ is the permutation matrix acting as $T_\tau|e_\pi\rangle = |e_{\tau\pi}\rangle$, then $(T_\tau^{-1}GT_\tau)_{\sigma, \pi} = G_{\tau\sigma, \tau\pi} = G_{\sigma, \tau}$. Inverting $G = T_\tau^{-1}GT_\tau$ yields $W = T_\tau^{-1}WT_\tau$ and thus $W_{\pi, \sigma} = W_{\tau\pi, \tau\sigma}$ for all τ , in particular $W_{\pi, \sigma} = W_{\text{id}, \pi^{-1}\sigma}$ for $\tau = \pi^{-1}$. For (b), we compute

$$\mathcal{G}_{k,d} = \sum_{\sigma} \text{tr}(R_{\pi^{-1}\sigma}) = \sum_{\sigma} \text{tr}(R_\sigma) = k! \text{tr}(P_{\text{Sym}, k, d}) = k! \binom{d+k-1}{k} = \frac{(d+k-1)!}{(d-1)!}. \quad (2.22)$$

Here, we used that the multiplication by π^{-1} can be absorbed into the sum (variable change), and the definition of the projector onto the symmetric subspace, $P_{\text{Sym}, k, d} = \frac{1}{k!} \sum_{\sigma} R_\sigma$, and the value of its trace, cf. Eqs. (1.54) and (1.59). For (c), we note that the definition of W as inverse of G implies

$$\sum_{\pi} W_{\sigma, \pi} G_{\pi, \tau} = \delta_{\sigma, \tau} \quad \Rightarrow \quad 1 = \sum_{\pi, \tau} W_{\sigma, \pi} G_{\pi, \tau} = \mathcal{G}_{k,d} \sum_{\pi} W_{\sigma, \pi} \quad \Rightarrow \quad \sum_{\pi} W_{\sigma, \pi} = \mathcal{G}_{k,d}^{-1}. \quad (2.23)$$

□

¹It is however not terribly complicated to make this work for $d < k$, see Sec. 2.3.

Some examples and exercises In the following, we will compute the Weingarten matrix for small values of k and illustrate the computations of Haar integrals using a number of examples. To this end, we use that the Gram matrix has a very simple form: It is the trace of a permutation $R_\tau = R_\pi^\dagger R_\sigma$ and we gave an expression for this in Eq. (1.46).

Example 2.2: Weingarten matrix for $k = 2$

Let us consider $k = 2$. Then we only have two permutations: the identity $\mathbb{1}$ and the flip/swap $\mathbb{F} = (2\ 1)$. There is only one non-trivial matrix element, namely $G_{\mathbb{1},\mathbb{F}} = \text{tr}(\mathbb{F}) = d$. Hence, the Gram and Weingarten matrices are

$$G = d^2 \begin{pmatrix} 1 & d^{-1} \\ d^{-1} & 1 \end{pmatrix}, \quad W = \frac{1}{d^2 - 1} \begin{pmatrix} 1 & -d^{-1} \\ -d^{-1} & 1 \end{pmatrix}. \quad (2.24)$$

Example 2.3: Average collision probability

The probability of obtaining the computational basis state x on $U|0\rangle$ is $p(x|U) = |\langle x|U|0\rangle|^2$. A measure of flatness of this distribution is the *collision probability*:

$$Z_U := \sum_x p(x|U)^2 = \sum_x |\langle x|U|0\rangle|^4. \quad (2.25)$$

Here, we are interested on how flat this distribution is *on average*, over Haar-random unitaries U . Due to the invariance of the Haar measure, we can simply absorb the X gates that prepare $|x\rangle = X^{x_1} \otimes \cdots \otimes X^{x_n}|0\rangle =: X(x)|0\rangle$ into the average:

$$Z := \int Z_U dU = \sum_x \int_U |\langle 0|X(x)U|0\rangle|^4 dU = d \int_U |\langle 0|U|0\rangle|^4 dU. \quad (2.26)$$

To compute the integral, we use second-order Weingarten calculus:

$$Z = d \int_U \text{tr}(|0\rangle\langle 0|^{\otimes 2} U^{\otimes 2} |0\rangle\langle 0|^{\otimes 2, \dagger}) dU \quad (2.27)$$

$$= d \sum_{\pi, \sigma \in S_2} W_{\pi, \sigma} \text{tr}(R_\sigma^\dagger |0\rangle\langle 0|^{\otimes 2}) \text{tr}(R_\pi |0\rangle\langle 0|^{\otimes 2}) \quad (2.28)$$

$$= d \sum_{\pi, \sigma \in S_2} W_{\pi, \sigma} \quad (2.29)$$

$$= 2d \mathcal{G}_{2,d}^{-1} = 2d \frac{(d-1)!}{(d+1)!} = \frac{2}{d+1}. \quad (2.30)$$

Here, we used Lem. 2.3. Note that we haven't actually used the exact form of the Weingarten matrix from Ex. 2.2. In fact, along the same lines we find that the average of $\sum_x p(x|U)^k$ is

$$I_k := \sum_x \int p(x|U)^k dU = k! d \mathcal{G}_{k,d}^{-1} = \frac{k! d!}{(d+k-1)!}. \quad (2.31)$$

Exercise 2.1. Using Weingarten calculus, compute the operator $S := d \int (U|0\rangle\langle 0|U^\dagger)^{\otimes 2} dU$.

2.2 Approximating the Haar measure: Unitary designs

In this chapter, we have seen how Haar integrals over the unitary group can be computed using Weingarten calculus. These tools will help us to design randomized protocols and algorithms for applications, and we will see an example for that already in the next chapter. In practice, Haar-random unitaries are however way too *expensive* to be useful: On a system with n qudits, a quantum




Figure 2.1: Caricature showing a “series expansion” of the Haar measure on the full unitary group (here depicted by a sphere) into finite subsets that agree with the k -th Haar moment. Taken from Kueng and Gross [1].

circuit implementing a Haar-random unitary requires a circuit depth that is exponential in n . This means that any quantum algorithm that utilizes Haar-random unitaries necessarily has exponential runtime.

In this section, we thus briefly discuss a concept introduced to lower the requirements for quantum randomness: *unitary k -designs*. These are sets of unitaries that mimic the Haar measure up to the k -th moment and can thus be used as a replacement in applications that rely on finite moments only. Importantly, unitary designs may be constructed using significantly less resources than Haar-random unitaries, in particular using polynomial-sized circuits.

2.2.1 Unitary designs

Formally, we define a unitary design as follows. Let $\mathcal{G} \subset \mathbf{U}(d)$ be a (finite) set of unitaries (we can extend this to infinite sets equipped with a suitable probability measure). Then, we call \mathcal{G} a *unitary k -design* if

$$\frac{1}{|\mathcal{G}|} \sum_{U \in \mathcal{G}} U^{\otimes k} A U^{\otimes k, \dagger} = M_k(A), \quad (2.32)$$

for all matrices A . We call the largest k for which Eq. (2.32) holds the *order* of the unitary design. Unitary designs can be seen as a set of points in the unitary group that are “sufficiently equally distributed” to reproduce the first moments of the Haar measure, cf. Fig. 2.1.

Note that the left hand side of Eq. (2.32) can be seen as the k -fold twirl over the set \mathcal{G} . This means that averages over the set \mathcal{G} can be computed using averages over the full unitary group – for instance using Weingarten calculus. Vice versa, in applications that only depend on k -fold twirls, Haar-random unitaries can be replaced by unitary k -designs, thereby enabling more resource-efficient and flexible implementations.

One might hope that one could find sufficiently symmetric subgroups of $\mathbf{U}(d)$ that gives natural candidates for unitary designs. However, it turns out that Eq. (2.32) together with the group structure imposes very strict conditions on such a subgroup, resulting in the fact these do not exist for $k \geq 4$ (and $d \geq 5$). Despite this, the most important example of a unitary design (with $k = 3$) is in fact a group, the *Clifford group*, which we will introduce in the next section.

Nevertheless, unitary k -designs exist for all k and dimensions d . Unfortunately, explicit constructions of general unitary k -designs are incredibly rare and the known ones are highly inefficient. To overcome these obstacles, it has been fruitful to demand that Eq. (2.32) holds only approximately – such *approximate unitary designs* can be realized much more easily and enable the efficient implementation of quantum randomness also for higher moments k . We will come back to this point in Ch. 4.

2.2.2 The Clifford group

The prototypical example of a unitary design is the *Clifford group*. Besides designs, the Clifford group plays a major role in quantum error correction, where it typically constitutes the set of “easy” gates in

fault-tolerant quantum computing. Clifford operations can also be efficiently simulated on a classical computer, making them the starting point for classical simulation algorithms and investigations of the “non-classicality” of quantum mechanics.

The simplest way to define the Clifford group is via its local generators: The single-qubit *phase gate* S and *Hadamard gate* H , as well as the two-qubit CNOT gate CX , given by

$$S|x\rangle = i^x|x\rangle \quad H|x\rangle = \frac{1}{\sqrt{2}}(|0\rangle + (-1)^x|1\rangle) \quad CX|x, y\rangle = |x, y \oplus x\rangle, \quad (2.33)$$

where $y \oplus x$ denotes addition modulo 2. The group that is generated by S and H on every qubit and CX on every pair of qubits is the n -qubit *Clifford group* Cl_n . It is a finite subgroup of $\text{U}(2^n)$ with $2^{O(n^2)}$ elements. Importantly, every Clifford unitary can be implemented using $O(n^2)$ generators S , H , and CX .

An important subgroup of the Clifford group is the *Pauli group*. Recall the Pauli operators

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad (2.34)$$

which we complement with the identity $\mathbb{1}$. Then, the multi-qubit Pauli operators on $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ are simply given by all possible tensor products of the single-qubit Pauli operators, in formula:

$$\sigma = \sigma_1 \otimes \cdots \otimes \sigma_n, \quad \sigma_i \in \{\mathbb{1}, X, Y, Z\}. \quad (2.35)$$

The n -qubit Pauli operators form a group up to phases, the so-called n -qubit Pauli group:

$$\text{P}_n := \{i^t \sigma_1 \otimes \cdots \otimes \sigma_n \mid t \in \mathbb{Z}_4, \sigma_i \in \{\mathbb{1}, X, Y, Z\}\}. \quad (2.36)$$

To see that P_n forms a subgroup of Cl_n , note that $Z = S^2$, $HZH = X$, and $Y = iXZ$. It turns out that we can characterize the Clifford group uniquely through the Pauli group: Clifford unitaries are exactly those unitaries that conjugate Paulis into Paulis, up to a phase.² In formula, we thus have

$$\text{Cl}_n \cdot \text{U}(1) = \{U \in \text{U}(2^n) \mid U \text{P}_n U^\dagger = \text{P}_n\}. \quad (2.37)$$

Note that we had to add arbitrary global phases ($\text{U}(1)$) to Cl_n as these are present on the right hand side as well.

The qudit case. The definition of the Clifford group can be readily extended to higher-dimensional qudits of dimension q . We will here focus on the case that $q > 2$ is **prime**, as this will matter for the design properties of the Clifford group.

We start by generalizing the generators of the qubit Clifford group. To this end, let $\omega_q := e^{2\pi i/q}$ be a primitive q -th root of unity, and let 2^{-1} be the inverse of 2 modulo q . Define

$$H_q|x\rangle := \frac{1}{\sqrt{q}} \sum_{y=0}^{q-1} \omega_q^{xy} |y\rangle \quad S_q|x\rangle := \omega_q^{2^{-1}x(x-1)} |x\rangle \quad CX_q|x, y\rangle := |x, y \oplus x\rangle, \quad (2.38)$$

where $y \oplus x$ denotes addition modulo q . Then, we again define the n -qudit Clifford group $\text{Cl}_n(q)$ to be the group generated by H_q , S_q on every qudit and CX_q on every pair of qudits. This is again a finite subgroup of $\text{U}(q^n)$ of order $q^{O(n^2)}$.

