Quantum Lazy Sampling and Path Recording for Any Group

Ben Foxman* Alex Lombardi† Fermi Ma‡ Barak Nehoran§ John Wright¶

Abstract

A central challenge in quantum algorithm analysis and cryptography is reasoning about algorithms with oracle access to a random group element (e.g. a random function, a random permutation, a random unitary). Can we efficiently simulate such algorithms? Can we determine what they know after t queries? Classically, an important tool for this is *lazy sampling*, where the oracle does not commit to the full group element at the beginning, but rather samples partial information about it on the fly. We study a quantum analog of lazy sampling: *compressed oracles* (or *recording oracles*), which are quantum data structures that allow such on-the-fly simulation for quantum queries. Compressed oracles were originally introduced by Zhandry (CRYPTO '19) for random functions, were generalized to random unitaries by Ma-Huang (STOC '25) and to permutations by Carolan (STOC '26), and have been employed to great effect in security proofs and query complexity lower bounds due to their *interpretability*.

In this work, we define and analyze a *general-purpose* and *interpretable* path-recording oracle, derived from first principles, that perfectly simulates random elements of any closed subgroup of $U(N)$. Our path-recording oracle stores superpositions of t input-output pairs $|(x_1, y_1), \ldots, (x_t, y_t)\rangle$, which encode a Feynman path explored by the algorithm and thus transparently records the information that the algorithm may have learned from its queries. Our compressed oracle builds on a recent work of Grinko and Yoshida (QIP '26), who proposed a different kind of general-purpose compressed oracle without clear interpretability. Crucially for applications, we derive an operationally useful mathematical description of our update procedure in terms of the commutant of the group's tensor power representation.

One powerful feature of our path-recording oracle is that it enables direct *comparisons* between compressed oracles for different groups, which gives a new technique for proving pseudorandomness results. For our main application, we formally relate the S_N and $U(N)$ compressed oracles, yielding what is arguably the simplest construction to date of pseudorandom unitaries: the product PC of a pseudorandom permutation and a random Clifford. This improves on the prior "PFC" construction of (Metger-Poremba-Sinha-Yuen, FOCS '24; Ma-Huang, STOC '25).

*Yale University, New Haven, CT, USA
†Princeton University, Princeton, NJ, USA
‡New York University, New York, NY, USA
§Columbia University, New York, NY, USA
¶University of California Berkeley, Berkeley, CA, USA

Contents

1 Introduction **4**
- 1.1 Prior Work . . . 6
- 1.2 Our Results . . . 7
  - 1.2.1 Tableau Recording . . . 8
  - 1.2.2 Path Recording . . . 10
  - 1.2.3 Approximations and Pseudorandomness of the PC Ensemble . . . 13

2 Technical Overview **14**
- 2.1 Tableau Recording . . . 14
  - 2.1.1 Other Query Types . . . 18
- 2.2 Path Recording . . . 20
- 2.3 Pseudorandomness of the PC Ensemble . . . 27

3 Preliminaries **30**
- 3.1 Notation for compact groups . . . 30
- 3.2 Purified Oracle Algorithms (or Adversaries) . . . 31
- 3.3 Peter-Weyl Theorem and the Fourier Transform . . . 31
  - 3.3.1 Properties of the Fourier Transform . . . 32
- 3.4 Complete Reducibility and Schur-Weyl Duality . . . 33
- 3.5 Diagram Algebras and the Schur Representation . . . 34
- 3.6 Trace forms and dual bases for semisimple algebras . . . 36
- 3.7 The (dual) Clebsch–Gordan transform and Clebsch–Gordan coefficients . . . 37
- 3.8 The Generalized Schur Transform . . . 39
- 3.9 Young Tableaux and Bratteli Paths . . . 40
- 3.10 Clebsch-Gordan Coefficients in the Gelfand-Tsetlin Basis of Unitaries . . . 43
- 3.11 The Fourier Basis of Semisimple Algebras . . . 45
- 3.12 Ratios of Irrep Dimensions for the Partition and Symmetric Group Algebras . . . 45

4 The Tableau Recording Oracle **46**
- 4.1 The Uncompressed Representation Oracle . . . 47
- 4.2 The Tableau Recording Oracle . . . 47
  - 4.2.1 Equivalence of the uncompressed and tableau recording oracles . . . 48
  - 4.2.2 Other query types . . . 51
  - 4.2.3 Efficient Implementation for the Unitary Group . . . 52

5 The Path Recording Oracle **59**
- 5.1 Defining the Path Recording Oracle . . . 60
- 5.2 Path Recording Theorem Statement . . . 62
- 5.3 Defining the partial isometry Γ_t . . . 63
- 5.4 Path Matrix Elements . . . 64
- 5.5 Proof of Theorem 5.3 . . . 66
- 5.6 Transpose queries . . . 67
- 5.7 Schur Orthogonality for the Commutant Algebras . . . 68
- 5.8 Interpretation of the Commutant EPR projector . . . 70
- 5.9 Adaptive Queries to the Update Rule . . . 71

6 Approximation and Specialization to Specific Groups **77**
6.1 The Unitary Path-Recording Oracle of [MH25] . . . 77
6.1.1 The Distinct Subspace . . . 77
6.1.2 The [MH25] oracle . . . 77
6.1.3 Proof of approximate equivalence . . . 78
6.2 The Unitary Haar Cipher . . . 82
6.2.1 Recording Oracle for the Unitary Haar Cipher . . . 86
6.2.2 Diagonal Unitaries . . . 89
6.3 Zhandry's Compressed Phase Oracle . . . 91

7 PC is a Pseudorandom Unitary **96**
7.1 The Distinct, Nonplussed Subspace . . . 96
7.1.1 Preliminaries . . . 96
7.1.2 Unitary 2-designs and the Distinct Nonplussed Subspace . . . 98
7.2 The Action of the Partition Algebra on the Distinct Nonplussed Subspace . . . 104
7.2.1 The Dual of a Permutation Element . . . 106
7.3 Additional Properties of the Update Operators . . . 108
7.4 Updates in the Distinct Nonplussed Subspace . . . 110
7.5 Putting it Together . . . 115

1 Introduction

In algorithms, complexity, cryptography, and quantum information science, it is common to encounter an algorithm or experiment that has access to some form of random transformation. In the context of quantum computing, the algorithm is typically quantum, and the transformation itself may be either classical (a function $x \mapsto f(x)$) or quantum (an n-qubit unitary transformation or a quantum channel). Although these transformations, such as random Boolean functions or Haar-random unitaries, are not computationally efficient, they serve several important purposes:

- They serve as an ideal model for justifying the correctness of algorithms. In practice, they can be replaced with an efficient transformation that is either provably pseudorandom or "scrambling" enough that the ideal model is accepted as a heuristic [BR93, SML+25]. This captures a variety of applications from the classical use of hashing in data structures [MU17] to theoretical justifications of quantum algorithms such as random circuit sampling [AAB+19].
- They can be used to prove *lower bounds* ruling out efficient black-box algorithms for computational problems (as "random instances" of black-box problems are often the hardest). For example, this can be used to rule out black-box quantum algorithms for NP that are faster than Grover search [BBBV97, Zha19].
- They can be used for analyzing the security of *cryptographic constructions* such as the Luby-Rackoff construction of pseudorandom permutations/block ciphers [LR88, Car26] and the Metger-Poremba-Sinha-Yuen construction (and Ma-Huang analysis) of pseudorandom unitaries [MPSY24, MH25]. Indeed, most such security proofs proceed by a hybrid argument: switching a pseudorandom object (such as a PRF) to a fully random one (a random function) and then arguing that an adversary cannot break the idealized variant of the construction.

In this work, we seek to understand how to analyze the behavior of such algorithms. Classically, the basic tool/data structure for such analysis is called "lazy sampling": instead of thinking of the exponential-size random object as being sampled "at the beginning," only sample the pieces of the object that are needed for each step of the algorithm. For example,

- Instead of sampling a random Boolean function $f : \{0,1\}^n \to \{0,1\}^n$, one can lazily sample its truth table. Whenever the algorithm wants to query $f(x)$, it (1) checks whether an input-output pair (x, y) has been recorded in the existing partial truth table (and uses y if so), and (2) otherwise samples a new random string y and appends (x, y) to the data structure.
- If instead the algorithm wants to query a random permutation $p : \{0,1\}^n \to \{0,1\}^n$, one can do exactly the same thing, except that each lazily sampled output y should be sampled uniformly from all output strings that do not already appear in the data structure.

In general, this kind of lazy sampling can be done for any distribution $\mathcal{D}$ of functions such that the conditional distribution

$$p(y \mid x, (x_1, y_1), \ldots, (x_t, y_t)) = \Pr_{f \leftarrow \mathcal{D}}\Big[f(x) = y \mid f(x_1) = y_1, \ldots, f(x_t) = y_t\Big]$$

is efficiently sampleable. For this efficient simulation, we think of f as undetermined outside of the partial truth table $(x_1, y_1), \ldots, (x_t, y_t)$.

Compressed Oracles and Recording Oracles. For quantum algorithms, the tool of choice for this problem has been *compressed oracles* [Zha19, MH25] (also known as *recording oracles*). Compressed oracles are *quantum data structures* that adapt the classical notion of lazy-sampling to allow quantum interactions. That is, they enable efficient simulation of quantum algorithms that interact (in a black-box manner) with a randomly chosen operation. In this work, this random operation will be a unitary sampled from some distribution $\mathcal{D}$. Specifically, we will consider the case in which $\mathcal{D}$ is chosen from the uniform measure over a finite group or, more generally, from the Haar measure over any compact Lie group.[1]

Informally, a compressed oracle records in superposition the partial information that the algorithm *would have discovered* about the group element g if it were interacting with the real oracle. Rather than fully sampling g, the compressed oracle is stateful, maintaining this partial information on-the-fly. But unlike classical lazy-sampling, the "random choices" are made coherently, in superposition. As the algorithm queries the oracle, its internal state and the data structure become entangled.

The functionality of compressed oracles is that a quantum query algorithm should not be able to distinguish whether it is interacting with the real oracle or with the stateful compressed oracle. This allows us to use the compressed oracle to analyze the behavior of the algorithm if it were to query the real oracle.

Remark 1.1 (Desiderata for compressed oracles)**.** We list some key desirable properties of a compressed oracle.

1. **Polynomial-space encoding of the data structure.** This is the bare minimum requirement to consider the oracle "compressed."
2. **Polynomial-time update algorithm.** This enables efficient simulation of quantum query algorithms accessing the oracle.
3. **Perfect simulation, or at least small simulation error.** That is, the algorithm should not be able to distinguish it from the real oracle. Ideally, this simulation is *perfect*, but if not, then we would like it to be valid for *as many queries* as possible.
4. **Interpretability.** This property is somewhat subjective. Given a quantum query algorithm $A^{(\cdot)}$ that we are simulating using a compressed oracle, we would like the data structure to encode, in some understandable way, useful information about the algorithm's computation. This is often crucial for proving any sort of lower bound (including cryptographic security proofs) against quantum query algorithms.

In the classical setting, lazy sampling algorithms indeed satisfy all of these properties (with property (2) requiring efficient conditional sampling). For example, after t classical queries to a random function f, lazy-sampling the function has the effect of recording exactly the information of a t-size partial truth table of f. Unlike classical algorithms, quantum algorithms can learn information *in superposition*, and even *forget* the information previously learned. Despite much recent effort, building and analyzing compressed oracles in the quantum setting has thus proved to be far more subtle. In this paper, we ask the following question:

Is there a general-purpose, interpretable *compressed oracle that allows one to implement such quantum lazy-sampling for every compact group?*

[1] Compact Lie groups are the groups that can be realized as closed subgroups of the unitary group in finite dimension, and include all finite groups.

1.1 Prior Work

We briefly discuss some prominent examples of existing compressed oracles:

Random Boolean functions. A seminal work of Zhandry [Zha19] built compressed oracles for simulating quantum queries to random Boolean functions. Zhandry considered two related query models:

- For a random function $f : \{0,1\}^n \to \{0,1\}$, queries to the n-qubit unitary $|x\rangle \mapsto (-1)^{f(x)} |x\rangle$.
- For a random function $f : \{0,1\}^n \to \{0,1\}^m$, queries to the $(n+m)$-qubit unitary $|x\rangle |y\rangle \mapsto |x\rangle |y \oplus f(x)\rangle$.

[Zha19] gives compressed oracles for both of these models. The compressed *phase* oracle (and its generalization to a controlled phase oracle) unambiguously satisfies all four of the above desiderata. Regarding (4), the compressed phase oracle simply records, for each input x that has been queried by the algorithm, the parity of the number of times that it has been queried. The compressed "standard" oracle is written to satisfy properties (1)–(3) and *mostly* satisfies (4) up to some small error terms (see, e.g., [Unr23, Section 3]).

These compressed oracles have been used to great effect in proving quantum query complexity lower bounds [Zha19, LZ19a, PCX25, TW25, JZ25] as well as in (post-)quantum cryptographic constructions and security analyses [Zha19, LZ19b, CFHL21, HLS24, Car26].

Haar-random unitaries. A recent work of Ma and Huang [MH25] constructs a compressed oracle that simulates Haar-random unitaries. Their construction achieves simulation error $O(t^2/N)$ for algorithms that make t queries to an N-dimensional unitary, and has an extremely natural interpretation: the compressed oracle records superpositions of *Feynman path information* $|x_1\rangle \mapsto |y_1\rangle, \ldots, |x_t\rangle \mapsto |y_t\rangle$ of the algorithm's queries, but *forgets the order* of the queries (so only the set $\{(x_1, y_1), \ldots, (x_t, y_t)\}$ is recorded). Thus their compressed oracle satisfies properties (1)–(4) except that it fails to achieve perfect simulation for (3).

Moreover, their compressed oracle has already proved quite useful in quantum cryptographic constructions [MH25, ABGL25b, ABGL25a, HY25, AGL25, SML+25, BG26], starting with the first provable construction of pseudorandom unitaries using post-quantum cryptography (analyzing the *"PFC"* construction proposed by [MPSY24]).

Random permutations. More recently, impressive progress has been made towards understanding compressed *permutations* [Unr23, MMW25, Car26].[2] Thus far, the model of choice has been a "standard" XOR oracle

$$|x\rangle |y\rangle |b\rangle \mapsto |x\rangle \ |y \oplus \pi^{(-1)^b}(x)\rangle |b\rangle$$

for a random permutation $\pi : \{0,1\}^n \to \{0,1\}^n$. The state-of-the-art result here is a compressed permutation oracle of Carolan [Car26] (as well as a compressed oracle of Unruh [Unr23] that is not an isometry) that provably simulates random permutations for up to $O(N^{1/12})$ queries [Car26]. Intriguingly, [Car26] proves quantum query lower bounds for several problems that are tight relative to his compressed permutation oracle, which would imply tight lower bounds relative to random permutations if the simulation error could be improved.

[2] See also [Ros21] for a related work on quantum query lower bounds for permutations. We discuss this related work in more detail in Section 1.2.1.

We emphasize that each of the above compressed oracles was arrived at via a (very clever) ad-hoc argument tailored to the specific simulation problem at hand. Indeed, it has proved especially challenging to identify the right kind of construction/argument for compressed permutations, let alone anything more sophisticated that might come along in the future.

Compressed Group Representations. A recent work of Grinko and Yoshida [GY25] proposed what we call a *tableau-recording oracle*: a general-purpose compressed oracle for simulating query access to a Haar-random element of a compact Lie group G that records the partial information in the Fourier basis of G. We describe the relevant formalism in more detail in Section 1.2, as it is essential to our work. Along the way to our main results, we also give a new proof of the main result of [GY25].

We emphasize that unlike prior compressed oracles, the tableau-recording oracle is not clearly interpretable (see Section 2.1 for more details). This presents a major obstruction to using it in applications such as query complexity lower bounds and cryptographic security proofs.

1.2 Our Results

In this work, we develop a general paradigm for designing and analyzing compressed oracles. In particular, we define an interpretable compressed oracle for a Haar-random element of an arbitrary compact Lie group G[3] that acts on n-qubit states via a representation $\rho : G \to U(N)$.[4] This setting easily captures the three prominent examples described above, as random functions, random permutations,[5] and Haar-random unitaries can all be described as Haar-random group representations.[6] Our framework also captures several other natural examples of interest, including:

- Haar-random *orthogonal* unitaries $O \leftarrow O(N)$ (unitaries with real entries).
- Random elements of the *hyperoctohedral* group of products PF, where P is a permutation unitary and F is a binary phase unitary. This group arises naturally in [MH25].
- Similarly, the *colored permutation* groups of products PF where F can be a q-ary or even continuous complex phase unitary.
- The Haar-random unitary cipher $|k\rangle \otimes |x\rangle \mapsto |k\rangle \otimes U_k |x\rangle$ for a collection of i.i.d. Haar-random unitaries $\{U_k\}$; this model represents an idealized pseudorandom unitary or encryption scheme.
- Random elements of *subgroups* of the symmetric group such as the alternating group A_N.

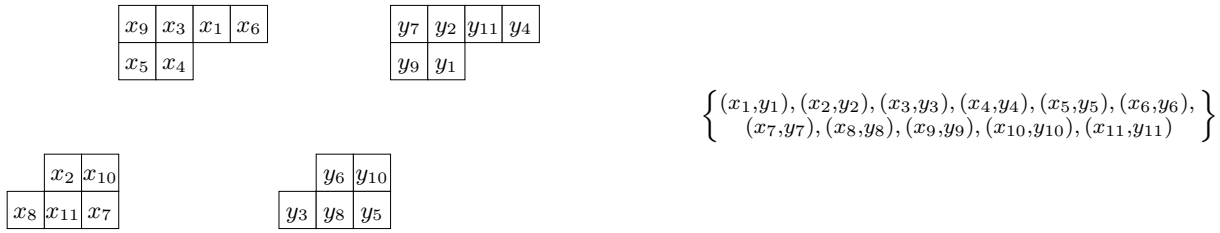
Since our compressed oracle can be seen as a generalization of the Ma-Huang oracle, we refer to our compressed oracle as the **path-recording oracle**.

[3] We note, again, that this includes all finite groups as a special case.

[4] A *unitary representation* of a group G is any homomorphism from the group to the unitaries. That is, it requires that $\rho(g)\,\rho(h) = \rho(gh)$ for all group elements $g, h \in G$.

[5] To capture permutations, we instead consider the *in-place* permutation oracle $|x\rangle \mapsto |\pi(x)\rangle$, which (along with its inverse) is algorithmically equivalent to the XOR oracle.

[6] We discuss later how to handle inverse queries, which are an *anti*-representation of the group.



(a) A basis element of the tableau recording.

(b) A basis element of the path recording.

Figure 1: In the two types of compressed oracles, the resulting data structures take different forms. The examples shown above correspond to the case in which the group is the unitary group.

In fact, we derive both the tableau-recording oracle and the path-recording oracle together from first principles, establishing the following:

- For every unitary representation ρ of a compact Lie group G, the tableau- and path-recording oracles perfectly simulate query access to a Haar-random $\rho(g)$ (and therefore, simulate each other). We identify simple Uhlmann transformations relating these compressed oracles to each other and to a naïve "uncompressed oracle" (see Section 2.1).
- In the special case $G = U(N)$, our framework easily recovers the Ma-Huang approximate path-recording oracle while avoiding their $O(t^2/N)$ simulation error.
- Other settings of our path-recording oracle automatically recover Zhandry's compressed oracle as well as other natural ideal models.

Finally, as our main application, we use our path-recording oracles to prove the security of a new, extremely simple construction of a pseudorandom unitary: the product PC of a random permutation P and a circuit C sampled from a unitary 2-design (such as a random Clifford circuit).

1.2.1 Tableau Recording

As a starting point, in Section 4, we give a new proof of the following theorem.

Theorem 1.2 ([Har05, GY25])**.** *Let ρ be an N-dimensional unitary representation of a compact Lie group G. Let $A^{\rho(g)}$ denote a quantum query algorithm with oracle access to $\rho(g)$ for a Haar-random $g \leftarrow G$ (for finite groups, the uniform distribution).*

Then, there is an isometry Q_ρ^{Tab} acting on the algorithm's query register A *and a recording register* R *with the following properties:*

1. *After t queries to Q_ρ^{Tab}, the dimension of* R *is at most N^{2t}.*
2. *Q_ρ^{Tab} has an efficient implementation relative to the* Clebsch-Gordan transform *[BCH05, Har05] and* dual Clebsch-Gordan transform *[Ngu23, GBO23] for G.*
3. *The mixed state obtained by running $A^{\rho(g)}$ for a Haar-random g is* identical *to that of running $A^{Q_\rho^{\mathrm{Tab}}}$ and tracing out the* R *register.*

Moreover, there is a variant of this isometry that also supports queries to the conjugate *representation $\rho(g)^*$, the inverse map $\rho(g)^\dagger$, the transpose map $\rho(g)^T$, and controlled versions of these unitaries.*

Thus, the Tableau-recording oracle *always* satisfies desiderata (1) and (3), and reduces property (2) (efficient updates) to a previously studied problem in algorithmic representation theory. In particular, applying this framework to the defining representation of the unitary group $U(N)$, this gives an efficient algorithm for *perfect* stateful simulation of Haar-random unitaries.

Related work to Theorem 1.2. The tableau recording isometry Q_ρ^{Tab} was originally studied in [Har05, Section 8.1.3] in the different context of "generalized phase estimation".[7] Its application to compressed oracles and simulation, including the first proof of Theorem 1.2 and the first efficient $U(N)$ implementation, is due to the recent work [GY25]. Our proof uses a different formalism/approach compared to [GY25], proceeding from first principles by purifying the group element in the (countably infinite-dimensional) Hilbert space of square integrable functions on the group, $L^2(G)$, and then viewing it in the Fourier basis to recover the tableau-recording oracle (see Section 2.1 for more details). This gives an explicit Uhlmann transformation (the Fourier transform) relating the compressed and naïve uncompressed representation oracles.

Additionally, a recent work of Rosmanis [Ros21] considers the special case of random permutations and proves a quantum query lower bound for permutation inversion. While [Ros21] uses many of the techniques typical of the compressed oracle literature, the work only considers an uncompressed oracle, with the purification register consisting of a structured superposition of permutations that agree with a partial assignment $\{x_i \to y_i\}_{i\in[t]}$. Such states are not orthogonal, and so it does not define or study a compressed oracle. It does take a representation-theoretic view of the projectors onto low and high success probability, which correspond to projectors onto certain natural subspaces in the Fourier/tableau basis.

Other Query Types. We briefly mention two improvements that we make to Theorem 1.2. First, we give simple black-box reductions between the four query types $\rho(g)$, $\rho(g)^*$, $\rho(g)^\dagger$, and $\rho(g)^T$.

Theorem 1.3. *Given black-box access*[8] *to* Q_ρ^{Tab}*, we can implement* $Q_{\rho^*}^{\mathrm{Tab}}$*,* $Q_{\rho^T}^{\mathrm{Tab}}$*, and* $Q_{\rho^\dagger}^{\mathrm{Tab}}$ *each with a single application of* Q_ρ^{Tab}*, and simple pre- and post-processing.*

In the case of transpose queries, $Q_{\rho^T}^{\mathrm{Tab}}$, this pre- and post-processing only involves exchanging two registers, which is, of course, simple and efficient. In the case of conjugate and inverse queries, this involves performing a bijective mapping between the labels and bases of two corresponding irreducible representations[9] (between an irrep and its *dual irrep*), and its efficiency depends on the group and how its irreps are encoded.

The Unitary Group. For the unitary case in particular, we show in Section 4.2.3 that this mapping is a simple signed flip on the tableau data structure (which consists of simply rotating the data structure, with an attached ± 1 phase depending on its parity; see Definition 4.2). This in particular allows us to show that the dual Clebsch-Gordan transform [Ngu23, GBO23] for the unitary group can be very simply reduced to the (standard) Clebsch-Gordan transform.

[7] Harrow was interested in an algorithm — called generalized phase estimation (GPE) — that allowed for measuring (and performing operations controlled on) the hidden irrep label associated with a state $|\psi\rangle \in V \simeq \bigoplus_\lambda V_\lambda \otimes M_\lambda$ belonging to a completely *reducible* representation of G, *without* having to explicitly decompose V. Rather than maintaining a recording register R for simulating queries to a Haar-random $\rho(g)$, Harrow used R as a temporary workspace for GPE, which only requires one query to what we call Q_ρ^{Tab}. In Section 8.1.3, Harrow observes the form of Q_ρ^{Tab} as part of a discussion on the relationship between the Clebsch-Gordan transform and the Fourier transform.

[8] For complex conjugate queries and transpose queries, we technically need the complex conjugate, $\overline{Q_\rho^{\mathrm{Tab}}}$, of Q_ρ^{Tab}. For all the cases we consider, it is possible to write Q_ρ^{Tab} using only real coefficients, which makes these the same. But even for general groups, this provides a useful reduction from conjugate and transpose queries to $\overline{Q_\rho^{\mathrm{Tab}}}$, which has a generic efficient implementation given any implementation of Q_ρ^{Tab}.

[9] A representation ρ_λ is *irreducible* if it has no non-trivial G-invariant subspaces. It turns out that for all compact Lie groups G, there is a countable list of all irreducible representations (or irreps) of G up to isomorphism, and for particular groups of interest enumerating irreps and their bases often has a simple combinatorial interpretation.

Lemma 1.4. *For the unitary group, the dual Clebsch-Gordan transform, in the same basis need for Theorem 1.2, is given by a* signed flip *(the* Flip *operation in Definition 4.2), followed by a Clebsch-Gordan transform, followed by another signed flip. That is,* $\mathrm{dCG} = \mathsf{Flip} \cdot \mathrm{CG} \cdot \mathsf{Flip}$.

Thus we can express the update of Theorem 1.2 purely in terms of the Clebsch-Gordan transform without requiring a separate dual Clebsch-Gordan transform. Moreover, the different query types are recovered by pre- and post-processing with a Swap, a Flip, or both.[10]

Query type	**Update**
Standard $\rho(g)$	Q_ρ^{Tab}
Transpose $\rho(g)^T$	$\mathsf{Swap} \cdot Q_\rho^{\mathrm{Tab}} \cdot \mathsf{Swap}$
Conjugate $\rho(g)^*$	$\mathsf{Flip} \cdot Q_\rho^{\mathrm{Tab}} \cdot \mathsf{Flip}$
Inverse $\rho(g)^\dagger$	$\mathsf{Swap} \cdot \mathsf{Flip} \cdot Q_\rho^{\mathrm{Tab}} \cdot \mathsf{Flip} \cdot \mathsf{Swap}$

Table 1: Summary of the four different query types for the unitary group. In each case, one must only pre- and post-process with one or both of two simple operations.

1.2.2 Path Recording

The primary drawback of Q_ρ^{Tab} is *interpretability*; given the contents of the recording register R, it is hard to intuit and argue what information about $\rho(g)$ has actually been recorded.

For our first main contribution, we give a *second* compressed oracle Q_ρ^{Path} that provides better interpretability. In order to motivate Q_ρ^{Path}, we briefly recall the [MH25] approximate path-recording oracle for the defining representation of $U(N)$. In the [MH25] oracle,

- The compressed oracle register R is spanned by *set states* $|\{(x_1, y_1), \ldots, (x_t, y_t)\}\rangle$, on which the restriction is imposed that the y_i are all distinct.
- When a new query is made on an input $|x\rangle$, the recording register $|S\rangle$ is updated to a uniform superposition of $|S \cup \{(x, y)\}\rangle$ over all $y \notin S$, while the state $|y\rangle$ (entangled with the recording register $|S \cup \{(x, y)\}\rangle$) is returned as the output.

The rough intuition for this oracle is that the compressed purification R is supposed to record Feynman path information $x_1 \mapsto y_1, \ldots, x_t \mapsto y_t$ about the oracle queries; however, only the *set* of pairs is recorded, which means that the *order* of the queries has been forgotten by the data structure.

We build a path-recording oracle that, unlike the one in [MH25], is *exact* (no simulation error), and moreover fully generalizes to work for any compact Lie group. While for $G = U(N)$ (and ρ the defining representation[11]), the path-recording oracle records Feynman paths up to a *permutation* symmetry (that is, a symmetry in the group algebra $\mathbb{C}[S_t]$), our general path-recording oracle records Feynman paths up to symmetries imposed by the *commutant algebra* of $\rho^{\otimes t}$. This is the algebra $\mathcal{A}_t$

[10] Here, we are writing Swap to swap between two sets of registers in the recording, and Flip to apply in parallel to both sets of registers. Note that Flip itself is simply rearranging the entries on each side, plus a sign that is trivial to compute based on the parity of all the entries.

[11] The *defining* or *fundamental representation* of $U(N)$ is the one that represents each unitary as its own matrix $\rho(U) = U$ for every group element $U \in U(N)$.

of all operators that commute with the tensor power of $\rho(g)$ for all group elements (see Table 2 for the commutant algebras of some important groups):

$$\mathcal{A}_t = \{A \in \mathrm{End}((\mathbb{C}^N)^{\otimes t}) : A \cdot \rho(g)^{\otimes t} = \rho(g)^{\otimes t} \cdot A \text{ for all } g \in G\}. =: \mathrm{End}_G((\mathbb{C}^N)^{\otimes t})$$

Group		Tensor space	Representation	Commutant algebra	
Unitary	$U(N)$	$V_N^{\otimes t}$	$U^{\otimes t}$	Symmetric Group Algebra	$\mathbb{C}[S_t]$
Unitary Product (Haar Cipher)	$U(N)^K$	$\left(V_N^{\oplus K}\right)^{\otimes t}$	$\left(\sum_{k\in[K]} \lvert k\rangle\langle k\rvert \otimes U_k\right)^{\otimes t}$	Colored Permutation Group Algebra	$\mathbb{C}[\mathbb{Z}_K \wr S_t]$
Orthogonal	$O(N)$	$V_N^{\otimes t}$	$U_g^{\otimes t}$ U_g is the standard embedding of $g \in O(N)$ as a unitary over $\mathbb{C}^t$	Brauer Algebra	$B_t(N)$
Colored Permutations	$\mathbb{Z}_r \wr S_t$	$V_N^{\otimes t}$	$(P_\pi F_f)^{\otimes t}$ P_π is the permutation matrix of $\pi \in S_N$, and F_f is diagonal with r'th roots of unity	Tanabe Algebra	$T_t(N, r)$
Symmetric	S_N	$(V_{N-1} \oplus V_1)^{\otimes t}$	$P_\sigma^{\otimes t}$ P_σ is the permutation matrix of $\sigma \in S_N$	Partition Algebra	$P_t(N)$
Boolean Functions	$\mathbb{Z}_2^N$	$\left(V_1^{\oplus N}\right)^{\otimes t}$	$\left(\sum_{x\in[N]} (\text{-}1)^{f(x)} \lvert x\rangle\langle x\rvert\right)^{\otimes t}$	Colored Even-Partition Algebra	

Table 2: Example instances of commutant relationships (Schur-Weyl dualities) between group actions on tensor spaces and their commutant algebras. The path-recording oracle for any compact Lie group can be described in terms of its corresponding commutant algebra.

Unlike the [MH25] oracle, the updates for our generalized path-recording also require reweighting different subspaces in an important representation-theoretic basis called the Schur basis (see Section 2.2 for details). For now, we informally state the form of our path-recording oracle to give a taste of the result.

Theorem 1.5 (informal, see Definition 5.1 and Theorem 5.3)**.** *Let ρ be an N-dimensional unitary representation of a compact Lie group G. Let $A^{\rho(g)}$ denote a quantum query algorithm with oracle access to $\rho(g)$ for a Haar-random $g \leftarrow G$.*

Then, there is an isometry Q_ρ^{Path} acting on the algorithm's query register A and a recording register R with the following properties:

1. *After t queries to Q_ρ^{Path}, the content of R is ordered Feynman path information $(x_1, \ldots, x_t, y_1, \ldots, y_t)$ for $x_i, y_i \in [N]$.*

2. *The t^{th} application of Q_ρ^{Path} can be expressed in the following form:*

$$Q_{\rho,t}^{\text{Path}} = \underbrace{\left(\Lambda_{\mathcal{A}_t}\right)_{\mathsf{R}}}_{\text{ReweightEnd}} \cdot \underbrace{\left(\Omega_{\mathcal{A}_t}\right)_{\mathsf{R}}}_{\text{Symmetrize}} \cdot \underbrace{\text{App}}_{\text{Append}} \cdot \underbrace{\left(\Lambda^+_{\mathcal{A}_{t-1}}\right)_{\mathsf{R}_{\leq t-1}}}_{\text{ReweightStart}}, \tag{1.1}$$

where:

- $\mathcal{A}_t = \text{End}_G((\mathbb{C}^N)^{\otimes t})$ *denotes the commutant algebra of* $\rho^{\otimes t}$.
- $\Lambda_{\mathcal{A}_t}, \Lambda^+_{\mathcal{A}_{t-1}}$ *denote certain* subspace reweighting *operators,*
- $\text{App} = \sum_{x_t,y_t} |y_t\rangle\langle x_t|_{\mathsf{A}} \otimes |(x_t,y_t)\rangle_{\mathsf{R}_t}$ *denotes the (un-normalized) initialization of an EPR pair shared between the algorithm and the recording.*
- $\Omega_{\mathcal{A}_t}$ *denotes a* symmetrization *projection that can be described in terms of* $\mathcal{A}_t$ *and represents* forgetting *partial Feynman path information.*

3. *The mixed state obtained by running $A^{\rho(g)}$ for a Haar-random g is* identical *to that of running $A^{Q_\rho^{\text{Path}}}$ and tracing out the* R *register.*

Moreover, essentially the same update also supports queries to $\rho(g)^T$, while another variant additionally supports queries to $\rho(g)^*$ and $\rho(g)^\dagger$.

We give more details in Sections 2.2 and 5, but we emphasize that the symmetrization operation $\Omega_{\mathcal{A}_t}$ has an extremely natural interpretation directly in the "path basis" $|(x_1,y_1),\ldots,(x_t,y_t)\rangle$. (Lemma 5.11): Choose any basis for the commutant algebra (for instance, for group algebras, simply take the basis of the group elements), and apply every such basis element in superposition.[12] For instance, for the unitary group, Ω_{S_t} amounts to applying all possible permutations, in superposition, among the t different (x_i,y_i) pairs. Thus, modulo understanding the subspace reweighting $\Lambda_{\mathcal{A}_t}$, the update rule is highly interpretable.

Equivalence of Tableau and Path Recording. We emphasize that as written, Q_ρ^{Path} is a mathematical description and is not necessarily algorithmic. On the other hand, we show that Q_ρ^{Path} is isometric to the algorithmic Q_ρ^{Tab} via two applications of the generalized Schur transform, special cases of which have been studied in [BCH05, Kro19, FTH23, Ngu23, GBO23, BFG+25], and which can be built from Clebsch-Gordan transforms. We thus have that the updates to our path recording oracle can likewise be made efficient with an efficient Clebsch-Gordan transform.

Concurrent Work on Path Recording. Concurrently to this work, we were made aware that the authors of [GY25] will update [GY25] to include a version of the path-recording oracle. Their path-recording oracle, like ours, is derived by use of a generalized Schur transform on the tableau-recording oracle. Their path recording update is defined by transporting the tableau update through this isometry. In the special case of $G = U(N)$, they also give a mathematical formula (different from ours) for the path basis update matrix coefficients. In this work, beyond defining the path-recording oracle for general compact Lie groups, we focus on providing an interpretable mathematical description of the path basis update rule (Theorem 1.5) and showing how to use it for security proofs and applications.