Equivalently, we can define $\text{Cl}_n(q)$ in terms of a qudit version of Pauli operators, defined by

$$Z_q|x\rangle := \omega_q^x |x\rangle, \quad X_q|x\rangle := |x \oplus 1\rangle, \quad Y_q := \omega_q^{2^{-1}} X_q^\dagger Z_q^\dagger. \quad (2.39)$$

The qudit Pauli group is then given as

$$\text{P}_n(q) := \{\omega_q^{t_0} \sigma_1^{t_1} \otimes \cdots \otimes \sigma_n^{t_n} \mid t_i \in \mathbb{Z}_q, \sigma_i \in \{X_q, Y_q, Z_q\}\}. \quad (2.40)$$

As in the qubit case, qudit Clifford unitaries map qudit Paulis to qudit Paulis, up to a phase.

²In group-theoretic terms, the Pauli group is a normal subgroup of the Clifford group, and the Clifford group is the normalizer of the Pauli group within the unitary group.

The Clifford group as a design. Finally, we have the following result on the design properties of the Clifford group.

Theorem 2.2. *Let q be prime. Then, the Clifford group $\text{Cl}_n(q)$ is a unitary 2-design but not a 3-design if q is odd, and a unitary 3-design but not a 4-design if $q = 2$.*

We will not prove this theorem here. It is not particularly difficult (see e.g. Ref. [2, Sec. 12.2] for a summary), but it requires a slightly deeper analysis of the structure of the Clifford group which shall not be the focus of this course.

2.3 Further reading

To be done.

MEASURING PROPERTIES OF MANY-BODY STATES USING CLASSICAL SHADOWS

Quantum experiments are getting better and better in coherently manipulating many-body quantum systems. We can use this to study interesting many-body phenomena by preparing exotic and weird quantum states. In quantum computers, a high level of control is necessary to manipulate quantum information and perform quantum computations.

However, even with perfect quantum control, we are left with the problem of extracting information from the prepared quantum states. An example for this would be the expectation value of a given Hamiltonian, because we are trying to find its ground state through a variational quantum algorithm. We could also be interested in more complicated properties, such as the entanglement of the state across some bipartition to verify or reject area laws.

In this chapter, we consider the problem of experimentally estimating an expectation value of the form $\text{tr}(O\rho)$, where ρ is the quantum state that is experimentally accessible and O is the observable of interest. Importantly, we do not assume that O can be measured directly. In the following, we will show that for some interesting classes of observables this problem can be solved using *classical shadows*. The main idea is to randomize measurement bases using a suitable ensemble of random unitaries, resulting in a partial classical representation of the quantum state (the “shadow”). The shadow can then be used by a classical computer to predict expectation values. Intriguingly, the same classical shadow can be used to predict expectation values of many observables at once, at a moderate (logarithmic) overhead in the size of the shadow (i.e. the number of measurements).

We will first review some basics on quantum measurements and the (partial) reconstruction of quantum states from measurement data. Afterwards, we introduce the idea of classical shadows and treat most common instances of the protocol based on Clifford unitaries. We will also hint briefly at other usecases discussed in the literature.

3.1 Reconstructing quantum states from measurements

Suppose we have access to many copies of a quantum state ρ and we want to obtain a classical description of ρ from measurement statistics, e.g., in the form of a density matrix. This reconstruction task is also called *quantum state tomography*. The central question is how should we choose our measurements such that this reconstruction succeeds (for any state), and how many measurements do we need in total?

Let us consider quantum measurements in a basis $|\varphi_1\rangle, \dots, |\varphi_d\rangle$ of the Hilbert space \mathcal{H} . The probability distribution over outcomes of this measurement follows Born’s rule:

$$\text{Born's rule: } p(i|\rho) := \langle \varphi_i | \rho | \varphi_i \rangle = \text{tr}(E_i \rho) \quad \text{where } E_i := |\varphi_i\rangle\langle \varphi_i|. \quad (3.1)$$

Clearly, measuring in a basis cannot be enough for state reconstruction because we can only assess the diagonal elements of ρ in the chosen basis. Hence, we are insensitive to any coherence in the basis: For instance, we cannot distinguish between any two states of the form $|0\rangle + e^{i\alpha}|1\rangle$ by looking at the diagonal elements of their density matrices only.

However, it turns out that if we combine measurements in sufficiently many bases, state reconstruction is possible and we will explicitly describe such a reconstruction algorithm in a moment. Before, we introduce some new notation that will simplify both the following computations as well as those in the upcoming chapters.

Operator bra-ket notation. Recall that the vector space $L(\mathcal{H})$ of linear operators on \mathcal{H} is equipped with the *trace inner product*:

$$(A|B) := \text{tr}(A^\dagger B). \quad (3.2)$$

With this notation, we can now express Born probabilities as

$$p(i|\rho) = (E_i|\rho), \quad (3.3)$$

with the interpretation of projecting the state ρ onto the pure state $E_i = |\varphi_i\rangle\langle\varphi_i|$. Furthermore, analogous to the ordinary bra-ket notation based on the inner product $\langle\cdot|\cdot\rangle$ on \mathcal{H} , we now introduce *operator bra-ket notation* by defining *operator kets* and *operator bras* as

$$|B\rangle := B, \quad (A| : B \equiv |B\rangle \mapsto (A|B). \quad (3.4)$$

Note that operator bras $(A|$ are linear forms on $L(\mathcal{H})$ (dual vectors), just as ordinary bras $\langle\psi|$ are dual vectors on \mathcal{H} . Again similar to the ordinary bra-ket notation, we introduce *operator bra-kets* as the outer products

$$|A\rangle(B| : C \equiv |C\rangle \mapsto |A\rangle(B|C), \quad (3.5)$$

As we will see in a moment, this gives us a convenient way to write down linear maps on operators, similar to expressions of the form $A = \sum_{i,j} A_{i,j} |i\rangle\langle j|$ for the ordinary bra-ket notation. For instance, the Weingarten expansion (2.19) attains the following appealing form in this notation:

$$M_k = \sum_{\pi, \sigma \in S_k} W_{\pi, \sigma} |R_\pi\rangle\langle R_\sigma|, \quad (3.6)$$

Following the quantum information language, we will refer to linear maps on operators, such as (3.5), as *superoperators* (as these are “operators on operators”). An important example of superoperators are *quantum channels*.

Reconstruction via linear inversion. Equipped with the new notation, let us come back to the reconstruction of quantum states from measurements. Consider the following superoperator:

$$S := \sum_{i=1}^d |E_i\rangle\langle E_i|. \quad (3.7)$$

We call S the *frame (super)operator* associated to the basis $(\varphi_i)_{i \in [d]}$. In the classical shadows literature S is also called the *measurement channel*. Note that we have $S(\rho) = \sum_{i=1}^d p(i|\rho) E_i$ for any state ρ and in this sense, the inability of reconstructing a state from a basis measurement is encoded in the non-injectivity of S .

However, it turns out that if we combine measurements in sufficiently many bases, we can perform successful state reconstruction. Suppose we are given bases $(\varphi_{i,j})_{i \in [d]}$ for $j = 1, \dots, m$ and we perform measurements in any of those. The frame operator of this combined measurement strategy is a convex combination of the single frame operators:

$$S = \frac{1}{m} \sum_{j=1}^m S_j = \frac{1}{m} \sum_{j=1}^m \sum_{i=1}^d |E_{i,j}\rangle\langle E_{i,j}|. \quad (3.8)$$

Suppose that this combined frame operator S is invertible, which is –as it turns out– necessary for reconstruction (cf. Ex. 3.1). Then, we can do the simply manipulation

$$\rho = S^{-1}S(\rho) = \frac{1}{m} \sum_{j=1}^m \sum_{i=1}^d S^{-1}|E_{i,j}\rangle\langle E_{i,j}|\rho = \frac{1}{m} \sum_{j=1}^m \sum_{i=1}^d p(i, j|\rho) \tilde{E}_{i,j}, \quad (3.9)$$

where $p(i, j|\rho) = (E_{i,j}|\rho)$ are again the Born probabilities and $\tilde{E}_{i,j} := S^{-1}(E_{i,j})$ are the *dual* measurement elements. This gives as a simply recipe to reconstruct the state ρ from the Born probabilities $p(i, j|\rho)$ obtained through measurements, which is typically called *linear inversion tomography*.

As it is generally the case for quantum state tomography, we need a lot of copies of ρ , i.e. measurements, to approximate ρ through the formula (3.9), namely at least $dr^2\varepsilon^{-2}$ many, where d is the Hilbert space dimension, $r = \text{rank}(\rho)$, and ε is the desired precision in trace distance. This means that the measurement effort of quantum state tomography scales exponentially with the number of qudits in the system and is thus limited to small systems only.

Exercise 3.1 (Informationally complete measurements). *In general, state reconstruction can only be successful if the Born probabilities differ for any two states $\rho \neq \rho'$. If this is the case, we call the measurement informationally complete (IC).*

- (a) *Show that measurements in several bases is informationally complete if and only any operator $X \in \mathbb{C}^{d \times d}$ can be written as a linear combination of the $E_{i,j}$, this is $X = \sum_{i,j} x_{i,j} E_{i,j}$ (i.e. the $E_{i,j}$ span the space of operators). Hint: Consider the linear map $V(x) := \sum_{i,j} x_{i,j} E_{i,j}$ from $\mathbb{C}^{d \times d}$ to $L(\mathcal{H})$ and its adjoint V^\dagger (bra-ket notation might be useful).*
- (b) *Show that measurements in several bases is informationally complete if and only if the frame operator (3.7) is invertible. Hint: Show that $S = \frac{1}{m} V V^\dagger$ where V is the map from (b).*

Exercise 3.2. *For those interested in quantum information: Show that the frame operator (3.7) is indeed a quantum channel by computing its Choi matrix.*

3.2 Shadow estimation with randomized measurements

In the previous section, we have seen that quantum state tomography can be realized based on simple linear inversion, but it requires exponentially many measurements rendering it very inefficient. But what if we are not interested in reconstructing the full state ρ , but only some of its *features*? Concretely, let us say that we want to reconstruct M linear functions of ρ , which we can write as $(O_s | \rho)$ ($s = 1, \dots, M$). We call this task *shadow tomography* as we do not observe the full state, but only some of features, similar to the shadow of an object that is illuminated from a certain direction, cf. Fig. 3.1. Can shadow tomography be done more efficiently than performing full quantum state tomography?

The answer is *yes, but it depends*, namely on the allowed measurement strategy (can we access only single copies of ρ , or some of them at once) and also on the observables O_s . In the following, we will focus on single-copy measurements and show how randomized measurements can be used to estimate some classes of expectation values very efficiently.




Figure 3.1: Shadow tomography is the task of determining only some of the features of a quantum state without performing full quantum state tomography. We do this by considering only few measurements of the state akin to the shadows of an object under few illumination angles as shown in this cartoon. Figure taken from the popular summary of Ref. [3].