[12] Technically, we apply each such basis element to the all the x_i's and its *dual transpose* to all the y_i's (see Lemma 5.11 for details). For group algebras, this amounts to applying the same group element across both the x and y entries.

1.2.3 Approximations and Pseudorandomness of the PC Ensemble

Our second main contribution is analyzing Q_ρ^{Path} in special cases of interest by introducing *approximations*. Our thesis is that for each group representation, knowing what the commutant algebra $\mathcal{A}_t$ looks like is the key to understanding what the exact updates to the Feynman path recording are effectively doing to the Feynman paths (when they record, when they forget, etc.), and how a more elementary approximate update can be derived.

As a proof of concept, in Section 6, we show how to derive the [MH25] path-recording oracle for Haar-random unitaries from our path-recording oracle (when specialized for $U(N)$) via a sequence of two approximations:

- restricting our oracle to appending *distinct* outputs $y_1, \ldots, y_t$ in the path basis, and
- replacing the subspace reweighting operators $\Lambda_{\mathcal{A}_t}, \Lambda^+_{\mathcal{A}_{t-1}}$ by *scalars* (i.e., no reweighting),

which results in the Ma-Huang oracle, rewritten as

$$Q_{MH}^{\mathrm{Path}} = \left(\sqrt{t} \cdot \Omega_{\mathbb{C}[S_t]}\right)_{\mathsf{R}} \cdot \frac{1}{\sqrt{N-t+1}} \left(\Pi_{\mathsf{Dist}}\right)_{\mathsf{R}_Y} \mathsf{App}_{\mathsf{AR}}.$$

Indeed, one can essentially read off the algorithmic Ma-Huang update from this expression: the key idea is identifying the Ma-Huang set states $\{(x_1, y_1), \ldots, (x_t, y_t)\}$ as S_t-symmetrized path states $\frac{1}{\sqrt{t!}} \sum_{\pi \in S_t} \pi^{\otimes 2} |x_1, \ldots, x_t, y_1, \ldots, y_t\rangle$.

Expanding on this, we show in Section 6 how the path-recording oracle can be used to *automatically* derive three other compressed oracles that have been either implicitly or explicitly studied before: Zhandry's original compressed (phase) oracle, the ideal Haar cipher $|k, x\rangle \mapsto |k\rangle \otimes U_k |x\rangle$, and Haar-random diagonal unitaries $|x\rangle \mapsto u_x |x\rangle$.

Finally, for our primary application, in Section 7 we analyze the path-recording oracle for the in-place representation of the permutation group S_N:

$$\rho(\pi) \cdot |x\rangle = |\pi(x)\rangle .$$

This compressed oracle is particularly well-behaved—it turns out to approximately match the compressed unitary oracle!—on recording states restricted to a dense "distinct, nonplussed" subspace $\Pi_{\mathsf{DNP}}^{X,Y}$. Through a careful analysis, we are able to leverage this to prove the following theorem:

Theorem 1.6 (see Theorem 7.34)**.** *For all t-query adversaries $\mathsf{Adv}^{(\cdot)}$, the oracle distribution $P \cdot C$ for a uniformly random in-place permutation P and C sampled from any unitary 2-design is $O(t^2/N)$-indistinguishable from a Haar-random unitary.*

This improves on the PFC construction of pseudorandom unitaries (PRUs) of Metger, Poremba, Sinha, and Yuen [MPSY24] (proved secure by Ma and Huang [MH25]): we prove that the random phase F is actually unnecessary, and PC is already a PRU! By using a pseudorandom permutation for P, this yields, in our opinion, the simplest PRU construction that has been analyzed to date.

Accordingly, the same proof gives a new, simple construction of unitary t-designs using $2t$-wise independent permutations:

Theorem 1.7 (See Corollary 7.35)**.** *Suppose that P is sampled from a family of non-adaptive (t, ε) post-quantum pseudorandom permutations (respectively, (t, ε)-adaptively secure post-quantum pseudorandom permutations) with* statistical security *and that C is sampled from an exact 2-design. Then, $P \cdot C$ is an $\varepsilon + O(t^2/N)$ approximate t-design (respectively, an $\varepsilon + O(t^2/N)$ approximate adaptive t-design).*

In particular (see [Zha12, MPSY24]), the conclusions hold if P is sampled from a family of $2t$-wise independent permutations, or $\frac{\varepsilon}{2} \cdot N^{-t}$-almost $2t$-wise independent permutations.

More broadly, because of its wide applicability, as well as its interpretability, we believe that our path-recording oracle will facilitate a variety of additional pseudorandomness results as well as new quantum query complexity lower bounds.

2 Technical Overview

In this section, we overview the main ideas behind our results and their proofs.

2.1 Tableau Recording

We begin with the tableau recording oracle. Again, we remark that the tableau recording oracle for finite groups can be seen implicitly in [Har05, Section 8.1.3] and was also defined and analyzed in [GY25]. We describe the construction for expository purposes, for some technical reasons (a different proof and some slight variations/improvements), and because it is connected to and will help explain our path-recording oracle.

First, given a representation $\rho : G \to U(N)$, why should we expect there to be a compressed oracle at all? The reason is a direct generalization of Zhandry's compressed oracle. Recall that in the case of the binary phase representation $\rho(f)\left|x\right\rangle = (-1)^{f(x)}\left|x\right\rangle$ of functions $f : \{0,1\}^n \to \{0,1\}$,[13] Zhandry's compressed oracle is constructed in two steps:

1. First, *purify* the random function f to a function register $\mathsf{F} \simeq \mathbb{C}^{\{0,1\}^{2^n}}$; here, F is initialized to the uniform superposition $\left|+\right\rangle$ and a function query is represented by the unitary
$$Q_\rho\left|x\right\rangle\left|f\right\rangle = (-1)^{f(x)}\left|x\right\rangle\left|f\right\rangle.$$
2. Second, apply the quantum Fourier transform (in this case, simply the Hadamard transform) to F: the initialization is then $\left|\emptyset\right\rangle$ rather than $\left|+\right\rangle$, while a function query is now written as
$$Q_\rho^{\mathsf{Fourier}}\left|x\right\rangle\left|D\right\rangle = \left|x\right\rangle\left|D \oplus \{x\}\right\rangle,$$
where D is interpreted as a subset of $\{0,1\}^n$, or "database", and $D \oplus \{x\}$ adds x to the database if $x \notin D$ and otherwise removes it.

The magic is that in the Fourier basis, after t queries, the database state $\left|D\right\rangle$ is always t-sparse. It can thus be stored in sparse form using $tn = t\log N$ space. (In addition, the update is extremely simple and can be implemented efficiently.)

How does this generalize to other group representations? The natural idea is to consider the same two-step process:

1. First, purify the random group element g to a "uniform superposition" on a register G. In the case of finite groups, we write $\mathsf{G} = \mathbb{C}[G]$ to be a Hilbert space with basis $\{\left|g\right\rangle, g \in G\}$. In this language, the initial state is indeed the uniform superposition $\left|+\right\rangle$, and the purified query can be written as
$$Q_\rho\left|x\right\rangle\left|g\right\rangle = \rho(g)\left|x\right\rangle \otimes \left|g\right\rangle.$$
In the case of infinite groups such as $U(N)$, this is not the right purification (it would have uncountable dimension); instead, we use the well-studied countable-dimension Hilbert space of square-integrable functions
$$L^2(G) = \left\{f : G \to \mathbb{C} \text{ such that } \int_G |f(g)|^2 dg < \infty\right\},$$

[13] We will later see that we should treat such functions as group elements of the group $\mathbb{Z}_2^N$, for $N = 2^n$.

where dg denotes the Haar measure on G. This Hilbert space comes equipped with commuting left- and right- actions of G (pre- and post- composition) that together define what is called the **regular representation** of G. This register is naturally initialized to the purified Haar measure ($f = \mathbf{1}$) and the purified query is analogous to above:

$$Q_\rho\Big(|x\rangle \otimes \int_G f(g)\,|g\rangle\,dg\Big) = \int_G \rho(g)\,|x\rangle \otimes f(g)\,|g\rangle\,dg.$$

2. Second, apply the quantum Fourier transform to G. When generalized from finite cyclic groups to arbitrary compact Lie groups, the Fourier transform is an isomorphism

$$F : L^2(G) \xrightarrow{\sim} \bigoplus_{\lambda\in\widehat{G}} V_\lambda \otimes V_\lambda^*, \tag{2.1}$$

where the index $\lambda \in \widehat{G}$ ranges over *all irreducible representations* (or *irreps*) of G (up to isomorphism)[14] and V_λ denotes the vector space underlying the irreducible representation ρ_λ.[15] In other words, there exists an orthonormal Fourier basis spanned by states

$$F^{-1}\,|\lambda, X, Y\rangle = \sqrt{\dim(V_\lambda)}\int_G \langle Y|\,\rho_\lambda(g)\,|X\rangle\,|g\rangle\,dg \in L^2(G) \tag{2.2}$$

for $\lambda \in \widehat{G}$, and $|X\rangle, |Y\rangle$ ranging over a basis of V_λ, such that the Fourier transform decomposes $L^2(G)$ in this basis.

We call this basis $\bigoplus_\lambda V_\lambda \otimes V_\lambda^*$ the **tableau basis** by analogy to the case $G = U(N)$: the irreducible representations of $U(N)$ are indexed by Young diagrams λ, and the basis elements $|X\rangle$ of V_λ correspond to semi-standard Young tableaux of shape λ (see Section 3.9 for details on Young diagrams and tableaux).

While this is a natural proposal, the big questions are:

- Is this Fourier basis representation compressible?
- What is the form of the query Q_ρ^{Tab} in this basis?

The first question is implicitly answered by the resolution to the second, but let us answer it separately anyway since it will be relevant for the path-recording oracle. One way to see why the basis is compressible is to observe that all recording states obtained by making t queries to Q_ρ are in the span of the following "t-power matrix element states" of $L^2(G)$:

$$\int_G \langle y_1, \ldots, y_t|\,\rho(g)^{\otimes t}\,|x_1, \ldots, x_t\rangle\,|g\rangle\,dg\,.$$

This already places a limit on the *dimension* of accessible states on the recording register, but it does not say how they appear in the Fourier basis. To understand this, we make use of the *Schur transform* [BCH05]

$$\mathsf{Schur} : (\mathbb{C}^N)^{\otimes t} \xrightarrow{\sim} \bigoplus_{\lambda\in\widehat{G}_t} V_\lambda \otimes V_\lambda^{\mathcal{A}_t}, \tag{2.3}$$

[14] For a compact Lie group, there are only countably many irreps.

[15] Note that we use $\rho(g)$ to refer to the queried representation, which may or may not be reducible. On the other hand, $\rho_\lambda(g)$ is the irreducible representation corresponding to $\lambda \in \widehat{G}$.

where $\lambda \in \widehat{G}_t$ ranges over a finite subset of the possible irreps[16] and $V_\lambda^{\mathcal{A}_t}$ is some multiplicity space (we will say more about $V_\lambda^{\mathcal{A}_t}$ later). In particular, this means that there is a Schur basis consisting of states $\mathsf{Schur}^\dagger |\lambda, X, S\rangle \in (\mathbb{C}^N)^{\otimes t}$ (such that the tensor power $\rho(g)^{\otimes t}$ acts as the λ irrep, mapping them to $\mathsf{Schur}^\dagger\Big(|\lambda\rangle \otimes \rho_\lambda(g) |X\rangle \otimes |S\rangle \Big)$), and such states together span $(\mathbb{C}^N)^{\otimes t}$.

We can thus insert the identity $\mathsf{Schur}^\dagger\mathsf{Schur}$ and write the t-power matrix element state as

$$\int_G \langle y_1, \ldots, y_t| \, \rho(g)^{\otimes t} \, |x_1, \ldots, x_t\rangle \, |g\rangle \, dg = \int_G \langle y_1, \ldots, y_t| \; \mathsf{Schur}^\dagger \, \mathsf{Schur} \; \rho(g)^{\otimes t} \; \mathsf{Schur}^\dagger \, \mathsf{Schur} \; |x_1, \ldots, x_t\rangle \, |g\rangle \, dg \,,$$

where $\mathsf{Schur} \; \rho(g)^{\otimes t} \; \mathsf{Schur}^\dagger = \sum_{\lambda \in \widehat{G}_t} |\lambda\rangle\langle\lambda| \otimes \rho_\lambda(G) \otimes \mathbf{I}$. This implies that each t-power matrix element state is in the span of a restricted set of Fourier basis states (Equation (2.2)),

$$\int_G \langle Y, T| \Big(\rho_\lambda(g) \otimes \mathbf{I}\Big) |X, S\rangle \, |g\rangle \, dg \;=\; \delta_{S,T} \int_G \langle Y| \, \rho_\lambda(g) \, |X\rangle \, |g\rangle \, dg \;\propto\; \delta_{S,T} \cdot F^{-1} \, |\lambda, X, Y\rangle \;,$$

such that λ comes from the restricted set of irreps $\widehat{G}_t$. This exactly says that its Fourier transform is supported on

$$\bigoplus_{\lambda \in \widehat{G}_t} V_\lambda \otimes V_\lambda^* ,$$

a small finite subset of the Fourier basis, of dimension at most $\dim((\mathbb{C}^N)^{\otimes 2t}) = N^{2t}$. So the purification is in fact sparse in the Fourier (tableau) basis! Moreover, it can be stored using space at most $O(t \log N)$.

Finally, we consider the update rule Q_ρ^{Tab}. For the $(t+1)$th query, this must be an isometry which, for each $\lambda \in \widehat{G}_t$, maps

$$|x\rangle \otimes |\lambda, X, Y\rangle \in \mathbb{C}^N \otimes V_\lambda \otimes V_\lambda^* \;\longrightarrow\; \mathbb{C}^N \otimes \bigoplus_{\lambda^+ \in \widehat{G}_{t+1}} V_{\lambda^+} \otimes V_{\lambda^+}^* \,,$$

whose basis elements look like

$$|y\rangle \otimes \left|\lambda^+, X^+, Y^+\right\rangle .$$

Intuitively, upon the algorithm's query, we want to grow X, Y "by one element each" to record an additional piece of information. It turns out that this update can be viewed as a composition of two instances of a well-studied map from representation theory: The Clebsch-Gordan transform [BCH05, Har05] is an isomorphism describing the irrep decomposition[17]

$$\mathbb{C}^N \otimes V_\lambda \xrightarrow{\sim} \bigoplus_{\lambda^+ \in \widehat{G}_{t+1}} V_{\lambda^+} \otimes M_{\lambda,\lambda^+}$$

for some multiplicity spaces M_{λ,λ^+}. For example, in the case $G = U(N)$ (where the M_{λ,λ^+} all have dimension 1), this describes the procedure for "inserting a new box" into a Young diagram/semi-standard Young tableau (as in Figure 2).

[16] For example, in the case $G = U(N)$ and ρ is the defining representation, $\widehat{G}_t$ is the set of all Young diagrams with exactly t boxes.

[17] Typically, the Clebsch-Gordan transform is reserved to the setting where both tensor factors are irreducible. In the body, we will often explicitly decompose ρ into irreps and then apply a controlled Clebsch-Gordan transform. However, it is also reasonable to consider a generalized definition of the Clebsch-Gordan transform allowing for reducible tensor factors. This will be useful, for example, for S_N, whose natural representation on $\mathbb{C}^N$ is reducible.

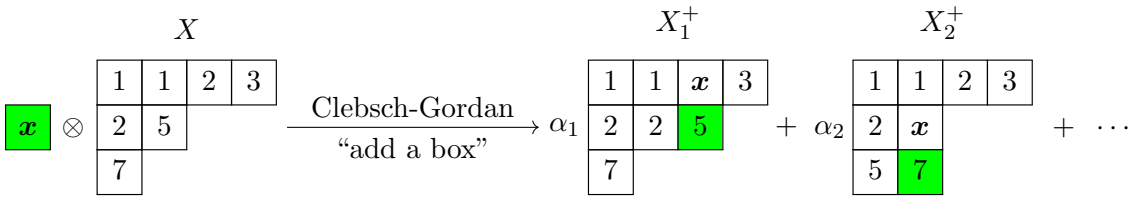


Figure 2: Example Clebsch-Gordan update for the unitary group, with the newly added box highlighted. Roughly speaking, each shape λ (Young diagram) that contains one extra box not in the original, and each valid filling (tableau) with the contents of X plus the new x value, receive some weight.

In general, applying this transformation to $|x, \lambda, X\rangle$ will produce a superposition of elements of the form

$$\left|\lambda, \lambda^+, X^+, m\right\rangle$$

with $|X^+\rangle \in V_{\lambda^+}$ and $|m\rangle \in M_{\lambda,\lambda^+}$. This (along with $|Y\rangle$), of course, does not match the type of vector that the update rule is supposed to output.

However, we can arrive at a valid element of $\bigoplus_{\lambda^+} V_{\lambda^+} \otimes V_{\lambda^+}^*$ by *inverting* another irrep decomposition

$$\bigoplus_{\lambda} V_\lambda^* \otimes M_{\lambda,\lambda^+} \to \mathbb{C}^N \otimes V_{\lambda^+}^*.$$

We refer to this map as an inverse *dual* Clebsch-Gordan transform.[18] Applying this map to $|\lambda^+, \lambda, Y, m\rangle$ yields an element of $\mathbb{C}^N \otimes \bigoplus_{\lambda^+ \in \widehat{G}_t} V_{\lambda^+} \otimes V_{\lambda^+}^*$. It turns out that this composition (see Fig. 3) is equal to Q_ρ^{Tab}!

In [Har05], Harrow gives an abstract algebraic proof that these maps are equivalent (in the case of finite groups). In Section 4, we give a proof via a direct calculation in the Fourier basis. This is somewhat necessary in order to reason about algorithmic efficiency, because the Clebsch-Gordan transform is only defined up to choices of basis for the irrep and multiplicity spaces.[19] We prove that the composition correctly computes Q_ρ^{Tab} whenever the Clebsch-Gordan and dual Clebsch-Gordan transforms are compatibly chosen to satisfy a certain identity on their matrix coefficients (see Lemma 3.21). Since, in the case of $G = U(N)$, a standard choice of Clebsch-Gordan transform on the Gelfand-Tsetlin basis of V_λ satisfies this matrix coefficient identity *and* has an efficient algorithmic implementation [BCH05, BFG+25], we conclude that $Q_{U(N)}^{\mathrm{Tab}}$ has an efficient algorithm.

[18] In the body, we technically define the dual Clebsch-Gordan transform to decompose $V_{\lambda^+} \otimes (\mathbb{C}^N)^*$ into (primal) irreps. Prior work on the dual Clebsch-Gordan transform assumes the representation of G on $\mathbb{C}^N$ is irreducible, implicitly identifies the dual irrep on $(\mathbb{C}^N)^*$ with a pre-defined irrep space V_{μ^*}, and decomposes $V_{\lambda^+} \otimes V_{\mu^*}$, which roughly forces a different choice of basis on the dual space. To align with prior work, we translate from our dual Clebsch-Gordan transform to the previously defined one at the level of matrix coefficients.

[19] Harrow implicitly forces a particular choice of basis for the multiplicity space, but algorithmically we are free to choose another basis.

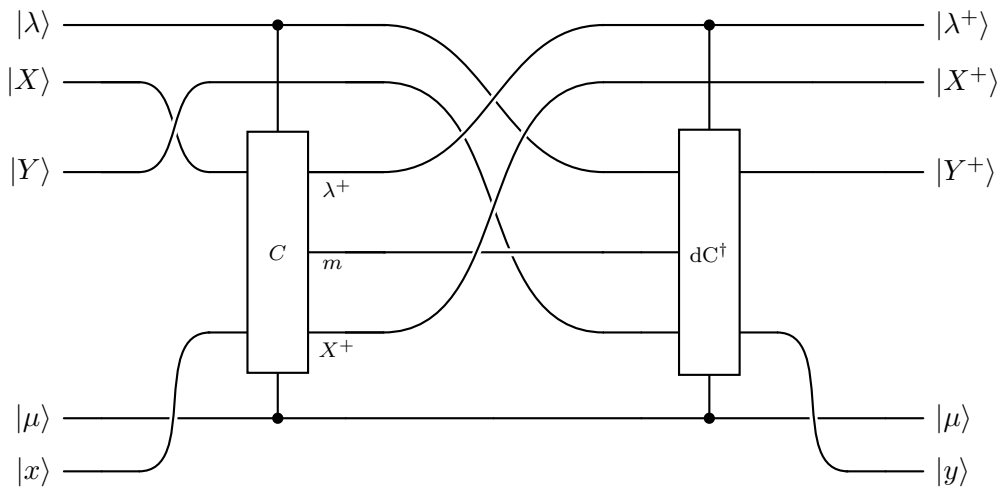


Figure 3: A single query can be decomposed into a Clebsch–Gordan transformation followed by the inverse dual-Clebsch-Gordan operation, though not on the same registers. The top three registers are those of the recording oracle's data structure, while the bottom two are the algorithm's registers. The algorithm's multiplicity register (in cases where it exists) is not shown here since it is not touched by the update. For unitaries, μ is the single-box fundamental irrep that acts as $\rho_{\square}(U) = U$.

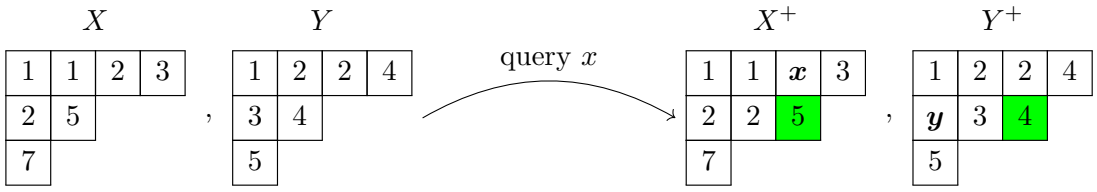


Figure 4: When updating the tableau recording, both tableaux grow in the same way, such that the resulting shapes will always be identical. For the unitary group, this (very roughly) means that we add a box in the same place (in superposition over valid locations) for both the X and Y tableaux.

In fact, rather than using direct implementations of a dual Clebsch-Gordan transform [Ngu23, GBO23], we observe a simple identity (Lemma 4.3) that efficiently reduces the $U(N)$ dual Clebsch-Gordan transform to the "primal" one, in the Gelfand-Tsetlin (or Young tableau) basis. This gives an efficient unitary compressed oracle in terms of only an efficient Clebsch-Gordan transform.

2.1.1 Other Query Types

What if we want to be able to query the complex conjugate $\rho(g)^*$, the inverse $\rho(g)^\dagger$, or the transpose $\rho(g)^T$ of the representation? And what if we want to be able to perform controlled queries?

Complex Conjugate Queries. We present two distinct ways to handle conjugate queries. We defer the second approach to the "transpose and inverse queries" discussion.

For the first approach, complex conjugate queries $\rho(g)^*$ are themselves another representation of the group, which, for unitary representations, is also known as the *dual representation* $\rho^*(g)$. So we can consider the same derivation, but for ρ^* instead of ρ. This has us switch the roles of the Clebsch-Gordan and dual Clebsch-Gordan transforms. Thus each $\rho^*(g)$ query can be implemented as in Figure 3, but with a dual Clebsch-Gordan followed by an inverse Clebsch-Gordan.

We also derive an alternative interpretation of this update for the special case $G = U(N)$: as referenced above, we show that a simple signed rotation (a Flip operation) of the Young tableaux is sufficient to implement a dual Clebsch-Gordan transform *in terms of* the primal Clebsch-Gordan transform. This means that (1) we can reduce conjugate queries to standard queries by simply flipping the Young tableaux $|X\rangle$ and $|Y\rangle$ before and after the query, and (2) even standard queries can be implemented relative to *just* a primal Clebsch-Gordan transform, without any explicit reference to a dual Clebsch-Gordan transform.

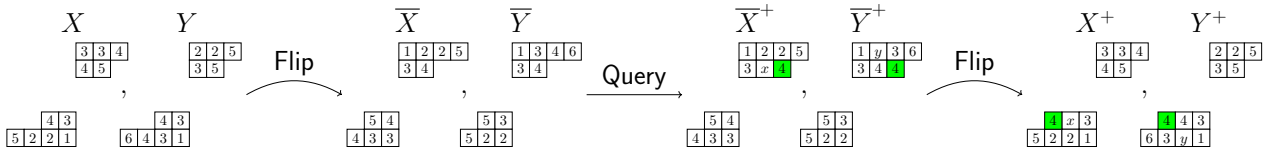


Figure 5: Example complex conjugate query for the unitary group.

Of course, we may want to make queries to both $\rho(g)$ and $\rho^*(g)$, or more generally to query both in superposition, as $|0\rangle\langle 0| \otimes \rho(g) + |1\rangle\langle 1| \otimes \rho^*(g)$. This is merely another representation of the group, and our framework above also works for this representation. The (generalized) Schur transform that results is closely related to the mixed Schur transform of [FTH23, Ngu23, GBO23] (more precisely, it is the direct sum of mixed Schur transforms for different sequences of $\rho(g)$ and $\rho^*(g)$). As before, each query is still a pair of Clebsch-Gordan transforms (the same Clebsch-Gordan transforms as for the standard query $\rho(g)$), depending on the control qubit.

Controlled Queries. Controlled queries can be handled similarly, taking the representation $\rho_{\mathsf{controlled}}(g) \coloneqq |0\rangle\langle 0| \otimes \mathbf{I} + |1\rangle\langle 1| \otimes \rho(g)$. Since this is also a representation of the group, we can similarly apply the framework to this representation. This will, as expected,[20] recover a tableau recording with an update that performs a normal $\rho(g)$ update on the $|1\rangle$ branch and does nothing on the $|0\rangle$ branch.

Transpose and Inverse Queries. Unlike complex conjugate queries, which are representations of the group, transpose and inverse queries are anti-representations,[21] and therefore do not fall directly into the framework above. However, we make use of the following simple idea: implementing a transpose query can be thought of as conjugating a *forward query* by the following (anti-linear) isometry on $L^2(G)$

$$\mathsf{Transpose} \cdot \int_G f(g)\,|g\rangle\,dg = \int_G f(g^{-1})^*\,|g\rangle\,dg.$$

In order to reason about this, we observe that Transpose implements the following simple involution on the Fourier basis (see Section 3.3.1):

$$\mathsf{Transpose} \cdot F^{-1}\,|\lambda, X, Y\rangle = F^{-1}\,|\lambda, Y, X\rangle\,.$$

Thus, performing a transpose query in the Tableau basis simply requires swapping the roles of $|X\rangle$ and $|Y\rangle$![22]

[20] Of course, since we already know what a controlled query *should* do, rerunning this framework for the controlled representation $\rho_{\mathsf{controlled}}(g)$ just to recover this behavior may be overkill. However, we say this to emphasize that we do not need any specialized ad hoc modifications for each new query type.

[21] Unlike representations, which have $\rho(g)\ \rho(h) = \rho(gh)$, anti-representations multiply in reverse order: $\rho^\dagger(g)\ \rho^\dagger(h) = \rho^\dagger(hg)$

[22] Since Transpose is anti-linear (or conjugate-linear) and is not a physically implementable operation, it turns out that rather than performing the update rule Q_ρ^{Tab} on the swapped registers, one needs to perform the complex

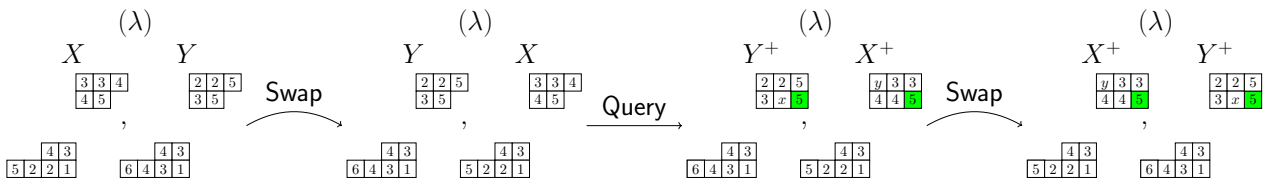


Figure 6: Example transpose query for the unitary group.

Similarly, we can perform an inverse query by implementing a conjugate query with $|X\rangle$ and $|Y\rangle$ reversed. In fact any anti-representation can be handled this way.

Finally, we remark that a similar idea can be used to derive a conjugate query in the Tableau basis, in terms of a (complex conjugated) standard query, using the anti-linear isometry

$$\mathsf{Conj} \cdot \int_G f(g) \, |g\rangle \, dg = \int_G f(g)^* \, |g\rangle \, dg.$$

We refer the reader to Section 4.2.2 for more details.

2.2 Path Recording

While the Tableau recording oracle is mathematically beautiful, satisfies perfect simulation, and can be algorithmically efficient, it seems quite difficult to prove operationally useful properties of this compressed oracle. In contrast, Zhandry's compressed oracle and the Ma-Huang path-recording oracle have easily interpretable "database states" on the recording register that can be used for proving quantum query lower bounds and pseudorandomness results. For example, in the Ma-Huang path-recording oracle,

- The recording register contains superpositions of "set states" of the form $|\{(x_1, y_1), \ldots, (x_t, y_t)\}\rangle$.
- To make a query on an input $|x\rangle$, and database state $|D\rangle$, one does the following:
  - Create an EPR pair in superposition over $y \notin \{y_1, \ldots, y_t\}$[23]
  - Swap one of the $|y\rangle$'s into the algorithm's register (replacing $|x\rangle$)
  - Add the pair (x, y) (using the second $|y\rangle$) to the database set D.

One might naturally ask: why are they stored as *sets* rather than ordered tuples? This is easy to understand in the case of **non-adaptive queries** applied to **distinct inputs**: if you query a (purified) Haar-random unitary on distinct inputs $|x_1, \ldots, x_t\rangle$, this results in the uncompressed

conjugate $\overline{Q_\rho^{\mathrm{Tab}}}$; but any circuit for Q_ρ^{Tab} can be generically converted into a circuit for $\overline{Q_\rho^{\mathrm{Tab}}}$; following [Zha25], we consider this to be a form of "black-box access" to the update. Moreover, in all the cases we consider, it is possible to write Q_ρ^{Tab} in a basis in which it is purely real, in which case $\overline{Q_\rho^{\mathrm{Tab}}} = Q_\rho^{\mathrm{Tab}}$ and this distinction becomes moot.

[23] We remark that while $y \notin \{y_1, \ldots, y_t\}$ is enforced to make this update procedure an isometry, this is *not* perfectly consistent with a Haar-random unitary (for example, if you query $U^{\otimes t} |x_1, x_2, \ldots, x_t\rangle$ you may observe the same output y multiple times). As we will show, this is one out of two sources of simulation error in the Ma-Huang oracle.

recording[24]

$$\int_U U\left|x_1\right\rangle \otimes \ldots \otimes U\left|x_t\right\rangle \otimes \left|U\right\rangle dU = \sum_{y_1,\ldots,y_t} \int_U \left\langle \vec{y}\right| U^{\otimes t}\left|\vec{x}\right\rangle \cdot \left|\vec{y}\right\rangle \otimes \left|U\right\rangle dU \tag{2.4}$$

$$= \sum_{y_1,\ldots,y_t} \left|\vec{y}\right\rangle \otimes \int_U \left\langle \vec{y}\right| U^{\otimes t}\left|\vec{x}\right\rangle \cdot \left|U\right\rangle dU \tag{2.5}$$

$$= \sum_{y_1,\ldots,y_t} \left|\vec{y}\right\rangle \otimes \int_U \prod_{i=1}^{t} \Big(\left\langle y_i\right| U\left|x_i\right\rangle \Big) \cdot \left|U\right\rangle dU\,, \tag{2.6}$$

and indeed the recording register has *forgotten* the ordering of $x_1, \ldots, x_t$ (but *remembers* the matching between x_i and y_i). Even more explicitly, let us try to compare the uncompressed $L^2(G)$ recording with a naïve path recording by defining the following map $\Theta_t : (\mathbb{C}^N)^{\otimes 2t} \to L^2(G)$.

$$\Theta_t \left|x_1, \ldots, x_t, y_1, \ldots, y_t\right\rangle = \int_U \left\langle \vec{y}\right| U^{\otimes t}\left|\vec{x}\right\rangle \cdot \left|U\right\rangle dU.$$

Then, if we think of a permutation $\pi \in S_t$ as acting on $\left|x_1, \ldots, x_t\right\rangle$ by permuting the registers

$$S(\pi)\left|x_1, \ldots, x_t\right\rangle = \left|x_{\pi^{-1}(1)}, \ldots, x_{\pi^{-1}(t)}\right\rangle,$$

we have that

$$\Theta_t \cdot \Big(S(\pi) \otimes S(\pi)\Big)\left|\vec{x}, \vec{y}\right\rangle = \int_U \left\langle \vec{y}\right| S(\pi)^\dagger U^{\otimes t} S(\pi)\left|\vec{x}\right\rangle \cdot \left|U\right\rangle dU \tag{2.7}$$

$$= \int_U \left\langle \vec{y}\right| U^{\otimes t}\left|\vec{x}\right\rangle \cdot \left|U\right\rangle dU \tag{2.8}$$

$$= \Theta_t \left|\vec{x}, \vec{y}\right\rangle. \tag{2.9}$$

This holds because every permutation $S(\pi)$ on the registers *commutes* with the tensor power operator $U^{\otimes t}$. In fact, let us write this identity slightly differently:

$$\Theta_t \cdot \Big(S(\pi) \otimes \mathbf{I}\Big)\left|\vec{x}, \vec{y}\right\rangle = \Theta_t \cdot \Big(\mathbf{I} \otimes S(\pi)^T\Big)\left|\vec{x}, \vec{y}\right\rangle, \tag{2.10}$$

which can be seen to hold by renaming $S(\pi)\left|\vec{y}\right\rangle \to \left|\vec{y}\right\rangle$. This means that the state $\left|\vec{x}, \vec{y}\right\rangle$ contains redundant information and should not be used as-is in a recording oracle.

However, if we instead used the state

$$\left|\{(x_1, y_1), \ldots, (x_t, y_t)\}\right\rangle := \frac{1}{\sqrt{t!}} \sum_{\pi \in S_t} \Big(S(\pi) \otimes S(\pi)\Big)\left|\vec{x}, \vec{y}\right\rangle,$$

then because we have symmetrized with respect to the symmetric group S_t, we have that

$$\left(S(\pi) \otimes S(\pi)\right)\left|\{(x_1, y_1), \ldots, (x_t, y_t)\}\right\rangle = \left|\{(x_1, y_1), \ldots, (x_t, y_t)\}\right\rangle,$$

so these states satisfy the same symmetries as the states $\Theta_t\left|x, y\right\rangle$. Indeed, the image of the projector[25]

$$\Omega_{S_t} = \frac{1}{t!} \sum_{\pi \in S_t} S(\pi) \otimes S(\pi), \tag{2.11}$$

[24] Here and throughout the paper, we use the notation $\vec{x} := (x_1, \ldots, x_t)$ to denote a tuple of entries.