We thus study measurements in random rotations of the computational basis, i.e. measurements are performed in the bases

$$|U, x\rangle := U^\dagger |x\rangle, \quad E_{U,x} := |U, x\rangle\langle U, x| = U^\dagger |x\rangle\langle x| U, \quad x \in [d], \quad (3.10)$$

where the unitary U is sampled uniformly at random from a finite set $\mathcal{G} \subset U(d)$ (here the adjoint comes from using the Heisenberg picture). The frame operator is then a probabilistic mixture of the frame operators in the rotated bases (3.10):

$$S = \frac{1}{|\mathcal{G}|} \sum_{U \in \mathcal{G}} \sum_{x=1}^d |E_{U,x}\rangle\langle E_{U,x}|. \quad (3.11)$$

The underlying observation is that we can use the linear inversion trick (3.9) to obtain the following identity for expectation values:

$$(O | \rho) = (O | S^{-1} S | \rho) = \frac{1}{|\mathcal{G}|} \sum_{U \in \mathcal{G}} \sum_{x=1}^d (O | S^{-1} | E_{U,x}\rangle\langle E_{U,x}| \rho) = \sum_{U \in \mathcal{G}} \sum_{x=1}^d (O | \tilde{E}_{U,x}) p(x|U, \rho) p(U), \quad (3.12)$$

where we set $p(x|U, \rho) = (E_{U,x}|\rho)$ and $p(U) = 1/|\mathcal{G}|$. We chose the later notation to highlight that we are sampling pairs (U, x) according to the joint distribution $p(x|U, \rho)p(U)$. This suggests to estimate the expectation value $(O|\rho)$ using the following protocol.

Protocol 3.1: Shadow estimation

Repeat the following steps N times:

- (i) Sample a unitary U uniformly at random from the set \mathcal{G} and apply it to ρ
- (ii) Measure in the computational basis resulting in outcome x
- (iii) Record the pair (U, x)

Estimate $(O|\rho)$ using the mean estimator of $\hat{\rho}_{U,x} := (O|\tilde{E}_{U,x})$ on the N samples.

Equation (3.12) then guarantees that shadow estimation 3.1 converges in expectation to $(O|\rho)$ (i.e. $\hat{\rho}$ is an unbiased estimator). But how fast does this estimator converge, i.e. how many measurements N do we have to perform? As we will see below, this depends a lot on the used ensemble \mathcal{G} and on the observable O (as well as on the state).

To answer the question on sample complexity, as well as other questions on the efficiency and practicability of the shadow estimation protocol 3.1, we have to analyze it in more depth. As a first step we will compute the frame operator (3.7) and its inverse, as it is needed in the construction of the estimator $\hat{\rho}_{U,x} = (O|S^{-1}|E_{U,x})$. The frame operator involves two invocations of U and U^\dagger , respectively, which we can explicitly see by considering the matrix elements

$$(A|S|B) = \sum_{x=1}^d \frac{1}{|\mathcal{G}|} \sum_{U \in \mathcal{G}} (A|E_{U,x})(E_{U,x}|B) = \sum_{x=1}^d \frac{1}{|\mathcal{G}|} \sum_{U \in \mathcal{G}} \text{tr}(A^\dagger \otimes B U^{\otimes 2}|x\rangle\langle x|^{\otimes 2} U^{\otimes 2,\dagger}) \quad (3.13)$$

Hence, if we take \mathcal{G} to be a unitary 2-design, we can replace the average over \mathcal{G} in Eq. (3.13) by an average over the unitary group, computable via Weingarten calculus. In the following, we will see that this brings S into a very simple form.

3.2.1 Clifford measurements

Computing the shadow estimator. Let us thus assume that we are randomizing over a unitary 2-design, or equivalently, over Haar-random unitaries, and compute the frame operator by $k = 2$ Weingarten calculus. To this end, we evaluate the integral (cf. Example 2.3):

$$(A|S|B) = \sum_{x=1}^d \int_{\mathcal{U}(d)} \text{tr}(A^\dagger \otimes B U^{\otimes 2}|x\rangle\langle x|^{\otimes 2} U^{\otimes 2,\dagger}) dU \quad (3.14)$$

$$= \sum_{x=1}^d \sum_{\pi, \sigma \in \mathbb{S}_2} W_{\pi, \sigma} \text{tr}(R_\sigma^\dagger A^\dagger \otimes B) \text{tr}(R_\pi |x\rangle\langle x|^{\otimes 2}) \quad (3.15)$$

$$= d \sum_{\pi, \sigma \in \mathbb{S}_2} W_{\pi, \sigma} \text{tr}(R_\sigma^\dagger A^\dagger \otimes B) \quad (3.16)$$

$$= d \frac{(d-1)!}{(d+1)!} \sum_{\sigma \in \mathbb{S}_2} \text{tr}(R_\sigma^\dagger A^\dagger \otimes B) \quad (3.17)$$

$$= \frac{1}{d+1} \left(\text{tr}(A^\dagger) \text{tr}(B) + \text{tr}(A^\dagger B) \right). \quad (3.18)$$

In the last step, we used that $\text{tr}(R_{(12)}A^\dagger \otimes B) = \text{tr}(A^\dagger B)$ which can be easily verified using the graphical notation (cf. Exercise Sheet 2). Next, we note that

$$\frac{1}{d+1} \left(\text{tr}(A^\dagger) \text{tr}(B) + \text{tr}(A^\dagger B) \right) = \frac{1}{d+1} \left((A|\mathbb{1})(\mathbb{1}|B) + (A|\text{id}|B) \right) \quad (3.19)$$

$$= \frac{1}{d+1} \left(d(A|\mathcal{D}|B) + (A|\text{id}|B) \right), \quad (3.20)$$

where $\mathcal{D} = \frac{1}{d}|\mathbb{1})(\mathbb{1}|$ is the completely depolarizing channel, acting as $\mathcal{D}(X) = \text{tr}(X)\mathbb{1}/d$. Hence, we have shown that

$$S = \frac{1}{d+1} (d\mathcal{D} + \text{id}) \equiv \mathcal{D}_{\frac{d}{d+1}}, \quad (3.21)$$

is a convex combination of the completely depolarizing channel and the identity, and thus a depolarizing channel of strength $d/(d+1)$. Moreover, as both channels are trace-preserving (TP), so is S . It is then straightforward to invert S :

$$Y = S(X) = \frac{\text{tr}(X)\mathbb{1} + X}{d+1} \Rightarrow X = (d+1)Y - \text{tr}(Y)\mathbb{1} \Rightarrow S^{-1} = (d+1)\text{id} - d\mathcal{D}, \quad (3.22)$$

where we used that S is TP and thus $\text{tr} X = \text{tr} S(X) = \text{tr} Y$.

Using the explicit form of S^{-1} , we can write the shadow estimator \hat{o} as follows

$$\hat{o}_{U,x} = (O|S^{-1}|E_{U,x}) = (d+1)(O|E_{U,x}) - \text{tr}(O) = (d+1) \text{tr}(OU^\dagger|x\rangle\langle x|U) - \text{tr}(O). \quad (3.23)$$

Thus, the evaluation of $\hat{o}_{U,x}$ requires us to classically compute the expectation value of O in the rotated basis states $U^\dagger|x\rangle$. We will comment on the efficiency of this computation later.

Number of measurements. A central question remains: *How efficient is shadow estimation in terms of the required number of measurements (sample complexity)?* This requires us to bound the convergence of the mean estimator associated to \hat{o} for which we will use Chebyshev's inequality.

Lemma 3.1 (Chebyshev's inequality). *Let X be a random variable and $\varepsilon > 0$. Then, we have*

$$\Pr[|X - \mathbb{E}[X]| > \varepsilon] \leq \frac{\text{Var}[X]}{\varepsilon^2}. \quad (3.24)$$

In particular, if X_1, \dots, X_N are independent and identically distributed (iid) random variables with mean μ and variance σ^2 , and $\bar{X} := \frac{1}{N} \sum_{i=1}^N X_i$, then

$$\Pr[|\bar{X} - \mu| > \varepsilon] \leq \frac{\sigma^2}{N\varepsilon^2}. \quad (3.25)$$

Chebyshev's inequality implies that we need $N \geq \varepsilon^{-2} \delta^{-1} \text{Var}[\hat{o}]$ many samples to get an ε -approximate estimate of $\text{tr}(O\rho)$ with probability at least $1 - \delta$. Hence, we need the variance of \hat{o} to be sufficiently small to get a decent bound on the sample complexity.

Before we compute $\text{Var}[\hat{o}]$, we note that

$$\text{Var}[\hat{o}] = \mathbb{E}[(\hat{o} - \mathbb{E}[\hat{o}])^2]. \quad (3.26)$$

but $\hat{o} - \mathbb{E}[\hat{o}]$ only depends on the traceless part of O , i.e. $O_0 = O - \text{tr}(O)\mathbb{1}/d$, since

$$\hat{o}_{U,x} - \mathbb{E}[\hat{o}] = (d+1) \text{tr}(OU^\dagger|x\rangle\langle x|U) - \text{tr}(O) - \text{tr}(O\rho) \quad (3.27)$$

$$= (d+1) \text{tr}(O_0U^\dagger|x\rangle\langle x|U) - \text{tr}(O_0\rho) + \text{tr}(O) \left(\frac{d+1}{d} - 1 - \frac{1}{d} \right) \quad (3.28)$$

$$= (d+1) \text{tr}(O_0U^\dagger|x\rangle\langle x|U) - \text{tr}(O_0\rho). \quad (3.29)$$

Thus, we can use the traceless observable O_0 instead of O in the future computations. Next, we use $\text{Var}[\hat{o}] = \mathbb{E}[\hat{o}^2] - \mathbb{E}[\hat{o}]^2$ and focus on the computation of $\mathbb{E}[\hat{o}^2]$.

Using the explicit form (3.23) of \hat{o} , we then compute

$$\mathbb{E}[\hat{o}_0^2] = \sum_{U \in \mathcal{G}} \sum_{x=1}^d (O_0 | \tilde{E}_{U,x})^2 p(x|U, \rho) p(U) \quad (3.30)$$

$$= \frac{1}{|\mathcal{G}|} \sum_{U \in \mathcal{G}} \sum_{x=1}^d (O_0 | E_{U,x})^2 (E_{U,x} | \rho) \quad (3.31)$$

$$= \frac{(d+1)^2}{|\mathcal{G}|} \sum_{U \in \mathcal{G}} \sum_{x=1}^d (O_0^{\otimes 2} \otimes \rho | E_{U,x}^{\otimes 3}) \quad (3.32)$$

$$= \frac{(d+1)^2}{|\mathcal{G}|} \sum_{U \in \mathcal{G}} \sum_{x=1}^d \text{tr} \left(O_0^{\otimes 2} \otimes \rho U^{\otimes 3} | x \rangle \langle x |^{\otimes 3} U^{\otimes 3, \dagger} \right). \quad (3.33)$$

To compute the latter twirl, we now assume that \mathcal{G} is a *unitary 3-design* such that we can apply Weingarten calculus. We then find, analogous to the computation of the frame operator (cf. Lem. 2.3):

$$\mathbb{E}[\hat{o}_0^2] = (d+1)^2 \sum_{x=1}^d \sum_{\pi, \sigma \in S_3} W_{\pi, \sigma} \text{tr} (O_0^{\otimes 2} \otimes \rho R_{\pi}) \text{tr} (R_{\sigma}^{\dagger} | x \rangle \langle x |^{\otimes 3}) \quad (3.34)$$

$$= d(d+1)^2 \frac{(d-1)!}{(d+2)!} \sum_{\pi \in S_3} \text{tr} (O_0^{\otimes 2} \otimes \rho R_{\pi}) \quad (3.35)$$

$$= \frac{d+1}{d+2} \text{tr} \left(\begin{array}{c} \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{\rho} \text{---} \end{array} + \begin{array}{c} \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{\rho} \text{---} \end{array} + \begin{array}{c} \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{\rho} \text{---} \end{array} + \begin{array}{c} \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{\rho} \text{---} \end{array} + \begin{array}{c} \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{\rho} \text{---} \end{array} + \begin{array}{c} \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{O_0} \text{---} \\ \text{---} \boxed{\rho} \text{---} \end{array} \right) \quad (3.36)$$

$$= \frac{d+1}{d+2} (\text{tr}(O_0^2) + 2 \text{tr}(O_0^2 \rho)) , \quad (3.37)$$

where we used that O_0 is traceless and thus the first, third, and fourth term in Eq. (3.36) vanishes. We thus get the following bound on the variance for a unitary 3-design:

$$\text{Var}[\hat{o}] \leq \mathbb{E}[\hat{o}_0^2] = \frac{d+1}{d+2} (\text{tr}(O_0^2) + 2 \text{tr}(O_0^2 \rho)) \leq \frac{d+1}{d+2} (\|O_0\|_2^2 + 2\|O_0\|_{\infty}^2) \leq 3\|O_0\|_2^2. \quad (3.38)$$

Here, $\|X\|_2^2 = (X|X) = \text{tr}(X^{\dagger}X)$ is the *Hilbert-Schmidt norm* and $\|X\|_{\infty} = \sup_{\psi} \|X|\psi\rangle\|$ is the *spectral* or *operator norm*. In the last step, we used Hölder's inequality, $\text{tr}(O_0^2 \rho) \leq \|O_0\|_{\infty}^2 \text{tr} \rho$, as well as the general inequality $\|X\|_{\infty} \leq \|X\|_2$.

Discussion of shadow estimation with Clifford unitaries. In the previous derivations we assumed that \mathcal{G} forms a unitary 3-design and for concreteness we will take $\mathcal{G} = \text{Cl}_n$ to be the n -qubit Clifford group.¹ Based on the variance (3.38), we can see that shadow estimation with Clifford unitaries is sample-efficient for observables O with bounded Hilbert-Schmidt norm, $\|O_0\|_2 \leq \|O\|_2 \leq \text{const.}$ (in the sense that the number of measurements N does only depend on the precision, not on the number of qubits). The most important example of such observables are *quantum states* for which $\|\sigma\|_2 = \sqrt{\text{tr}(\sigma^2)} \leq 1$. The associated expectation values then have the form $\text{tr}(\rho\sigma) = (\rho|\sigma)$ and can be interpreted as *overlaps between the states*. Moreover, if $\sigma = |\psi\rangle\langle\psi|$, then $\text{tr}(\rho\sigma) = \langle\psi|\rho|\psi\rangle \equiv F(\rho, \sigma)$ coincides with the usual definition of fidelity between two quantum states (for mixed states, the formula is a bit more complicated). Hence, we can use shadow estimation with Clifford unitaries to efficiently perform *fidelity estimation* with pure target states.

While shadow estimation is sample-efficient, the *classical post-processing* of the measurement data requires the evaluation of $\hat{o}_{U,x}$ using Eq. (3.23). This typically involves the classical simulation of the

¹See Sec. 3.3 for the qudit case.

unitary evolution of $|x\rangle$ under U^\dagger . For the considered Clifford unitaries, it turns out that $U^\dagger|x\rangle$ can be efficiently computed on a classical computer (in time $O(n^2)$) and has an efficient classical description ($O(n^2)$ bits) that can be stored (*Gottesman-Knill theorem*).

However, the evaluation of $\hat{o}_{U,x}$ also involves taking the inner product with O which is, in general, inefficient, i.e. we have to expect that this scales exponentially with the number of qubits n if O does not possess a special structure that we can exploit.

But how does this compare to other methods of fidelity estimation? The *direct fidelity estimation* protocol achieves the same task, but requires a number of measurements that depend on the target state $|\psi\rangle$ and typically scale exponentially in the number of qubits n , while the demands on the classical computer are negligible. In this sense, shadows move the complexity of fidelity estimation from measurements to the classical post-processing, which may be advantageous in near-term devices where taking measurements is more costly than running computations on a powerful computer. For a detailed comparison of fidelity estimation based on classical shadows and direct fidelity estimation, see also Leone, Oliviero, and Hamma [4].

Classical post-processing with Clifford unitaries for stabilizer state fidelity estimation Fidelity estimation is sample efficient, but the classical post-processing requires the evaluation of

$$\hat{o}_{U,x} = (d+1)|\langle\phi|U^\dagger|x\rangle|^2 - 1. \quad (3.39)$$

This is efficiently implementable with Clifford gates.

First, we note that all the states $|x\rangle$ in the computational basis are $+1$ eigenstates of N independent and commuting observables $(-1)^{x_i}Z_i$ where $Z_i = I^{\otimes i-1} \otimes Z \otimes I^{\otimes n-i}$, namely $(-1)^{x_i}Z_i|x\rangle = |x\rangle$. Recall by definition the Clifford group maps individual Pauli strings into individual Pauli strings. This means that $|\psi\rangle = U|x\rangle$ for some $|x\rangle$ in the computational basis and $U \in \text{Cl}_N$ Clifford unitary fulfills

$$|\psi\rangle = U|x\rangle = U(-1)^{x_i}Z_i|x\rangle = U(-1)^{x_i}Z_iU^\dagger U|x\rangle = \tilde{Z}_i|\psi\rangle. \quad (3.40)$$

Thus, the state $|\psi\rangle$ also is the $+1$ eigenstate of new N independent and commuting Pauli strings $\tilde{Z}_i \equiv U(-1)^{x_i}Z_iU^\dagger$. State fulfilling this property are denoted stabilizer states and their associated \tilde{Z}_i the stabilizers. An equivalent characterization is that stabilizer states are given by the action of a random Clifford unitary on a computational basis state. The knowledge of these Pauli strings is sufficient to fix the knowledge of the state. Indeed, as a density matrix

$$|\psi\rangle\langle\psi| = \prod_{i=1}^N \left(\frac{I + \tilde{Z}_i}{2} \right) = \frac{1}{2^N} \sum_{P \in \mathcal{S}(\psi)} P, \quad (3.41)$$

where $\mathcal{S}(\psi) = \text{span}(\tilde{Z}_1, \dots, \tilde{Z}_N)$ is the so-called stabilizer group of $|\psi\rangle$. (It is easy to see that $\tilde{Z}_i\tilde{Z}_j|\psi\rangle = |\psi\rangle$ and that $[\tilde{Z}_i, \tilde{Z}_j] = 0$).

Every Pauli string can be univoquely determined by

$$P = (i)^\phi (X^{a_1}Z^{b_1}) \otimes (X^{a_2}Z^{b_2}) \otimes \dots \otimes (X^{a_N}Z^{b_N}), \quad (3.42)$$

where $a_j, b_j \in \mathbb{Z}_2$ and $\phi \in \{0, 1, 2, 3\}$.

Predicting many expectation values at once. We note that the above estimation strategy can also be extended to estimate several expectation values $(O_1|\rho), \dots, (O_M|\rho)$ at once. To make this precise, note that for the simultaneous estimation of all expectation values we have to choose the failure probability for the i -th mean estimator to be δ/M such that the joint failure probability is uniformly bounded by δ (union bound). Hence, we need the following number of measurements in total:

$$N \geq \frac{M}{\varepsilon^2 \delta} \max_{i=1, \dots, M} \text{Var}[\hat{o}_i], \quad (\text{mean estimators}), \quad (3.43)$$

Depending on the size of M , it may be beneficial to replace the mean estimator for each observable by a so-called *median-of-means estimator*. The idea behind the median-of-means estimator is that the data set is decomposed into K equally sized batches of size N' for any of which an independent mean estimator $\hat{\delta}_i^{(j)}(N')$ is computed. Finally, we take the median over all these mean estimators, in formula

$$\hat{\delta}_i(N', K) = \text{median} \left\{ \hat{\delta}_i^{(1)}(N'), \dots, \hat{\delta}_i^{(K)}(N') \right\}. \quad (3.44)$$

The advantage of this approach is that it is more robust against deviations from the mean. Indeed, a statistical analysis of the median-of-means estimator shows that the dependency on the failure probability of estimating all M expectation values is improved from M/δ to $\log(M/\delta)$ at the cost of larger constants. More precisely, we obtain the following bound on the number of samples

$$N \geq \frac{68}{\varepsilon^2} \log \left(\frac{2M}{\delta} \right) \max_{i=1, \dots, M} \text{Var}[\hat{\delta}_i], \quad (\text{median-of-means estimators}), \quad (3.45)$$

with batch size $K = 2 \log(2M/\delta)$, to guarantee an ε -approximate estimation of all M expectation values with probability at least $1 - \delta$.

Finally, we note that for the median-of-means estimator to be beneficial, we need $68 \log(2M/\delta) \leq M/\delta$ which is true for $M/\delta \geq 464.76$. Assuming $\delta = 0.01$ (1% failure probability), we thus obtain $M \geq 5$.

3.2.2 Pauli measurements

In the following, we consider the ensemble given by *local* Clifford unitaries, these are unitaries of the form $U = U_1 \otimes \dots \otimes U_n$ where $U_i \in \text{Cl}_1$ are single-qubit Clifford unitaries. We denote this set by $\text{LCl}_n := \text{Cl}_1^{\otimes n}$. The effect of local Clifford unitaries is to locally change the basis: Instead of measuring each qubit in the Z basis, each qubit is measured in either the X , Y , or Z basis, depending on the Clifford unitary U_i . This is why these are also referred to as *Pauli measurements*. The hope in this choice of ensemble is that it may be better adjusted to measuring *local observables* of the form $O = O_1 \otimes \dots \otimes O_r \otimes \mathbb{1}^{\otimes(n-r)}$ which only act non-trivially on r qubits.

Computing the shadow estimator. Because all of the expressions factorize, we can re-use our computations from Sec. 3.2.1. For instance,

$$E_{U,x} = (U_1^\dagger \otimes \dots \otimes U_n^\dagger) |x_1, \dots, x_n\rangle \langle x_1, \dots, x_n| (U_1 \otimes \dots \otimes U_n) \bigotimes_{i=1}^n U_i^\dagger |x_i\rangle \langle x_i| U_i = \bigotimes_{i=1}^n E_{U_i, x_i}. \quad (3.46)$$

Hence, the frame operator becomes

$$S = \frac{1}{|\text{Cl}_1|^n} \sum_{U_1, \dots, U_n \in \text{Cl}_1} \sum_{x \in \{0,1\}^n} \bigotimes_{i=1}^n |E_{U_i, x_i}\rangle \langle E_{U_i, x_i}| = S_1^{\otimes n}, \quad (3.47)$$

where S_1 is the local frame operator for which we can use Eq. (3.21) with $d = 2$:

$$S_1 = \frac{1}{|\text{Cl}_1|} \sum_{U \in \text{Cl}_1} \sum_{x \in \{0,1\}} |E_{U,x}\rangle \langle E_{U,x}| = \frac{1}{3}(2\mathcal{D} + \text{id}). \quad (3.48)$$