[25] It is straightforward to check from the definition that this operator is idempotent (that is, that $(\Omega_{S_t})^2 = \Omega_{S_t}$) and Hermitian.

is precisely the set of states $|\psi\rangle$ satisfying the identity

$$(S(\pi) \otimes \mathbf{I}) |\psi\rangle = (\mathbf{I} \otimes S(\pi)^T) |\psi\rangle . \tag{2.12}$$

Although our argument so far has not proved it, it turns out that the image of Ω_{S_t} is *precisely* the subspace of valid recording states for *arbitrary* quantum query algorithms, even those that are adaptive and may query on non-distinct inputs.

What is Ω_{S_t}? Let us pause and try to better understand the Ω_{S_t} operator. The right tool for this job is **Schur-Weyl duality**. Schur-Weyl duality for the unitary representation $U^{\otimes t}$ on $(\mathbb{C}^N)^{\otimes t}$ amounts to employing the following facts:

1. The *only* operators on $(\mathbb{C}^N)^{\otimes t}$ that commute with every $U^{\otimes t}$ are linear combinations of the permutation elements $S(\pi)$. In other words, the $U(N)$-commutant algebra of $(\mathbb{C}^N)^{\otimes t}$ is the group algebra $\mathbb{C}[S_t]$.
2. Conversely, the *only* operators on $(\mathbb{C}^N)^{\otimes t}$ that commute with every $S(\pi)$ are linear combinations of the diagonal action operators $U^{\otimes t}$.

According to Schur-Weyl duality, this turns out to imply that the space $(\mathbb{C}^N)^{\otimes t}$ decomposes in the following way:

$$(\mathbb{C}^N)^{\otimes t} \simeq \bigoplus_{\lambda} V_\lambda^{U(N)} \otimes V_\lambda^{S_t}, \tag{2.13}$$

where $V_\lambda^{U(N)}$ is the space of an irreducible representation of $U(N)$, $V_\lambda^{S_t}$ is a *corresponding* irreducible representation space of S_t. The isometry implementing this mapping is the *Schur transform* of Equation (2.3). In this decomposition, $U^{\otimes t}$ acts as

$$U^{\otimes t} \,|_{V_\lambda^{U(N)} \otimes V_\lambda^{S_t}} = \rho_\lambda(U) \otimes \mathbf{I}.$$

That is, $U^{\otimes t}$ acts properly as the irrep transformation $\rho_\lambda(U)$ on $V_\lambda^{U(N)}$ and ignores the commutant register $V_\lambda^{S_t}$. On the other hand, the permutation elements $S(\pi)$ act on the other tensor factor:

$$S(\pi) \,|_{V_\lambda^{U(N)} \otimes V_\lambda^{S_t}} = \mathbf{I} \otimes \rho_\lambda(\pi).$$

With this decomposition, we can understand the projection Ω_{S_t}: since states $|\psi\rangle \in (\mathbb{C}^N)^{\otimes t} \otimes (\mathbb{C}^N)^{\otimes t}$ in the image of Ω_{S_t} satisfy Eq. (2.12) and the $S(\pi)$ span the commutant algebra, we conclude that $(A \otimes \mathbf{I}) |\psi\rangle = (\mathbf{I} \otimes A^T) |\psi\rangle$ for every $A \in \mathbb{C}[S_t]$. These operators A are precisely those that act only on the commutant irreps:

$$A = \bigoplus_{\lambda} \mathbf{I}_{V_\lambda^{U(N)}} \otimes A_\lambda, \qquad A_\lambda \in \mathrm{End}(V_\lambda^{S_t})$$

Thus, for any state $|\psi\rangle \in (\mathbb{C}^N)^{\otimes t} \otimes (\mathbb{C}^N)^{\otimes t}$, which we can generically write (in the Schur-transformed basis) as $|\psi\rangle = \sum_{\lambda,\lambda'} \alpha_{\lambda,\lambda'} |\lambda, \lambda'\rangle \otimes |\psi_{\lambda,\lambda'}\rangle$, for $|\psi_{\lambda,\lambda'}\rangle \in V_\lambda^{U(N)} \otimes V_{\lambda'}^{U(N)} \otimes V_\lambda^{S_t} \otimes V_{\lambda'}^{S_t}$, we have that

$$\begin{aligned} (A \otimes \mathbf{I}) |\psi\rangle &= \sum_{\lambda,\lambda'} \alpha_{\lambda,\lambda'} |\lambda, \lambda'\rangle \otimes (\mathbf{I}_{V_\lambda^{U(N)}} \otimes \mathbf{I}_{V_{\lambda'}^{U(N)}} \otimes A_\lambda \otimes \mathbf{I}_{V_{\lambda'}^{S_t}}) |\psi_{\lambda,\lambda'}\rangle \\ = (\mathbf{I} \otimes A^T) |\psi\rangle &= \sum_{\lambda,\lambda'} \alpha_{\lambda,\lambda'} |\lambda, \lambda'\rangle \otimes (\mathbf{I}_{V_\lambda^{U(N)}} \otimes \mathbf{I}_{V_{\lambda'}^{U(N)}} \otimes \mathbf{I}_{V_\lambda^{S_t}} \otimes A_{\lambda'}^T) |\psi_{\lambda,\lambda'}\rangle \end{aligned}$$

We can see, therefore, that any state satisfying Eq. (2.12) must be such that:

- only the blocks $\lambda = \lambda'$ have nonzero weight (consider $A = \Pi_\lambda$ that projects onto a fixed λ-block), and
- each component $|\psi_\lambda\rangle$ is an EPR state on the commutant registers $V_\lambda^{S_t} \otimes V_\lambda^{S_t}$ in tensor product with an arbitrary state on the $U(N)$ registers (because the EPR state is the unique state, up to overall phase, satisfying $(\mathbf{I} \otimes A_\lambda^T)\,|\psi_\lambda\rangle = (A_\lambda \otimes \mathbf{I})\,|\psi_\lambda\rangle$ for all linear maps A_λ).

Thus, we refer to Ω_{S_t} as the **commutant EPR projector**.

Our Path-Recording Oracle. With this thought process in mind, we are ready to introduce our exact path-recording compressed oracle. As before, we assume that our goal is to simulate query-access to a single unitary representation of G

$$Q_\rho(|x\rangle \otimes \int_G f(g)\,|g\rangle\,dg) = \int_G \rho(g)\,|x\rangle \otimes f(g)\,|g\rangle\,dg.$$

By our earlier argument, we know that by applying the Fourier transform F on the $L^2(G)$ register, after t queries the recording register will be supported on

$$\bigoplus_{\lambda \in \widehat{G}_t} V_\lambda \otimes V_\lambda^*,$$

where $\widehat{G}_t$ denotes exactly the set of irreps that appear in the Schur-Weyl decomposition of $(\mathbb{C}^N)^{\otimes t}$:

$$(\mathbb{C}^N)^{\otimes t} \simeq \bigoplus_{\lambda \in \widehat{G}_t} V_\lambda^G \otimes V_\lambda^{\mathcal{A}_t}.$$

Compared to the case of $G = U(N)$ (and ρ the defining representation) discussed earlier, this Schur-Weyl decomposition is defined with respect to the commutant algebra

$$\mathcal{A}_t = \{A \in \mathrm{End}((\mathbb{C}^N)^{\otimes t}) : A \cdot \rho(g)^{\otimes t} = \rho(g)^{\otimes t} \cdot A \text{ for all } g \in G\}.$$

For the defining representation of $U(N)$, we have that $\mathcal{A}_t = \mathbb{C}[S_t]$, while for other group representations, the commutant will be another semi-simple diagram algebra (see Table 2 for more examples). For a general representation (just as in the special case before), one can observe that uncompressed recording states obey symmetries imposed by $\mathcal{A}_t$:

$$\Theta_t \cdot (A \otimes \mathbf{I}) = \Theta_t \cdot (\mathbf{I} \otimes A^T) \text{ for all } A \in \mathcal{A}_t.$$

Therefore, we should aim for a path-recording oracle with database states in the subspace in which these two commutant actions have been identified, which is exactly the commutant EPR subspace that is the image of $\Omega_{\mathcal{A}_t}$ in the Schur basis.

We now derive a path-recording oracle from the tableau-recording oracle. To begin, we will define an isometry

$$\Gamma_t : \bigoplus_{\lambda \in \widehat{G}_t} V_\lambda^G \otimes (V_\lambda^G)^* \to (\mathbb{C}^N)^{\otimes t} \otimes ((\mathbb{C}^N)^*)^{\otimes t}$$

that maps from the tableau basis to the path basis. Let $|\lambda, X, Y\rangle$ be a basis vector in the tableau basis, where $\lambda \in \widehat{G}_t$. The isometry Γ_t acts as follows:

1. It first creates an EPR state

$$|\mathrm{EPR}_\lambda\rangle = \frac{1}{\sqrt{\dim(V_\lambda^{\mathcal{A}_t})}} \sum_{S \in \mathcal{B}(V_\lambda^{\mathcal{A}_t})} |S\rangle\,|S^*\rangle\,,$$

   where $\mathcal{B}(V_\lambda^{\mathcal{A}_t})$ is a basis for $V_\lambda^{\mathcal{A}_t}$, and we interpret $|S^*\rangle \in (V_\lambda^{\mathcal{A}_t})^*$ as a dual vector, and appends it to the tableau basis vector. This produces the state $|\lambda, X, Y\rangle\,|\mathrm{EPR}_\lambda\rangle$, which is an element of the space $V_\lambda^G \otimes (V_\lambda^G)^* \otimes V_\lambda^{\mathcal{A}_t} \otimes (V_\lambda^{\mathcal{A}_t})^*$.

2. Next, it applies the inverse Schur transform $\mathsf{Schur}^\dagger$ to $V_\lambda^G \otimes V_\lambda^{\mathcal{A}_t}$ and the inverse dual Schur transform Schur^T to $(V_\lambda^G \otimes V_\lambda^{\mathcal{A}_t})^*$. The resulting state is an element of $(\mathbb{C}^N)^{\otimes t} \otimes ((\mathbb{C}^N)^*)^{\otimes t}$, and can therefore be written as a superposition over t-tuples of pairs $|(x_1, y_1), \ldots, (x_t, y_t)\rangle$.

Having defined Γ_t, it is straightforward to state the path-recording oracle as[26]

$$Q_\rho^{\mathrm{Path}} \coloneqq (\mathbf{I} \otimes \Gamma_{t+1})\, Q_\rho^{\mathrm{Tab}}\, (\mathbf{I} \otimes \Gamma_t^\dagger) = (\mathbf{I} \otimes (\Gamma_{t+1} \cdot F))\, Q_\rho\, (\mathbf{I} \otimes (F^\dagger \cdot \Gamma_t^\dagger))\,.$$

In other words, if the purification register is in the path basis, Q_ρ^{Path} simply applies $\Gamma_t^\dagger$ to map it to the tableau basis. It then applies the tableau-basis oracle, and then applies Γ_{t+1} to convert the purification register back to the path basis.

Interpreting Q_ρ^{Path}. Of course, the main reason for defining the path-recording oracle is to be able to interpret the recording states and the effect of an oracle query. To do this, we prove that Q_ρ^{Path} has the following form

Theorem 2.1 (See Definition 5.1 and Theorem 5.3)**.** *The path-recording oracle Q_ρ^{Path} can be expressed as*

$$Q_{\rho,t+1}^{\mathrm{Path}} = \sum_{x,y \in [N]} |y\rangle\langle x| \otimes \Lambda_{t+1} \cdot \Omega_{\mathcal{A}_{t+1}} \cdot \mathrm{App}_{x,y}^{(t)} \cdot \Lambda_t^+ \cdot \Omega_{\mathcal{A}_t},$$

where:

- *$\mathcal{A}_t$ denotes the G-commutant algebra of $(\mathbb{C}^N)^{\otimes t}$.*
- *$\Omega_{\mathcal{A}_t}$ denotes the (controlled-λ) projection onto $|\mathrm{EPR}_\lambda\rangle$ in the Schur basis,*
- *$\mathrm{App}_{x,y}^{(t)}$ denotes the map appending $|(x,y)\rangle$ to the database state.*[27]
- *$\Lambda_t = \bigoplus_{\lambda \in \widehat{G}_t} \sqrt{\frac{\dim(V_\lambda^{\mathcal{A}_t})}{\dim(V_\lambda^G)}} \cdot (\Pi_\lambda \otimes \mathbf{I})$, where Π_λ denotes projecting onto the λ-component in the Schur basis.*[28]
- *$\Lambda_t^+ = \bigoplus_{\lambda \in \widehat{G}_t} \sqrt{\frac{\dim(V_\lambda^G)}{\dim(V_\lambda^{\mathcal{A}_t})}} \cdot (\Pi_\lambda \otimes \mathbf{I})$ is (pseudo-)inverse to Λ_t*

[26] Here, again, F is the Fourier transform on $L^2(G)$ as before (Equation (2.1)).

[27] Note the database here contains an *ordered* list of (x, y) pairs, and this operator merely appends to that list. As we will see, the operator $\Omega_{\mathcal{A}_t}$ resymmetrizes the list, removing this ordering (and anything else that is not to be recorded).

[28] Since Ω_{A_t} has already projected onto equal λ-labels on the two Schur bases, it is equivalent here to project onto the λ block on either half or both.

In other words, assuming that the recording register was already in the image of $\Omega_{\mathcal{A}_t}$ (which will be true if it was obtained by running an algorithm that queries Q_ρ^{Path}), the update is the following sequence of operations.

- Reweight subspaces in the Schur basis.
- Initialize an EPR state $\sum_{y\in[N]} |y\rangle\,|y\rangle$,
- Append (x, y) to the database using the query input and half of the EPR state. The other half of the EPR state will be returned as the response.
- Re-impose the symmetry $\Omega_{\mathcal{A}_{t+1}}$ on the recording register.[29]
- Reweight the (new) subspaces in the Schur basis again.

We refer the reader to Section 5 for the formal proof of Theorem 2.1. The high-level idea is to relate the isometry $\Gamma_t \cdot F$ with the map Θ_t from earlier (which aimed to analogize the path state $|x, y\rangle$ with the recording state $\int_G \langle y|\,\rho(g)^{\otimes t}\,|x\rangle \cdot |g\rangle\,dg$) by proving that

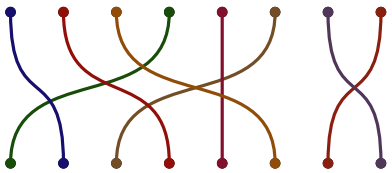
$$\Gamma_t \cdot F \cdot \Theta_t = \Lambda_t \cdot \Omega_{\mathcal{A}_t}\,.$$

At this point, there are two remaining questions to answer regarding interpretation of this formula for Q_ρ^{Path} for general groups:

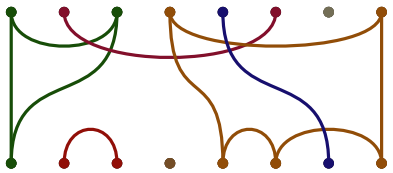
1. The commutant EPR projector $\Omega_{\mathcal{A}_t}$ may appear somewhat mysterious in the path basis outside of the case of $G = U(N)$ and $\mathcal{A}_t = \mathbb{C}[S_t]$ in the earlier example.
2. The subspace reweighting operators Λ_t are applied in the Schur basis, which has no simple/obvious interpretation in the path basis.

For the first question, we show that $\Omega_{\mathcal{A}_t}$ has a clear general-purpose form directly in the path basis, generalizing the formula $\Omega_{S_t} = \frac{1}{t!}\sum_{\pi\in S_t} S(\pi)\otimes S(\pi)$ for the case $G = U(N)$. This allows us to think of Ω_{S_t} as an operation in the standard ("path") basis rather than one the Schur basis.

This is accomplished as follows: let $\mathcal{B}(\mathcal{A}_t)$ denote a basis for the commutant algebra. For example, this basis can consist of permutations $\pi \in S_t$ for the case $G = U(N)$ and $\mathcal{A} = \mathbb{C}[S_t]$. For the case $G = S_N$, the commutant algebra $\mathcal{A}_t = P_t(N)$ is the *partition algebra* [HR05], which has a basis consisting of *set-partition diagrams* (see Figure 7b).



(a) An example basis element of the symmetric group algebra, $\mathbb{C}[S_t]$, for $t = 8$.



(b) An example basis element of the partition algebra, $P_t(N)$, for $t = 8$.

Figure 7: The symmetric group algebra, $\mathbb{C}[S_t]$, the partition algebra, $P_t(N)$, and other diagram algebras can be described in terms of a basis of partition diagrams that partition t top vertices and t bottom vertices into components. See Section 3.5 for more details.

[29] We remark that the recording register was previously in the image of $\Omega_{\mathcal{A}_t}$, so intuitively $\Omega_{\mathcal{A}_{t+1}}$ is only imposing additional symmetries involving the last register.

Using a variant of the Schur orthogonality relations we show (see Lemma 5.11) that in general,

$$\Omega_{\mathcal{A}_t} = \sum_{a \in \mathcal{B}(\mathcal{A}_t)} S(a) \otimes S(a^*)^T,$$

where a^* denotes a *dual basis element* with respect to the trace pairing $\mathrm{tr}_{\mathcal{A}_t}(ab)$.[30] As an example, for $G = U(N)$, the elements of $\mathcal{B}(\mathbb{C}[S_t])$ are orthogonal under this trace pairing, except that $\mathrm{tr}\big([\pi] \cdot [\pi^{-1}]\big) = t!$. We thus have that $[\pi]^* = \frac{1}{t!}[\pi^{-1}]$, recovering the form of Ω_{S_t} in Equation (2.11). Of course, this operator $\Omega_{\mathcal{A}_t}$ depends on the group G and its representation, but understanding it simply amounts to understanding the commutant algebra $\mathcal{A}_t$ (in other words, the algebraic symmetries of $\rho^{\otimes t}$) and does not require the Schur basis. For example, we study this operator closely for the case $G = S_N$ in Section 7, as well as some simpler cases in Section 6.

Thus, what remains in order to understand the exact path-recording oracle is to understand the reweighting operator $\Lambda_{\mathcal{A}_t}$. As we discuss next, if one is willing to make *approximations*, this challenge can be surmounted as well.

[MH25] as an approximation of our oracle. How does our path-recording oracle (for the case $G = U(N)$) compare with the Ma-Huang oracle? Here is a direct comparison:

- The [MH25] oracle and our oracle both begin by creating an EPR state over outputs $|y\rangle$; however, [MH25] restricts y to not collide (in the standard basis) with the y values in the existing recording state.
- The [MH25] oracle then adds the pair $|(x, y)\rangle$ to an unordered *set*, while our path-recording oracle first appends $|(x, y)\rangle$ to an *ordered list* or *t-tuple* and then symmetrizes with respect to S_t (permuting the pairs into all possible positions in the list). On states with distinct y_i, these operations are the same up to a multiplicative scalar.
- Finally, our oracle contains two subspace reweightings (one at the beginning and one at the end) that the [MH25] oracle does not have.

Thus, the fact that the [MH25] oracle approximates ours (for $G = U(N)$) follows from two facts that we prove in Section 6:

1. For any t-query adversary $\mathcal{A}^{Q_\rho^{\mathrm{Path}}}$, the probability that the y_i in the recording register are all distinct is at least $1 - O(t^2/N)$.
2. The rescaling operators Λ_t are *scalar multiples of the identity*, up to multiplicative error $(1 \pm O(t^2/N))$.

Conversely, this gives a method of deriving the Ma-Huang approximate path-recording oracle from first principles:

1. Theorem 2.1 describes a general exact path-recording oracle.
2. For a simpler form with easier interpretability, approximate the Λ_{S_t} reweightings with scalar multiples of the identity, resulting in an operator $\widetilde{Q}_\rho^{\mathrm{Path}}$.

[30] $\mathrm{tr}_{\mathcal{A}_t}(ab)$ is defined to be the trace of the "multiplication by ab map" on $\mathcal{A}_t$. The dual basis element a^* corresponding to a basis element $a \in \mathcal{B}(\mathcal{A}_t)$ is the unique element of $\mathcal{A}_t$ such that $\mathrm{tr}_{\mathcal{A}_t}(a^*\, b) = \delta_{ab}$ where b runs over the basis of the algebra, $\mathcal{B}(\mathcal{A}_t)$.

3. Without the subspace reweighting, $\widetilde{Q}_\rho^{\text{Path}}$ is not an isometry; however, it *is* an isometry when inserted y values are restricted to be distinct. This gives the interpretable—and, upon inspection, efficiently implementable—approximate path-recording oracle

$$Q_\rho^{MH} = \frac{1}{\sqrt{N-t}} \sum_{x,y\in[N]} |y\rangle\langle x| \otimes \sqrt{t+1} \cdot \Omega_{\mathcal{A}_{t+1}} \cdot \Pi_{\text{Dist}}^Y \cdot \text{App}_{x,y}^{(t)},$$

which is precisely the path-recording oracle of [MH25].

See Section 6, for additional examples of how to derive path recording oracles for other groups in this way.

2.3 Pseudorandomness of the PC Ensemble

Finally, we describe our approach for proving that our PC ensemble is a pseudorandom unitary (Theorem 1.6). To gain some intuition, we can start by asking: what is the difference between querying a random permutation unitary[31] P and querying a Haar-random unitary U? There are two simple, detectable differences:

- If you query any permutation unitary P on a uniform superposition[32] $|+\rangle \in \mathbb{C}^N$, the result is $|+\rangle$. However, for Haar-random unitary U, the state $U|+\rangle$ is a Haar-random state in $\mathbb{C}^N$.
- If you query any permutation $P^{\otimes 2}$ on two copies the same computational basis state $|x\rangle|x\rangle$, the resulting state $|P(x)\rangle|P(x)\rangle$ has the same computational basis state $y = P(x)$ appearing on both registers. However, for Haar-random U, the state $U^{\otimes 2}|x\rangle|x\rangle$ is two copies of a Haar-random state, which will almost never agree in the computational basis.

However, conveniently, prepending P with a random Clifford circuit (or any unitary 2-design) C thwarts both of these attacks:

- If you query $P \cdot C$ on an arbitrary pure state $|\psi\rangle$ (independent of P, C), $C|\psi\rangle$ has almost no overlap with $|+\rangle$.
- If you query $(P \cdot C)^{\otimes 2}$ on an arbitrary state $|\psi\rangle \in (\mathbb{C}^N)^{\otimes 2}$, the state $C^{\otimes 2}|\psi\rangle$ with high probability has no collisions in the computational basis. That is, it is almost completely contained within the *distinct subspace* $\Pi_{\text{Dist}} = \text{Span}\{|x_1, x_2\rangle : x_1 \neq x_2\}$.

Of course, just because these two attacks are thwarted does not necessarily mean that *all* possible attacks are thwarted. Perhaps a more clever attack, which makes queries on something more sophisticated, could distinguish a random permutation from a random unitary?

We show that this is not the case: At a high level, we use our path-recording formalism, specialized to both groups S_N and $U(N)$, to prove that these are in fact the *only obstructions* to a random permutation P looking like a random unitary, and so multiplying P by C does in fact result in a pseudorandom unitary. More formally, for any t, we define the "distinct, nonplussed" subspace $\mathsf{DNP} \subset (\mathbb{C}^N)^{\otimes t}$ to be the intersection of the distinct subspace with the "nonplussed subspace" $\mathsf{NoPlus} = \text{Im}((\mathbf{I} - |+\rangle\langle+|)^{\otimes t})$. We must argue both that it suffices to consider only the distinct

[31] The permutation here is an element of $G = S_N$ (not to be confused with the commutant algebra $\mathbb{C}[S_t]$ that we considered earlier for the case where $G = U(N)$). For each such permutation, it acts on $\mathbb{C}^N$ by permuting the N basis vectors in the standard basis.

[32] We will always use the plus state to denote the uniform positive superposition in the computational basis, $|+\rangle \coloneqq \frac{1}{\sqrt{N}} \sum_{z\in[N]} |z\rangle$.

nonplussed subspace, and that permutations and unitaries are statistically indistinguishable on this subspace. We sketch the main points of the argument below (and see Section 7 for the proof in detail).

We know from Theorem 1.5 that queries to a Haar random unitary $U \leftarrow U(N)$ are perfectly indistinguishable from queries the the path-recording oracle $Q^{\mathrm{Path}}_{U(N)}$, and similarly, that queries to a random permutation $P \leftarrow S_N$ are perfectly indistinguishable from those to $Q^{\mathrm{Path}}_{S_N}$. To prove that $P \cdot C$ is indistinguishable from a Haar random unitary, we can thus take any t-query adversary Adv_t which may attempt to distinguish them, and *directly compare* the path-recording purifications[33]

$$\left|\mathsf{Adv}_t^{Q^{\mathrm{Path}}_{S_N}\cdot C}\right\rangle \qquad \text{and} \qquad \left|\mathsf{Adv}_t^{Q^{\mathrm{Path}}_{U(N)}\cdot C}\right\rangle . \tag{2.14}$$

Crucially, the two path-recording oracles are acting on the *same* Hilbert space $((\mathbb{C}^N)^{\otimes t})^{\otimes 2}$ in their recording registers, which is what allows us to make the direct comparison. We proceed in three steps:

1. **Projecting onto the Distinct Nonplussed Subspace.** First, we prove that for any possible adversary making t queries to either of the two path-recording oracles, the recording register will only ever have at most $O(t^2/N)$ mass outside the distinct nonplussed subspace $\mathsf{DNP}^{\otimes 2} \subset ((\mathbb{C}^N)^{\otimes t})^{\otimes 2}$.

   The distinct nonplussed subspace DNP, as the intersection of two simpler subspaces (the distinct subspace, Dist, and the nonplussed subspace, NoPlus) is somewhat tricky to reason about on its own. But by bounding the *Friedrichs angle* between the two subspaces, we show that the projection $\Pi^{X,Y}_{\mathsf{DNP}}$ onto the distinct nonplussed subspace can be approximated by the product $\Pi^{X,Y}_{\mathsf{Dist}} \cdot \Pi^{X,Y}_{\mathsf{NoPlus}}$ of projections onto the two larger subspaces.

   We then show a *ricochet property* of the path-recording oracles. This says, roughly, that any unitary C that is applied before each path-recording oracle query can instead be thought of as if it were applied in parallel as $C^{\otimes t}$ to the X part of the recording register. We can then use the 2-design property of random Cliffords to argue that the X part of the recording register will, with high probability, be in the distinct subspace, and separately in the nonplussed subspace. A direct calculation also shows the same for the Y part of the recording register. We then argue by a quantum union bound that the recording register must, with similarly high probability, be in the subspaces simultaneously.

   Thus, we may apply the projection $\Pi^{X,Y}_{\mathsf{DNP}}$ onto the distinct nonplussed subspace, while incurring error at most $O(t^2/N)$.

2. **The Symmetrization Operators $\Omega_{\mathcal{A}_t}$ on the Distinct Nonplussed Subspace.** Second, we prove that the commutant EPR projectors Ω_{S_t} and $\Omega_{P_t(N)}$ are in fact *equal* on the subspace $\mathsf{DNP}^{\otimes 2}$, meaning

   $$\Omega_{S_t}\, \Pi^{X,Y}_{\mathsf{DNP}} = \Omega_{P_t(N)}\, \Pi^{X,Y}_{\mathsf{DNP}} .$$

   In our opinion, this is the most important step of the analysis. Consider the operators Ω_{S_t} and $\Omega_{P_t(N)}$ written in the diagram bases of the two commutant algebras (see Figure 7):

   $$\Omega_{S_t} = \frac{1}{t!}\sum_{a\in S_t} S(a)\otimes S(a) \qquad\qquad \Omega_{P_t(N)} = \sum_{a\in\mathcal{B}(P_t(N))} S(a)\otimes S(a^*)^T \tag{2.15}$$

[33] The (right-) invariance of the Haar measure of $U(N)$ says that for a unitary U drawn from the Haar measure and any other unitary C, $U \cdot C$ is identically distributed to U. Thus queries to $Q^{\mathrm{Path}}_{U(N)} \cdot C$ are perfectly indistinguishable from queries to $U \cdot C$ (by Theorem 1.5), and consequently from those to the Haar random unitary U.

Since the partition diagrams of $P_t(N)$ can be arbitrary partitions of the $2t$ vertices (t top/output vertices and t bottom/input vertices), they range over a much larger set than the permutation diagrams of $\mathbb{C}[S_t]$ (which require a perfect matching between the t input vertices and the t output vertices in order to be a permutation in S_t). Thus $\Omega_{P_t(N)}$ ranges over a much larger basis than does Ω_{S_t}. How could they be equal?

But observe something important about how these set-partition diagrams act on $(\mathbb{C}^N)^{\otimes t}$: Suppose a diagram $a \in \mathcal{B}(P_t(N))$ has a component with at least two vertices on the bottom row (say, vertices numbered i and j when counting from the left). Then its representation $S(a)$ on the space $(\mathbb{C}^N)^{\otimes t}$ has the form $\sum_{z \in [N], \ldots} |\ldots\rangle\langle\ldots, z, \ldots, z, \ldots|$ with the two z's appearing (at least) on the i'th and j'th positions. But such an operator acts as exactly 0 unless the state it acts on has a collision (in the computational basis) on those two registers. So we know that on the distinct subspace, such diagrams must drop out of the sum on the right of Equation (2.15). Now suppose instead that a diagram $a \in \mathcal{B}(P_t(N))$ has some component consisting of only a single vertex on the bottom row (say, at vertex numbered i from the left). Then its representation $S(a)$ on $(\mathbb{C}^N)^{\otimes t}$ has the form $\sum_{z \in [N], \ldots} |\ldots\rangle\langle\ldots, z, \ldots|$ with the z appearing only at position i in the bra of the operator. Such an operator acts as exactly 0 unless the state it acts on has a plus state $|+\rangle$ in position i. So we know that on the *nonplussed subspace*, such diagrams must drop out of the right-hand sum of Equation (2.15). On the *distinct nonplussed subspace*, *both* kinds of set-partition diagrams are annihilated. It is not hard at this point to convince yourself that the only diagrams remaining are those in which each component has exactly one vertex on the bottom row, and one vertex on the top row of the diagram, in other words, *permutation diagrams*.

We are not quite done yet, however, since we are now left with

$$\sum_{\substack{a \in \mathcal{B}(P_t(N)) \\ a \text{ is a permutation diagram}}} S(a) \otimes S(a^{*_{P_t(N)}})^T$$

where the dual relationship $a^{*_{P_t(N)}}$ is still with respect to the trace pairing of the partition algebra $P_t(N)$ (and not of the symmetric group algebra $\mathbb{C}[S_t]$). Nevertheless, we prove that for every permutation $\pi \in S_t$, the dual element $[\pi]^{*_{P_t(N)}}$ is a linear combination of $\frac{1}{t!}[\pi^{-1}]$ (matching the S_t case) and *non-permutation* elements of $\mathcal{B}(P_t(N))$ (which, as we have argued, act as zero on the distinct nonplussed subspace DNP).

Thus, after our restriction to $\mathsf{DNP}^{\otimes 2}$, the path-recording updates are *identical* up to the subspace reweighting operators.

3. **The Reweighting Operators $\Lambda_{\mathcal{A}_t}$ on the Distinct Nonplussed Subspace.** At this point, we already know that Λ_{S_t} is approximately a constant. We additionally show that $\Lambda_{P_t(N)}$ is approximately constant (indeed, the same constant) on the subspace $\mathsf{DNP}^{\otimes 2}$.

   This conclusion comes from viewing the distinct nonplussed subspace from the standpoint of the representation theory of S_N: it turns out that in the Schur basis of S_N and $P_t(N)$, the projector Π_{DNP} projects $\bigoplus_\lambda V_\lambda^{S_N} \otimes V_\lambda^{P_t(N)}$ precisely onto the "full-box" irrep labels λ. In more detail, each irrep label $\lambda \in \widehat{P_t}(N)$ corresponds to a Young diagram with at most t boxes below the top row.[34] We show that the distinct nonplussed projector Π_{DNP} projects onto the irreps λ with *exactly* t boxes below the top row. Moreover, we can compute that the

[34] by the common convention of indexing irreps of S_N with Young diagrams of exactly N boxes. Often, we ignore the top row of boxes, and instead count only the boxes that are moved to the second row or further.

dimensions of such irreps are very close (up to a multiplicative $(1 \pm O(t^2/N))$) between S_N and U_N and between $P_t(N)$ and $\mathbb{C}[S_t]$, showing that Λ_{S_t} and $\Lambda_{P_t(N)}$ are approximately the same, up to error $O(t^2/N)$.

We can thus conclude that the two purified states in Equation (2.14), when querying the path-recording oracles for either a random unitary or a random permutation (preceded with a random Clifford), have distance at most $O(t^2/N)$, which is negligible when $t = \text{poly}(n)$ and $N = 2^n$. Thus the t-query adversary cannot distinguish $P \cdot C$ from a Haar random unitary with anything better than negligible advantage. We finally conclude by replacing the random permutation P with either a pseudorandom permutation to get a pseudorandom unitary or a $2t$-wise independent permutation to get a unitary t-design.

3 Preliminaries

Throughout this paper, we consider the following setting. Let G be a compact Lie group. Without loss of generality, G can be thought of as a closed subgroup of the group of unitary matrices on a finite-dimensional Hilbert space. Compact Lie groups admit a unique normalized Haar measure that allows us to define uniform averages over G. We use $\int_G dg$ to denote integration under the Haar measure on G. This setting captures the following special cases of interest to quantum computation:

- $G = U(N)$ is the group of $N \times N$ unitary matrices.
- $G = O(N)$ is the group of $N \times N$ real orthogonal matrices.
- $G = (U(1))^N$ is the group of $N \times N$ diagonal unitary matrices.
- Any finite group G, in which case $\int_G f(g)dg = \frac{1}{|G|}\sum_g f(g)$. This includes the permutation group S_N, which can be embedded as $N \times N$ permutation matrices, as well as the abelian group $\{\pm 1\}^N$ embedded as diagonal sign matrices.

We will consider a finite-dimensional representation $\rho : G \mapsto U(N)$ of the group G. Our goal is to perform efficient on-the-fly simulation of a quantum algorithm A with "oracle access to $\rho(g)$ for a random g", where g is sampled (once) from the Haar measure and used for all queries. By efficient on-the-fly simulation, we mean a stateful simulation such that the t-th query that A makes can be simulated by a $\text{poly}(t, \log(N))$-time algorithm operating on a $\text{poly}(t, \log(N))$-qubit Hilbert space.

3.1 Notation for compact groups

Definition 3.1 (Haar measure)**.** Let G be a compact Lie group. We say that a (regular Borel) measure μ is left-invariant if for any measureable subset of the group, $S \subseteq G$, and for every $g \in G$, $\mu(gS) = \mu(S)$. Similarly, we say that μ is right-invariant if for all $S \subseteq G$, and for every $g \in G$, $\mu(Sg) = \mu(S)$. The *Haar measure* on G is the unique such measure (up to a rescaling) which is both left-invariant and right-invariant.