In particular, $S^{-1} = (3\text{id} - 2\mathcal{D})^{\otimes n}$ by Eq. (3.22). For our r -local observable $O = O_1 \otimes \cdots \otimes O_r \otimes \mathbb{1}^{\otimes (n-r)}$ we can then use that each factor $(3\text{id} - 2\mathcal{D})$ preserves trace such that

$$\hat{\delta}_{U,x} = (O | S^{-1} | E_{U,x}) \quad (3.49)$$

$$= \prod_{i=1}^r (O_i | 3\text{id} - 2\mathcal{D} | E_{U_i, x_i}) \prod_{i=r+1}^n (\mathbb{1} | E_{U_i, x_i}) \quad (3.50)$$

$$= \prod_{i=1}^r \hat{\delta}_{i, U_i, x_i} \quad (3.51)$$

$$= \prod_{i=1}^r \left(3 \text{tr}(O_i U_i^\dagger | x_i \rangle \langle x_i | U_i) - \text{tr}(O_i) \right). \quad (3.52)$$

Number of measurements. Next, we compute the variance of $\hat{\delta}$. Because the estimator only depends on the first r qubits we can resum the remaining local terms inside the Born probability using that $\frac{1}{|\text{Cl}_1|} \sum_{V \in \text{Cl}_1} \sum_{y \in \{0,1\}} V^\dagger |y\rangle \langle y| V = \mathbb{1}$

$$\mathbb{E}[\hat{\delta}^2] = \frac{1}{|\text{Cl}_1|^n} \sum_{U \in \text{Cl}_1^n} \sum_{x \in \{0,1\}^n} \prod_{i=1}^r \hat{\delta}_{i, U_i, x_i}^2 \text{tr} \left(\bigotimes_{i=1}^n U_i^\dagger |x_i\rangle \langle x_i| U_i \rho \right) \quad (3.53)$$

$$= \frac{1}{|\text{Cl}_1|^r} \sum_{U \in \text{Cl}_1^r} \sum_{x \in \{0,1\}^r} \prod_{i=1}^r \hat{\delta}_{i, U_i, x_i}^2 \text{tr} \left(\bigotimes_{i=1}^r U_i^\dagger |x_i\rangle \langle x_i| U_i \otimes \mathbb{1}^{\otimes (n-r)} \rho \right) \quad (3.54)$$

$$= \frac{1}{|\text{Cl}_1|^r} \sum_{U \in \text{Cl}_1^r} \sum_{x \in \{0,1\}^r} \prod_{i=1}^r \hat{\delta}_{i, U_i, x_i}^2 \text{tr} \left(\bigotimes_{i=1}^r U_i^\dagger |x_i\rangle \langle x_i| U_i \rho' \right), \quad (3.55)$$

where $\rho' = \text{tr}_{r+1, \dots, n} \rho$ is the reduced state on the first r qubits.

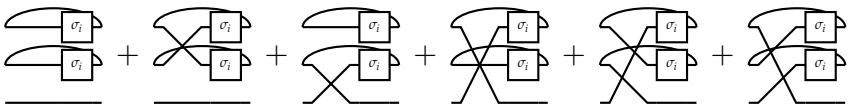
In the following, we concentrate on the case that $O = \sigma_1 \otimes \cdots \otimes \sigma_r \otimes \mathbb{1}^{\otimes (n-r)}$ is a Pauli operator supported on r qubits (here $\sigma_i \in \{X, Y, Z\}$) and refer to Ref. [5] for general r -local observables. Using the explicit form of the shadow estimator (3.52), we can then compute in analogy to Sec. 3.2.1:

$$\mathbb{E}[\hat{\delta}^2] = \frac{1}{|\text{Cl}_1|^r} \sum_{U \in \text{Cl}_1^r} \sum_{x \in \{0,1\}^r} \prod_{i=1}^r 9 \text{tr}(\sigma_i U_i^\dagger |x_i\rangle \langle x_i| U_i)^2 \text{tr} \left(\bigotimes_{i=1}^r U_i^\dagger |x_i\rangle \langle x_i| U_i \rho' \right) \quad (3.56)$$

$$= \text{tr} \left[\mathbb{1}_r^{\otimes 2} \otimes \rho' \bigotimes_{i=1}^r \left(\frac{9}{|\text{Cl}_1|} \sum_{U \in \text{Cl}_1} \sum_{x \in \{0,1\}} (U_i^\dagger |x_i\rangle \langle x_i| U_i)^{\otimes 3} \sigma_i^{\otimes 2} \otimes \mathbb{1}_1 \right) \right] \quad (3.57)$$

$$= \text{tr} \left[\mathbb{1}_r^{\otimes 2} \otimes \rho' \bigotimes_{i=1}^r \left(\frac{3}{4} \sum_{\pi \in \mathbb{S}_3} R_\pi \sigma_i^{\otimes 2} \otimes \mathbb{1}_1 \right) \right]. \quad (3.58)$$

In the last step, we used the previous computation (3.35) (for $d = 2$). We can now perform the partial trace over the first two systems based on the previous graphical calculus (3.36):



$$= \text{tr}(\sigma_i^2) \mathbb{1} + 2\sigma_i^2 = 4\mathbb{1}. \quad (3.59)$$

Here, we used that Pauli operators square to the identity: $X^2 = Y^2 = Z^2 = \mathbb{1}$. Hence, we obtain the following variance from Eq. (3.58):

$$\text{Var}[\hat{\delta}] \leq \mathbb{E}[\hat{\delta}^2] = \text{tr} \left[\rho' \bigotimes_{i=1}^r (3\mathbb{1}) \right] = 3^r. \quad (3.60)$$

Discussion of shadow estimation with local Clifford unitaries. It is worth highlighting that, when using local Clifford unitaries for shadow estimation, the postprocessing become particularly efficient. At the same time, the sampling complexity for $r = O(1)$ is constant, meaning that the whole shadow tomography performance is efficient. In this setting, each measurement outcome corresponds to projecting onto a local Pauli basis element, and the corresponding classical shadow can be reconstructed with computational cost scaling only linearly with the system size. Moreover, for observables that are themselves local (i.e., supported on a constant number of qubits), the postprocessing requires evaluating simple expectation values over these classical shadows. This makes the overall procedure highly scalable in practice, since both the sample complexity and the computational overhead remain independent of the global system size, provided the observables are local.

3.2.3 Beyond linear observables: the purity

Consider the problem of estimating the purity $\mathcal{P}(\rho) \equiv \text{tr}(\rho^2)$ of a state ρ , for instance the reduced density matrix of a global pure state $|\Psi\rangle$. How do we estimate this quantity? Consider we have M snapshot $\tilde{\rho}^r \equiv (U^{(r)})^\dagger |x^{(r)}\rangle \langle x^{(r)}| U^{(r)}$. Since U are identically, independent, randomly distributed from an ensemble \mathcal{G} , and $\{x^{(r)} | r = 1, \dots, M\}$ are i.i.d. from $U^{(r)} \rho (U^{(r)})^\dagger$, the stochastic objects $\tilde{\rho}^{(r)}$ and $\tilde{\rho}^{(j)}$ are uncorrelated for $r \neq j$. Thus, we can construct an estimator of the system purity by computing the monte-carlo sampling

$$\mathcal{P}^{(e)} = \frac{1}{M(M-1)} \sum_{r \neq j, j, r=1}^M \text{tr}(\tilde{\rho}^{(j)} \tilde{\rho}^{(r)}) = \frac{2}{M(M-1)} \sum_{1 \leq j < r \leq M} \text{tr}(\tilde{\rho}^{(j)} \tilde{\rho}^{(r)}) \quad (3.61)$$

We want to compute the variance of the purity estimator $\mathcal{P}^{(e)}$ for global Clifford estimator. We have

$$\text{Var} [\mathcal{P}^{(e)}] = \binom{M}{2}^{-2} \sum_{r < r'} \sum_{s < s'} \left(\mathbb{E} [\text{tr}(\tilde{\rho}^{(r)} \tilde{\rho}^{(r')}) \text{tr}(\tilde{\rho}^{(s)} \tilde{\rho}^{(s')})] - \mathcal{P}(\rho)^2 \right)$$

Let us expand these terms. When $s = r$ and $r' = s'$, we have

$$\binom{M}{2}^{-2} \sum_{r < r'} \mathbb{E} [\text{tr}(\tilde{\rho}^{(r)} \tilde{\rho}^{(r')})^2] = \binom{M}{2}^{-1} \mathbb{E} [\text{tr}(\tilde{\rho}^{(1)} \tilde{\rho}^{(2)})^2], \quad (3.62)$$

where we identified two reference copies via Haar invariance. If all the terms are different, one gets via statistical independence of the copies and the replica invariance

$$\mathbb{E} [\text{tr}(\tilde{\rho}^{(1)} \tilde{\rho}^{(2)}) \text{tr}(\tilde{\rho}^{(3)} \tilde{\rho}^{(4)})] = \mathcal{P}(\rho)^2. \quad (3.63)$$

Finally, if at least one term $l = i$ or $l = j$, then we have using Haar invariance and identically independent distributed

$$2 \binom{M}{2}^{-2} \sum_{i < j} \sum_{k \neq i, j} \mathbb{E} [\text{tr}(\tilde{\rho}^{(1)} \tilde{\rho}^{(2)}) \text{tr}(\tilde{\rho}^{(1)} \tilde{\rho}^{(3)})] = 2 \binom{N}{2}^{-1} (N-2) \mathbb{E} [\text{tr}(\tilde{\rho}^{(1)} \rho)^2]. \quad (3.64)$$

Putting these terms together, we get the final result

$$\text{Var} [\mathcal{P}^{(e)}] = \binom{M}{2}^{-1} \text{Var} [\text{tr}(\tilde{\rho}^{(1)} \tilde{\rho}^{(2)})] + \binom{M}{2}^{-1} 2(M-2) \text{Var} [\text{tr}(\tilde{\rho}^{(1)} \rho)].$$

The first term can be written explicitly as

$$\begin{aligned}
\text{Var} \left[\text{tr}(\tilde{\rho}^{(1)} \tilde{\rho}^{(2)}) \right] &= \mathbb{E}_{U, V \in \text{Cl}(2^N)} \mathbb{E}_{b, d} \text{tr} \left[\left((2^N + 1) U^\dagger |b\rangle \langle b| U - \mathbb{I} \right) \left((2^N + 1) V^\dagger |d\rangle \langle d| V - \mathbb{I} \right) \right]^2 - \mathcal{P}(\rho)^2 \\
&= \int d\mu_H(U) d\mu_H(V) \left\{ 2^{2N} \left[(2^N + 1)^2 |\langle 0 | UV^\dagger | 0 \rangle|^2 - (2^N + 2) \right]^2 \right. \\
&\quad \times \langle 0 | U \rho U^\dagger | 0 \rangle \langle 0 | V \rho V^\dagger | 0 \rangle \left. \right\} - \mathcal{P}(\rho)^2 \\
&= 2^{2N} (2^N + 1)^4 \int d\mu_H(U) d\mu_H(V) |\langle 0 | UV^\dagger | 0 \rangle|^4 \langle 0 | U \rho U^\dagger | 0 \rangle \langle 0 | V \rho V^\dagger | 0 \rangle \\
&\quad - (2^N + 2 + \mathcal{P}(\rho))^2, \tag{3.65}
\end{aligned}$$

where in the second line we invoked the 3-design property of Clifford unitaries to replace the sum with an integral over Haar measure $d\mu_H(U)$. The first term in the last line can then be rewritten as

$$2^{2N} (2^N + 1)^4 \int d\mu_H(U) \langle 0 | U \rho U^\dagger | 0 \rangle \text{tr} \left\{ \left((U^\dagger)^{\otimes 2} |0\rangle \langle 0|^{\otimes 2} U^{\otimes 2} \otimes \rho \right) \int d\mu_H(V) (V^\dagger)^{\otimes 3} |0\rangle \langle 0|^{\otimes 3} V^{\otimes 3} \right\}, \tag{3.66}$$

and now the last integral can be computed using the Weingarten calculus, yielding

$$\begin{aligned}
&= \frac{2^N (2^N + 1)^3}{2^N + 2} \int d\mu_H(U) \langle 0 | U \rho U^\dagger | 0 \rangle (2 + 4 \langle 0 | U \rho U^\dagger | 0 \rangle) \\
&= \frac{2^N (2^N + 1)^3}{2^N + 2} \left(\frac{2}{2^N} + 4 \frac{1 + \mathcal{P}(\rho)}{2^N (2^N + 1)} \right) = \frac{2(2^N + 1)^2}{2^N + 2} (2^N + 3 + 2\mathcal{P}(\rho)), \tag{3.67}
\end{aligned}$$

thus arriving at

$$\text{Var} \left[\text{tr}(\tilde{\rho}^{(1)} \tilde{\rho}^{(2)}) \right] = \frac{2(2^N + 1)^2}{2^N + 2} (2^N + 3 + 2\mathcal{P}(\rho)) - (2^N + 2 + \mathcal{P}(\rho))^2. \tag{3.68}$$

Following similar steps, the second term in Eq. (3.65) can be written as

$$\begin{aligned}
\text{Var} \left[\text{tr}(\tilde{\rho}^{(1)} \rho) \right] &= \mathbb{E}_{U \in \text{Cl}(2^N)} \mathbb{E}_b \text{tr} \left[\left((2^N + 1) U^\dagger |b\rangle \langle b| U - \mathbb{I} \right) \rho \right]^2 - \mathcal{P}(\rho)^2 \\
&= 2^N \int d\mu_H(U) \left((2^N + 1)^2 \langle 0 | U \rho U^\dagger | 0 \rangle^2 - 2(2^N + 1) \langle 0 | U \rho U^\dagger | 0 \rangle + 1 \right) \langle 0 | U \rho U^\dagger | 0 \rangle - \mathcal{P}(\rho)^2 \\
&= \frac{2^N + 1}{2^N + 2} (1 + 3\mathcal{P}(\rho) + 2 \text{tr}(\rho^3)) - 2(1 + \mathcal{P}(\rho)) + 1 - \mathcal{P}(\rho)^2 \\
&= \frac{2^N + 1}{2^N + 2} (1 + 3\mathcal{P}(\rho) + 2 \text{tr}(\rho^3)) - (1 + \mathcal{P}(\rho))^2. \tag{3.69}
\end{aligned}$$

Putting everything together, we finally arrive at

$$\begin{aligned}
\text{Var} \left[\mathcal{P}^{(e)} \right] &= \binom{M}{2}^{-1} \left\{ \frac{2(2^N + 1)^2}{2^N + 2} (2^N + 3 + 2\mathcal{P}(\rho)) - (2^N + 2 + \mathcal{P}(\rho))^2 \right. \\
&\quad \left. + 2(M - 2) \left[\frac{2^N + 1}{2^N + 2} (1 + 3\mathcal{P}(\rho) + 2 \text{tr}(\rho^3)) - (1 + \mathcal{P}(\rho))^2 \right] \right\}. \tag{3.70}
\end{aligned}$$

Thus the leading term is $\text{Var}[\mathcal{P}_\infty^{(e)}] \sim \frac{2(2^{2N})}{M(M-1)}$ for large system size N .

3.3 Further reading

To be done.

As argued previously in Sec. 2.2, Haar-random unitaries are notoriously expensive to implement in that they need exponentially deep quantum circuits. For this reason, unitary designs have been introduced as a way to mimic the Haar measure up to a certain moment. *The* prototypical example of a unitary design is the Clifford group which is thus extensively used throughout the literature, and also in our treatment of classical shadows in Ch. 3.

Nevertheless, there are two major problems with unitary designs: First, there are no practical examples of unitary designs beyond the Clifford group and we are thus limited to third moments. However, there are interesting properties in many-body physics that are described by much higher moments. Second, Clifford unitaries need linear-depth circuits to be implemented and this is likely the case for any unitary design. However, near-term devices are noisy and thus effectively limited in the depth of the circuits they can execute. We would thus prefer ensembles with a more fine-grained control on the circuit depth and, ideally, sub-linearly sized circuits.

From a practical point of view, it thus seems to be a good idea to take a *constructive* approach to random unitaries by considering families of quantum circuits with variable depth, where the (local) gates are drawn at random. We call the so-constructed ensemble a *random quantum circuit* (RQC).

4.1 Random quantum circuits

Intuitively, RQCs should become ‘more and more random’ with increasing circuit depth. However, RQCs will typically never be a unitary design in the sense of the defining equation (2.32), this is

$$M_k(A) \neq M_k^{\text{RQC}}(A) := \mathbb{E}_{U \sim \text{RQC}} \left[U^{\otimes k} A U^{\otimes k, \dagger} \right], \quad (4.1)$$

where the average on the right hand side is over all unitaries in the RQC ensemble. Instead, (4.1) will be fulfilled up to an error ε that becomes smaller with circuit depth. We then say that the RQC forms an ε -approximate unitary design.




Figure 4.1: Random quantum circuits are circuits in which the gates are drawn randomly. Typically, these gates are local and arranged in a specific layout, for instance in the popular *brickwork layout* shown here. These circuits are typically composed of *layers* (shown by the dashed box) and gates in each layer are drawn independently.

In this chapter, we will consider a specific construction of random quantum circuits which we call *layered* (we will encounter other examples later). For this, we assume that the circuit is composed of repeated and independent *layers*. These layers typically have some specific structure or layout, for instance, gates could be arranged in a *brickwork* pattern as in Fig. 4.1. Importantly, the gates within every layer are drawn independently at random such that there are no correlations between the gates or the layers.

Let us consider a single random layer distributed according to a probability measure ν on $U(d)$ and the associated k -fold twirl, typically called the $(k$ -th) *moment operator* in the RQC literature:

$$M_k^\nu := \int U^{\otimes k}(\cdot) U^{\otimes k,\dagger} d\nu(U). \quad (4.2)$$

Then, a random circuit with L layers has the form $U = U_L \cdots U_1$ where $U_i \sim \nu$.¹ The moment operator of the L -layer RQC is then given by

$$M_k^{\nu_L} = \int U_L^{\otimes k} \cdots U_1^{\otimes k}(\cdot) U_1^{\otimes k,\dagger} \cdots U_L^{\otimes k,\dagger} d\nu(U_1) \cdots d\nu(U_L) \quad (4.3)$$

$$= \int U_L^{\otimes k} \left(\cdots \left(\int U_2^{\otimes k} \left(\int U_1^{\otimes k}(\cdot) U_1^{\otimes k,\dagger} d\nu(U_1) \right) d\nu(U_2) \right) \cdots \right) U_L^{\otimes k,\dagger} d\nu(U_L) \quad (4.4)$$

$$= (M_k^\nu)^L. \quad (4.5)$$

We also emphasize that the (local) structure of each layer is also reflected in the moment operator M_k^ν . For instance, if we have a brickwork layout as in Fig. 4.1, then we can decompose every layer unitary as $U = U_{\text{even}} U_{\text{odd}}$ where $U_{\text{odd}} = U_{12} \otimes U_{34} \otimes \cdots \otimes U_{n-1,n}$ and $U_{\text{even}} = \mathbb{1} \otimes U_{23} \otimes U_{45} \otimes \cdots \otimes U_{n-2,n-1} \otimes \mathbb{1}$. Let us for simplicity assume that all local gates U_{ij} are distributed identically, for instance according to the 2-qudit Haar measure. Then, $M_k^\nu = M_k^{\text{even}} M_k^{\text{odd}}$ by the same argument as above, and

$$M_k^{\text{even}} = M_k^{12} \otimes M_k^{34} \otimes \cdots \otimes M_k^{n-1,n} = \left(M_k^{\text{loc}} \right)^{\otimes n/2}. \quad (4.6)$$

A similar argument holds for the odd layers. We can visualize this by taking the brickwork circuit 4.1 and stack the k copies of each local unitary (channel) on top of each other, cf. Fig. 4.2. The averages are now taken over the all bricks individually.

Following the intuition presented at the beginning of this section, the moment operator of a RQC should converge to the Haar-random one with increasing circuit depth, this is

$$M_k^{\nu_L} = (M_k^\nu)^L \rightarrow M_k \quad \text{for } L \rightarrow \infty. \quad (4.7)$$

One way of making this precise is by considering the eigenvalues of M_k^ν . Since the eigenvalues of $M_k^{\nu_L}$ are simply powers of those, it is sufficient to study a single layer. But how different can the single-layer moment operator M_k^ν be from the Haar-random moment operator M_k which we have studied so far?

Before we delve into this question, let us briefly recall that the adjoint ϕ^\dagger of a linear map ϕ is defined by the equation

$$(A|\phi(B)) = (\phi^\dagger(A)|B). \quad (4.8)$$

Applied to the map $\phi = V(\cdot)V^\dagger$, we thus find that $\phi^\dagger = V^\dagger(\cdot)V = \phi^{-1}$. Hence,

$$(M_k^\nu)^\dagger = \int U^{\otimes k,\dagger}(\cdot) U^{\otimes k} d\nu(U) = \int U^{\otimes k}(\cdot) U^{\otimes k,\dagger} d\nu(U^\dagger) = M_k^{\tilde{\nu}}, \quad (4.9)$$

where $\tilde{\nu}(U) = \nu(U^\dagger)$ is the measure under inversion. For M_k^ν to be self-adjoint, it is thus sufficient that ν is invariant under inversion which is however not always the case. For instance, the brickwork example in Fig. 4.1 is not: Inverting a layer means exchanging the even and odd sublayers which

¹Note that the product is distributed according to the convolution measure ν^{*L} .




Figure 4.2: Visualization of the k -fold action on *replicas* by stacking the unitaries on top of each other. The local structure of the underlying circuit is preserved. Any k -th moment of the RQC can be written in terms of ‘boundary conditions’ A and B to the circuit, and subsequent ensemble averages.

clearly leads to a different measure. Thus, we can generally not guarantee that the eigenvalues of M_k^ν are real or that M_k^ν is even diagonalizable.

However, there is a simple observation about M_k^ν : Since all unitaries of the form $U^{\otimes k}$ commute with permutations R_π , we still have

$$M_k^\nu(R_\pi) = U^{\otimes k, \dagger} R_\pi U^{\otimes k} d\nu(U) = R_\pi, \quad (4.10)$$

no matter what distribution ν we choose. As M_k^ν might not be a self-adjoint superoperator, the right eigenoperators R_π are not automatically left eigenoperators as well. However, by Eq. (4.9), we find that

$$(M_k^\nu)^\dagger(R_\pi) = \int U^{\otimes k, \dagger} R_\pi U^{\otimes k} d\nu(U) = R_\pi. \quad (4.11)$$

Hence, all permutations R_π are left and right eigenoperators of eigenvalue 1, just as in the Haar-random case. This means that M_k^ν has the following matrix form in a suitable basis:

$$M_k^\nu \simeq \begin{pmatrix} \mathbb{1}_k! & 0 \\ 0 & * \end{pmatrix}, \quad (4.12)$$

where $*$ is defined on the orthocomplement of the unitary commutant Comm_k . Note that the Haar-random moment operator M_k has the same structure, but with $*$ = 0. Now, the L -layer moment operator $(M_k^\nu)^L$ is obtained by taking L -th powers of the blocks. Thus, to obtain the desired convergence (4.7), we need that $*^L \rightarrow 0$. We here analyze the convergence in *spectral norm*. Recall that the spectral norm $\|\phi\|_\infty$ is defined as the largest singular value of ϕ , or equivalently, the largest eigenvalue of the self-adjoint operator $\sqrt{\phi^\dagger \phi}$.