Definition 3.2 (Group element kets)**.** Let G be a compact Lie group with normalized Haar measure dg. We write $\{|g\rangle : g \in G\}$ for kets labeled by group elements, with $\langle g|g'\rangle = \delta_{g,g'}$ (so $\langle g|g\rangle = 1$). To make this notation consistent between the finite and infinite cases, we write dg to denote a "half-measure". Then for any (square-integrable) function $f : G \mapsto \mathbb{C}$, we define

$$|f\rangle = \int_G f(g)\,|g\rangle\ dg$$

with the convention that

$$\langle f_1 | f_2 \rangle = \int_G \overline{f_1(g)} f_2(g)\, dg.$$

Thus, $\langle f | f \rangle = 1$ when $\int_G |f(g)|^2\, dg = 1$. For finite G, this means that $|f\rangle = \frac{1}{|G|} \sum_{g \in G} f(g)\, |g\rangle$ and we identify unit-norm basis vectors $|g\rangle = \sqrt{|G|} \cdot |f_g\rangle$, where $f_g(g') = \delta_{g,g'}$.

Definition 3.3 (L^2 function space)**.** $L^2(G)$ is defined to be the space of square-integrable functions $f : G \to \mathbb{C}$. Each function is represented by $|f\rangle$ as in Definition 3.2.

Definition 3.4 (Regular representation)**.** The *left regular representation*

$$L : G \to U(L^2(G))$$

is defined by

$$(L_g f)(x) = f(g^{-1}x), \qquad f \in L^2(G),\ g, x \in G.$$

This gives $L^2(G)$ the structure of a unitary representation of G. We note that unlike all other representations considered in this work, the regular representation may be infinite-dimensional.

3.2 Purified Oracle Algorithms (or Adversaries)

We formalize the notion of an algorithm, or adversary, with oracle access to a unitary $\mathcal{O}_{\mathsf{AR}}$ acting on the algorithm's query register A along with a purification register R that is hidden from the algorithm. This formalism enables all of our compressed oracle definitions and will be used explicitly in Sections 5.9, 7.4 and 7.5.

Definition 3.5 (Purified Oracle Adversaries)**.** A t-query adversary, denoted $\mathsf{Adv}_t^{\mathcal{O}}$, is specified by a list of $n + m$ qubit unitaries $(A_1, \ldots, A_t)$. A_i acts on a n-qubit *query register* A, along with a m-qubit *workspace register* B. The purified oracle $\mathcal{O}$ acts on both the adversary's query register, as well as a *database register* R, which is inaccessible to the adversary. After t queries to $\mathcal{O}$, the state of the adversary is[35]

$$\mathrm{tr}_{\mathsf{R}}\left[\mathsf{Adv}_t^{\mathcal{O}}\left(|0\rangle_{\mathsf{AB}} \otimes |\emptyset\rangle_{\mathsf{R}}\right)\right] = \mathrm{tr}_{\mathsf{R}}\left[\left(\prod_{i=1}^{t} \mathcal{O}_{\mathsf{AR}} \cdot A_{i\,\mathsf{AB}}\right)\left(|0\rangle_{\mathsf{AB}} \otimes |\emptyset\rangle_{\mathsf{R}}\right)\right] \tag{3.1}$$

In this paper, we study purified oracle adversaries $\mathsf{Adv}_t^{\mathcal{O}}$ that have access to an oracle $\mathcal{O}$ corresponding to a purified random group representation $\rho(g)$.

3.3 Peter-Weyl Theorem and the Fourier Transform

The first step towards understanding how to simulate group representations is defining the Fourier transform, which is an algorithmic form of the Peter–Weyl theorem.

Theorem 3.6 (Peter–Weyl)**.** *There is a unitary isomorphism of representations*

$$L^2(G) \;\simeq\; \bigoplus_{\lambda \in \widehat{G}} V_\lambda \otimes V_\lambda^*,$$

where $L^2(G)$ denotes the left regular representation and g acts as $\rho_\lambda(g) \otimes \mathbf{I}_{V_\lambda^*}$ on the right hand side. Here, $\widehat{G}$ refers to the set of isomorphism classes of irreducible representations of G.

[35] We are using the conventions that $\mathrm{tr}[|\psi\rangle] := \mathrm{tr}[\,|\psi\rangle\langle\psi|\,]$ and that $\prod_{i=1}^{t} X_i := X_t\, X_{t-1} \cdots X_2\, X_1$.

Definition 3.7 (Fourier Transform)**.** We define the Fourier transform as the Peter-Weyl decomposition isomorphism in a particular basis. For every irreducible representation λ, choose an orthonormal basis $\{\,|X\rangle : X \in \mathcal{I}_\lambda\,\}$ of V_λ, with dual basis $\{|Y^*\rangle : Y \in \mathcal{I}_\lambda\,\}$ of V_λ^*, where $|Y^*\rangle \in V_\lambda^*$ acts as $\langle Y|$ on V_λ. This leads to a basis $|\lambda, X, Y\rangle$ of the right hand side $\bigoplus_{\lambda\in\widehat{G}} V_\lambda \otimes V_\lambda^*$ of the form

$$|\lambda, X, Y\rangle = |X, Y^*\rangle \in V_\lambda \otimes V_\lambda^*.$$

In this basis, it turns out that the isomorphism F in the Peter-Weyl theorem can be written as

$$F^{-1}|\lambda, X, Y\rangle = \sqrt{d_\lambda} \int_G \langle Y|\, \rho_\lambda(g)\, |X\rangle \;\; |g\rangle\, dg.$$

We define this isomorphism F to be the Fourier transform. By the unitarity of the Fourier transform, the equation above implies the **Schur orthogonality relations**:

Lemma 3.8 (Schur orthogonality for compact groups)**.** *For all $\lambda, \lambda', X, X', Y, Y'$, we have*

$$\sqrt{d_\lambda}\sqrt{d_{\lambda'}} \int_G \overline{\langle Y|\, \rho_\lambda(g)\, |X\rangle}\, \langle Y'|\, \rho_{\lambda'}(g)\, |X'\rangle\, dg = \delta_{\lambda,\lambda'} \cdot \delta_{X,X'} \cdot \delta_{Y,Y'}.$$

In Section 5.7, we derive generalized Schur orthogonality relations for arbitrary semisimple algebras.

3.3.1 Properties of the Fourier Transform

In this section, we consider the following simple isometries defined on $L^2(G)$:

$$\mathsf{Conj} : |f\rangle \mapsto \int_G f(g)^*\, |g\rangle\, dg \tag{3.2}$$

$$\mathsf{Inv} : |f\rangle \mapsto \int_G f(g^{-1})\, |g\rangle\, dg \tag{3.3}$$

$$\mathsf{Transpose} : |f\rangle \mapsto \int_G f(g^{-1})^*\, |g\rangle\, dg \tag{3.4}$$

where the third map is the composition of the first two. Note that Conj and Transpose are anti-linear (or conjugate-linear) rather than linear maps. We now describe how these maps act on the Fourier basis. For the first,

$$\mathsf{Conj}\cdot F^{-1}\,|\lambda, X, Y\rangle = \sqrt{d_\lambda} \int_G \langle Y|\, \rho_\lambda(g)\, |X\rangle^* \;\; |g\rangle\, dg = \sqrt{d_\lambda} \int_G \langle Y|\, \rho_\lambda(g)^*\, |X\rangle \;\; |g\rangle\, dg = F^{-1}\,|\lambda^*, X, Y\rangle\,,$$

where λ^* denotes a dual representation acting as $\rho_{\lambda^*}(g) = \rho_\lambda(g)^*$ on an appropriate dual basis, and we have identified $|X\rangle$ non-canonically as a basis element of V_{λ^*}. For the second, we have

$$\mathsf{Inv}\cdot F^{-1}\,|\lambda, X, Y\rangle = \sqrt{d_\lambda} \int_G \langle Y|\, \rho_\lambda(g^{-1})\, |X\rangle \;\; |g\rangle\, dg = \sqrt{d_\lambda} \int_G \langle Y|\, \rho_\lambda(g)^\dagger\, |X\rangle \;\; |g\rangle\, dg = F^{-1}\,|\lambda^*, Y, X\rangle\,,$$

via the identity $\langle Y|\, U^\dagger\, |X\rangle = \langle X|\, U^*\, |Y\rangle$, where $|Y\rangle$ and $|X\rangle$ are identified as elements of V_λ^* and $(V_\lambda^*)^*$, respectively. Finally, this tells us that the third map is given by

$$\mathsf{Transpose}\cdot F^{-1}\,|\lambda, X, Y\rangle = F^{-1}\,|\lambda, Y, X\rangle\,.$$

3.4 Complete Reducibility and Schur-Weyl Duality

In order to simulate oracle access to ρ, we will reduce to the case where ρ is an *irreducible* representation of G. To do this, we make use of the following result. Let $\widehat{G}$ denote the set of all finite-dimensional irreducible representations of G.

Theorem 3.9 (Complete reducibility for compact Lie groups, [Hal13, Section 4.10])**.** *Let G be a compact Lie group and let ρ be a finite-dimensional unitary representation. Then ρ decomposes as a finite direct sum of irreducible representations*

$$\rho \simeq \bigoplus_{i=1}^{k} \rho_{\lambda_i}.$$

for $\lambda_i \in \widehat{G}$ (not necessarily distinct). We write $\mathrm{Irr}(\rho)$ to be the subset of $\widehat{G}$ which appears in the decomposition of ρ with non-zero multiplicity.

Finally, of particular interest to us is the decomposition of *tensor power representations*

$$\rho^{\otimes t}(g) = \rho(g)^{\otimes t}, \qquad g \in G$$

given a fixed "base" representation ρ. It turns out that the decomposition of $\rho^{\otimes t}$ into irreducible representations can be understood in terms of the commutant algebra of the (tensor power) representation, defined below.

Definition 3.10 (Commutant algebra)**.** Let $\rho : G \to U(N)$ be a unitary representation of a compact Lie group G. For any integer $t \geq 1$, consider the tensor power representation $\rho^{\otimes t}$ acting on $(\mathbb{C}^N)^{\otimes t}$. The *commutant algebra* $\mathcal{A}_t$ of $\rho^{\otimes t}$ is the set of all $N^t \times N^t$ matrices that commute with every $\rho^{\otimes t}(g)$:

$$\mathcal{A}_t \;=\; \mathrm{End}_G((\mathbb{C}^N)^{\otimes t}) = \{\, X \in \mathbb{C}^{N^t \times N^t} : X\rho^{\otimes t}(g) = \rho^{\otimes t}(g)X \text{ for all } g \in G \,\}.$$

By definition, the commutant algebra $\mathcal{A}_t$ acts on $(\mathbb{C}^N)^{\otimes t}$ in a way that commutes with the action of G. Throughout the paper, we make use of Schur-Weyl duality, which explains the structure of the representation $\rho^{\otimes t}$ using the commutant algebra $\mathcal{A}_t$.

Theorem 3.11 (Schur–Weyl Duality)**.** *Let G be a compact Lie group and ρ be a finite-dimensional unitary representation $\rho : G \to U(V)$. Let $\mathcal{A}_t \simeq \mathrm{End}_G(V^{\otimes t})$ denote the commutant algebra.*

Then, the space $V^{\otimes t}$ decomposes as

$$V^{\otimes t} \;\simeq\; \bigoplus_{\lambda} V_G^{\lambda} \otimes V_{\mathcal{A}_t}^{\lambda},$$

where each V_G^{λ} is an irreducible representation of G and each $V_{\mathcal{A}_t}^{\lambda}$ is an irreducible representation of $\mathcal{A}_t$. Under this decomposition, G and $\mathcal{A}_t$ act as

$$\rho^{\otimes t}(g) = \bigoplus_{\lambda} \rho_\lambda(g) \otimes \mathbf{I}_{V_{\mathcal{A}_t}^{\lambda}}, \qquad X = \bigoplus_{\lambda} \mathbf{I}_{V_G^{\lambda}} \otimes \rho_\lambda(X), \quad X \in \mathcal{A}_t.$$

In particular, the $\mathcal{A}_t$-commutant of $V^{\otimes t}$ can be described as the following finite-dimensional algebra:

$$\mathrm{End}_{\mathcal{A}_t}(V^{\otimes t}) = \mathrm{Span}\{\rho(g)^{\otimes t}\}_{g \in G}.$$

This theorem is a standard consequence of complete reducibility along with the double centralizer theorem [GW09, Section 4.1.5].

Example 3.12. Take $G = U(N)$, and let ρ be the fundamental representation, satisfying $\rho(U) = U$. Classical Schur-Weyl duality [Sch27, Wey39] proves that the commutant algebra $\mathcal{A}_t$ of $\rho^{\otimes t}$ is the algebra spanned by all *permutation matrices* on the t registers:

$$\mathcal{A}_t = \operatorname{span}\{X_\pi : \pi \in S_t\}, \quad X_\pi \left|i_1, i_2, \ldots, i_t\right\rangle = \left|i_{\pi^{-1}(1)}, i_{\pi^{-1}(2)}, \ldots, i_{\pi^{-1}(t)}\right\rangle \tag{3.5}$$

which is itself a representation of the symmetric group algebra $\mathbb{C}[S_t]$ known as the *permutation representation*.

3.5 Diagram Algebras and the Schur Representation

Throughout the paper, we consider other instances of Schur-Weyl duality which generalize the classical $(U(N), \mathbb{C}[S_t])$-duality in Example 3.12. In these settings, the relevant commutant algebras are obtained by extending the basis of $\mathbb{C}[S_t]$ (which consists of all permutations in S_t) to more general *partition diagrams*:

Definition 3.13 (Partition Diagram)**.** A *partition diagram* is a set partition of the $2n$ vertices $\{1, \ldots, n\} \cup \{1', \ldots, n'\}$. Partition diagrams are typically visualized using a $n \times 2$ grid. We refer to the top row of the diagram as the *output* row, and the bottom row of the diagram as the *input* row.

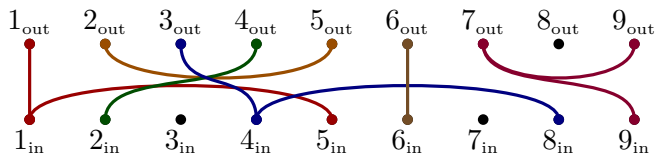
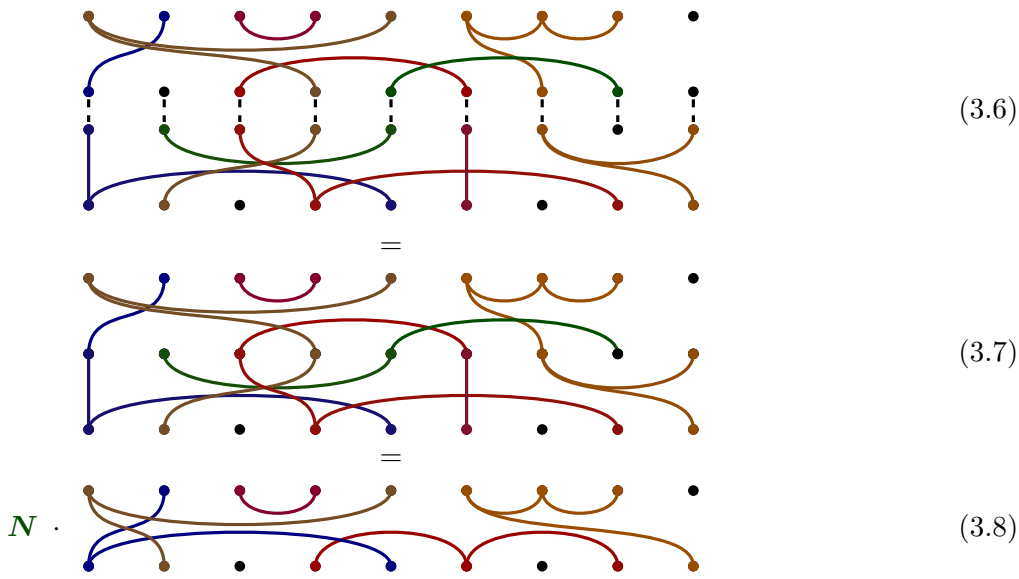


Figure 8: Example of a partition diagram.
Note that the connected components are important, but the specific graph connecting them is not.

Definition 3.14 (Multiplication of Diagrams)**.** Given two partition diagrams D_1 and D_2, the product diagram $D_1 D_2$ is defined by identifying the vertices in the bottom row of D_1 with the vertices in the top row of D_2 (see Fig. 9). The vertices in the identified middle rows are then removed—if c connected components are removed this way, the product is multiplied by N^c. Here, N is a free parameter, although we will typically identify N with the dimension of the group in Table 3.



(3.6)

(3.7)

(3.8)

Figure 9: Multiplication of partition diagrams. The green component is confined to the middle layer and drops out to become an extra factor of N.

Note that when D_1 and D_2 are *permutation diagrams*, where each component in the input row is connected to exactly one component in the output row, Definition 3.14 coincides with the usual composition rule for permutations in S_t.

Definition 3.15 (Diagram Algebras [KS08])**.** Fix an integer $t \geq 0$. Let S be a subset of partition diagrams with t columns. The *diagram algebra generated by* S is the $\mathbb{C}$-algebra

$$\text{Alg}_{\mathbb{C}}(S) = \text{span}_{\mathbb{C}}\{D_1 D_2 \cdots D_k : k \geq 0,\ D_i \in S\}, \tag{3.9}$$

using the multiplication rule defined in Definition 3.14.

For example, the symmetric group algebra $\mathbb{C}[S_t]$ is a diagram algebra with a basis given by the set of all *permutation diagrams* with t columns, i.e. diagrams whose connected components all contain exactly one vertex in the input and output rows. Other examples of diagram algebras include the *Brauer algebra* $B_t(N)$ ([Bra37]), spanned by partition diagrams where every connected component has size 2, and the *partition algebra* $P_t(N)$, ([Mar96]), which is spanned by the set of all partition diagrams. In Section 6.2, we also consider a generalization of partition diagrams where each connected component is assigned a "color" $k \in [K]$. For more details, see Section 6.2.

Just as diagram algebras generalize the symmetric group algebra $\mathbb{C}[S_t]$, the *Schur representation* generalizes the permutation representation (Example 3.12):

Definition 3.16 (The Schur Representation)**.** Let A be a diagram algebra. The *Schur representation*[36] $S(\cdot)$ is an N^t-dimensional representation, defined on a diagram D as follows:

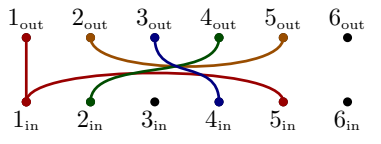
$$S(D) = \sum_{\vec{x},\vec{y}\in[N]^t} \left(\prod_{\substack{u,v\in\{1,\dots,n,1',\dots,n'\} \\ u,\, v \text{ in the same block in } D}} \delta_{z_u,z_v} \right) |\vec{x}\rangle \langle\vec{y}|\,, \tag{3.10}$$

with

$$z_r = \begin{cases} x_r & r \in \{1,\dots,n\} \\ y_r & r \in \{1',\dots,n'\} \end{cases} \tag{3.11}$$

In other words, $S(D)_{\vec{x},\vec{y}} = 1$ if the labeling of vertices given by $\vec{x},\vec{y}$ is constant on every connected component of D, and 0 otherwise.

For example,

$S\Bigg($ 

$\Bigg) = \sum_{\vec{x}\in[N]^7} |x_1,x_2,x_3,x_4,x_2,x_6\rangle\langle x_1,x_4,x_5,x_3,x_1,x_7|$

3.6 Trace forms and dual bases for semisimple algebras

These preliminaries are adapted from [HR05, Section 5].

In this section, we work over a finite-dimensional semisimple algebra A over the complex numbers; this means that A is isomorphic to a direct product of finite-dimensional complex matrix algebras.

Let A be any such algebra, and let $a \in A$. We define the *regular trace* of a by

$$\mathrm{tr}_A(a) := \mathrm{tr}(L_a),$$

where tr denotes the trace of a linear operator and L_a denotes left multiplication by a. This gives a symmetric bilinear pairing

$$\langle a,b\rangle_A := \mathrm{tr}_A(ab).$$

The pairing is nondegenerate because A is semisimple.

Definition 3.17 (Dual basis)**.** Given a basis $\mathcal{B}(A)$ of A, we write a^* for the dual basis element corresponding to $a \in B(A)$; that is,

$$\mathrm{tr}_A(ab^*) = \delta_{a,b} \qquad \text{for all } a,b \in \mathcal{B}(A).$$

Throughout the paper, dual bases are taken with respect to the regular trace tr_A. We will later use this dual basis to interpret the commutant EPR projector Ω_{A_t}.

[36] This representation is sometimes also called the "permutation representation" in the literature, since it generalizes the permutation representation of the symmetric group. For more general diagram algebras, however, it goes beyond permuting registers. Furthermore, even for the symmetric group, there is a chance for confusion with the representation as *permutation matrices* (in-place permutations of the standard basis vectors). Other sources call this the "natural representation", but of course, which representation is natural depends on the algebra and its presentation. To disambiguate, we therefore follow [FND26] in calling it the Schur representation.

Group		Tensor space	Representation	Commutant algebra	
Unitary	$U(N)$	$V_N^{\otimes t}$	$U^{\otimes t}$	Symmetric Group Algebra	$\mathbb{C}[S_t]$
Unitary Product (Haar Cipher)	$U(N)^K$	$\left(V_N^{\oplus K}\right)^{\otimes t}$	$\left(\sum_{k\in[K]} \|k\rangle\langle k\| \otimes U_k\right)^{\otimes t}$	Colored Permutation Group Algebra	$\mathbb{C}[\mathbb{Z}_K \wr S_t]$
Orthogonal	$O(N)$	$V_N^{\otimes t}$	$U_g^{\otimes t}$ U_g is the standard embedding of $g \in O(N)$ as a unitary over $\mathbb{C}^t$	Brauer Algebra	$B_t(N)$
Colored Permutations	$\mathbb{Z}_r \wr S_t$	$V_N^{\otimes t}$	$(P_\pi F_f)^{\otimes t}$ P_π is the permutation matrix of $\pi \in S_N$, and F_f is diagonal with r'th roots of unity	Tanabe Algebra	$T_t(N, r)$
Symmetric	S_N	$(V_{N-1} \oplus V_1)^{\otimes t}$	$P_\sigma^{\otimes t}$ P_σ is the permutation matrix of $\sigma \in S_N$	Partition Algebra	$P_t(N)$
Boolean Functions	$\mathbb{Z}_2^N$	$\left(V_1^{\oplus N}\right)^{\otimes t}$	$\left(\sum_{x\in[N]} (-1)^{f(x)} \|x\rangle\langle x\|\right)^{\otimes t}$	Colored Even-Partition Algebra	

Table 3: Schur-Weyl dualities between different compact Lie groups and their commutant algebras. In all cases, the commutant algebra is the Schur representation (Definition 3.16) of the given diagram algebra. See Section 6.2 for the definition of the Schur representation of $\mathbb{C}[\mathbb{Z}_K \wr S_t]$. In many of these cases, the given commutant algebra relation only holds when t is sufficiently small relative to N. For larger t, the commutant instead becomes a quotient algebra of the one listed; see Section 3.9 for more details and Section 6.2.2 for an example in which this becomes important.

Example 3.18 (Group algebras)**.** Let $A = \mathbb{C}[G]$ be the group algebra of a finite group G, with basis $\{[g] : g \in G\}$. Left multiplication by $[g]$ permutes this basis. Hence

$$\mathrm{tr}_{\mathbb{C}[G]}([g]) = \begin{cases} |G|, & g = e, \\ 0, & g \neq e. \end{cases} \tag{3.12}$$

Therefore, for $a = \sum_{g\in G} a_g[g]$,

$$\mathrm{tr}_{\mathbb{C}[G]}(a) = |G|\, a_e.$$

The dual basis is then given by

$$[g]^* = \frac{1}{|G|}[g^{-1}].$$

3.7 The (dual) Clebsch–Gordan transform and Clebsch–Gordan coefficients

Let $(\rho_\lambda, V_G^\lambda)$ and (ρ_μ, V_G^μ) be irreducible unitary representations of G. By complete reducibility (Theorem 3.9), we have a decomposition

$$V_G^\lambda \otimes V_G^\mu \simeq \bigoplus_{\lambda^+ \in \widehat{G}} \left(M_{\lambda,\mu}^{\lambda^+} \otimes V_G^{\lambda^+}\right),$$

where G acts as $\mathbf{I}_{M^{\lambda^+}_{\lambda,\mu}} \otimes \rho_{\lambda^+}$ on the right hand side. In this setting, we let $m^{\lambda^+}_{\lambda,\mu} = \dim(M^{\lambda^+}_{\lambda,\mu})$ and define a standard basis $\{|r\rangle\}$ for $M^{\lambda^+}_{\lambda,\mu}$.

Definition 3.19 (Clebsch–Gordan transform)**.** A *Clebsch–Gordan transform* is any family of unitary, G-respecting isomorphisms

$$C_{\lambda,\mu} : V_G^\lambda \otimes V_G^\mu \xrightarrow{\simeq} \bigoplus_{\lambda^+} \big(M^{\lambda^+}_{\lambda,\mu} \otimes V_G^{\lambda^+}\big),$$

meaning that for all $g \in G$,

$$C_{\lambda,\mu}\Big(\rho_\lambda(g) \otimes \rho_\mu(g)\Big)C^\dagger_{\lambda,\mu} = \bigoplus_{\lambda^+ \in \widehat{G}} \big(\mathbf{I}_{M^{\lambda^+}_{\lambda,\mu}} \otimes \rho_{\lambda^+}(g)\big)$$

The Clebsch–Gordan transform is unique up to independent unitary changes of bases on each multiplicity space $M^{\lambda^+}_{\lambda,\mu}$. We will use $C = \sum_{\lambda,\mu} |\lambda\rangle\langle\lambda| \otimes |\mu\rangle\langle\mu| \otimes C_{\lambda,\mu}$ to denote the Clebsch–Gordan transform defined on all $V_G^\lambda \otimes V_G^\mu$ simultaneously.

Clebsch–Gordan Coefficients. Next, we give notation for the matrix elements of C in fixed bases for all of the underlying spaces. Under our conventions from earlier, we use

$$\{\,|X\rangle\}_X, \quad \{\,|x\rangle\}_x, \quad \{\,|r, X^+\rangle\}_{r,X^+}$$

as bases for V_G^λ, V_G^μ, and $M^{\lambda^+}_{\lambda,\mu} \otimes V_G^{\lambda^+}$. We can then write C as an explicit matrix:

$$C = \sum_{\substack{\lambda,\mu,\lambda^+ \\ X,x \\ r,X^+}} \langle \lambda, \mu, \lambda^+, r, X^+|\, C\, |\lambda, \mu, X, x\rangle \cdot\, |\lambda, \mu, \lambda^+, r, X^+\rangle\langle \lambda, \mu, X, x|\,,$$

where we call the matrix elements $\langle \lambda, \mu, \lambda^+, r, X^+|\, C\, |\lambda, \mu, X, x\rangle$ the **Clebsch–Gordan coefficients** for G.

In the event that $m^{\lambda^+}_{\lambda,\mu} = 1$ for all λ^+, we may drop the $|r\rangle$ from our notation.

Dual Clebsch–Gordan Transform and Coefficients. Finally, we define a related *dual* Clebsch–Gordan transform and state an identity involving Clebsch–Gordan and dual Clebsch–Gordan coefficients that we make use of to analyze our compressed representations. Departing slightly from prior work [Ngu23, GBO23] but more in line with [Har05], we first define the dual Clebsch-Gordan transform using dual vector spaces. Then, we fix an irreducible representation labeling convention so that for every λ, there is a "dual label" λ^* such that $V_G^{\lambda^*} \simeq \big(V_G^\lambda\big)^*$ and the action of G on $\big(V_G^\lambda\big)^*$ is the dual representation $(\rho_{\lambda^*}(g) \cdot v^*)(w) = v^*(\rho_\lambda(g)^\dagger \cdot w) = (\rho_\lambda(g)v)^*(w)$. Thus, in an explicit dual basis, $\rho_{\lambda^*}(g)$ acts as $\rho_\lambda(g)^*$.

Under these conventions, we then define "dual Clebsch-Gordan coefficients" on the underlying irrep spaces (with no explicit dual spaces).

Definition 3.20 (Dual Clebsch–Gordan transform)**.** A *Dual Clebsch–Gordan transform* is any family of unitary, G-respecting isomorphisms

$$dC_{\lambda,\mu} : V_G^\lambda \otimes (V_G^\mu)^* \xrightarrow{\simeq} \bigoplus_{\lambda^-} \big(M^{\lambda^-}_{\lambda,\mu,*} \otimes V_G^{\lambda^-}\big),$$

meaning that for all $g \in G$,

$$dC_{\lambda,\mu}\Big(\rho_\lambda(g) \otimes \rho^*_\mu(g)\Big)dC^\dagger_{\lambda,\mu} = \bigoplus_{\lambda^- \in \widehat{G}} \big(\mathbf{I}_{M^{\lambda^-}_{\lambda,\mu,*}} \otimes \rho_{\lambda^-}(g)\big)$$

To be compatible with prior work, given a collection of isomorphisms $(V_G^\mu)^* \simeq V_G^{\mu^*}$ for irrep label pairs (μ, μ^*) (suppressed for simplicity), we abuse notation and identify this dual Clebsch–Gordan transform with a map

$$dC_{\lambda,\mu} : V_G^\lambda \otimes V_G^{\mu^*} \to \bigoplus_{\lambda^- \in \widehat{G}} (M_{\lambda,\mu^*}^{\lambda^-} \otimes V_G^{\lambda^-}).$$

In other words, we have

$$dC = \sum_{\substack{\lambda,\mu,\lambda^- \\ Y,y \\ r,Y^-}} \langle \lambda, \mu, \lambda^-, r, Y^- | \, dC \, | \lambda, \mu, Y, y \rangle \cdot \, | \lambda, \mu, \lambda^-, r, Y^- \rangle\langle \lambda, \mu, Y, y | \,,$$

with

$$\langle \lambda, \mu, \lambda^-, r, Y^- | \, dC \, | \lambda, \mu, Y, y \rangle = \, \langle \lambda, \mu^*, \lambda^-, r, Y^- | \, C \, | \lambda, \mu^*, Y, y \rangle$$

and

$$dC_{\lambda,\mu}\Big(\rho_\lambda(g) \otimes \rho_{\mu^*}(g)\Big) dC_{\lambda,\mu}^\dagger = \bigoplus_{\lambda^- \in \widehat{G}} \big(\mathbf{I}_{M_{\lambda,\mu^*}^{\lambda^-}} \otimes \rho_{\lambda^-}(g)\big).$$

Up to an identification of bases of $\big(V_G^\mu\big)^*$ and $V_G^{\mu^*}$, this leaves us with the same degrees of freedom in choosing a dual Clebsch–Gordan transform as we have for choosing a Clebsch–Gordan transform.

An additional property. For our compressed representations, we work with choices of Clebsch–Gordan and dual Clebsch–Gordan whose coefficients satisfy an additional identity.

Lemma 3.21. *There exist choices of C, dC satisfying the identities*

$$\langle \lambda^+, \mu, X^+, x | \, dC^\dagger \, | \lambda^+, \mu, \lambda, r, X \rangle = \sqrt{\frac{\dim\big(V_G^\lambda\big)}{\dim\big(V_G^{\lambda^+}\big)}} \, \langle \lambda, \mu, \lambda^+, r, X^+ | \, C \, | \lambda, \mu, X, x \rangle^* \tag{3.13}$$

In fact, [VK92] give the identity

$$\langle \lambda^+, \mu, X^+, x | \, dC^\dagger \, | \lambda^+, \mu, \lambda, r, X \rangle = \sqrt{\frac{\dim\big(V_G^\lambda\big)}{\dim\big(V_G^{\lambda^+}\big)}} \, \langle \lambda, \mu, \lambda^+, r, X^+ | \, C \, | \lambda, \mu, X, x \rangle$$

for choices of C, dC with real coefficients. We state this relaxed condition in case this allows for additional efficient instantiations.

3.8 The Generalized Schur Transform

We will also make use of a generalization of the Schur transform, defined below.

Definition 3.22 (Schur transform)**.** Let (ρ, V) be a finite-dimensional (not necessarily irreducible) unitary representation of G. A *Schur transform* is any family of unitary, G-respecting isomorphisms

$$\mathsf{Schur}_{\rho,t} : \; V^{\otimes t} \xrightarrow{\simeq} \bigoplus_\lambda \big(V_G^\lambda \otimes V_{\mathcal{A}_t}^\lambda\big),$$

meaning that for all $g \in G$,

$$\mathsf{Schur}_{\rho,t}\left(\rho(g)^{\otimes t}\right) \mathsf{Schur}_{\rho,t}^\dagger = \bigoplus_{\lambda \in \widehat{G}} \big(\rho_\lambda(g) \otimes \mathbf{I}_{V_{\mathcal{A}_t}^\lambda}\big)$$

We will use $\mathsf{Schur}_\rho = \sum_{t\geq 0} |t\rangle\langle t| \otimes \mathsf{Schur}_{\rho,t}$ to denote the Schur transform defined on all t simultaneously, extending $\mathsf{Schur}_{\rho,t}$ for smaller t to act as identity for all but the first t registers. By convention, we take $\mathsf{Schur}_{\rho,0}$ to be the isometry that takes no input and produces the label for the trivial 1-dimensional irrep of G, with no multiplicity.

Schur Transform for a Locally Block-Diagonal Representation. We will often consider cases in which ρ is already block-diagonal, with $V = \bigoplus_{\mu\in\mathrm{Irr}(\rho)} V_G^\mu \otimes M^\mu$ and $\rho(g) = \sum_{\mu\in\mathrm{Irr}(\rho)} |\mu\rangle\langle\mu| \otimes \rho_\mu(g) \otimes \mathbf{I}_{M^\mu(g)}$. This Schur transform is equivalent to the one described above up to the local transformation that block-diagonalizes each copy of V (which will often be efficient, but not always).

The local multiplicity registers $\{M_i^\mu\}_{i\in[t]}$ will end up being collected as part of the total multiplicity $V_{\mathcal{A}_t}^\lambda$. We will sometimes ignore the local multiplicity registers for notational simplicity, since at least in principle, a Schur transform that considers them can be built from one that does not by tacking them onto the $V_{\mathcal{A}_t}^\lambda$ register at the end (though this may not yield the desired basis for $V_{\mathcal{A}_t}^\lambda$).

The representation ρ determines which subset of irreps will appear (or equivalently, the set of available irreps will determine the representation). Furthermore, in this basis, we can view the Schur transform as not explicitly depending on the representation ρ, since we can view it as accepting any irrep in $\widehat{G}$, not just those that appear in ρ (though of course, if $\widehat{G}$ has infinite cardinality, any actual implementation must choose some subset of $\widehat{G}$ to allow). We will then just write $\mathsf{Schur}_{G,t}$ and $\mathsf{Schur}_G = \sum_t |t\rangle\langle t| \otimes \mathsf{Schur}_{G,t}$, or just Schur when the group is clear from context.

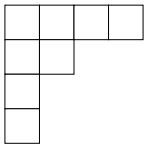
The Schur transform can then be written as a map from sequences of the form

$$(\vec{\mu}, \vec{x}) := ((\mu_1, x_1), (\mu_2, x_2), \ldots, (\mu_t, x_t)) \tag{3.14}$$

to superpositions of states of the form $|\lambda, T_{\mathcal{A}_t}, T_G\rangle$, where $\lambda \in \widehat{G}$, $|T_{\mathcal{A}_t}\rangle \in V_{\mathcal{A}_t}^\lambda$, and $|T_G\rangle \in V_G^\lambda$, with coefficients $\langle\lambda, T_{\mathcal{A}_t}, T_G| \mathsf{Schur} |\vec{\mu}, \vec{x}\rangle$.

3.9 Young Tableaux and Bratteli Paths

The irreducible representations of the groups we study (and their corresponding commutant algebras) are indexed by integer partitions, typically visualized with *Young diagrams*. For example, the following Young diagram corresponds to the partition $(5, 2, 2, 1)$ of $N = 10$:



We use $|\lambda|$ to denote the total number of boxes in λ, λ_i to denote the number of boxes in the ith row, and λ'_i to denote the number of boxes in the ith column.

The Symmetric Group.

Definition 3.23 (Standard Young Tableau)**.** Let λ be a Young diagram. A *standard Young tableau* (SYT) corresponding to λ is a filling of the Young diagram with the numbers $\{1, 2, \ldots |\lambda|\}$, such that all rows of λ are filled in increasing order from left to right, and all columns of λ are filled in increasing order from top to bottom.

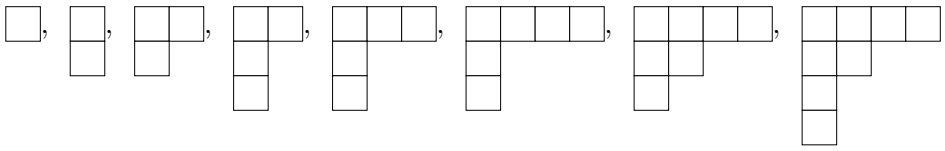
For example, the following is one valid standard Young tableau corresponding to the Young diagram above:

1	3	5	6
2	7		
4			
8			

An equivalent characterization of an SYT T is as a chain of Young diagrams

$$\emptyset, A_1, A_2, \ldots, A_{|\lambda|} = \lambda \tag{3.15}$$

where A_i is the shape formed by the boxes of T whose labels are at most i. For the SYT above, the corresponding chain is given by

$\emptyset$,  (3.16)

Lemma 3.24 (Hook Length Formula)**.** *Given a Young diagram λ, the hook length formula f^λ is*

$$f^\lambda = \frac{|\lambda|!}{\prod_{b\in\lambda} h(b)} \tag{3.17}$$

where $h(b)$ is the hook length of the box b. The hook length is equal to the number of boxes in the same row strictly to the right of b, plus the number of boxes in the same column strictly below b, plus one.

Lemma 3.25 ([Sag01])**.** *For any $\lambda \in \widehat{\mathbb{C}[S_t]}$, f^λ counts the number of standard Young tableaux corresponding to λ. Equivalently,* $\dim\left(V^\lambda_{S_{|\lambda|}}\right) = f^\lambda$.

The Unitary Group.

Definition 3.26 (Semi-Standard Young Tableau)**.** Let $\lambda = (\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N)$ be a weakly decreasing sequence of integers, not necessarily nonnegative. Write $\lambda_{(+)} = (\lambda_1, \ldots, \lambda_p)$ for the positive entries of λ (so $\lambda_p > 0 \geq \lambda_{p+1}$), and $\lambda_{(-)} = (\lambda_{N+1-q}, \ldots, \lambda_N)$ for the negative entries of λ (so $\lambda_{N-q} \geq 0 > \lambda_{N+1-q}$), and let $\mu = (\mu_1 \geq \cdots \geq \mu_q)$, where $\mu_j := -\lambda_{N+1-j}$, denote the partition formed from the negative entries of λ, reversed and negated.

A *semi-standard Young tableau* (SSYT)[37] of shape λ with entries in $[N] := \{1, \ldots, N\}$ consists of:

- a filling of the Young diagram of $\lambda_{(+)}$ (the *positive tableau*) with entries in $[N]$, weakly increasing along each row (left to right) and strictly increasing down each column (top to bottom); and
- a filling of the Young diagram of μ (the *negative tableau*) with entries in $[N]$, weakly decreasing along each row (left to right) and strictly decreasing down each column (top to bottom),

[37] Semi-standard Young tableaux are often defined only with positive rows and boxes, while tableaux with negative rows are called *rational* semistandard tableaux [Kin70, Ste87, Kwo08]. We generally allow negative rows, and specify explicitly when negative rows are not present. We will thus often drop the word "rational" and refer to both as simply semistandard Young tableau.

subject to the following compatibility condition. For $x \in [N]$, let $r^+(x)$ denote the largest row index $i \in \{1, \ldots, p\}$ (rows of $\lambda_{(+)}$ numbered from the top) containing the entry x in the positive tableau, or $r^+(x) := 0$ if x does not appear there. Define $r^-(x)$ analogously, as the largest row index $j \in \{1, \ldots, q\}$ (rows of μ numbered from the top, or equivalently, rows of $\lambda_{(-)}$) containing x in the negative tableau, or 0 if x does not appear there. We require

$$r^+(x) + r^-(x) \;\leq\; x \qquad \text{for all } x \in [N]. \tag{3.18}$$

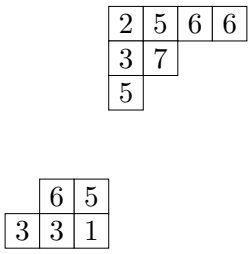


Figure 10: Example of a semistandard tableau with shape $(4, 2, 1, 0, 0, -2, -3)$, where $N = 7$.

Let the *complement* of a negative column be a positive column containing all the entries of $[N]$ besides those in the original column (arranged in order). Then the condition in Equation (3.18) is equivalent to saying that shifting the tableau by one column (by converting the first negative column to its complement) produces a valid ordering in the resulting positive tableau.

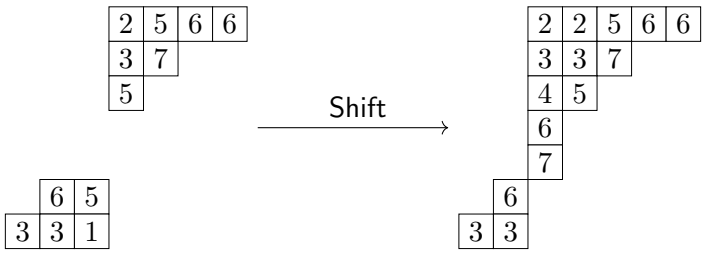


Figure 11: Shifting the semistandard Young tableau by a column maintains a valid order in the positive tableau.

The Weyl Dimension Formula [Wey39] counts the number of semi-standard Young tableaux of a certain shape λ and thus gives the dimension of the corresponding irreducible representation.

Lemma 3.27. *Let λ be a Young diagram. The number of semi-standard Young tableaux of shape λ is*

$$\dim\left(V^{\lambda}_{U(N)}\right) = \prod_{i=1}^{N} \prod_{j=i+1}^{N} \frac{(\lambda_i - i) - (\lambda_j - j)}{j - i}$$

The Partition Algebra. The partition algebra $P_t(N)$ is semisimple whenever $N > 2t - 2$ [HJ20]. In this regime, the irreps of $P_t(N)$ are indexed by N-box Young diagrams such that all but at most

t boxes are in the first row [[HR05], Theorem 2.24b], i.e.

$$\widehat{P_t(N)} = \{\lambda : |\lambda| = N,\ \lambda_0 \geq N - t\} \tag{3.19}$$

We will adopt the convention that when λ refers to an irrep of the partition algebra, the first row is ignored. With this convention, the irreps of the partition algebra are instead indexed by all Young diagrams with at most t boxes, i.e.

$$\widehat{P_t(N)} = \{\lambda : |\lambda| \leq t\} \tag{3.20}$$

With this convention, the dimension of an irrep $\lambda \in \widehat{P_t(N)}$ is [HJ20] given as follows:

Lemma 3.28.

$$\dim\left(V^\lambda_{P_t(N)}\right) = f^\lambda \sum_{i=|\lambda|}^{t} \left\{ {t \atop i} \right\} \binom{i}{|\lambda|} \tag{3.21}$$

where f^λ is the hook length formula in Lemma 3.24, and $\left\{ {t \atop i} \right\}$ is a Stirling number of the second kind. Note that if $|\lambda| = t$, then $\dim\left(V^\lambda_{P_t(N)}\right) = f^\lambda$.

We will sometimes say an irrep λ is a *full box* irrep of $P_t(N)$ if $|\lambda| = t$.

3.10 Clebsch-Gordan Coefficients in the Gelfand-Tsetlin Basis of Unitaries

Lemma 3.29 ([VK92])**.** *Let G be the unitary group $U(N)$. Let X be a semi-standard Young tableau of shape λ which may have positive and/or negative rows, and let $\boxed{x}$ be a Young tableau of shape $\mu = \square$, the standard irrep of $U(N)$. We can write the Young tableau as a Gelfand-Tsetlin pattern, that is, a sequence of N shapes $(X^{\leq \ell})_{\ell \in [N]}$, where each $X^{\leq \ell}$ is the shape of the X tableau when it is restricted to entries at most ℓ (and removing $N - \ell$ empty rows such that it has ℓ rows). For each shape $X^{\leq \ell}$, let $X^{\leq \ell}_k$ be the number of boxes on the k^{th} row (which may be positive or negative).*

In the Gelfand-Tsetlin basis, the Clebsch-Gordan coefficient $\langle \lambda, \mu, \lambda^+, r, X^+ | \, C \, | \lambda, \mu, X, x \rangle$ for tensoring a single box Young tableau $\boxed{x}$ has the following properties:

1. *The transition from shape λ to λ^+ has no multiplicity, so r is redundant.*
2. *It has non-zero value only when λ^+ differs from λ by adding a single box (or removing a negative box)*
3. *It has non-zero value only when X^+ has the form*

   $$X^+ = (X^{\leq 1}, X^{\leq 2}, \ldots, X^{\leq x-1}, X^{\leq x}, X^{\leq x+1}, \ldots, X^{\leq N}) + (0, 0, \ldots, 0, \square_{\Delta_x}, \square_{\Delta_{x+1}}, \ldots, \square_{\Delta_N}),$$

   where "$+\square_{\Delta_z}$" means "add a box to the $\Delta_z{}^{th}$ row" (or remove a negative box), for some $(\Delta_x \in [x], \ldots, \Delta_N \in [N])$. That is, every shape starting with the x^{th} one will have a net increase of one box.
4. *In this case, the Clebsch-Gordan coefficients can be written as a product of reduced Wigner*

coefficients, or scalar factors. Namely,

$$\langle \lambda, \mu, \lambda^+, r, X^+ |\, C \,| \lambda, \mu, X, x \rangle = \left| \frac{\prod_{k=1}^{x-1} \left(X_k^{\leq x-1} - X_{\bar{\Delta}_x}^{\leq x} + \Delta_x - k - 1 \right)}{\prod_{\substack{k=1 \\ k \neq \Delta_x}}^{x} \left(X_k^{\leq x} - X_{\bar{\Delta}_x}^{\leq x} + \Delta_x - k \right)} \right|^{\frac{1}{2}}$$

$$\cdot \prod_{z>x} \left| \frac{\prod_{\substack{k=1 \\ k \neq \Delta_{z-1}}}^{z-1} \left(X_k^{\leq z-1} - X_{\bar{\Delta}_z}^{\leq z} + \Delta_z - k - 1 \right)}{\prod_{\substack{k=1 \\ k \neq \Delta_z}}^{z} \left(X_k^{\leq z} - X_{\bar{\Delta}_z}^{\leq z} + \Delta_z - k \right)} \right|^{\frac{1}{2}} \left| \frac{\prod_{\substack{k=1 \\ k \neq \Delta_z}}^{z} \left(X_k^{\leq z} - X_{\bar{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k \right)}{\prod_{\substack{k=1 \\ k \neq \Delta_{z-1}}}^{z-1} \left(X_k^{\leq z-1} - X_{\bar{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k - 1 \right)} \right|^{\frac{1}{2}} (-1)^{\Delta_z > \Delta_{z-1}}$$

Similarly, in the Gelfand-Tsetlin basis, the dual Clebsch-Gordan coefficients

$$\langle \lambda, \mu, \lambda^-, r, X^- |\, dC \,| \lambda, \mu, X, x \rangle = \langle \lambda, \bar{\mu}, \lambda^-, r, X^- |\, C \,| \lambda, \bar{\mu}, X, x \rangle$$

(where $\bar{\mu}$ is the dual of the standard representation) have the properties:

1. *The transition from shape λ to λ^- has no multiplicity.*
2. *It has non-zero value only when λ^- differs from λ by removing a single box (or adding a negative box)*
3. *It has non-zero value only when X^- has the form*

   $$X^- = (X^{\leq 1}, X^{\leq 2}, \ldots, X^{\leq x-1}, X^{\leq x}, X^{\leq x+1}, \ldots, X^{\leq N}) - (0, 0, \ldots, 0, \square_{\Delta_x}, \square_{\Delta_{x+1}}, \ldots, \square_{\Delta_N}),$$

   where "$-\square_{\Delta_z}$" means "remove a box from the $\Delta_z{}^{th}$ row" (or add a negative box), for some $(\Delta_x \in [x], \ldots, \Delta_N \in [N])$. That is, every shape starting with the x^{th} one will have a net decrease of one box.
4. *In this case, as before, the Clebsch-Gordan coefficients can be written as a product of reduced Wigner coefficients, or scalar factors. Namely,*

$$\langle \lambda, \mu, \lambda^-, r, X^- |\, dC \,| \lambda, \mu, X, x \rangle = \left| \frac{\prod_{k=1}^{x-1} \left(X_k^{\leq x-1} - X_{\bar{\Delta}_x}^{\leq x} + \Delta_x - k \right)}{\prod_{\substack{k=1 \\ k \neq \Delta_x}}^{x} \left(X_k^{\leq x} - X_{\bar{\Delta}_x}^{\leq x} + \Delta_x - k \right)} \right|^{\frac{1}{2}}$$

$$\cdot \prod_{z>x} \left| \frac{\prod_{\substack{k=1 \\ k \neq \Delta_{z-1}}}^{z-1} \left(X_k^{\leq z-1} - X_{\bar{\Delta}_z}^{\leq z} + \Delta_z - k \right)}{\prod_{\substack{k=1 \\ k \neq \Delta_z}}^{z} \left(X_k^{\leq z} - X_{\bar{\Delta}_z}^{\leq z} + \Delta_z - k \right)} \right|^{\frac{1}{2}} \left| \frac{\prod_{\substack{k=1 \\ k \neq \Delta_z}}^{z} \left(X_k^{\leq z} - X_{\bar{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k + 1 \right)}{\prod_{\substack{k=1 \\ k \neq \Delta_{z-1}}}^{z-1} \left(X_k^{\leq z-1} - X_{\bar{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k + 1 \right)} \right|^{\frac{1}{2}} (-1)^{\Delta_z > \Delta_{z-1}}$$

3.11 The Fourier Basis of Semisimple Algebras

Definition 3.30 (The Fourier Basis)**.** Let $\mathcal{A}$ be a finite-dimensional semisimple algebra, i.e.

$$\mathcal{A} \overset{f}{\simeq} \bigoplus_{\lambda \in \widehat{\mathcal{A}}} \mathrm{End}(V^\lambda) \tag{3.22}$$

A *Fourier basis* $\left\{E^\lambda_{i,j}\right\}_{\lambda,i,j}$ for $\mathcal{A}$ is the preimage of the natural matrix unit basis $\{|i\rangle\langle j|_\lambda\}_{\lambda,i,j}$ under the isomorphism f.

In general, the isomorphism in Eq. (3.22) is not unique, and so the Fourier basis is only defined relative to f. Moreover, the matrix unit basis $\{|i\rangle\langle j|_\lambda\}_{\lambda,i,j}$ implicitly assumes a choice of basis for each irrep space V^λ.[38]

Lemma 3.31 ([FND26])**.** *Let $\mathcal{A}_k$ be a finite-dimensional semisimple algebra, and let $\left\{E^{k,\lambda}_{i,j}\right\}_{\substack{\lambda \in \widehat{\mathcal{A}}_k \\ i,j \in \mathcal{B}(V^\lambda_{\mathcal{A}_k})}}$ be a Fourier basis. Then, we have that*

1. *[FND26, Lemma 2.13b]: $\left(E^{k,\lambda}_{i,j}\right)^* = \frac{1}{\dim(V^\lambda_{\mathcal{A}_k})} E^{k,\lambda}_{j,i}$, where $*$ denotes the dual basis element (Definition 3.17).*
2. *[FND26, Section 4.5.2]: $S\left(E^{k,\lambda}_{i,j}\right)^T = S\left(E^{k,\lambda}_{j,i}\right)$, where $S(\cdot)$ denotes the Schur representation (Definition 3.16).*
3. *[FND26, Lemma 2.13a]: $E^{k,\lambda}_{i,j} E^{k,\mu}_{\ell,m} = \delta_{\lambda\mu}\, \delta_{j\ell}\, E^{k,\lambda}_{i,m}$.*

3.12 Ratios of Irrep Dimensions for the Partition and Symmetric Group Algebras

Lemma 3.32. *Consider the $(S_N, P_t(N))$-Schur duality, where $t^2 \leq N/4$ (in this range, $P_t(N)$ is always semisimple). For any Young diagram λ with a full number of boxes $|\lambda| = t$,*

$$\frac{\dim\left(V^\lambda_{P_t(N)}\right)}{\dim\left(V^{(N-t,\lambda)}_{S_N}\right)} = \frac{t!}{N^t}\left(1 \pm O\left(\frac{t^2}{N}\right)\right), \tag{3.23}$$

where $(N-t,\lambda)$ is the Young diagram with an extra first row of $N-t$ boxes.

Similarly, for the $(U(N), \mathbb{C}[S_t])$-Schur duality, we have that for all irreps,[39]

$$\frac{\dim\left(V^\lambda_{\mathbb{C}[S_t]}\right)}{\dim\left(V^\lambda_{U(N)}\right)} = \frac{t!}{N^t}\left(1 \pm O\left(\frac{t^2}{N}\right)\right) \tag{3.24}$$

[38] In algorithmic representation theory, a standard choice of basis for each irrep space is a *subalgebra adapted basis* [DR90, MRR03, BCH05, MRW16, FND26], which is constructed according to the branching rules when restricting an irrep of $\mathcal{A}$ to an irrep of a subalgebra $\mathcal{B}$. We will not require this assumption.

[39] Note that in this case, *all* irreps correspond to Young diagrams with a full number of boxes ($|\lambda| = t \ \ \forall \lambda$).

Proof. When $|\lambda| = t$, we have that

$$\dim\left(V_{S_N}^{(N-t,\lambda)}\right) = \frac{N!}{\prod_{(i,j)\in\lambda} h_\lambda(i,j) \cdot \prod_{i=1}^{N-t}(N-t-i+1+\lambda'_i)}, \qquad \text{(Lemma 3.25)}$$

$$\dim\left(V_{P_t(N)}^{\lambda}\right) = \dim\left(V_{S_t}^{\lambda}\right) = f^\lambda = \frac{t!}{\prod_{(i,j)\in\lambda} h_\lambda(i,j)}, \qquad \text{(Lemma 3.28, Lemma 3.25)}$$

$$\dim\left(V_{U(N)}^{\lambda}\right) = \prod_{(i,j)\in\lambda} \frac{N+j-i}{h_\lambda(i,j)}. \qquad \text{(Lemma 3.27)}$$

where λ'_i is the i'th column length of λ. Therefore, the irrep dimension ratio for the full box irreps of the $(S_N, P_t(N))$ duality is

$$\frac{\dim\left(V_{P_t(N)}^{\lambda}\right)}{\dim\left(V_{S_N}^{(N-t,\lambda)}\right)} = \frac{t! \cdot \prod_{i=1}^{N-t}(N-t-i+1+\lambda'_i)}{N!} \tag{3.25}$$

$$= \frac{t! \cdot (N-t)! \cdot \prod_{i=1}^{N-t}(N-t-i+1+\lambda'_i)}{N! \cdot \prod_{i=1}^{N-t}(N-t-i+1)} \tag{3.26}$$

$$= \frac{1}{\binom{N}{t}} \cdot \prod_{i=1}^{N-t}\left(\frac{N-t-i+1+\lambda'_i}{N-t-i+1}\right) \tag{3.27}$$

$$= \frac{1}{\binom{N}{t}} \cdot \prod_{i=1}^{t}\left(\frac{N-t-i+1+\lambda'_i}{N-t-i+1}\right) \qquad (\lambda'_i = 0 \text{ for } i > t)$$

$$= \frac{1}{\binom{N}{t}} \cdot \left(1 \pm O\left(\frac{t}{N}\right)\right)^t \qquad (\lambda'_i \leq |\lambda|)$$

$$= \frac{1}{\binom{N}{t}} \cdot \left(1 \pm O\left(\frac{t^2}{N}\right)\right) \qquad (t^2 = O(N))$$

$$= \frac{t!}{N^t} \cdot \left(1 \pm O\left(\frac{t^2}{N}\right)\right) \tag{3.28}$$

For the $(U(N), \mathbb{C}[S_t])$ duality,

$$\frac{\dim\left(V_{S_t}^{\lambda}\right)}{\dim\left(V_{U(N)}^{\lambda}\right)} = \frac{t!}{\prod_{(i,j)\in\lambda}(N+j-i)} \tag{3.29}$$

$$= \frac{t!}{N^t} \cdot \left(1 \pm O\left(\frac{t^2}{N}\right)\right) \qquad (|j-i| < |\lambda|)$$

□

4 The Tableau Recording Oracle

Let G be a compact Lie group and let $\rho : G \to U(N)$ be a finite-dimensional unitary representation of G. We wish to give an efficient simulation of query access to $\rho(g)$ for g sampled from the Haar measure of G (where g is re-used for all queries). In fact, we wish to simulate (controlled) access to all of $\rho(g), \rho(g)^\dagger, \rho(g)^*$, and $\rho(g)^\top$.

4.1 The Uncompressed Representation Oracle

We first consider the following purification of the above process:

- The adversary's register is $V = \mathbb{C}^N$, containing the input on which they wish to evaluate $\rho(g)$.
- The recording register is the Hilbert space $L^2(G)$ (which may be countably infinite-dimensional), initialized to the Haar measure $dg = |\mathbf{1}\rangle \in L^2(G)$ using the constant $f(g) = 1$ function.
- A query is represented by the unitary

$$Q_\rho\Big(\,|x\rangle \otimes |f\rangle\,\Big) = \int_G f(g)\cdot \rho(g)\,|x\rangle \otimes |g\rangle\,dg. \tag{4.1}$$

  Since each $\rho(g)$ is a unitary, this is in fact a norm-preserving transformation. We also use the following notation to denote Q_ρ:

$$Q_\rho = \int_G \rho(g) \otimes |g\rangle\langle g|\,dg. \tag{4.2}$$

In fact, using Q_ρ, we can also easily describe a process that has query access to $\rho(g), \rho(g)^\dagger, \rho(g)^*$, and $\rho(g)^\top$:[40]

- Controlled access to $\rho(g)^\dagger$ is simulated by (controlled) $\Big(\mathbf{I}_V \otimes \mathsf{Inv}\Big)Q_\rho\Big(\mathbf{I}_V \otimes \mathsf{Inv}\Big)$,
- Controlled access to $\rho(g)^*$ is simulated by (controlled) $\Big(\mathbf{I}_V \otimes \mathsf{Conj}\Big)Q_\rho\Big(\mathbf{I}_V \otimes \mathsf{Conj}\Big)$,
- Controlled access to $\rho(g)^\top$ is simulated by (controlled) $\Big(\mathbf{I}_V \otimes \mathsf{Transpose}\Big)Q_\rho\Big(\mathbf{I}_V \otimes \mathsf{Transpose}\Big)$.

In this section, our goal is to give an equivalent simulation of this oracle using the Clebsch-Gordan transform. To do this, we must first decompose ρ into irreducible representations, i.e., $V \simeq \bigoplus_\mu M^\mu \otimes V_G^\mu$, where each $\mu \in \widehat{G}$ is an irreducible representation that appears in V with multiplicity $\dim(M^\mu)$. We represent vectors $|\psi\rangle \in V$ using the basis $|\mu\rangle \otimes |m\rangle \otimes |x\rangle$ for $|m\rangle \in M^\mu$ and $|x\rangle \in V_G^\mu$.

In this basis, we can re-write Q_ρ as

$$Q_\rho = \sum_\mu \Big(\,|\mu\rangle\langle\mu| \otimes \mathbf{I}_{M_\mu} \otimes \int_G \rho_\mu(g) \otimes\, |g\rangle\langle g|\,dg\Big) = \sum_\mu |\mu\rangle\langle\mu| \otimes \mathbf{I} \otimes Q_{\rho_\mu}. \tag{4.3}$$

4.2 The Tableau Recording Oracle

The uncompressed recording oracle uses an extremely large, potentially even infinite-dimensional, Hilbert space in its purification, and is thus certainly not efficient in general. However, it turns out that in the Fourier basis (that is, considering the image of $L^2(G)$ under the Fourier transform F) we get a recording oracle that is efficient whenever the Clebsch–Gordan transform C is efficient! Indeed, when zero queries have been made, the recording register is in the trivial irrep state $F \cdot |\mathbf{1}\rangle = |\lambda = \mathsf{triv}, X = 1, Y = 1\rangle$, because $|\mathbf{1}\rangle \in L^2(G)$ is the unique normalized state that is invariant under G-translation.

Formally, we define the tableau recording oracle Q_ρ^{Tab} [GY25] as follows.

[40] See Section 3.3.1 for the definitions of Inv, Conj, and Transpose.

- As before, the adversary's register contains linear combinations of $|\mu\rangle \otimes |m\rangle \otimes |x\rangle$ for $|m\rangle \in M^\mu$ and $|x\rangle \in V_G^\mu$.
- The recording register now contains the following sub-registers:
  - A register containing a label $\lambda \in \widehat{G}$ of an irreducible representation ρ_λ. While there are infinitely many such λ, it turns out that after t queries to the tableau recording oracle, at most N^t such λ are "reachable," resulting in a finite-dimensional Hilbert space after any finite number of queries, which can be encoded in at most $t \log(N)$ qubits.
  - Two registers X and Y containing elements of V_λ and V_λ^* controlled on $|\lambda\rangle$. We also non-canonically identify V_λ^* with V_λ by choosing a fixed basis of V_λ with corresponding dual basis of V_λ^*.
    V_λ has dimension at most N^t for any λ reachable after t queries, and can also therefore be encoded in at most $t \log(N)$ qubits.
- The update rule is given by the composition of the following two maps (see Figure 12):
  - First, given $|\mu, m, x\rangle \otimes |\lambda, X, Y\rangle$, apply $C\,|\lambda, \mu, X, x\rangle$. This results in a superposition of states of the form $|\lambda, \mu, \lambda^+, r, X^+\rangle$ (see Definition 3.19).
  - Then, on the registers containing $\lambda^+, \mu, \lambda, r, Y$, apply $dC^\dagger$, the inverse of the dual Clebsch–Gordan transform. This results in a superposition of states of the form $|\lambda^+, \mu, Y^+, y\rangle$.
  - The registers containing $|\mu, m, y\rangle$ are returned to the adversary, while the registers containing $|\lambda^+, X^+, Y^+\rangle$ are the updated state of the recording oracle.

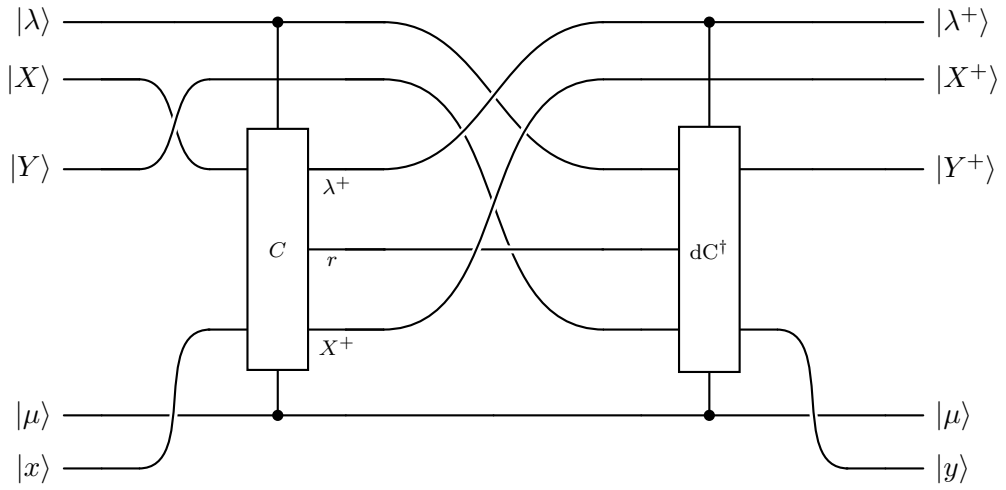


Figure 12: A single query can be decomposed into a Clebsch–Gordan transformation followed by an inverse dual-Clebsch–Gordan operation, though not on the same registers. The adversary's multiplicity register is not shown since it is not touched by the update.

4.2.1 Equivalence of the uncompressed and tableau recording oracles

Having defined our two recording oracles, we are ready to state the main result of this section.

Theorem 4.1 (Clebsch–Gordan Updates for Tableau Recordings)**.** *Let $F : L^2(G) \to \bigoplus_\lambda V_\lambda \otimes V_\lambda^*$ denote the Fourier transform. Then, the uncompressed representation oracle is isometric to the tableau recording oracle via the map F: in particular, $(\mathbf{I} \otimes F) \cdot Q_\rho \cdot (\mathbf{I} \otimes F)^{-1}$ is equal to the tableau recording oracle update rule of a Clebsch–Gordan transform followed by an inverse dual Clebsch–Gordan transform.*

This theorem is similar to the main result of [GY25], but additionally proves that the Fourier transform is an explicit Uhlmann transformation relating the tableau oracle to the uncompressed oracle.

Proof. To show this, we expand out

$$\left\langle \mu, m, y, \lambda^+, X^+, Y^+ \right| (\mathbf{I} \otimes F) \cdot Q_\rho \cdot (\mathbf{I} \otimes F)^\dagger \left| \mu, m, x, \lambda, X, Y \right\rangle \tag{4.4}$$

using the following expressions:

$$(\mathbf{I} \otimes F)^\dagger \left| \mu, x, \lambda, X, Y \right\rangle = \left| \mu, x \right\rangle \otimes \left(\sqrt{\dim(V_G^\lambda)} \int \left\langle Y \right| \rho_\lambda(g) \left| X \right\rangle \cdot \left| g \right\rangle \right) \tag{4.5}$$

$$(\mathbf{I} \otimes F)^\dagger \left| \mu, y, \lambda^+, X^+, Y^+ \right\rangle = \left| \mu, y \right\rangle \otimes \left(\sqrt{\dim(V_G^{\lambda^+})} \int \left\langle Y^+ \right| \rho_{\lambda^+}(g) \left| X^+ \right\rangle \cdot \left| g \right\rangle \right). \tag{4.6}$$

Plugging in the definition of Q, these imply that

$$\begin{aligned}
(4.4) = &\int_G \left\langle \mu, m, y \right| \Big(\left| \mu, m \right\rangle\!\left\langle \mu, m \right| \otimes \rho_\mu(g) \Big) \left| \mu, m, x \right\rangle && (4.7)\\
&\cdot \left(\sqrt{\mathrm{d}(V_G^\lambda) \cdot \mathrm{d}(V_G^{\lambda^+})} \left\langle Y^+ \right| \overline{\rho_{\lambda^+}(g)} \left| X^+ \right\rangle \left\langle Y \right| \rho_\lambda(g) \left| X \right\rangle \right) dg && (4.8)\\
=&\sqrt{\mathrm{d}(V_G^\lambda) \cdot \mathrm{d}(V_G^{\lambda^+})} \int_G \left\langle Y^+ \right| \overline{\rho_{\lambda^+}(g)} \left| X^+ \right\rangle \underbrace{\left\langle y \right| \rho_\mu(g) \left| x \right\rangle \left\langle Y \right| \rho_\lambda(g) \left| X \right\rangle}_{(*)} dg && (4.9)
\end{aligned}$$

Let us expand out

$$(*) = \langle \lambda, Y, \mu, y | \Big(|\lambda\rangle\langle\lambda| \otimes \rho_\lambda(g) \otimes |\mu\rangle\langle\mu| \otimes \rho_\mu(g) \Big) |\lambda, X, \mu, x\rangle \tag{4.10}$$

$$= \langle \lambda, \mu, Y, y | C^\dagger \cdot \Big(C \Big(|\lambda\rangle\langle\lambda| \otimes |\mu\rangle\langle\mu| \otimes \rho_\lambda(g) \otimes \rho_\mu(g) \Big) C^\dagger \Big) \cdot C |\lambda, \mu, X, x\rangle \tag{4.11}$$

$$= \langle \lambda, \mu, Y, y | C^\dagger \cdot \Bigg(|\lambda\rangle\langle\lambda| \otimes |\mu\rangle\langle\mu| \otimes \sum_{\lambda^\#, r \in [m_{\lambda,\mu}^{\lambda^\#}]} |r\rangle\langle r| \otimes |\lambda^\#\rangle\langle\lambda^\#| \otimes \rho_{\lambda^\#}(g) \Bigg) \cdot C |\lambda, \mu, X, x\rangle \tag{4.12}$$

$$= \langle \lambda, \mu, Y, y | C^\dagger \cdot \Bigg(|\lambda\rangle\langle\lambda| \otimes |\mu\rangle\langle\mu| \otimes \sum_{\lambda^\#, r \in [m_{\lambda,\mu}^{\lambda^\#}]} |r\rangle\langle r| \otimes |\lambda^\#\rangle\langle\lambda^\#| \tag{4.13}$$

$$\otimes \sum_{X^\#, Y^\#} |Y^\#\rangle\langle Y^\#| \, \rho_{\lambda^\#}(g) \, |X^\#\rangle\langle X^\#| \Bigg) \cdot C |\lambda, \mu, X, x\rangle \tag{4.14}$$

$$= \sum_{\substack{\lambda^\#, X^\#, Y^\#, \\ r \in [m_{\lambda,\mu}^{\lambda^\#}]}} \langle \lambda, \mu, Y, y | C^\dagger |\lambda, \mu, \lambda^\#, r, Y^\#\rangle \cdot \langle \lambda, \mu, \lambda^\#, r, X^\# | C |\lambda, \mu, X, x\rangle \cdot \langle Y^\# | \rho_{\lambda^\#}(g) |X^\#\rangle \tag{4.15}$$

Plugging this back into (4.9), we have

$$(4.4) = \sqrt{\frac{\mathrm{d}(V_G^\lambda)}{\mathrm{d}(V_G^{\lambda^+})}} \sum_{\substack{\lambda^\#, X^\#, Y^\#, \\ r \in [m_{\lambda,\mu}^{\lambda^\#}]}} \langle \lambda, \mu, Y, y | C^\dagger |\lambda, \mu, \lambda^\#, r, Y^\#\rangle \, \langle \lambda, \mu, \lambda^\#, r, X^\# | C |\lambda, \mu, X, x\rangle \tag{4.16}$$

$$\cdot \underbrace{\Big(\mathrm{d}(V_G^{\lambda^+}) \cdot \int_G \langle Y^+ | \, \overline{\rho_{\lambda^+}(g)} \, |X^+\rangle \, \langle Y^\# | \, \rho_{\lambda^\#}(g) \, |X^\#\rangle \, dg \Big)}_{1 \text{ if } (\lambda^+, X^+, Y^+) = (\lambda^\#, X^\#, Y^\#), \text{ else } 0} \tag{4.17}$$

$$= \sqrt{\frac{\mathrm{d}(V_G^\lambda)}{\mathrm{d}(V_G^{\lambda^+})}} \sum_{r \in [m_{\lambda,\mu}^{\lambda^+}]} \langle \lambda, \mu, Y, y | C^\dagger |\lambda, \mu, \lambda^+, r, Y^+\rangle \, \langle \lambda, \mu, \lambda^+, r, X^+ | C |\lambda, \mu, X, x\rangle \, . \tag{4.18}$$

Now, we invoke the dagger of Lemma 3.21 to obtain

$$(4.4) = \sum_{r \in [m_{\lambda^+,\lambda,\mu}]} \langle \lambda^+, \mu, Y^+, y | \, dC^\dagger \, |\lambda^+, \mu, \lambda, r, Y\rangle \, \langle \lambda, \mu, \lambda^+, r, X^+ | C |\lambda, \mu, X, x\rangle \, . \tag{4.19}$$

Finally, we observe that these matrix entries correspond to the composition of the Clebsch–Gordan and inverse dual Clebsch–Gordan unitaries applied to the appropriate sub-registers indicated below:

$$C = \sum_{\lambda, \lambda^+, \mu, r, X, X^+, x} \langle \lambda, \mu, \lambda^+, r, X^+ | C |\lambda, \mu, X, x\rangle \cdot |\lambda, \mu, \lambda^+, r, X^+\rangle\langle\lambda, \mu, X, x| \tag{4.20}$$

$$dC^\dagger = \sum_{\lambda,\lambda^+,\mu,r,Y,Y^+,y} \langle \lambda^+,\mu,Y^+,y|\, dC^\dagger\, |\lambda^+,\mu,\lambda,r,Y\rangle \cdot |\lambda^+,\mu,Y^+,y\rangle\langle\lambda^+,\mu,\lambda,r,Y|\,. \tag{4.21}$$

This completes the derivation of the Clebsch–Gordan update rule. □

4.2.2 Other query types

The derived update rule implies that we can implement compressed conjugate, transpose, and inverse queries with the (linear[41]) isometries defined below

- $\Big(\mathbf{I}\otimes F\Big)\Big(\mathbf{I}_V\otimes \mathsf{Conj}\Big)\Big(\mathbf{I}\otimes F^\dagger\Big)\cdot Q_\rho^{\mathrm{Tab}}\cdot\Big(\mathbf{I}\otimes F\Big)\Big(\mathbf{I}_V\otimes \mathsf{Conj}\Big)\Big(\mathbf{I}\otimes F^\dagger\Big)$,
- $\Big(\mathbf{I}\otimes F\Big)\Big(\mathbf{I}_V\otimes \mathsf{Transpose}\Big)\Big(\mathbf{I}\otimes F^\dagger\Big)\cdot Q_\rho^{\mathrm{Tab}}\cdot\Big(\mathbf{I}\otimes F\Big)\Big(\mathbf{I}_V\otimes \mathsf{Transpose}\Big)\Big(\mathbf{I}\otimes F^\dagger\Big)$,
- $\Big(\mathbf{I}\otimes F\Big)\Big(\mathbf{I}_V\otimes \mathsf{Inv}\Big)\Big(\mathbf{I}\otimes F^\dagger\Big)\cdot Q_\rho^{\mathrm{Tab}}\cdot\Big(\mathbf{I}\otimes F\Big)\Big(\mathbf{I}_V\otimes \mathsf{Inv}\Big)\Big(\mathbf{I}\otimes F^\dagger\Big)$.