We say that M_k^ν is *gapped* if $\|*\|_\infty \leq 1 - \Delta_k$ with the *spectral gap* $\Delta_k > 0$. Again, this guarantees that $(M_k^\nu)^L \rightarrow M_k$ and, moreover,

$$\|(M_k^\nu)^L - M_k\|_\infty = \|(M_k^\nu - M_k)^L\|_\infty \leq \|M_k^\nu - M_k\|_\infty^L \leq (1 - \Delta_k)^L. \quad (4.13)$$

Here, we used that $M_k^\nu M_k = M_k = M_k M_k^\nu$ by the invariance of the Haar measure and that the spectral norm is submultiplicative. Hence, Eq. (4.13) states that the *rate of convergence* is determined by the spectral gap.

It turns out that RQCs are generally gapped, at least if ν is a universal measure, i.e. its support contains a universal gate set. We will give a proof of this statement in Sec. 4.2, but since it is not of

major importance for the remaining course, we consider this section optional and not treat the proof in the lecture.

What does the spectral gap imply for the convergence of concrete moments to their Haar-random value, i.e. for expressions of the form $\text{tr}[B(M_k^\nu)^L(A)]$? These moments correspond to certain boundary conditions on the k -replica circuit, cf. Fig. 4.2. To bound the approximation error of a RQC, we use that the spectral norm is an operator norm, meaning that

$$\|\phi\|_\infty = \sup_{X \in L(\mathcal{H})} \frac{\|\phi(X)\|_2}{\|X\|_2}. \quad (4.14)$$

Hence, by using Cauchy-Schwarz inequality:

$$\begin{aligned} \text{tr} \left[B((M_k^\nu)^L - M_k)(A) \right] &\leq \|B\|_2 \|((M_k^\nu)^L - M_k)(A)\|_2 \\ &\leq \|A\|_2 \|B\|_2 (\|(M_k^\nu)^L - M_k\|_\infty) \leq (1 - \Delta_k)^L \|A\|_2 \|B\|_2 \end{aligned} \quad (4.15)$$

Hence, the approximation error is lower than ε if

$$L \geq \frac{\log(\varepsilon) - \log(\|A\|_2 \|B\|_2)}{\log(1 - \Delta_k)}. \quad (4.16)$$

Using that $\log(1 + x) \leq x$ for $x \geq -1$ this is implied for

$$L \geq \Delta_k^{-1} (\log(1/\varepsilon) + \log(\|A\|_2 \|B\|_2)). \quad (4.17)$$

Note that $\log(1/\varepsilon) \geq 0$ since $0 < \varepsilon < 1$. Hence, the required number of layers is determined by the size of the gap and the approximation error.

Now one may wonder that if $\|A\|_2 = \|B\|_2 = 1$ (for instance for quantum states) and Δ_k is constant, we would obtain convergence we have a constant number of layers. While it is true that we would then achieve a constant error ε with a constant number of layers, this would not be enough as the moment we are trying to approximate, $\text{tr}[BM_k(A)]$, is typically exponentially small in n and k . Concretely, we typically need an *additive* error of the order of q^{-2nk} to get a good approximation. We can directly aim for a *relative* error ε by modifying our bound as follows

$$L \geq \Delta_k^{-1} (\log(1/\varepsilon) + \log(\|A\|_2 \|B\|_2 / \text{tr}[BM_k(A)])) = \Delta_k^{-1} (\log(1/\varepsilon) + \log(1 / \text{tr}[\hat{B}M_k(\hat{A})])) , \quad (4.18)$$

where $\hat{A} = A/\|A\|_2$ and $\hat{B} = B/\|B\|_2$. Now inserting $\text{tr}[\hat{B}M_k(\hat{A})] \approx q^{-2nk}$ we get

$$L \geq \Delta_k^{-1} (\log(1/\varepsilon) + 2 \log(q)nk) . \quad (4.19)$$

This already suggests that we need a number of layers that is linear in nk to get a good approximation to the Haar-random value of our moment.

Perhaps not surprisingly, the spectral gap depends strongly on the choice of RQC and will generally be a function of the number of qudits n and copies k . For instance, if we apply only a single 2-qudit gate in every layer instead of $O(n)$ as in the brickwork circuit, Fig. 4.1, we should expect a rescaling of $O(1/n)$ of the spectral gap since we need $O(n)$ more layers to achieve the same overall number of gates. It is generally expected that the spectral gap of RQCs like the brickwork circuit is constant in both n and k , at least for $k \leq O(d)$. Only recently, this was proven (up to logarithmic factors): The spectral gap of brickwork RQCs is $\Delta_k \geq c \log(k)^{-7}$ for some constant $c > 0$.

Conclusion. Combining the spectral gap with the error discussion, we can generally say that

$$L \geq c^{-1} \log(k)^7 (\log(1/\varepsilon) + 2 \log(q)nk) , \quad (4.20)$$

many brickwork layers are sufficient to approximate arbitrary Haar moments. This involves circuits that are of linear depth in nk . Hence, *random quantum circuits of linear depth behave as Haar-random unitaries for (almost) all practical purposes.*

4.2 More on spectral gaps*

We have the following structural result on the spectrum of M_k^ν .

Theorem 4.1 (Spectral properties of M_k^ν). *We have the following general properties:*

- (i) All eigenvalues of M_k^ν have absolute values ≤ 1 ,
- (ii) The 1-eigenspace of M_k^ν is at least $k!$ -fold degenerate, containing the unitary commutant Comm_k ,
- (iii) If ν is universal (its support generates a dense subgroup of $U(d)$), then the 1-eigenspace is exactly $k!$ -dimensional. Thus, M_k^ν is gapped.

Proof of Theorem 4.1. (i) Let us assume that A is an eigenoperator with eigenvalue $\lambda \in \mathbb{C}$, i.e. $M_k^\nu(A) = \lambda A$. We denote by $\|A\|_2^2 = (A|A)$ the Hilbert-Schmidt or Frobenius norm (the norm associated to the trace inner product). Note that this norm is invariant under unitaries: $\|VA\|_2^2 = \text{tr}((VA)^\dagger VA) = \text{tr}(A^\dagger A) = \|A\|_2^2$. Then, triangle and Cauchy-Schwarz inequality yield

$$|\lambda| \|A\|_2^2 = |(A|M_k^\nu(A))| \leq \int |(A|U^{\otimes k}AU^{\otimes k,\dagger})| d\nu(U) \leq \int \|A\|_2 \|U^{\otimes k}AU^{\otimes k,\dagger}\|_2 d\nu(U) = \|A\|_2^2. \quad (4.21)$$

Hence, all eigenvalues lie in the unit disk.

(ii) was already proven in Sec. 4.1.

(iii) Let A be an eigenoperator of eigenvalue 1 and consider the map $\mathcal{X}_U = \text{id} - U^{\otimes k}(\cdot)U^{\otimes k,\dagger}$. Then:

$$\int (A|\mathcal{X}_{U^\dagger}\mathcal{X}_U|A) d\nu(U) = 2\|A\|_2^2 - \int (A|U^{\otimes k}AU^{\otimes k,\dagger} + U^{\otimes k,\dagger}AU^{\otimes k}) d\nu(U) \quad (4.22)$$

$$= 2\|A\|_2^2 - \int 2 \text{Re} \left((A|U^{\otimes k}AU^{\otimes k,\dagger}) \right) d\nu(U) \quad (4.23)$$

$$= 2\|A\|_2^2 - 2 \text{Re}(A|M_k^\nu(A)) = 0. \quad (4.24)$$

However, we also have

$$(A|\mathcal{X}_{U^\dagger}\mathcal{X}_U|A) = \text{tr} \left(A^\dagger (\mathcal{X}_U(A) - U^{\otimes k,\dagger}\mathcal{X}_U(A)U^{\otimes k}) \right) \quad (4.25)$$

$$= \text{tr} \left((A - U^{\otimes k}AU^{\otimes k,\dagger})^\dagger \mathcal{X}_U(A) \right) = (\mathcal{X}_U(A)|\mathcal{X}_U(A)) \geq 0. \quad (4.26)$$

Hence, the only way that the integral above can be zero is if

$$0 = (\mathcal{X}_U(A)|\mathcal{X}_U(A)) = \|A - U^{\otimes k}AU^{\otimes k,\dagger}\|_2^2 \Leftrightarrow A = U^{\otimes k}AU^{\otimes k,\dagger}, \quad (4.27)$$

holds ν -almost everywhere (i.e. everywhere except on a subset of unitaries that have ν -measure zero). Hence, if the support of ν contains a set of generators, A has to commute with generators of $U(d)$ and hence lies in the commutant of a dense subgroup of $U(d)$. By continuity, this commutant is the same as the one of $U(d)$ and thus the 1-eigenspace of M_k^ν is exactly Comm_k . To be done: show that M_k^ν is gapped (could still have an additional singular value 1). \square

4.3 Further reading

To be done.

BIBLIOGRAPHY

- [1] Richard Kueng and David Gross. *Qubit stabilizer states are complex projective 3-designs*. 2015. [arXiv: 1510.02767](#).
- [2] Markus Heinrich. *On stabiliser techniques and their application to simulation and certification of quantum devices*. PhD thesis. Universität zu Köln, 2021.
- [3] Srilekha Gandhari, Victor V. Albert, Thomas Gerrits, Jacob M. Taylor, and Michael J. Gullans. *Precision Bounds on Continuous-Variable State Tomography Using Classical Shadows*. In: *PRX Quantum* 5.1 (2024), p. 010346.
- [4] Lorenzo Leone, Salvatore F. E. Oliviero, and Alioscia Hama. *Nonstabilizerness determining the hardness of direct fidelity estimation*. In: *Physical Review A* 107.2 (2023), p. 022429.
- [5] Hsin-Yuan Huang, Richard Kueng, and John Preskill. *Predicting many properties of a quantum system from very few measurements*. In: *Nature Physics* 16.10 (2020), pp. 1050–1057.

SOME LINEAR ALGEBRA

This section gives a basic introduction to the linear algebraic concepts used in this course. Most of this should already be known from linear algebra and quantum mechanics lectures. At this point, the lectures notes are more detailed than the lecture to achieve a certain self-containment of the notes and provide a reference for later stages of the course.

A.1 States, operators, superoperators

State space As usual, quantum mechanics is modeled on a *Hilbert space* \mathcal{H} , which we take, in good quantum info tradition, to be *finite-dimensional* for the remainder of this course. Hence, we can simply think of $\mathcal{H} = \mathbb{C}^d$ with the standard basis $|x\rangle$ labeled by integers $x = 0, 1, \dots, d-1$, and the standard inner product

$$\langle \psi | \varphi \rangle = \sum_{x=0}^{d-1} \bar{\psi}_x \varphi_x, \quad (\text{A.1})$$

where $\psi_x = \langle x | \psi \rangle$ and $\varphi_x = \langle x | \varphi \rangle$ are the coefficients in the standard basis. Typically, we take vectors $\psi \in \mathcal{H}$ to be normalized: $\langle \psi | \psi \rangle = 1$.

The notation of the inner product as a ‘bracket’ motivates the popular *Dirac* or *bra-ket notation* which we adopt here: In this context, vectors $\psi \in \mathcal{H}$ are called *kets* and written as $|\psi\rangle$. The corresponding *bra* is a dual vector $\langle \psi | \in \mathcal{H}^*$ and given by the linear form $\mathcal{H} \ni \varphi \mapsto \langle \psi | \varphi \rangle$.¹ While the pairing between a bra and ket yields the inner product (the ‘bracket’), the pairing between a ket and bra forms a so-called *outer product* $|\psi\rangle\langle \varphi|$ which is a linear operator on \mathcal{H} that acts as $\mathcal{H} \ni \chi \mapsto |\psi\rangle\langle \varphi | \chi \rangle$.