Using the three identities from Section 3.3.1, we can see that

- $\Big(\mathbf{I}\otimes F\Big)\Big(\mathbf{I}_V\otimes \mathsf{Conj}\Big)\Big(\mathbf{I}\otimes F^\dagger\Big)$ performs the map

$$\sum_{\lambda,X,Y}\alpha_{\lambda,X,Y}\,|\lambda,X,Y\rangle \mapsto \sum_{\lambda,X,Y}\overline{\alpha}_{\lambda,X,Y}\,|\lambda^*,X,Y\rangle\,.$$

- $\Big(\mathbf{I}\otimes F\Big)\Big(\mathbf{I}_V\otimes \mathsf{Transpose}\Big)\Big(\mathbf{I}\otimes F^\dagger\Big)$ performs the map

$$\sum_{\lambda,X,Y}\alpha_{\lambda,X,Y}\,|\lambda,X,Y\rangle \mapsto \sum_{\lambda,X,Y}\overline{\alpha}_{\lambda,X,Y}\,|\lambda,Y,X\rangle\,.$$

- $\Big(\mathbf{I}\otimes F\Big)\Big(\mathbf{I}_V\otimes \mathsf{Inv}\Big)\Big(\mathbf{I}\otimes F^\dagger\Big)$ performs the map

$$\sum_{\lambda,X,Y}\alpha_{\lambda,X,Y}\,|\lambda,X,Y\rangle \mapsto \sum_{\lambda,X,Y}\alpha_{\lambda,X,Y}\,|\lambda^*,Y,X\rangle\,.$$

Thus, we conclude that

- Conjugate queries can be implemented using the *complex conjugate* $\overline{Q_\rho^{\mathrm{Tab}}}$ of Q_ρ^{Tab} conjugated by the linear isometry $|\lambda,X,Y\rangle \mapsto |\lambda^*,X,Y\rangle$.

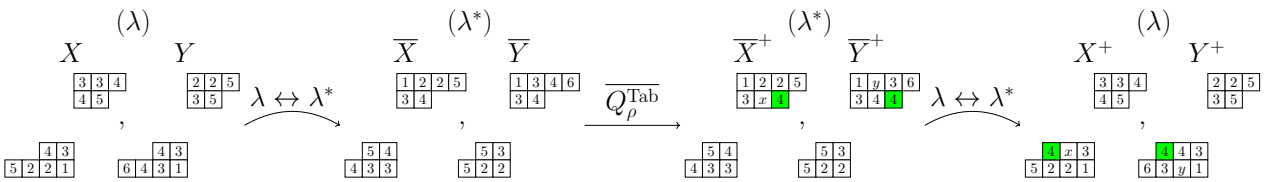


Figure 13: Example basis vectors in a complex conjugate query for the unitary group.

[41] as opposed to anti-linear (or conjugate linear)

- Transpose queries can be implemented using $\overline{Q_\rho^{\mathrm{Tab}}}$ conjugated by the unitary $|\lambda, X, Y\rangle \mapsto |\lambda, Y, X\rangle$

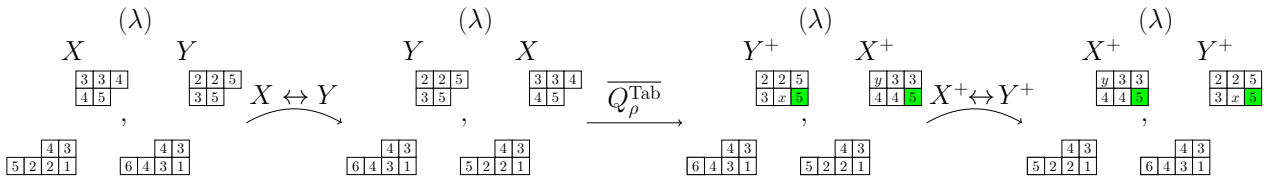


Figure 14: Example basis vectors in a transpose query for the unitary group.

- Inverse queries can be implemented using Q_ρ^{Tab} conjugated by the linear isometry $|\lambda, X, Y\rangle \mapsto |\lambda^*, Y, X\rangle$.

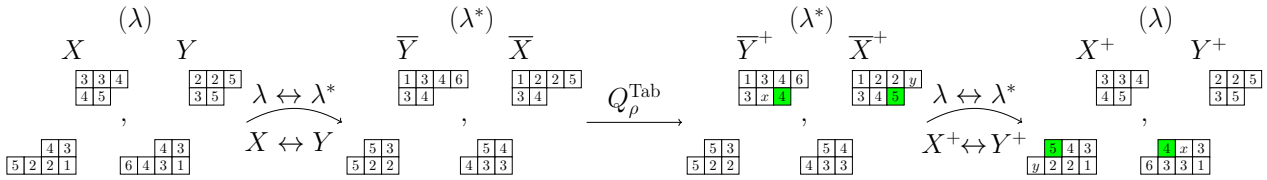


Figure 15: Example basis vectors in an inverse query for the unitary group.

Thus, given any polynomial-size circuit for Q_ρ^{Tab} and efficient isomorphisms between V_{λ^*} and V_λ^*, we have corresponding polynomial-size circuits for these other query types.

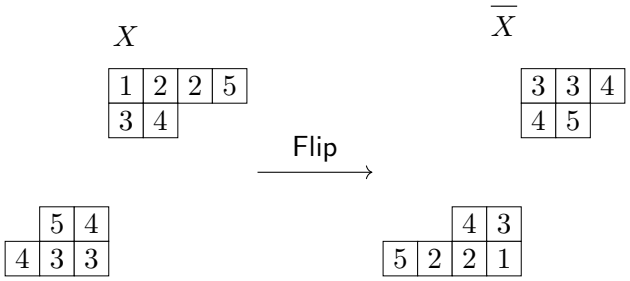
4.2.3 Efficient Implementation for the Unitary Group

We now show how to efficiently implement the update rule, Q_ρ^{Tab}, in Theorem 4.1 for the case of Haar random unitaries, $G = U(N)$. The Clebsch–Gordan transform for unitaries was shown to have an efficient algorithm in [BCH05, Har05, BFG+25], and this algorithm was extended to the mixed unitary case, in which $\rho(U)$ could be either the standard irrep or the dual irrep, by [Ngu23, FTH23, GBO23]. For the case of the irrep λ consisting of a tableau containing at most t boxes (either positive or negative), and the oracle representation being a single-box irrep (either the standard or dual irrep), they show how to implement the Clebsch–Gordan transform with error at most ε in time $\mathrm{poly}(\log(N), t, \log(\frac{1}{\varepsilon}))$.

We show here that the Clebsch–Gordan transform implemented by these algorithms, or any similar algorithm that works in the Gelfand-Tsetlin basis, implies a corresponding dual Clebsch–Gordan algorithm with the same complexity which satisfies Lemma 3.21 and can be used in Theorem 4.1.

Definition 4.2 (Signed Flipping of a Young Tableau)**.** For every basis state $|X\rangle$ given by a semi-standard Young tableau X (allowing negative boxes), let $\mathsf{Flip}\,|X\rangle$ be the operation that, up to a ± 1

sign, flips the Young tableau X by rotating it 180 degrees, for example, sending



(4.22)

with a phase of $(-1)^{Nt+p}$ where p is the *total parity* (even or odd) of all the entries in the tableau. (Equivalently, the phase is $(-1)^{\sum_{x_i \in X}(N-x_i)}$ where the sum runs over all the entries in the boxes of the tableau.)

For instance, in the example in Equation (4.22) above, there are 6 odd entries (namely 1, 5, 3, 5, 3, and 3), so the sum is even, and $t = 11$ is odd, so in this case we would apply a -1 phase whenever N is odd.

When working with a semi-standard Young tableau X, we always apply a sign when flipping. Therefore, when clear from context, we refer to the signed flipping of a tableau as simply the "flipping operation."

Lemma 4.3. *Let $G = U(N)$, and let $\mu = \square$ be the standard irrep of $U(N)$. The dual Clebsch–Gordan transform implied by Lemma 3.29 is equivalent to applying the signed flip of Definition 4.2 to the Young tableau X (and a corresponding unsigned flip to corresponding shape λ), applying the standard Clebsch–Gordan transform of Lemma 3.29, and then (signed) flipping it back.*

That is, we have that

$$\langle \lambda, \mu, \lambda^-, X^- | \, dC \, | \lambda, \mu, X, x \rangle = \langle \lambda, \mu, \lambda^-, X^- | \, \mathsf{Flip} \cdot C \cdot \mathsf{Flip} \, | \lambda, \mu, X, x \rangle \, . \tag{4.23}$$

Where the Flip *applies to the X register, and also flips λ and λ^- (without sign) accordingly.*

This means that every efficient algorithm for the Clebsch–Gordan transform implies a corresponding efficient dual Clebsch–Gordan transform algorithm where we Flip*, apply Clebsch–Gordan, and then* Flip *back.*

Proof. We can show that the reduced Wigner coefficients, or scalar factors, of Lemma 3.29 match up term by term (up to a ± 1 phase, which we will handle separately). Let $Z = \overline{X}$ be the *unsigned*[42] flip of X. Then $Z_k^{\leq z} = -X_{z-k+1}^{\leq z}$ (a positive row turns into a negative row and vice versa, and the row counting is reversed to go from z down to 1 rather than 1 to z). We can see in Lemma 3.29 that the Clebsch–Gordan scalar factors for Z are zero in the same places as for the dual Clebsch–Gordan, so it suffices to consider nonzero scalar factors.

[42] We are starting with the unsigned flip, which we then use to compute the necessary sign. This is the only time we take the flip to be unsigned.

First scalar factor. The first scalar factor of the Clebsch–Gordan coefficient of adding $(0, 0, \ldots, 0, \square_{\Delta_x}, \square_{\Delta_{x+1}}, \ldots, \square_{\Delta_N})$ to Z is

$$\left| \frac{\prod_{k=1}^{x-1} \left(Z_k^{\leq x-1} - Z_{\overline{\Delta}_x}^{\leq x} + \Delta_x - k - 1 \right)}{\prod_{\substack{k=1 \\ k \neq \Delta_x}}^{x} \left(Z_k^{\leq x} - Z_{\overline{\Delta}_x}^{\leq x} + \Delta_x - k \right)} \right|^{\frac{1}{2}} \tag{4.24}$$

$$= \left| \frac{\prod_{k=1}^{x-1} \left(-X_{x-k}^{\leq x-1} + X_{x-\Delta_x+1}^{\leq x} + \Delta_x - k - 1 \right)}{\prod_{\substack{k=1 \\ k \neq \Delta_x}}^{x} \left(-X_{x-k+1}^{\leq x} + X_{x-\Delta_x+1}^{\leq x} + \Delta_x - k \right)} \right|^{\frac{1}{2}} \tag{4.25}$$

$$= \left| \frac{\prod_{k=1}^{x-1} \left(-X_k^{\leq x-1} + X_{x-\Delta_x+1}^{\leq x} + \Delta_x - x + k - 1 \right)}{\prod_{\substack{k=1 \\ x-k+1 \neq \Delta_x}}^{x} \left(-X_k^{\leq x} + X_{x-\Delta_x+1}^{\leq x} + \Delta_x - x + k - 1 \right)} \right|^{\frac{1}{2}}$$

(by re-indexing $k \to x - k$ in numerator and $k \to x - k + 1$ in denominator)

$$= \left| \frac{\prod_{k=1}^{x-1} \left(X_k^{\leq x-1} - X_{x-\Delta_x+1}^{\leq x} + (x - \Delta_x + 1) - k \right)}{\prod_{\substack{k=1 \\ k \neq (x-\Delta_x+1)}}^{x} \left(X_k^{\leq x} - X_{x-\Delta_x+1}^{\leq x} + (x - \Delta_x + 1) - k \right)} \right|^{\frac{1}{2}} . \tag{4.26}$$

When we now relabel $\Delta_x \to (x - \Delta_x + 1)$ (since each row of the Gelfand-Tsetlin pattern is reversed when we apply the flipping operation), we can see that this becomes identical to the first scalar factor of the dual Clebsch–Gordan coefficient of subtracting $(0, 0, \ldots, 0, \square_{\left(x-\Delta_x+1\right)}, \square_{\left((x+1)-\Delta_{(x+1)}+1\right)}, \ldots, \square_{\left(N-\Delta_N+1\right)})$ from X.

Remaining scalar factors. A similar argument holds for the other scalar factors. For the Clebsch–Gordan coefficient of adding $(0, 0, \ldots, 0, \square_{\Delta_x}, \square_{\Delta_{x+1}}, \ldots, \square_{\Delta_N})$ to the tableau $Z = \overline{X}$, they are

$$\left|\frac{\prod_{\substack{k=1\\k\neq\Delta_{(z-1)}}}^{z-1}\left(Z_{\overline{k}}^{\leq z-1} - Z_{\overline{\Delta}_z}^{\leq z} + \Delta_z - k - 1\right)}{\prod_{\substack{k=1\\k\neq\Delta_z}}^{z}\left(Z_{\overline{k}}^{\leq z} - Z_{\overline{\Delta}_z}^{\leq z} + \Delta_z - k\right)}\right|^{\frac{1}{2}}\left|\frac{\prod_{\substack{k=1\\k\neq\Delta_z}}^{z}\left(Z_{\overline{k}}^{\leq z} - Z_{\overline{\Delta}_{(z-1)}}^{\leq z-1} + \Delta_{(z-1)} - k\right)}{\prod_{\substack{k=1\\k\neq\Delta_{(z-1)}}}^{z-1}\left(Z_{\overline{k}}^{\leq z-1} - Z_{\overline{\Delta}_{(z-1)}}^{\leq z-1} + \Delta_{(z-1)} - k - 1\right)}\right|^{\frac{1}{2}}{}^{(-1)^{\Delta_z > \Delta_{(z-1)}}} \tag{4.27}$$

$$= \left|\frac{\prod_{\substack{k=1\\k\neq\Delta_{(z-1)}}}^{z-1}\left(-X_{z-k}^{\leq z-1} + X_{z-\Delta_z+1}^{\leq z} + \Delta_z - k - 1\right)}{\prod_{\substack{k=1\\k\neq\Delta_z}}^{z}\left(-X_{z-k+1}^{\leq z} + X_{z-\Delta_z+1}^{\leq z} + \Delta_z - k\right)}\right|^{\frac{1}{2}}\left|\frac{\prod_{\substack{k=1\\k\neq\Delta_z}}^{z}\left(-X_{z-k+1}^{\leq z} + X_{z-\Delta_{(z-1)}}^{\leq z-1} + \Delta_{(z-1)} - k\right)}{\prod_{\substack{k=1\\k\neq\Delta_{(z-1)}}}^{z-1}\left(-X_{z-k}^{\leq z-1} + X_{z-\Delta_{(z-1)}}^{\leq z-1} + \Delta_{(z-1)} - k - 1\right)}\right|^{\frac{1}{2}}{}^{(-1)^{\Delta_z > \Delta_{(z-1)}}} \tag{4.28}$$

$$= \left|\frac{\prod_{\substack{k=1\\z-k\neq\Delta_{(z-1)}}}^{z-1}\left(-X_{\overline{k}}^{\leq z-1} + X_{z-\Delta_z+1}^{\leq z} + \Delta_z - z + k - 1\right)}{\prod_{\substack{k=1\\z-k+1\neq\Delta_z}}^{z}\left(-X_{\overline{k}}^{\leq z} + X_{z-\Delta_z+1}^{\leq z} + \Delta_z - z + k - 1\right)}\right|^{\frac{1}{2}}\left|\frac{\prod_{\substack{k=1\\z-k+1\neq\Delta_z}}^{z}\left(-X_{\overline{k}}^{\leq z} + X_{z-\Delta_{(z-1)}}^{\leq z-1} + \Delta_{(z-1)} - z + k - 1\right)}{\prod_{\substack{k=1\\z-k\neq\Delta_{(z-1)}}}^{z-1}\left(-X_{\overline{k}}^{\leq z-1} + X_{z-\Delta_{(z-1)}}^{\leq z-1} + \Delta_{(z-1)} - z + k - 1\right)}\right|^{\frac{1}{2}}{}^{(-1)^{\Delta_z > \Delta_{(z-1)}}} \tag{4.29}$$

$$= \left|\frac{\prod_{\substack{k=1\\k\neq z-\Delta_{(z-1)}}}^{z-1}\left(X_{\overline{k}}^{\leq z-1} - X_{z-\Delta_z+1}^{\leq z} + (z - \Delta_z + 1) - k\right)}{\prod_{\substack{k=1\\k\neq z-\Delta_z+1}}^{z}\left(X_{\overline{k}}^{\leq z} - X_{z-\Delta_z+1}^{\leq z} + (z - \Delta_z + 1) - k\right)}\right|^{\frac{1}{2}}\left|\frac{\prod_{\substack{k=1\\k\neq z-\Delta_z+1}}^{z}\left(X_{\overline{k}}^{\leq z} - X_{z-\Delta_{(z-1)}}^{\leq z-1} + (z - \Delta_{(z-1)}) - k + 1\right)}{\prod_{\substack{k=1\\z-k\neq\Delta_{(z-1)}}}^{z-1}\left(X_{\overline{k}}^{\leq z-1} - X_{z-\Delta_{(z-1)}}^{\leq z-1} + (z - \Delta_{(z-1)}) - k + 1\right)}\right|^{\frac{1}{2}}{}^{(-1)^{\Delta_z > \Delta_{(z-1)}}}. \tag{4.30}$$

As before, we now relabel $\Delta_z \to (z - \Delta_z + 1)$ and $\Delta_{z-1} \to (z - \Delta_{z-1})$, and, up to a ± 1 phase, it becomes identical to the corresponding scalar factor of the dual Clebsch–Gordan coefficient of subtracting $(0, \ldots, \square_{\left(x-\Delta_x+1\right)}, \square_{\left((x+1)-\Delta_{(x+1)}+1\right)}, \ldots, \square_{\left(N-\Delta_N+1\right)})$ from X.

It now remains to handle the sign. For each scalar factor besides the first (which does not come with a sign), the attached sign (after relabeling $\Delta_z \to (z - \Delta_z + 1)$ and $\Delta_{z-1} \to (z - \Delta_{z-1})$) is $(-1)^{(z-\Delta_z+1)>(z-\Delta_{(z-1)})} = (-1)^{\Delta_z<(\Delta_{(z-1)}+1)} = (-1)^{\Delta_z \leq \Delta_{(z-1)}}$ for the (unsigned) flipped Clebsch–Gordan, but $(-1)^{\Delta_z > \Delta_{(z-1)}}$ for the dual Clebsch–Gordan coefficient. Since the conditions $\Delta_z \leq \Delta_{(z-1)}$ and $\Delta_z > \Delta_{(z-1)}$ are mutually exclusive, they always disagree, meaning that the scalar factor for our unsigned flipped Clebsch–Gordan always has the wrong sign.

The product of the scalar factors that makes up the Clebsch–Gordan coefficient has $N - x$ such

scalar factors, which means that we are off by a phase of $(-1)^{N-x} = (-1)^{N+x}$. This is exactly the sign that we get by applying a phase of $(-1)^{\sum_{x_i \in X}(N-x_i)}$ both times that we flip. The phase from each value x_i in the tableau that is not added/removed will cancel between the flip before and after the Clebsch–Gordan transform. The only phase remaining will be from the entries that are added or removed in between.

The Clebsch–Gordan transform of inserting x either adds a positive box containing x, removes a negative box containing x, or replaces a negative box containing x with y and adds positive box containing y. In any of these cases, when performing the signed flip before and after the Clebsch–Gordan transform, the resulting phase is $(-1)^{N-x}$ as desired, completing the proof. □

Claim 4.4. *Let C be the Clebsch–Gordan transform implied by Lemma 3.29, and let dC be the corresponding dual Clebsch–Gordan transform, given by applying C on the flipped tableau as in Lemma 4.3. Then C and dC satisfy Lemma 3.21. Namely,*

$$\langle \lambda^+, \lambda, X| \, dC \, |\lambda^+, X^+, x\rangle = \sqrt{\frac{\dim\left(V^{\lambda}_{U(N)}\right)}{\dim\left(V^{\lambda^+}_{U(N)}\right)}} \, \langle \lambda, \lambda^+, X^+| \, C \, |\lambda, X, x\rangle \tag{4.31}$$

This implies that an efficient algorithm for C alone is enough to give an efficient perfect update algorithm for the tableau recording oracle for random unitaries.

Proof. We first observe that each Clebsch–Gordan coefficient and its dual Clebsch–Gordan coefficient are non-zero on the same support. Where they are both non-zero, we proceed by computing the ratio

$$\frac{\langle \lambda, \lambda^+, X^+| \, C \, |\lambda, X, x\rangle}{\langle \lambda^+, \lambda, X| \, dC \, |\lambda^+, X^+, x\rangle} \tag{4.32}$$

by evaluating the ratios of each scalar factor individually.

First scalar factor. The ratio of the first scalar factors of the two Clebsch–Gordan coefficients is

$$\left|\frac{\prod_{k=1}^{x-1}\left(X_k^{\leq x-1} - X_{\Delta_x}^{\leq x} + \Delta_x - k - 1\right)}{\prod_{\substack{k=1\\k\neq\Delta_x}}^{x}\left(X_k^{\leq x} - X_{\Delta_x}^{\leq x} + \Delta_x - k\right)}\right|^{\frac{1}{2}} \left|\frac{\prod_{\substack{k=1\\k\neq\Delta_x}}^{x}\left((X^+)_k^{\leq x} - (X^+)_{\Delta_x}^{\leq x} + \Delta_x - k\right)}{\prod_{k=1}^{x-1}\left((X^+)_k^{\leq x-1} - (X^+)_{\Delta_x}^{\leq x} + \Delta_x - k\right)}\right|^{\frac{1}{2}} \tag{4.33}$$

$$= \left|\frac{\prod_{k=1}^{x-1}\left(X_k^{\leq x-1} - X_{\Delta_x}^{\leq x} + \Delta_x - k - 1\right)}{\prod_{\substack{k=1\\k\neq\Delta_x}}^{x}\left(X_k^{\leq x} - X_{\Delta_x}^{\leq x} + \Delta_x - k\right)}\right|^{\frac{1}{2}} \left|\frac{\prod_{\substack{k=1\\k\neq\Delta_x}}^{x}\left(X_k^{\leq x} - X_{\Delta_x}^{\leq x} + \Delta_x - k - 1\right)}{\prod_{k=1}^{x-1}\left(X_k^{\leq x-1} - X_{\Delta_x}^{\leq x} + \Delta_x - k - 1\right)}\right|^{\frac{1}{2}} \tag{4.34}$$

$$= \left| \frac{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_x}}^{x} \left(X_{\overline{k}}^{\leq x} - X_{\overline{\Delta}_x}^{\leq x} + \Delta_x - k - 1 \right)}{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_x}}^{x} \left(X_{\overline{k}}^{\leq x} - X_{\overline{\Delta}_x}^{\leq x} + \Delta_x - k \right)} \right|^{\frac{1}{2}} \tag{4.35}$$

$$= \left| \prod_{\substack{k=1\\k\neq\Delta_x}}^{x} \frac{X_{\overline{k}}^{\leq x} - X_{\overline{\Delta}_x}^{\leq x} + \Delta_x - k - 1}{X_{\overline{k}}^{\leq x} - X_{\overline{\Delta}_x}^{\leq x} + \Delta_x - k} \right|^{\frac{1}{2}} \tag{4.36}$$

Remaining scalar factors. The ratios of the remaining scalar factors of the two Clebsch–Gordan coefficients are

$$\left| \frac{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_{z-1}}}^{z-1} \left(X_{\overline{k}}^{\leq z-1} - X_{\overline{\Delta}_z}^{\leq z} + \Delta_z - k - 1 \right)}{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_z}}^{z} \left(X_{\overline{k}}^{\leq z} - X_{\overline{\Delta}_z}^{\leq z} + \Delta_z - k \right)} \right|^{\frac{1}{2}} \left| \frac{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_z}}^{z} \left(X_{\overline{k}}^{\leq z} - X_{\overline{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k \right)}{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_{z-1}}}^{z-1} \left(X_{\overline{k}}^{\leq z-1} - X_{\overline{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k - 1 \right)} \right|^{\frac{1}{2}} (-1)^{\Delta_z > \Delta_{z-1}} \tag{4.37}$$

$$\cdot \left| \frac{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_z}}^{z} \left((X^+)_{\overline{k}}^{\leq z} - (X^+)_{\overline{\Delta}_z}^{\leq z} + \Delta_z - k \right)}{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_{z-1}}}^{z-1} \left((X^+)_{\overline{k}}^{\leq z-1} - (X^+)_{\overline{\Delta}_z}^{\leq z} + \Delta_z - k \right)} \right|^{\frac{1}{2}} \left| \frac{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_{z-1}}}^{z-1} \left((X^+)_{\overline{k}}^{\leq z-1} - (X^+)_{\overline{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k + 1 \right)}{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_z}}^{z} \left((X^+)_{\overline{k}}^{\leq z} - (X^+)_{\overline{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k + 1 \right)} \right|^{\frac{1}{2}} (-1)^{\Delta_z > \Delta_{z-1}} \tag{4.38}$$

$$= \left| \frac{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_{z-1}}}^{z-1} \left(X_{\overline{k}}^{\leq z-1} - X_{\overline{\Delta}_z}^{\leq z} + \Delta_z - k - 1 \right)}{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_z}}^{z} \left(X_{\overline{k}}^{\leq z} - X_{\overline{\Delta}_z}^{\leq z} + \Delta_z - k \right)} \right|^{\frac{1}{2}} \left| \frac{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_z}}^{z} \left(X_{\overline{k}}^{\leq z} - X_{\overline{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k \right)}{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_{z-1}}}^{z-1} \left(X_{\overline{k}}^{\leq z-1} - X_{\overline{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k - 1 \right)} \right|^{\frac{1}{2}} \tag{4.39}$$

$$\cdot \left| \frac{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_z}}^{z} \left(X_{\overline{k}}^{\leq z} - X_{\overline{\Delta}_z}^{\leq z} + \Delta_z - k - 1 \right)}{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_{z-1}}}^{z-1} \left(X_{\overline{k}}^{\leq z-1} - X_{\overline{\Delta}_z}^{\leq z} + \Delta_z - k - 1 \right)} \right|^{\frac{1}{2}} \left| \frac{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_{z-1}}}^{z-1} \left(X_{\overline{k}}^{\leq z-1} - X_{\overline{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k \right)}{\displaystyle\prod_{\substack{k=1\\k\neq\Delta_z}}^{z} \left(X_{\overline{k}}^{\leq z} - X_{\overline{\Delta}_{z-1}}^{\leq z-1} + \Delta_{z-1} - k \right)} \right|^{\frac{1}{2}} \tag{4.40}$$

$$= \left| \prod_{\substack{k=1 \\ k\neq\Delta_z}}^{z} \frac{\left(X_k^{\leq z} - X_{\Delta_z}^{\leq z} + \Delta_z - k - 1\right)}{\left(X_k^{\leq z} - X_{\Delta_z}^{\leq z} + \Delta_z - k\right)} \right|^{\frac{1}{2}} \left| \prod_{\substack{k=1 \\ k\neq\Delta_{z-1}}}^{z-1} \frac{\left(X_k^{\leq z-1} - X_{\Delta_{z-1}}^{\leq z-1} + \Delta_{z-1} - k\right)}{\left(X_k^{\leq z-1} - X_{\Delta_{z-1}}^{\leq z-1} + \Delta_{z-1} - k - 1\right)} \right|^{\frac{1}{2}} \tag{4.41}$$

$$= \left| \prod_{\substack{k=1 \\ k\neq\Delta_z}}^{z} \frac{X_k^{\leq z} - X_{\Delta_z}^{\leq z} + \Delta_z - k - 1}{X_k^{\leq z} - X_{\Delta_z}^{\leq z} + \Delta_z - k} \right|^{\frac{1}{2}} \left| \prod_{\substack{k=1 \\ k\neq\Delta_{z-1}}}^{z-1} \frac{X_k^{\leq z-1} - X_{\Delta_{z-1}}^{\leq z-1} + \Delta_{z-1} - k}{X_k^{\leq z-1} - X_{\Delta_{z-1}}^{\leq z-1} + \Delta_{z-1} - k - 1} \right|^{\frac{1}{2}} \tag{4.42}$$

Putting them together. We now combine these to get the ratio of the Clebsch–Gordan coefficients as the product of the ratios of the scalar factors:

$$\frac{\langle \lambda, \lambda^+, X^+ | \, C \, | \lambda, X, x \rangle}{\langle \lambda^+, \lambda, X | \, dC \, | \lambda^+, X^+, x \rangle} \tag{4.43}$$

$$= \left| \prod_{\substack{k=1 \\ k\neq\Delta_x}}^{x} \frac{X_k^{\leq x} - X_{\Delta_x}^{\leq x} + \Delta_x - k - 1}{X_k^{\leq x} - X_{\Delta_x}^{\leq x} + \Delta_x - k} \right|^{\frac{1}{2}} \tag{4.44}$$

$$\cdot \prod_{z=x+1}^{N} \left| \prod_{\substack{k=1 \\ k\neq\Delta_z}}^{z} \frac{X_k^{\leq z} - X_{\Delta_z}^{\leq z} + \Delta_z - k - 1}{X_k^{\leq z} - X_{\Delta_z}^{\leq z} + \Delta_z - k} \right|^{\frac{1}{2}} \left| \prod_{\substack{k=1 \\ k\neq\Delta_{z-1}}}^{z-1} \frac{X_k^{\leq z-1} - X_{\Delta_{z-1}}^{\leq z-1} + \Delta_{z-1} - k}{X_k^{\leq z-1} - X_{\Delta_{z-1}}^{\leq z-1} + \Delta_{z-1} - k - 1} \right|^{\frac{1}{2}} \tag{4.45}$$

This is a telescoping product, since in each factor of the outside product, the second part cancels with the first part of the previous factor, ultimately giving

$$= \left| \prod_{\substack{k=1 \\ k\neq\Delta_N}}^{N} \frac{X_k^{\leq N} - X_{\Delta_N}^{\leq N} + \Delta_N - k - 1}{X_k^{\leq N} - X_{\Delta_N}^{\leq N} + \Delta_N - k} \right|^{\frac{1}{2}} \tag{4.46}$$

$$= \left| \prod_{\substack{k=1 \\ k\neq\Delta_N}}^{N} \frac{\lambda_k - (\lambda_{\Delta_N} + 1) + \Delta_N - k}{\lambda_k - \lambda_{\Delta_N} + \Delta_N - k} \right|^{\frac{1}{2}} \tag{4.47}$$

$$= \left| \prod_{\substack{k=1 \\ k\neq\Delta_N}}^{N} \frac{\lambda_k^+ - \lambda_{\Delta_N}^+ + \Delta_N - k}{\lambda_k - \lambda_{\Delta_N} + \Delta_N - k} \right|^{\frac{1}{2}} \tag{4.48}$$

$$= \left| \prod_{\substack{k=1 \\ k\neq\Delta_N}}^{N} \frac{(\lambda_k^+ - k) - (\lambda_{\Delta_N}^+ - \Delta_N)}{(\lambda_k - k) - (\lambda_{\Delta_N} - \Delta_N)} \right|^{\frac{1}{2}} \tag{4.49}$$

$$= \left| \prod_{i=1}^{\Delta_N - 1} \frac{(\lambda_i^+ - i) - (\lambda_{\Delta_N}^+ - \Delta_N)}{(\lambda_i - i) - (\lambda_{\Delta_N} - \Delta_N)} \right|^{\frac{1}{2}} \left| \prod_{j=\Delta_N+1}^{N} \frac{(\lambda_{\Delta_N}^+ - \Delta_N) - (\lambda_j^+ - j)}{(\lambda_{\Delta_N} - \Delta_N) - (\lambda_j - j)} \right|^{\frac{1}{2}} \cdot \underbrace{\left| \prod_{\substack{i<j \\ i \neq \Delta_N \\ j \neq \Delta_N}} \frac{(\lambda_i^+ - i) - (\lambda_j^+ - j)}{(\lambda_i - i) - (\lambda_j - j)} \right|^{\frac{1}{2}}}_{=1 \text{ since } \lambda \text{ and } \lambda^+ \text{ only differ on row } \Delta_N} \tag{4.50}$$

$$= \left| \prod_{i=1}^{N} \prod_{j=i+1}^{N} \frac{(\lambda_i^+ - i) - (\lambda_j^+ - j)}{(\lambda_i - i) - (\lambda_j - j)} \right|^{\frac{1}{2}} \tag{4.51}$$

$$= \left| \frac{\dim\left(V_{U(N)}^{\lambda^+}\right)}{\dim\left(V_{U(N)}^{\lambda}\right)} \right|^{\frac{1}{2}} . \qquad \square$$

5 The Path Recording Oracle

The tableau-recording oracle gives an exact simulation of queries to a Haar-random group representation by storing the purification in the Fourier basis of $L^2(G)$. Unfortunately, the tableau-recording oracle can be challenging to interpret operationally: from the contents of the tableau register, it can be difficult to characterize in simple terms what information about the queried representation has been learned or fixed by the algorithm.

The goal of this section is to construct an equivalent recording in a path basis. Informally, after t queries, we would like the recording register to store a superposition of Feynman path information $|x_1, \ldots, x_t, y_1, \ldots, y_t\rangle$. For a fixed group element g, this path corresponds to the matrix element

$$\prod_{i=1}^{t} \langle y_i | \rho(g) | x_i \rangle = \langle y_1, \ldots, y_t | \rho(g)^{\otimes t} | x_1, \ldots, x_t \rangle,$$

and indeed, a non-adaptive algorithm that queries the uncompressed oracle Q_ρ on inputs $|x_1, \ldots, x_t\rangle$ would result in the state

$$\sum_{y_1, \ldots, y_t} |y_1, \ldots, y_t\rangle \otimes \int_G \langle \vec{y}| \, \rho^{\otimes t} \, |\vec{x}\rangle \cdot |g\rangle \, dg.$$

How does this compare with recording the path $|\vec{x}, \vec{y}\rangle$? To make this precise, we introduce the matrix element map

$$\Theta_t : V^{\otimes t} \otimes \overline{V}^{\otimes t} \longrightarrow L^2(G)$$

defined as

$$\Theta_t |\vec{x}, \vec{y}\rangle := \int_G \langle \vec{y} | \rho(g)^{\otimes t} | \vec{x} \rangle |g\rangle \, dg, \tag{5.1}$$

so that

$$Q_\rho^{\otimes t} \, |\vec{x}\rangle = \sum_{\vec{y} \in \mathcal{B}(V)} |\vec{y}\rangle \otimes \Theta_t(|\vec{x}\rangle \, |\vec{y}\rangle). \tag{5.2}$$

Using this, we can see that for any particular representation ρ, path states contain some redundancies. Let

$$\mathcal{A}_t = \mathrm{End}_G(V^{\otimes t})$$

be the commutant algebra of the tensor-power representation. As discussed in the technical overview, since every $A \in \mathcal{A}_t$ commutes with $\rho(g)^{\otimes t}$, the matrix-element map satisfies

$$\Theta_t(A \otimes \mathbf{I}) = \Theta_t(\mathbf{I} \otimes A^T).$$

Therefore, it should be the case that $(A \otimes \mathbf{I}) |\psi\rangle = (\mathbf{I} \otimes A^T) |\psi\rangle$ on valid path-recording states. This identification is implemented by the commutant EPR projector $\Omega_{\mathcal{A}_t}$.

Furthermore, path states may require different relative normalizations. These factors are encoded in the reweighting operator $\Lambda_{\mathcal{A}_t}$ defined later.

With these ingredients, the path-recording update has the following form: append a new input-output pair (x_t, y_t), project the recording register onto the commutant-EPR subspace $\Omega_{\mathcal{A}_t}$, and apply an appropriate Schur-Weyl subspace reweighting. The remainder of this section defines these operators precisely, proves that the resulting path-recording oracle is isometric to the tableau recording oracle, gives an interpretation of $\Omega_{\mathcal{A}_t}$ in terms of the algebra $\mathcal{A}_t$ (when it is semisimple), and derives forms of the update that are useful for adaptive query algorithms.

5.1 Defining the Path Recording Oracle

Let

$$\rho : G \to U(V)$$

be a finite-dimensional unitary representation of a compact Lie group G. Fix an orthonormal basis $\mathcal{B}(V) = \{|x\rangle\}$ of V. We use the same ket notation $|y\rangle$ for the corresponding conjugate basis of $\overline{V}$, specifying registers when needed.
For each t, let

$$\mathsf{Schur}_t : V^{\otimes t} \longrightarrow \bigoplus_{\lambda} V_G^{\lambda} \otimes V_{\mathcal{A}_t}^{\lambda}$$

be the generalized Schur transform for $V^{\otimes t}$, so that

$$\rho^{\otimes t}(g) = \bigoplus_{\lambda} \rho_{\lambda}(g) \otimes \mathbf{I}_{V_{\mathcal{A}_t}^{\lambda}}.$$

Define the ordered path space

$$\mathsf{Path}_t \coloneqq V^{\otimes t} \otimes \overline{V}^{\otimes t}.$$

A path basis vector is written

$$|\vec{x}, \vec{y}\rangle \coloneqq |x_1, \ldots, x_t\rangle \otimes |y_1, \ldots, y_t\rangle .$$

Note that $|y_1, \ldots, y_t\rangle$ implicitly corresponds to the dual basis element $|y_1^*, \ldots, y_t^*\rangle \in \overline{V}^{\otimes t}$. Next, we define the *append* map

$$\mathrm{App}_{x,y}^{(t)} : \mathsf{Path}_t \longrightarrow \mathsf{Path}_{t+1}$$

by

$$\mathrm{App}_{x,y}^{(t)} |\vec{x}, \vec{y}\rangle = |x_1, \ldots, x_t, x\rangle \otimes |y_1, \ldots, y_t, y\rangle . \tag{5.3}$$

Equivalently,

$$\mathrm{App}_{x,y}^{(t)} = \mathbf{I} \otimes |x, y\rangle .$$

We relate the path space $V^{\otimes t} \otimes \overline{V}^{\otimes t}$ to the (Fourier/tableau) compressed oracle via the double Schur transform

$$\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t} : V^{\otimes t} \otimes \overline{V}^{\otimes t} \longrightarrow \bigoplus_{\lambda,\lambda'} \left(V_G^{\lambda} \otimes V_{\mathcal{A}_t}^{\lambda} \right) \otimes \overline{\left(V_G^{\lambda'} \otimes V_{\mathcal{A}_t}^{\lambda'} \right)}.$$

We write double-Schur basis vectors as

$$|\lambda, X, T\rangle \otimes \left|\lambda', Y, R\right\rangle,$$

where

$$X \in \mathcal{B}(V_G^\lambda), \qquad Y \in \mathcal{B}(V_G^{\lambda'}), \qquad T \in \mathcal{B}(V_{\mathcal{A}_t}^\lambda), \qquad R \in \mathcal{B}(V_{\mathcal{A}_t}^{\lambda'}).$$

This is again shorthand for the vector $|\lambda, X, T\rangle \otimes \overline{|\lambda', Y, R\rangle}$.

Let $\widehat{\Pi}_{\lambda,t}$ be the projector in the Schur basis onto block λ, that is, onto $V_G^\lambda \otimes V_{\mathcal{A}_t}^\lambda$. We then define its path basis version as

$$\Pi_{\lambda,t} := \mathsf{Schur}_t^\dagger \cdot \widehat{\Pi}_{\lambda,t} \cdot \mathsf{Schur}_t. \tag{5.4}$$

In a slight abuse of notation, we use $\widehat{\Pi}_{\lambda,t}^{\otimes 2}$ to denote $\widehat{\Pi}_{\lambda,t} \otimes \overline{\widehat{\Pi}_{\lambda,t}}$, which is the projector in the double-Schur basis onto the (λ, λ) block

$$\left(V_G^\lambda \otimes V_{\mathcal{A}_t}^\lambda\right) \otimes \overline{\left(V_G^\lambda \otimes V_{\mathcal{A}_t}^\lambda\right)}.$$

Equivalently,

$$\widehat{\Pi}_{\lambda,t}^{\otimes 2}\left(|\lambda_1, X, T\rangle \otimes |\lambda_2, Y, R\rangle\right) = \begin{cases} |\lambda, X, T\rangle \otimes |\lambda, Y, R\rangle\,, & \lambda_1 = \lambda_2 = \lambda, \\ 0, & \text{otherwise.} \end{cases}$$

The corresponding path basis projector is then

$$\Pi_{\lambda,t}^{\otimes 2} := \left(\mathsf{Schur}_t^\dagger \otimes \mathsf{Schur}_t^\top\right) \widehat{\Pi}_{\lambda,t}^{\otimes 2} \left(\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t}\right).$$

For each λ, define the normalized EPR state on the pair of $\mathcal{A}_t$-registers by

$$\left|\mathsf{EPR}_{\mathcal{A}_t}^\lambda\right\rangle := \frac{1}{\sqrt{\dim(V_{\mathcal{A}_t}^\lambda)}} \sum_{T \in \mathcal{B}(V_{\mathcal{A}_t}^\lambda)} |T\rangle \otimes |T\rangle\,.$$

Let $\widehat{\Pi}_{\mathrm{EPR},t}^{\mathcal{A}_t}$ be the double-Schur-basis projector that acts as zero on blocks $\lambda \neq \lambda'$, and on the $\lambda = \lambda'$ block projects the two $\mathcal{A}_t$-registers onto $\left|\mathsf{EPR}_{\mathcal{A}_t}^\lambda\right\rangle$ while acting as identity on the two G-registers. Equivalently,

$$\widehat{\Pi}_{\mathrm{EPR},t}^{\mathcal{A}_t}\left(|\lambda, X, T\rangle \otimes \left|\lambda', Y, R\right\rangle\right) = \begin{cases} \dfrac{\delta_{T,R}}{\dim(V_{\mathcal{A}_t}^\lambda)} \displaystyle\sum_{S \in \mathcal{B}(V_{\mathcal{A}_t}^\lambda)} |\lambda, X, S\rangle \otimes |\lambda, Y, S\rangle\,, & \lambda = \lambda', \\ 0, & \lambda \neq \lambda'. \end{cases}$$

Via the double Schur transform, this induces the following projector on the path space.

$$\Omega_{\mathcal{A}_t} := \left(\mathsf{Schur}_t^\dagger \otimes \mathsf{Schur}_t^\top\right) \widehat{\Pi}_{\mathrm{EPR},t}^{\mathcal{A}_t} \left(\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t}\right). \tag{5.5}$$

Finally, define the subspace-reweighting operator

$$W_t := \sum_\lambda \sqrt{\frac{\dim(V_{\mathcal{A}_t}^\lambda)}{\dim(V_G^\lambda)}} \Pi_{\lambda,t}^{\otimes 2}. \tag{5.6}$$

This operator reweights each double-Schur subspace with matching irrep label λ on the two sides.

On the equal-irrep subspace

$$\Pi_{\mathrm{diag},t} := \sum_{\lambda} \Pi_{\lambda,t}^{\otimes 2},$$

the inverse reweighting is

$$W_t^+ := \sum_{\lambda} \sqrt{\frac{\dim(V_G^\lambda)}{\dim(V_{\mathcal{A}_t}^\lambda)}} \Pi_{\lambda,t}^{\otimes 2}. \tag{5.7}$$

Thus

$$W_t^+ W_t = W_t W_t^+ = \Pi_{\mathrm{diag},t}.$$

Since

$$\Omega_{\mathcal{A}_t} \preceq \Pi_{\mathrm{diag},t},$$

this is a valid inverse restricted to the legal path subspace.

5.2 Path Recording Theorem Statement

With these definitions set up, we are ready to state and analyze our path recording oracle.

Definition 5.1 (Path Recording Oracle)**.** The path recording oracle is defined to be the following linear operator. The input space of the operator is $V \otimes \bigoplus_t \mathsf{Path}_t$.

On subspace $V \otimes \mathsf{Path}_t$, the operator is defined as

$$Q_{\rho,t+1}^{\mathrm{Path}} = \sum_{x,y \in \mathcal{B}(V)} |y\rangle\langle x| \otimes W_{t+1} \cdot \Omega_{\mathcal{A}_{t+1}} \cdot \mathrm{App}_{x,y}^{(t)} \cdot W_t^+ \cdot \Omega_{\mathcal{A}_t}, \tag{5.8}$$

which outputs an element of $V \otimes \mathsf{Path}_{t+1}$.

When restricted to the legal path-recording subspace $\Omega_{\mathcal{A}_t}$, the operator is given by

$$Q_{\rho,t+1}^{\mathrm{Path}} = \sum_{x,y \in \mathcal{B}(V)} |y\rangle\langle x| \otimes W_{t+1} \Omega_{\mathcal{A}_{t+1}} \, \mathrm{App}_{x,y}^{(t)} \, W_t^+. \tag{5.9}$$

Remark 5.2 (Equivalent definition)**.** Another definition of the path recording oracle is the following:

$$Q_{\rho,t+1}^{\mathrm{Path}} = \sum_{x,y \in \mathcal{B}(V)} |y\rangle\langle x| \otimes \Lambda_{\mathcal{A}_{t+1}} \cdot \Omega_{\mathcal{A}_{t+1}} \cdot \mathrm{App}_{x,y}^{(t)} \cdot \Lambda_{\mathcal{A}_t}^+ \cdot \Omega_{\mathcal{A}_t}, \tag{5.10}$$

where

$$\Lambda_{\mathcal{A}_t} = \sum_{\lambda} \sqrt{\frac{\dim(V_{\mathcal{A}_t}^\lambda)}{\dim(V_G^\lambda)}} \cdot \Big(\Pi_{\lambda,t} \otimes \mathbf{I}\Big) \tag{5.11}$$

and $\Lambda_{\mathcal{A}_t}^+$ is inverse to $\Lambda_{\mathcal{A}_t}$ on $\mathrm{Im}(\mathsf{Schur}_t)$. This is equivalent because $\Omega_{\mathcal{A}_t}$ in particular projects onto $\Pi_{\mathrm{diag},t}$, on which $\Lambda_{\mathcal{A}_t}$ and W_t are equivalent (as are their pseudoinverses).

We will prove that $Q_{\rho,t+1}^{\mathrm{Path}}$ is equivalent to Q_ρ on an appropriate subspace of legal "reachable" states. Let

$$\widehat{G}_t := \mathrm{Irr}(\rho^{\otimes t}) = \Big\{\lambda \in \widehat{G} : \dim(V_{\mathcal{A}_t}^\lambda) > 0\Big\}$$

be the set of irreps appearing in the Schur–Weyl decomposition of $V^{\otimes t}$, and let

$$\mathsf{Tab}_t := \bigoplus_{\lambda \in \widehat{G}_t} V_G^\lambda \otimes \overline{V_G^\lambda}$$

be the t-query reachable tableau space.

Theorem 5.3 (Correctness of the path-recording oracle)**.** *There is a sequence of partial isometries* $\Gamma_t : \bigoplus_\lambda V_G^\lambda \otimes \overline{V}_G^\lambda \to \mathsf{Path}_t$ *such that the following holds. For every* t*,*

$$Q^{\mathrm{Path}}_{\rho,t+1} \cdot (\mathbf{I}_V \otimes \Gamma_t \cdot F) = (\mathbf{I}_V \otimes \Gamma_{t+1} \cdot F) \cdot Q_\rho, \tag{5.12}$$

when inputs are restricted to $V \otimes F^{-1}(\mathsf{Tab}_t)$*.*

Therefore, for every t*-query quantum algorithm* A *whose query register is* V*, for all input states* $|\psi\rangle$*, we have that*

$$A^{Q_\rho}(|\psi\rangle) \simeq A^{Q^{\mathrm{Path}}_\rho}(|\psi\rangle)$$

via a fixed partial isometry $\Gamma_t \cdot F$ *on the recording register.*

Setup. With this notation, the path-recording update is given by

$$Q^{\mathrm{Path}}_{\rho,t+1} = \sum_{x,y \in \mathcal{B}(V)} |y\rangle\langle x| \otimes W_{t+1}\Omega_{\mathcal{A}_{t+1}} \mathrm{App}^{(t)}_{x,y} W_t^+ \Omega_{\mathcal{A}_t}.$$

We will prove

$$Q^{\mathrm{Path}}_{\rho,t+1}(\mathbf{I}_V \otimes \Gamma_t F) = (\mathbf{I}_V \otimes \Gamma_{t+1} F) Q_\rho$$

on the appropriately restricted input states.

5.3 Defining the partial isometry Γ_t

To begin proving Theorem 5.3, we first define the partial isometries

$$\Gamma_t : \bigoplus_\lambda V_G^\lambda \otimes \overline{V_G^\lambda} \longrightarrow \mathsf{Path}_t.$$

For every irrep λ, we write the basis states of $V_G^\lambda \otimes \overline{V}_G^\lambda$ as $|\lambda, X, Y\rangle$.

With this setup, we have that

$$\mathrm{Im}\left(\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t}\right) = \bigoplus_{\lambda,\lambda' \in \widehat{G}_t} \left(V_G^\lambda \otimes V_{\mathcal{A}_t}^\lambda\right) \otimes \overline{\left(V_G^{\lambda'} \otimes V_{\mathcal{A}_t}^{\lambda'}\right)}.$$

Since $\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t}$ is an isometry, we thus conclude that

$$\left(\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t}\right)^\dagger \left(\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t}\right) = \mathbf{I}_{\mathsf{Path}_t}, \qquad \left(\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t}\right)\left(\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t}\right)^\dagger = \Pi_{\mathrm{Im}(\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t})}.$$

Next, we define the $\mathcal{A}_t$ EPR isometry

$$E_t : \bigoplus_\lambda V_G^\lambda \otimes \overline{V}_G^\lambda \longrightarrow \bigoplus_\lambda \left(V_G^\lambda \otimes V_{\mathcal{A}_t}^\lambda\right) \otimes \overline{\left(V_G^\lambda \otimes V_{\mathcal{A}_t}^\lambda\right)} \subset \bigoplus_{\lambda,\lambda'} \left(V_G^\lambda \otimes V_{\mathcal{A}_t}^\lambda\right) \otimes \overline{\left(V_G^{\lambda'} \otimes V_{\mathcal{A}_t}^{\lambda'}\right)}$$

by

$$E_t\,|\lambda, X, Y\rangle := \frac{1}{\sqrt{\dim(V_{\mathcal{A}_t}^\lambda)}} \sum_{R \in \mathcal{B}(V_{\mathcal{A}_t}^\lambda)} |\lambda, X, R\rangle \otimes |\lambda, Y, R\rangle\,, \qquad \lambda \in \widehat{G}_t$$

and $E_t\,|\lambda, X, Y\rangle = 0$ otherwise. Since this operation is simply tensoring with a normalized EPR state on the two $\mathcal{A}_t$-registers, E_t is an isometry when restricted to Tab_t. Moreover, by construction,

$$\mathrm{Im}(E_t) \subset \mathrm{Im}(\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t})$$

The adjoint of E_t is given on double-Schur basis states by

$$E_t^\dagger\left(|\lambda, X, T\rangle \otimes \left|\lambda', Y, R\right\rangle\right) = \begin{cases} \dfrac{1}{\sqrt{\dim(V_{\mathcal{A}_t}^\lambda)}} |\lambda, X, Y\rangle, & \lambda = \lambda' \in \widehat{G}_t \text{ and } T = R, \\ 0, & \text{otherwise.} \end{cases}$$

With these definitions, it follows directly that

$$E_t^\dagger E_t = \Pi_{\mathsf{Tab}_t}, \qquad E_t E_t^\dagger = \widehat{\Pi}_{\mathrm{EPR},t}^{\mathcal{A}_t}.$$

Finally, we are ready to define the operators

$$\Gamma_t := \left(\mathsf{Schur}_t^\dagger \otimes \mathsf{Schur}_t^\top\right) E_t.$$

To understand the domain and image on which Γ_t is an isometry, we calculate

$$\Gamma_t^\dagger \Gamma_t = E_t^\dagger \Pi_{\mathrm{Im}(\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t})} E_t = E_t^\dagger E_t = \Pi_{\mathsf{Tab}_t},$$

while

$$\begin{aligned} \Gamma_t \Gamma_t^\dagger &= (\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t})^\dagger E_t E_t^\dagger (\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t}) \\ &= (\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t})^\dagger \widehat{\Pi}_{\mathrm{EPR},t}^{\mathcal{A}_t} (\mathsf{Schur}_t \otimes \overline{\mathsf{Schur}_t}) \\ &= \Omega_{\mathcal{A}_t}, \end{aligned}$$

where the last equation holds by definition of $\Omega_{\mathcal{A}_t}$. Thus, we have proved the following fact.

Claim 5.4. *Γ_t is an isometry between Tab_t and the legal path subspace $\mathrm{Im}(\Omega_{\mathcal{A}_t}) \subset \mathsf{Path}_t$.*

As a result, we see that $\Gamma_t \cdot F$ is an isometry between $F^{-1}(\mathsf{Tab}_t)$ and $\mathrm{Im}(\Omega_{\mathcal{A}_t})$. We next give a characterization of the subspace $F^{-1}(\mathsf{Tab}_t)$.

5.4 Path Matrix Elements

To properly understand Γ_t, we consider the matrix element map

$$\Theta_t : \mathsf{Path}_t \longrightarrow L^2(G)$$

defined as

$$\Theta_t |\vec{x}, \vec{y}\rangle := \int_G \langle \vec{y}| \rho^{\otimes t}(g) |\vec{x}\rangle |g\rangle \, dg. \tag{5.13}$$

We make the following claims about the map Θ_t.

Lemma 5.5. $\mathrm{Im}(F \cdot \Theta_t \cdot \Omega_{\mathcal{A}_t}) = \mathrm{Im}(F \cdot \Theta_t) = \mathsf{Tab}_t$. *Moreover, as operators defined on Path_t,*

$$\Gamma_t \cdot F \cdot \Theta_t = W_t \Omega_{\mathcal{A}_t}. \tag{5.14}$$

Proof. We first prove that $\mathsf{Tab}_t \subset \mathrm{Im}(F \cdot \Theta_t \cdot \Omega_{\mathcal{A}_t})$. To prove this, let $|\lambda, X, Y\rangle \in \mathsf{Tab}_t$ be a tableau basis vector. Let $|T\rangle \in V_{\mathcal{A}_t}^\lambda$ be an arbitrary basis state, which we also identify with a basis state in

$\overline{V}^{\lambda}_{\mathcal{A}_t}$. Then, we observe that

$$\begin{aligned}\Theta_t\Big(\mathsf{Schur}^\dagger\,|\lambda,X,T\rangle\otimes\mathsf{Schur}^\top\,|\lambda,Y,T\rangle\Big) &= \int_G \langle\lambda,Y,T|\,\mathsf{Schur}\cdot\rho^{\otimes t}(g)\cdot\mathsf{Schur}^\dagger\,|\lambda,X,T\rangle\cdot|g\rangle\,dg\\ &= \int_G \langle\lambda,Y,T|\,\rho_\lambda(g)\,|\lambda,X,T\rangle\cdot|g\rangle\,dg \qquad (5.15)\\ &= \frac{1}{\sqrt{\dim(V_G^\lambda)}}\cdot F^{-1}\,|\lambda,X,Y\rangle\,. \qquad (5.16)\end{aligned}$$

Since this holds for every basis state $|T\rangle$, we can initialize the $V^{\lambda}_{\mathcal{A}_t}\otimes\overline{V}^{\lambda}_{\mathcal{A}_t}$ to be an EPR state (instead of a fixed basis vector) and also conclude that $F^{-1}\,|\lambda,X,Y\rangle\in\mathrm{Im}(\Theta_t\Omega_{\mathcal{A}_t})$.

Next, we prove that $\mathrm{Im}(F\cdot\Theta_t)\subset\mathsf{Tab}_t$ as well as Eq. (5.14). It suffices to prove both of these claims on basis vectors $|\vec{x},\vec{y}\rangle\in\mathsf{Path}_t$. Expand the Schur transform as

$$\mathsf{Schur}_t\,|\vec{x}\rangle=\sum_{\lambda,X,T}S_{\vec{x}}^{\lambda,X,T}\,|\lambda,X,T\rangle\,,$$

where

$$X\in\mathcal{B}(V_G^\lambda),\qquad T\in\mathcal{B}(V_{\mathcal{A}_t}^\lambda).$$

Since Schur-Weyl duality exactly tells us that the representation $\rho^{\otimes t}$ decomposes as

$$\rho^{\otimes t}(g)=\bigoplus_\lambda\rho_\lambda(g)\otimes\mathbf{I}_{V^\lambda_{\mathcal{A}_t}}\,,$$

we have that

$$\langle\vec{y}|\,\rho^{\otimes t}(g)\,|\vec{x}\rangle=\sum_{\lambda,X,Y,T}\overline{S_{\vec{y}}^{\lambda,Y,T}}S_{\vec{x}}^{\lambda,X,T}\,\langle Y|\,\rho_\lambda(g)\,|X\rangle\,.$$

Therefore, by the definition of the Fourier transform,

$$F\Theta_t\,|\vec{x},\vec{y}\rangle=\sum_{\lambda,X,Y,T}\frac{\overline{S_{\vec{y}}^{\lambda,Y,T}}S_{\vec{x}}^{\lambda,X,T}}{\sqrt{\dim(V_G^\lambda)}}\,|\lambda,X,Y\rangle\,.$$

Since the coefficient of $|\lambda,X,Y\rangle$ is nonzero only if both $|\lambda,X,T\rangle$ and $|\lambda,Y,T\rangle$ are in the image of the Schur transform for some T, this tells us that $F\Theta_t\,|\vec{x},\vec{y}\rangle\in\mathsf{Tab}_t$.

Next, applying $\Gamma_t=(\mathsf{Schur}_t^\dagger\otimes\mathsf{Schur}_t^\top)E_t$ gives

$$\Gamma_tF\Theta_t\,|\vec{x},\vec{y}\rangle=\sum_{\lambda,X,Y,T,R}\frac{\overline{S_{\vec{y}}^{\lambda,Y,T}}S_{\vec{x}}^{\lambda,X,T}}{\sqrt{\dim(V_G^\lambda)\dim(V_{\mathcal{A}_t}^\lambda)}}\Big(\mathsf{Schur}_t^\dagger\otimes\mathsf{Schur}_t^\top\Big)\left(|\lambda,X,R\rangle\otimes|\lambda,Y,R\rangle\right).\qquad(5.17)$$

On the other hand,

$$\left(\mathsf{Schur}_t\otimes\overline{\mathsf{Schur}_t}\right)|\vec{x},\vec{y}\rangle=\sum_{\lambda,\lambda',X,Y,T,R}S_{\vec{x}}^{\lambda,X,T}\overline{S_{\vec{y}}^{\lambda',Y,R}}\,|\lambda,X,T\rangle\otimes\left|\lambda',Y,R\right\rangle.$$

Projecting onto the EPR subspace gives

$$\Omega_{\mathcal{A}_t}\,|\vec{x},\vec{y}\rangle=\Big(\mathsf{Schur}_t^\dagger\otimes\mathsf{Schur}_t^\top\Big)\sum_{\lambda,X,Y,T,R}\frac{S_{\vec{x}}^{\lambda,X,T}\overline{S_{\vec{y}}^{\lambda,Y,T}}}{\dim(V_{\mathcal{A}_t}^\lambda)}\,|\lambda,X,R\rangle\otimes|\lambda,Y,R\rangle\,.$$

The operator W_t multiplies the λ-subspace by

$$\sqrt{\frac{\dim(V_{\mathcal{A}_t}^{\lambda})}{\dim(V_G^{\lambda})}}.$$

Hence

$$W_t\Omega_{\mathcal{A}_t}\left|\vec{x},\vec{y}\right\rangle = \sum_{\lambda,X,Y,T,R} \frac{\overline{S_{\vec{y}}^{\lambda,Y,T}} S_{\vec{x}}^{\lambda,X,T}}{\sqrt{\dim(V_G^{\lambda})\dim(V_{\mathcal{A}_t}^{\lambda})}} \left(\mathsf{Schur}_t^{\dagger} \otimes \mathsf{Schur}_t^{\top}\right)\left(|\lambda,X,R\rangle \otimes |\lambda,Y,R\rangle\right). \tag{5.18}$$

We conclude that

$$\Gamma_t F \Theta_t \left|\vec{x},\vec{y}\right\rangle = W_t\Omega_{\mathcal{A}_t}\left|\vec{x},\vec{y}\right\rangle,$$

as desired. □

5.5 Proof of Theorem 5.3

To prove Theorem 5.3, we derive an expression that is simultaneously equal to both $(\mathbf{I}_V \otimes \Gamma_{t+1}F)Q_\rho$ and $Q_{\rho,t+1}^{\mathrm{Path}}(\mathbf{I}_V \otimes \Gamma_t F)$.

Beginning with the former, we claim the following operator identity holds on inputs $|x\rangle \otimes |\vec{x}\rangle\,|\vec{y}\rangle$.

Claim 5.6.

$$Q_\rho \cdot (|x\rangle \otimes \Theta_t) = \sum_y |y\rangle \otimes \Theta_{t+1}\,\mathrm{App}_{x,y}^{(t)}. \tag{5.19}$$

To see this, we simply observe that the left-hand side is equal to

$$Q_\rho\left(|x\rangle \otimes \int_G \langle\vec{y}|\,\rho^{\otimes t}(g)\,|\vec{x}\rangle\,|g\rangle\;dg\right) = \sum_{y\in\mathcal{B}(V)} |y\rangle \otimes \int_G \langle y|\,\rho(g)\,|x\rangle\,\langle\vec{y}|\,\rho^{\otimes t}(g)\,|\vec{x}\rangle\,|g\rangle\;dg.$$

and the second tensor factor exactly contains the $|\vec{x},x,\vec{y},y\rangle$ matrix element function.

Applying $\mathbf{I}_V \otimes \Gamma_{t+1}F$ to the recording register and invoking Eq. (5.14), we see that

$$\begin{aligned}(\mathbf{I}_V \otimes \Gamma_{t+1}F)Q_\rho\,(|x\rangle \otimes \Theta_t) &= \sum_y |y\rangle \otimes \Gamma_{t+1}F\Theta_{t+1}\,\mathrm{App}_{x,y}^{(t)} \\ &= \sum_y |y\rangle \otimes W_{t+1}\Omega_{\mathcal{A}_{t+1}}\,\mathrm{App}_{x,y}^{(t)}.\end{aligned}$$

Next, for every $x \in V$, we calculate (again invoking Eq. (5.14))

$$\begin{aligned}Q_{\rho,t+1}^{\mathrm{Path}}(\mathbf{I}_V \otimes \Gamma_t F)\,(|x\rangle \otimes \Theta_t) &= Q_{\rho,t+1}^{\mathrm{Path}}\,(|x\rangle \otimes W_t\Omega_{\mathcal{A}_t}) &(5.20)\\ &= \sum_y |y\rangle \otimes W_{t+1}\Omega_{\mathcal{A}_{t+1}}\,\mathrm{App}_{x,y}^{(t)}\,W_t^{+}\Omega_{\mathcal{A}_t}W_t\Omega_{\mathcal{A}_t} &(5.21)\\ &= \sum_y |y\rangle \otimes W_{t+1}\Omega_{\mathcal{A}_{t+1}}\,\mathrm{App}_{x,y}^{(t)}\,\Omega_{\mathcal{A}_t}, &(5.22)\end{aligned}$$

where the last equation holds because $\Omega_{\mathcal{A}_t}$ and W_t^{+} commute and $W_t^{+}W_t\Omega_{\mathcal{A}_t} = \Omega_{\mathcal{A}_t}$.

Inserting an additional $\Omega_{\mathcal{A}_t}$ on the right, we conclude that

$$Q^{\mathrm{Path}}_{\rho,t+1}(\mathbf{I}_V \otimes \Gamma_t F)(\mathbf{I}_V \otimes \Theta_t \Omega_{\mathcal{A}_t}) = (\mathbf{I}_V \otimes \Gamma_{t+1} F) Q_\rho (\mathbf{I}_V \otimes \Theta_t \Omega_{\mathcal{A}_t}),$$

and thus

$$Q^{\mathrm{Path}}_{\rho,t+1}(\mathbf{I}_V \otimes \Gamma_t F) = (\mathbf{I}_V \otimes \Gamma_{t+1} F) Q_\rho$$

when restricted to inputs in $V \otimes \mathrm{Im}(\Theta_t \Omega_{\mathcal{A}_t}) = V \otimes F^{-1}(\mathsf{Tab}_t)$, which is the first claim of Theorem 5.3.

Quantum query algorithms. To obtain the second claim of Theorem 5.3, we note that all recording states obtainable by making t queries to Q_ρ lie in $F^{-1}(\mathsf{Tab}_t)$.

Lemma 5.7 (Reachable states are in $\mathrm{Im}(\Theta_t)$)**.** *Fix any adaptive quantum algorithm whose query register is V and queries Q_ρ. Then, its recording register always remains in $\mathrm{Im}(\Theta_t) = F^{-1}(\mathsf{Tab}_t)$.*

Proof. This follows immediately from Eq. (5.19) and an induction on t. □

Finally, we observe that in this context, the one-query identity implies the t-query statement by induction. In particular, suppose that immediately before the $(t+1)$-st query, the path-recorded simulation state is obtained from the uncompressed representation oracle by applying

$$\mathbf{I} \otimes \Gamma_t F$$

to the recording register. The adversary's next local unitary commutes with the recording-register isometry. The next oracle call is handled by the identity

$$Q^{\mathrm{Path}}_{\rho,t+1}(\mathbf{I}_V \otimes \Gamma_t F) = (\mathbf{I}_V \otimes \Gamma_{t+1} F) Q_\rho.$$

Therefore the invariant is preserved with $t+1$ in place of t.

Starting from the trivial depth-0 recording register, induction over the t queries gives

$$A^{Q_\rho}(|\psi\rangle) \simeq A^{Q^{\mathrm{Path}}_\rho}(|\psi\rangle)$$

via the final isometry $\Gamma_t F$. This proves the second claim of Theorem 5.3.

5.6 Transpose queries

The same argument also gives a path-recording update for transpose queries. Here transpose is taken with respect to the fixed basis $\mathcal{B}(V)$.