Linear operators The vector space of all linear operators $A : \mathcal{H} \rightarrow \mathcal{H}$ is denoted by $L(\mathcal{H})$. For any $A \in L(\mathcal{H})$, its *adjoint* A^\dagger is the linear operator for which

$$\langle \psi | A \varphi \rangle = \langle A^\dagger \psi | \varphi \rangle, \quad \forall \psi, \varphi \in \mathcal{H}. \quad (\text{A.2})$$

If represented in an orthonormal basis, such as the standard basis, the adjoint operator is the conjugate transpose matrix, $A^\dagger = \bar{A}^\top$.

Definition A.1. In the following, we define some relevant classes of operators:

- **Hermitian (or self-adjoint) operator:** $A \in L(\mathcal{H})$ such that $A^\dagger = A$. Hermitian operators have only real eigenvalues and an orthonormal eigenbasis.
- **Unitary operator:** $U \in L(\mathcal{H})$ such that $U^\dagger U = U U^\dagger = \mathbb{1}$.
- **Positive semi-definite (psd) operator:** Hermitian $A \in L(\mathcal{H})$ with only non-negative eigenvalues. We write $A \geq 0$.
- **Projector:** Hermitian $P \in L(\mathcal{H})$ such that $P^2 = P$.
- **Quantum state:** $\rho \in L(\mathcal{H})$ such that $\rho \geq 0$ and $\text{tr } \rho = 1$. ρ is called *pure* if it is a projector: $\rho^2 = \rho$. Pure states have rank one and are of the form $\rho = |\psi\rangle\langle \psi|$.

Finally, the unitaries on \mathcal{H} form the unitary group $U(\mathcal{H}) = U(d)$.

¹In mathematics, this is called the *Riesz representation theorem*.

The vector space $\text{End}(\mathcal{H})$ of linear operators on \mathcal{H} forms a Hilbert space of dimension $(\dim \mathcal{H}) = d^2$ in its own right with the *Hilbert-Schmidt* or *trace inner product*:

$$(X|Y) := \text{tr}(X^\dagger Y). \quad (\text{A.3})$$

In particular, we can introduce an orthonormal operator basis as a set of operators A_1, \dots, A_{d^2} such that $(A_i|A_j) = \delta_{ij}$. We will now introduce an import example of such a basis, the *Pauli basis*.

Example A.1: Pauli basis

Recall the Pauli operators

$$\sigma_{0,1} \equiv X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_{1,1} \equiv Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_{1,0} \equiv Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \quad (\text{A.4})$$

which we complement with the identity $\sigma_{0,0} = \mathbb{1}$. Then, the multi-qubit Pauli operators on $\mathcal{H} = (\mathbb{C}^2)^{\otimes n}$ are simply given by all possible tensor products of the single-qubit Pauli operators, in formula:

$$\sigma_a := \sigma_{a_1, a_2} \otimes \dots \otimes \sigma_{a_{2n-1}, a_{2n}}, \quad a \in \mathbb{Z}_2^{2n}. \quad (\text{A.5})$$

Pauli operators are orthogonal, $(\sigma_a|\sigma_b) = 2^n \delta_{a,b}$. In particular, the normalized Pauli operators $\hat{\sigma}_a = 2^{-n/2} \sigma_a$ form an orthonormal operator basis. Note that Pauli operators can be generalized to arbitrary dimensions and they give rise to an orthonormal operator basis in any of those.

We leave it as an exercise to show some basic properties of Pauli operators.

Exercise A.1 (Properties of Pauli operators). *Using the definition of Pauli operators, Eq. (A.5), show the following properties:*

- (a) $\sigma_a^\dagger = \sigma_a$ and $\sigma_a^2 = \mathbb{1}$, i.e. the multi-qubit Pauli operators are both Hermitian and unitary.
- (b) $\sigma_a \sigma_b \propto \sigma_{a+b}$, where addition is in \mathbb{Z}_2^{2n} , i.e. modulo two.
- (c) $\sigma_a \sigma_b = (-1)^{[a,b]} \sigma_b \sigma_a$ where $[a,b] := \sum_{i=1}^n a_i b_{n+i} + a_{n+i} b_i$.
- (d) $(\sigma_a|\sigma_b) = 2^n \delta_{a,b}$.

Superoperators and quantum channels Following a common nomenclature, we refer to linear maps $\phi : \mathcal{L}(\mathcal{H}) \rightarrow \mathcal{L}(\mathcal{H})$ as *superoperators* (on \mathcal{H}). We call ϕ *positivity-preserving* or simply *positive* iff $\phi(A) \geq 0$ for all $A \geq 0$. As it turns out, positive maps are not necessarily positive when we let them act on a subsystem of a composite system, i.e. if we consider $\phi \otimes \text{id}_{\mathcal{A}}$ for some auxillary system \mathcal{A} . Thus, we say that ϕ is *completely positive* iff $\phi \otimes \text{id}_{\mathcal{A}}$ is positive for any auxillary system \mathcal{A} . Completely positive maps are the ones which we consider ‘physical’, as the map quantum states to quantum states. This leads us to the definition of a quantum channel:

Definition A.2 (Quantum channel). *A quantum channel is a superoperator ϕ that is completely positive and trace-preserving, this is $\phi \otimes \text{id}_{\mathcal{A}}$ is positive for any auxillary system \mathcal{A} and $\text{tr}(\phi(A)) = \text{tr}(A)$ for all $A \in \mathcal{L}(\mathcal{H})$. We call ϕ unital iff $\phi(\mathbb{1}) = \mathbb{1}$.*

Example A.2: Quantum channels

Some examples of quantum channels are the following:

- *Unitary channels*: $\phi(X) = UXU^\dagger$ for $U \in \mathcal{U}(\mathcal{H})$.
- *Mixed-unitary channels*: $\phi(X) = \sum_i \lambda_i U_i X U_i^\dagger$ for $U_i \in \mathcal{U}(\mathcal{H})$, $\lambda_i \geq 0$, and $\sum_i \lambda_i = 1$. These are convex combinations of unitary channels.
- *Dephasing channel*: $\phi(X) = \sum_x \langle x | X | x \rangle |x\rangle \langle x|$.
- *Reset channels*: $\phi(X) = \text{tr}(X)\rho$ for a fixed quantum state ρ .

To denote superoperators, it is handy to introduce an ‘operator Dirac notation’ as follows: In analogy to the usual Dirac notation, we use the Hilbert-Schmidt inner product to define *operator kets and bras* by $|Y\rangle \equiv Y$ and $\langle X| : Y \mapsto (X|Y)$. Likewise, we can define outer products $|X\rangle\langle Y|$ which are now linear maps on $L(\mathcal{H})$, i.e. superoperators, acting as $A \mapsto (Y|A)X$.

The ‘operator bra-ket notation’ is especially useful to expand a superoperator in an operator basis, i.e. write down its matrix representation. Typically, we will use the (normalized) Pauli basis in this context, but any orthonormal basis works similarly. To this end, we observe that $\text{id} = \sum_a |\hat{\sigma}_a\rangle\langle \hat{\sigma}_a|$ and thus

$$\phi = \sum_{a,b} |\hat{\sigma}_a\rangle\langle \hat{\sigma}_a| \phi |\hat{\sigma}_b\rangle\langle \hat{\sigma}_b| =: \sum_{a,b} \phi_{a,b} |\hat{\sigma}_a\rangle\langle \hat{\sigma}_b| \quad (\text{A.6})$$

The matrix $(\phi_{a,b})_{a,b}$ is the representation of ϕ in the Pauli basis. For quantum channels, this matrix has certain properties, which we here leave as an exercise:

Exercise A.2. Let ϕ be a multi-qubit quantum channel and let $(\phi_{a,b})_{a,b}$ be its matrix representation in the Pauli basis. Show that

(a) $(\phi_{a,b})_{a,b}$ is real.

(b) $\phi_{a,0} = \delta_{a,0}$. If ϕ is unital, it also holds $\phi_{0,b} = \delta_{0,b}$, hence $\phi \simeq \begin{pmatrix} 1 & 0 \\ 0 & * \end{pmatrix}$.

(c) Suppose ϕ is a Pauli channel, this is $\phi(X) = \sum_a \lambda_a \sigma_a X \sigma_a^\dagger$ (for convex coefficients λ_a). Then, $(\phi_{a,b})_{a,b}$ is diagonal (use Ex. A.1).

Norms Throughout this paper, we use *Schatten p -norms* which are defined for any linear map $X \in L(\mathcal{V}, \mathcal{W})$ between Hilbert spaces \mathcal{V} and \mathcal{W} and $p \in [1, \infty]$ as

$$\|X\|_p := \left(\text{tr}|X|^p \right)^{\frac{1}{p}} = \left(\sum_{i=1}^d \sigma_i^p \right)^{\frac{1}{p}}, \quad (\text{A.7})$$

where $|X| := \sqrt{X^\dagger X} \in L(\mathcal{V})$ and $\sigma_i \geq 0$ are the singular values of X , i.e. the square roots of the eigenvalues of the positive semidefinite operator $X^\dagger X$. In particular, we use the *trace norm* $p = 1$, the *Hilbert-Schmidt norm* $p = 2$, and the *spectral norm* $p = \infty$. The definition of Schatten norms only relies on the Hilbert space structure of the underlying vector space, thus these norms can be defined for operators and superoperators alike.

A.2 Non-orthonormal bases

Let $(f_i)_{i \in [d]}$ be a basis of a Hilbert space \mathcal{V} . Thus, every $v \in \mathcal{V}$ has a unique expansion $v = \sum_i v_i f_i$. If $(f_i)_i$ is orthonormal, then the coefficients v_i can be simply expressed as $v_i = \langle f_i | v \rangle$. This can be

generalized to arbitrary bases by introducing the concept of a *dual basis* $(\tilde{f}_i)_i$ which is defined by the linear system of equations

$$\langle \tilde{f}_i | f_j \rangle = \delta_{i,j}. \quad (\text{A.8})$$

As $(f_i)_i$ is a basis, this system has a unique solution. It is now straightforward to verify that

$$\langle \tilde{f}_i | v \rangle = \sum_j v_j \langle \tilde{f}_i | f_j \rangle = v_i. \quad (\text{A.9})$$

Moreover, this implies that

$$\left(\sum_i |f_i\rangle \langle \tilde{f}_i| \right) (v) = \sum_i v_i f_i = v, \quad (\text{A.10})$$

for all $v \in \mathcal{V}$ and hence $\sum_i |f_i\rangle \langle \tilde{f}_i| = \text{id}_{\mathcal{V}}$.

The dual basis can be computed using the *Gram matrix*

$$G_{i,j} := \langle v_i | v_j \rangle. \quad (\text{A.11})$$

One can show that G is generally positive semi-definite and since the v_i are linearly independent, the eigenvalues are in fact strictly larger than zero. Hence, it is invertible and we define its inverse as $W := G^{-1}$. Then, the dual basis can be expressed as

$$\tilde{v}_i := \sum_j W_{i,j} v_j. \quad (\text{A.12})$$

Indeed:

$$\langle \tilde{v}_i | v_j \rangle = \sum_k W_{i,k} \langle v_k | v_j \rangle = \sum_k W_{i,k} G_{k,j} = \delta_{i,j}. \quad (\text{A.13})$$