Theorem 5.8 (Transpose path-recording update)**.** *Let $Q_{\rho^\top}$ denote the purified oracle which applies $\rho(g)^\top$ to the adversary's query register:*

$$Q_{\rho^\top}\left(|x\rangle \otimes \int_G f(g)|g\rangle\, dg\right) = \int_G f(g)\, \rho(g)^\top |x\rangle \otimes |g\rangle\, dg.$$

Then, for every $x \in \mathcal{B}(V)$,

$$Q_{\rho^\top} \cdot (|x\rangle \otimes \Theta_t) = \sum_{y \in \mathcal{B}(V)} |y\rangle \otimes \Theta_{t+1}\, \mathrm{App}^{(t)}_{y,x}\,.$$

Consequently, if we define

$$Q^{\mathrm{Path},\top}_{\rho,t+1} := \sum_{x,y \in \mathcal{B}(V)} |y\rangle\langle x| \otimes W_{t+1} \Omega_{\mathcal{A}_{t+1}}\, \mathrm{App}^{(t)}_{y,x}\, W_t^+ \Omega_{\mathcal{A}_t},$$

then

$$Q_{\rho,t+1}^{\mathrm{Path},\top} \cdot (\mathbf{I}_V \otimes \Gamma_t F) = (\mathbf{I}_V \otimes \Gamma_{t+1} F) \cdot Q_{\rho^\top}$$

on inputs restricted to $V \otimes F^{-1}(\mathsf{Tab}_t)$.

Proof. First we prove the matrix-element identity. By definition,

$$\Theta_t |\vec{x}, \vec{y}\rangle = \int_G \langle \vec{y}|\rho(g)^{\otimes t}|\vec{x}\rangle |g\rangle \, dg.$$

Therefore

$$Q_{\rho^\top}\left(|x\rangle \otimes \Theta_t |\vec{x}, \vec{y}\rangle\right) = \sum_{y \in \mathcal{B}(V)} |y\rangle \otimes \int_G \langle y|\rho(g)^\top |x\rangle \langle \vec{y}|\rho(g)^{\otimes t}|\vec{x}\rangle |g\rangle \, dg \tag{5.23}$$

$$= \sum_{y \in \mathcal{B}(V)} |y\rangle \otimes \int_G \langle x|\rho(g)|y\rangle \langle \vec{y}|\rho(g)^{\otimes t}|\vec{x}\rangle |g\rangle \, dg \tag{5.24}$$

$$= \sum_{y \in \mathcal{B}(V)} |y\rangle \otimes \Theta_{t+1} \operatorname{App}_{y,x}^{(t)} |\vec{x}, \vec{y}\rangle. \tag{5.25}$$

Thus a transpose query appends the reversed pair (y, x). We now prove the path-recording identity. By Eq. (5.14),

$$\Gamma_t F \Theta_t = W_t \Omega_{\mathcal{A}_t}, \qquad \Gamma_{t+1} F \Theta_{t+1} = W_{t+1} \Omega_{\mathcal{A}_{t+1}}.$$

Applying $\mathbf{I}_V \otimes \Gamma_{t+1} F$ to the transpose matrix-element identity gives

$$(\mathbf{I}_V \otimes \Gamma_{t+1} F) Q_{\rho^\top}(|x\rangle \otimes \Theta_t) = \sum_y |y\rangle \otimes W_{t+1} \Omega_{\mathcal{A}_{t+1}} \operatorname{App}_{y,x}^{(t)}.$$

On the other hand,

$$Q_{\rho,t+1}^{\mathrm{Path},\top}(|x\rangle \otimes W_t \Omega_{\mathcal{A}_t}) = \sum_y |y\rangle \otimes W_{t+1} \Omega_{\mathcal{A}_{t+1}} \operatorname{App}_{y,x}^{(t)} W_t^+ \Omega_{\mathcal{A}_t} W_t \Omega_{\mathcal{A}_t} \tag{5.26}$$

$$= \sum_y |y\rangle \otimes W_{t+1} \Omega_{\mathcal{A}_{t+1}} \operatorname{App}_{y,x}^{(t)} \Omega_{\mathcal{A}_t}, \tag{5.27}$$

using $W_t^+ W_t \Omega_{\mathcal{A}_t} = \Omega_{\mathcal{A}_t}$. Therefore the two expressions agree when restricted to inputs in

$$V \otimes \operatorname{Im}(\Theta_t \Omega_{\mathcal{A}_t}) = V \otimes F^{-1}(\mathsf{Tab}_t),$$

which proves the theorem. □

Thus arbitrary adaptive algorithms making forward and transpose queries are simulated by using $\operatorname{App}_{x,y}^{(t)}$ for a forward query and $\operatorname{App}_{y,x}^{(t)}$ for a transpose query.

5.7 Schur Orthogonality for the Commutant Algebras

Our next step is to *interpret* the commutant EPR projector $\Omega_{\mathcal{A}_t}$. To do this, we first show that whenever the commutant algebra is semisimple, it satisfies the Schur orthogonality relations for the matrix elements of its simple modules.

Theorem 5.9 (Schur Orthogonality)**.** *Let $\mathcal{A}$ be a semisimple algebra, and let $\mathcal{B}(\mathcal{A})$ be a basis for $\mathcal{A}$. For $a \in \mathcal{A}$, let $\mathrm{tr}(a)$ be the trace of the regular representation of $\mathcal{A}$. For each $a \in \mathcal{A}$, let a^* be its dual element under this trace. Let $\lambda, \mu \in \Lambda(\mathcal{A})$ be the labels of two simple modules, $V_{\mathcal{A}}^{\lambda}$ and $V_{\mathcal{A}}^{\mu}$, of $\mathcal{A}$. Let $|a\rangle, |b\rangle \in \mathcal{B}(V_{\mathcal{A}}^{\lambda})$ and $|c\rangle, |d\rangle \in \mathcal{B}(V_{\mathcal{A}}^{\mu})$. Then*[43]

$$\sum_{\sigma \in \mathcal{B}(\mathcal{A})} \langle a| \rho_\lambda(\sigma) |b\rangle \ \langle c| \rho_\mu(\sigma^*) |d\rangle = \frac{1}{\dim(V_{\mathcal{A}}^{\lambda})} \delta_{\lambda\mu} \delta_{ad} \delta_{bc} \tag{5.28}$$

Proof. Consider the matrix

$$T_{bc} \coloneqq \sum_{\sigma \in \mathcal{B}(\mathcal{A})} \rho_\lambda(\sigma) \cdot |b\rangle\langle c| \cdot \rho_\mu(\sigma^*).$$

Our goal is to show that

$$T_{bc} = \begin{cases} \frac{1}{\dim(V_{\mathcal{A}}^{\lambda})} \mathbf{I}_{V_{\mathcal{A}}^{\lambda}} & \text{if } \lambda = \mu \text{ and } b = c, \\ 0 & \text{otherwise.} \end{cases}$$

If we can show this, then

$$\sum_{\sigma \in \mathcal{B}(\mathcal{A})} \langle a| \rho_\lambda(\sigma) |b\rangle \ \langle c| \rho_\mu(\sigma^*) |d\rangle = \langle a| T_{bc} |d\rangle$$

is equal to $\frac{1}{\dim(V_{\mathcal{A}}^{\lambda})} \langle a| \mathbf{I}_{M_\lambda} |d\rangle = \frac{1}{\dim(V_{\mathcal{A}}^{\lambda})} \langle a|d\rangle = \frac{1}{\dim(V_{\mathcal{A}}^{\lambda})} \delta_{ad}$ if $\lambda = \mu$ and $b = c$ and is equal to 0 otherwise. This gives the desired equality.

To prove this, we observe that by [HR05, Proposition 5.2(b)], applied with $M = V_{\mathcal{A}}^{\mu}$, $N = V_{\mathcal{A}}^{\lambda}$, and $\phi = |b\rangle\langle c|$, the map $T_{b,c}$ is a $\mathcal{A}$-module homomorphism, meaning that for every $a \in \mathcal{A}$,

$$T_{bc} \rho_\mu(a) = \rho_\lambda(a) T_{bc}.$$

We now consider two cases.

Case 1: $\lambda \neq \mu$**.** If $\lambda \neq \mu$, Schur's lemma ([HR05, Proposition 5.3(2)]) states that $T_{bc} = 0$.

Case 2: $\lambda = \mu$**.** In this case, Schur's lemma ([HR05, Proposition 5.3(1)]) states that $T_{bc} = \alpha \cdot \mathbf{I}_{V_{\mathcal{A}}^{\lambda}}$ for some $\alpha \in \mathbb{C}$. Taking the (linear algebraic) trace of both sides, we have that

$$\mathrm{tr}(T_{bc}) = \alpha \cdot \dim(V_{\mathcal{A}}^{\lambda}),$$

from which we see that $\alpha = \mathrm{tr}(T_{bc}) / \dim(V_{\mathcal{A}}^{\lambda})$. It therefore suffices to compute the trace of T_{bc}.

To compute this, we write

$$\begin{aligned} \mathrm{tr}(T_{bc}) &= \sum_{\sigma \in \mathcal{B}(\mathcal{A})} \mathrm{tr}(\rho_\lambda(\sigma) \cdot |b\rangle\langle c| \cdot \rho_\lambda(\sigma^*)) \\ &= \sum_{\sigma \in \mathcal{B}(\mathcal{A})} \mathrm{tr}(|b\rangle\langle c| \cdot \rho_\lambda(\sigma^*) \cdot \rho_\lambda(\sigma)) \\ &= \mathrm{tr}\left(|b\rangle\langle c| \cdot \rho_\lambda\Big(\sum_{\sigma \in \mathcal{B}(\mathcal{A})} \sigma^* \sigma \Big) \right) \\ &= \mathrm{tr}(|b\rangle\langle c|) \\ &= \delta_{bc}, \end{aligned}$$

[43] Note that compared with the Schur orthogonality theorem for irrep matrix elements of *groups*, the second matrix element is transposed rather than complex conjugated. As we mention below in Remark 5.10, this is because for groups, σ^* is proportional to its inverse, which of course conjugate-transposes the irrep.

where we have invoked the identity $\sum_{\sigma\in\mathcal{B}(\mathcal{A})}\sigma^*\sigma = 1$ [HR05, Proof of Theorem 5.8]. □

An alternative proof of Theorem 5.9 also appears in [FND26, Corollary 2.12].

Remark 5.10. Note that this generalizes the Schur orthogonality for groups, which states that

$$\frac{1}{|G|}\sum_{g\in G}\langle a|\,\rho_\lambda(g)\,|b\rangle\ \overline{\langle d|\,\rho_\mu(g)\,|c\rangle} = \sum_{g\in\mathcal{B}(\mathbb{C}[G])}\langle a|\,\rho_\lambda(g)\,|b\rangle\ \langle c|\,\rho_\mu\left(\frac{1}{|G|}g^{-1}\right)|d\rangle = \frac{1}{\dim(V_{\mathcal{A}}^\lambda)}\delta_{\lambda\mu}\delta_{ad}\delta_{bc} \tag{5.29}$$

since note that whenever $\mathcal{A}$ is a group algebra $\mathbb{C}[G]$, then we have that, $g^* = \frac{1}{|G|}g^{-1}\quad\forall g\in G$ (since $\frac{1}{|G|}\operatorname{tr}\left(h^{-1}g\right) = \delta_{hg}$).

5.8 Interpretation of the Commutant EPR projector

Lemma 5.11. *Let $\rho : G\to\operatorname{End}(V)$ be a representation, and consider the tensor product representation $\rho^{\otimes t}$ on $V^{\otimes t}$. Let $\mathcal{A}_t$ denote a semisimple algebra that is isomorphic to the commutant algebra $\operatorname{End}_G(V^{\otimes t})$; for any $a\in\mathcal{A}_t$, let $S(a)\in\operatorname{End}_G(V^{\otimes t})$ denote its action.*

Finally, let Schur *be the corresponding generalized Schur transform. Let $V_{\mathcal{A}_t}^\lambda$ and V_G^λ be the corresponding simple $\mathcal{A}_t$-modules and G-representations with bases $\mathcal{B}(V_{\mathcal{A}_t}^\lambda)$ and $\mathcal{B}(V_G^\lambda)$ implied by the Schur transform.*

Then,

$$\Omega_{\mathcal{A}_t} = \sum_{\sigma\in\mathcal{B}(\mathcal{A}_t)} S(\sigma)\otimes S(\sigma^*)^T. \tag{5.30}$$

Proof. We calculate

$$\left(\mathsf{Schur}\otimes\overline{\mathsf{Schur}}\right)\left(\sum_{\sigma\in\mathcal{B}(\mathcal{A}_t)} S(\sigma)\otimes S(\sigma^*)^T\right)\left(\mathsf{Schur}^\dagger\otimes\mathsf{Schur}^T\right) \tag{5.31}$$

$$= \sum_{\lambda,\mu,\in\widehat{\mathcal{A}}_t} |\lambda,\mu\rangle\langle\lambda,\mu|\otimes\left(\sum_{\sigma\in\mathcal{B}(\mathcal{A}_t)}\rho_\lambda(\sigma)\otimes\rho_\mu(\sigma^*)^T\right)\otimes\mathbf{I}_{V_G^\lambda}\otimes\mathbf{I}_{V_G^\mu}$$

(definition of Schur transform (Definition 3.22))

$$= \sum_{\lambda,\mu,\in\widehat{\mathcal{A}}_t} |\lambda,\mu\rangle\langle\lambda,\mu|\otimes\left(\sum_{|a\rangle,|b\rangle|c\rangle,|d\rangle,\in\mathcal{B}(V_{\mathcal{A}_t}^\lambda)}|a\rangle\langle b|\otimes|d\rangle\langle c|\right)\otimes\mathbf{I}_{V_G^\lambda\otimes V_G^\mu}\cdot\sum_{\sigma\in\mathcal{B}(\mathcal{A}_t)}\langle a|\,\rho_\lambda(\sigma)\,|b\rangle\cdot\langle d|\,\rho_\mu(\sigma^*)^T\,|c\rangle$$

($\sum_a |a\rangle\langle a| = \mathbf{I}$)

$$= \sum_{\lambda,\mu,\in\widehat{\mathcal{A}}_t} |\lambda,\mu\rangle\langle\lambda,\mu|\otimes\left(\sum_{|a\rangle,|b\rangle|c\rangle,|d\rangle,\in\mathcal{B}(V_{\mathcal{A}_t}^\lambda)}|a\rangle\langle b|\otimes|d\rangle\langle c|\right)\otimes\mathbf{I}_{V_G^\lambda}\otimes\mathbf{I}_{V_G^\mu}\cdot\frac{1}{\dim(V_\lambda)}\delta_{\lambda\mu}\delta_{ad}\delta_{bc}$$

(Theorem 5.9)

$$= \sum_{\lambda \in \widehat{\mathcal{A}}_t} \frac{1}{\dim(V_{\mathcal{A}_t}^\lambda)} |\lambda, \lambda\rangle\langle\lambda, \lambda| \otimes \left(\sum_{|a\rangle, |b\rangle \in \mathcal{B}(V_{\mathcal{A}_t}^\lambda)} |a\rangle\langle b| \otimes |a\rangle\langle b| \right) \otimes \mathbf{I}_{V_G^\lambda} \otimes \mathbf{I}_{V_G^\mu} \tag{5.32}$$

$$= \sum_{\substack{\lambda \in \widehat{\mathcal{A}}_t \\ |a\rangle, |b\rangle \in \mathcal{B}(V_{\mathcal{A}_t}^\lambda)}} \frac{1}{\dim(V_{\mathcal{A}_t}^\lambda)} |\lambda, \lambda\rangle\langle\lambda, \lambda| \otimes |a, a\rangle\langle b, b| \otimes \mathbf{I}_{V_G^\lambda} \otimes \mathbf{I}_{V_G^\mu} \tag{5.33}$$

$$= \left(\mathsf{Schur} \otimes \overline{\mathsf{Schur}}\right) \Omega_{\mathcal{A}_t} \left(\mathsf{Schur}^\dagger \otimes \mathsf{Schur}^T\right) \qquad \square$$

Example 5.12 (Group Algebra Commutant EPR Projector)**.** In the case of $\mathcal{A}_t = \mathbb{C}[H]$ for some finite group H, we have that $[h]^* = \frac{1}{|H|}[h^{-1}]$ for all $h \in H$ (Example 3.18); since $S([h^{-1}])^T = \overline{S([h])}$ for all h, this means that the commutant EPR projector formula simplifies to

$$\Omega_{\mathbb{C}[H]} = \frac{1}{|H|} \sum_{h \in H} S([h]) \otimes \overline{S([h])}.$$

This further simplifies to

$$\Omega_{\mathbb{C}[H]} = \frac{1}{|H|} \sum_{h \in H} S([h]) \otimes S([h])$$

whenever the Schur representation $S(\cdot)$ has real entries, such as in the case of $H = S_t$.

5.9 Adaptive Queries to the Update Rule

In this section, we derive an equivalent form for the general update rule in Section 5.1, which we will use for the main argument in Section 7.5. Specifically, we are interested in analyzing (potentially adaptive) quantum query algorithms that have oracle access to one of two different isometries:

- The path basis update Q_ρ^{Path}, or
- $Q_\rho^{\text{Path}} \cdot C$, which applies Q_ρ^{Path} after first applying a unitary C to the adversary's query register.

We will use A, B, and R to denote the query, workspace, and database registers as defined in Definition 3.5. We will also subdivide R as

$$\mathsf{R} = (\mathsf{R}_0, \mathsf{R}_1 = (\mathsf{R}_{\mathsf{X},1}, \mathsf{R}_{\mathsf{Y},1}), \mathsf{R}_2 = (\mathsf{R}_{\mathsf{X},2}, \mathsf{R}_{\mathsf{Y},2}), \ldots, \mathsf{R}_{\mathsf{t}} = (\mathsf{R}_{\mathsf{X},\mathsf{t}}, \mathsf{R}_{\mathsf{Y},\mathsf{t}})) \tag{5.34}$$

where R_{i} denotes the register to which the update operator adds the ith query, and $\mathsf{R}_{\mathsf{X},\mathsf{i}}$ and $\mathsf{R}_{\mathsf{Y},\mathsf{i}}$ denote the registers for recording the input and output for the i'th query. We also write $\mathsf{R}_{\mathsf{X}} := (\mathsf{R}_{\mathsf{X},1}, \mathsf{R}_{\mathsf{X},2}, \ldots \mathsf{R}_{\mathsf{X},\mathsf{t}})$ and $\mathsf{R}_{\mathsf{Y}} := (\mathsf{R}_{\mathsf{Y},1}, \mathsf{R}_{\mathsf{Y},2}, \ldots \mathsf{R}_{\mathsf{Y},\mathsf{t}})$. We will also use R_0 to denote register containing the initial database, before making any queries (which we will assume to be properly formatted as a list of pairs of registers like the rest of R). Most often, we will take R_0 to be the empty (0 qubit) register (with associated Hilbert space $\mathbb{C}^{2^0} = \mathbb{C}$). We write such an empty register as $|\emptyset\rangle_{\mathsf{R}}$. However it can be non-empty if, for example, we are not starting with a clean slate and a previous algorithm has already made queries. We will also write $\mathsf{R}_{\leq i} := (\mathsf{R}_0, \mathsf{R}_1, \ldots, \mathsf{R}_i)$.

We will use $\mathcal{A}_i$ to denote the commutant algebra of the k'th tensor product of the group representation, where $k = |\mathsf{R}_{\leq i}| = i + |\mathsf{R}_0|$ (the size of R_0 being the number of pairs in the initial database). When the database is empty, then this is the commutant algebra of the trivial representation of the group, which is simply the one-dimensional algebra $\mathbb{C}$.

Notation for the Update Rule. We will frequently drop the subscript t from the Schur transform Schur_t when the tensor power is unambiguous. For ease of working with the update rule from Section 5.1, we will write it in terms of the following operators:

Let

$$\Pi_{\lambda,i} := \mathsf{Schur}^\dagger \Big(|\lambda\rangle\langle\lambda| \otimes \mathbf{I}_{V^\lambda_{\mathcal{A}_i}} \otimes \mathbf{I}_{V^\lambda_G} \Big) \mathsf{Schur} \otimes \mathbf{I}_{\overline{V}^{\otimes t}} \,, \tag{5.35}$$

$$\Lambda_{\lambda,i} := \beta_{\lambda,i}\, \Pi_{\lambda,i} \,, \qquad \Lambda^+_{\lambda,i} := \beta^{-1}_{\lambda,i}\, \Pi_{\lambda,i} \,, \qquad \text{where } \beta_{\lambda,i} := \sqrt{\frac{\dim(V^\lambda_{\mathcal{A}_i})}{\dim(V^\lambda_G)}} \,, \tag{5.36}$$

$$\Lambda_{\mathcal{A}_i} := \sum_{\lambda \in \widehat{\mathcal{A}}_i} \Lambda_{\lambda,i} \,, \qquad \Lambda^+_{\mathcal{A}_i} := \sum_{\lambda \in \widehat{\mathcal{A}}_i} \Lambda^+_{\lambda,i} \tag{5.37}$$

$$\Omega_{\mathcal{A}_i} := \Big(\mathsf{Schur}^\dagger \otimes \mathsf{Schur}^T \Big) \left(\sum_{\substack{\lambda \in \widehat{\mathcal{A}}_i \\ |a\rangle, |b\rangle \in \mathcal{B}(V^\lambda_{\mathcal{A}_i})}} \frac{1}{\dim(V^\lambda_{\mathcal{A}_i})} |\lambda,\lambda\rangle\langle\lambda,\lambda| \otimes |a,a\rangle\langle b,b| \otimes \mathbf{I}^{\otimes 2}_{V^\lambda_G} \right) \big(\mathsf{Schur} \otimes \overline{\mathsf{Schur}} \big) \tag{5.38}$$

$$= \sum_{\sigma \in \mathcal{B}(\mathcal{A}_i)} S(\sigma) \otimes S(\sigma^*)^T \,. \qquad \text{(by Lemma 5.11, where } S \text{ is the Schur representation)}$$

$$\mathrm{App}_{\mathsf{A}\,\mathsf{R}_i} := \sum_{x,y \in [N]} (|y\rangle\langle x|)_{\mathsf{A}} \otimes |x,y\rangle_{\mathsf{R}_i} \tag{5.39}$$

When the database register is of size $i-1$, we can write the path recording update rule (Remark 5.2), when restricted to input databases in $\mathrm{Im}(\Omega_{\mathcal{A}_{i-1}})$, as

$$Q^{\mathrm{Path}}_\rho = V := \Big(\Lambda_{\mathcal{A}_i} \Big)_{\mathsf{R}_{\le i}} \Big(\Omega_{\mathcal{A}_i} \Big)_{\mathsf{R}_{\le i}} \Big(\mathrm{App} \Big)_{\mathsf{A}\,\mathsf{R}_i} \Big(\Lambda^+_{\mathcal{A}_{i-1}} \Big)_{\mathsf{R}_{\le i-1}} \tag{5.40}$$

Rewriting the Update Operator for Adaptive Queries.

Lemma 5.13 (Telescoping product of irrep ratio reweighting)**.** *Let V be the isometry defined in Eq. (5.40), and define the reduced update operator (which without the rescaling is not an isometry) to be*

$$\widetilde{V} := \Big(\Omega_{\mathcal{A}_i} \Big)_{\mathsf{R}_{\le i}} \Big(\mathrm{App} \Big)_{\mathsf{A}\mathsf{R}_i} \tag{5.41}$$

$$= \sum_{x_i, y_i \in [N]} |y_i\rangle\langle x_i|_{\mathsf{A}} \otimes \Omega_{\mathcal{A}_i} \Big(\mathbf{I}_{\mathsf{R}_{<i}} \otimes |x_i, y_i\rangle_{\mathsf{R}_i} \Big) \,. \tag{5.42}$$

whenever the database register R is of size $i-1$. Let Adv_t be any t-query algorithm with workspace register A, ancilla register B, and that makes queries to the update rule V on workspace register A and database register R. Then

$$\mathsf{Adv}^V_t = \Big(\Lambda_{\mathcal{A}_t} \Big)_{\mathsf{R}_{\le t}} \mathsf{Adv}^{\widetilde{V}}_t \Big(\Lambda^+_{\mathcal{A}_0} \Big)_{\mathsf{R}_0} \tag{5.43}$$

In other words, if we make t queries to the path recording oracle, then this is the same as if we were to apply the path recording update rule without the reweighting, and then apply a reweighting just once at the end.[44]

[44] As described in Remark 5.19, when beginning with an empty database, the first reweighting operator always occurs on the one-dimensional irreps corresponding to $|\lambda| = 0$, and is thus just the identity operator.

Proof. Let $(A_1)_{\mathsf{AB}}, \dots, (A_t)_{\mathsf{AB}}$ be the unitaries that Adv applies in between queries. We can write

$$\mathsf{Adv}_t^V \tag{5.44}$$

$$= \prod_{i=1}^{t} (V_{\mathsf{AR}} \cdot A_{i\,\mathsf{AB}}) \tag{5.45}$$

$$= \prod_{i=1}^{t} \left(\left(\Lambda_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\mathrm{App}\right)_{\mathsf{AR}_i} \left(\Lambda^+_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\le i-1}} (A_i)_{\mathsf{AB}} \right) \qquad \text{(by the definition of } V\text{)}$$

$$= \prod_{i=1}^{t} \left(\left(\sum_{\lambda^+ \in \widehat{\mathcal{A}}_i} \beta_{\lambda^+,i}\, \Pi_{\lambda^+,i} \right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\mathrm{App}\right)_{\mathsf{AR}_i} \left(\sum_{\lambda \in \widehat{\mathcal{A}}_{i-1}} \beta^{-1}_{\lambda,i-1}\, \Pi_{\lambda,i-1} \right)_{\mathsf{R}_{\le \mathrm{i}-1}} (A_i)_{\mathsf{AB}} \right)$$

(by the definition of $\Lambda_{\mathcal{A}_i}$)

$$= \sum_{\substack{\lambda_1,\dots,\lambda_t \\ \lambda_1^+,\dots,\lambda_t^+ \\ \lambda_i \in \widehat{\mathcal{A}}_{i-1},\, \lambda_i^+ \in \widehat{\mathcal{A}}_i}} \prod_{i=1}^{t} \left(\left(\beta_{\lambda_i^+,i}\, \Pi_{\lambda_i^+,i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\mathrm{App}\right)_{\mathsf{AR}_i} \left(\beta^{-1}_{\lambda_i,i-1}\, \Pi_{\lambda_i,i-1}\right)_{\mathsf{R}_{\le \mathrm{i}-1}} (A_i)_{\mathsf{AB}} \right) \tag{5.46}$$

$$= \sum_{\substack{\lambda_1,\dots,\lambda_t \\ \lambda_1^+,\dots,\lambda_t^+ \\ \lambda_i \in \widehat{\mathcal{A}}_{i-1},\, \lambda_i^+ \in \widehat{\mathcal{A}}_i}} \beta_{\lambda_t^+,t} \prod_{i=1}^{t} \left(\left(\beta^{-1}_{\lambda_{i+1},i}\, \Pi_{\lambda_{i+1},i}\right)_{\mathsf{R}_{\le i}} \left(\beta_{\lambda_i^+,i}\, \Pi_{\lambda_i^+,i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\mathrm{App}\right)_{\mathsf{AR}_i} (A_i)_{\mathsf{AB}} \right) \left(\beta^{-1}_{\lambda_1,0}\, \Pi_{\lambda_1,0}\right)_{\mathsf{R}_0}$$

(by moving each $\beta^{-1}_{\lambda_i,i-1}\, \Pi_{\lambda_i,i-1}$ to the previous factor in the product[45])

$$= \sum_{\substack{\lambda_1,\dots,\lambda_t \\ \lambda_1^+,\dots,\lambda_t^+ \\ \lambda_i \in \widehat{\mathcal{A}}_{i-1},\, \lambda_i^+ \in \widehat{\mathcal{A}}_i}} \prod_{i=1}^{t} \left(\frac{\beta_{\lambda_i^+,i}}{\beta_{\lambda_{i+1},i}}\, \left(\Pi_{\lambda_{i+1},i}\right)_{\mathsf{R}_{\le i}} \left(\Pi_{\lambda_i^+,i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\mathrm{App}\right)_{\mathsf{AR}_i} (A_i)_{\mathsf{AB}} \right) (\Pi_{\lambda_1,0})_{\mathsf{R}_0}\, \frac{\beta_{\lambda_t^+,t}}{\beta_{\lambda_1,0}} \tag{5.47}$$

$$= \sum_{\substack{\lambda_1,\dots,\lambda_t \\ \lambda_1^+,\dots,\lambda_t^+ \\ \lambda_i \in \widehat{\mathcal{A}}_{i-1},\, \lambda_i^+ \in \widehat{\mathcal{A}}_i}} \prod_{i=1}^{t} \left(\frac{\beta_{\lambda_i^+,i}}{\beta_{\lambda_{i+1},i}}\, \delta_{\lambda_{i+1},\lambda_i^+}\, \left(\Pi_{\lambda_i^+,i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\mathrm{App}\right)_{\mathsf{AR}_i} (A_i)_{\mathsf{AB}} \right) (\Pi_{\lambda_1,0})_{\mathsf{R}_0}\, \frac{\beta_{\lambda_t^+,t}}{\beta_{\lambda_1,0}}$$

$$\left(\text{since } \Pi_{\lambda_{i+1}} \Pi_{\lambda_i^+} = \delta_{\lambda_{i+1},\lambda_i^+}\, \Pi_{\lambda_i^+}\right)$$

$$= \sum_{\substack{\lambda_1,\lambda_1^+,\dots,\lambda_t^+ \\ \lambda_1 \in \widehat{\mathcal{A}}_0,\, \lambda_i^+ \in \widehat{\mathcal{A}}_i}} \prod_{i=1}^{t} \left(\left(\Pi_{\lambda_i^+,i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\mathrm{App}\right)_{\mathsf{AR}_i} (A_i)_{\mathsf{AB}} \right) (\Pi_{\lambda_1,0})_{\mathsf{R}_0}\, \frac{\beta_{\lambda_t^+,t}}{\beta_{\lambda_1,0}} \tag{5.48}$$

$$= \prod_{i=1}^{t} \left(\left(\sum_{\lambda_i^+ \in \widehat{\mathcal{A}}_i} \Pi_{\lambda_i^+,i} \right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\mathrm{App}\right)_{\mathsf{AR}_i} (A_i)_{\mathsf{AB}} \right) \left(\sum_{\lambda_1 \in \widehat{\mathcal{A}}_0} \Pi_{\lambda_1,0} \right)_{\mathsf{R}_0} \frac{\beta_{\lambda_t^+,t}}{\beta_{\lambda_1,0}} \tag{5.49}$$

$$= \left(\sum_{\lambda_t^+ \in \widehat{\mathcal{A}}_t} \beta_{\lambda_t^+,t}\, \Pi_{\lambda_t^+,t} \right)_{\mathsf{R}_{\le t}} \prod_{i=1}^{t} \left(\left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\mathrm{App}\right)_{\mathsf{AR}_i} (A_i)_{\mathsf{AB}} \right) \left(\sum_{\lambda_1 \in \widehat{\mathcal{A}}_0} \beta^{-1}_{\lambda_1,0}\, \Pi_{\lambda_1,0} \right)_{\mathsf{R}_0}$$

(Using that $\sum_\lambda \Pi_\lambda = \mathbf{I}$ in all except the last factor of the product)

$$= \left(\Lambda_{\mathcal{A}_t}\right)_{\mathsf{R}_{\leq t}} \mathsf{Adv}_t^{\widetilde{V}} \left(\Lambda^+_{\mathcal{A}_0}\right)_{\mathsf{R}_0} \qquad \text{(by the definitions of } \Lambda_{\mathcal{A}_i},\ \widetilde{V})$$

□

Lemma 5.14. *The same holds for queries to $V \cdot C$ for any unitary $C \in U(N)$ acting on the* A *register. That is,*

$$\mathsf{Adv}_t^{V\cdot C} = \left(\Lambda_{\mathcal{A}_t}\right)_{\mathsf{R}_{\leq t}} \mathsf{Adv}_t^{\widetilde{V}\cdot C} \left(\Lambda^+_{\mathcal{A}_0}\right)_{\mathsf{R}_0}. \tag{5.50}$$

The proof is the same as that of Lemma 5.13 above, but where each $A_{i\,\mathsf{AB}}$ has an extra C attached to the left.

Lemma 5.15 (Ricochet Property of the Append Operator)**.** *For any N-dimensional unitary $C \in U(N)$,*

$$\mathrm{App}_{\mathsf{A}\,\mathsf{R_i}} \cdot C_{\mathsf{A}} = C_{\mathsf{R_{X,i}}} \cdot \mathrm{App}_{\mathsf{A}\,\mathsf{R_i}} \tag{5.51}$$

Proof.

$$\begin{aligned}
\mathrm{App}_{\mathsf{A}\,\mathsf{R_i}} \cdot C_{\mathsf{A}} &= \sum_{x,y\in[N]} |y\rangle\langle x|_{\mathsf{A}}\, C_{\mathsf{A}} \otimes |x,y\rangle_{\mathsf{R}_i} && \text{(Definition of } \mathrm{App}_{\mathsf{A}\,\mathsf{R_i}}) \\
&= \sum_{x,y,z\in[N]} |y\rangle_{\mathsf{A}} \langle x|C|z\rangle \langle z|_{\mathsf{A}} \otimes |x,y\rangle_{\mathsf{R}_i} && (\textstyle\sum_z |z\rangle\langle z| = \mathbf{I}) \\
&= \sum_{y,z\in[N]} |y\rangle\langle z|_{\mathsf{A}} \otimes \sum_{x\in[N]} \langle x|C|z\rangle\, |x,y\rangle_{\mathsf{R}_i} && (5.52) \\
&= \sum_{y,z\in[N]} |y\rangle\langle z|_{\mathsf{A}} \otimes \sum_{x\in[N]} |x\rangle_{\mathsf{R_{X,i}}} \langle x|C|z\rangle\, |y\rangle_{\mathsf{R_{Y,i}}} && (5.53) \\
&= \sum_{y,z\in[N]} |y\rangle\langle z|_{\mathsf{A}} \otimes C\,|z\rangle_{\mathsf{R_{X,i}}} \otimes |y\rangle_{\mathsf{R_{Y,i}}} && (\textstyle\sum_x |x\rangle\langle x| = \mathbf{I}) \\
&= C_{\mathsf{R_{X,i}}} \cdot \mathrm{App}_{\mathsf{A}\,\mathsf{R_i}} && \square
\end{aligned}$$

Lemma 5.16 (Alternate Form of $\mathsf{Adv}_t^{\widetilde{V}\cdot C}$)**.**

$$\mathsf{Adv}_t^{\widetilde{V}\cdot C} = \left(\prod_{i=1}^{t} (\Omega_{\mathcal{A}_i})_{\mathsf{R}_{\leq i}}\right) \cdot \left(C^{\otimes t}\right)_{\mathsf{R_X}} \cdot \mathsf{Adv}_t^{\mathrm{App}} \tag{5.54}$$

In other words, the following process is identical to querying $\widetilde{V} \cdot C$: during each query, only the App operator is applied. Then, after all queries have been made, $C^{\otimes t}$ is applied to $\mathsf{R_X}$, the input half of the database in the oracle register, and then for every $i \in [t]$, $\Omega_{\mathcal{A}_i}$ is applied to the first i registers of R.

Proof.

$$\begin{aligned}
\mathsf{Adv}_t^{\widetilde{V}\cdot C} &= \prod_{i=1}^{t} \left(\widetilde{V}_{\mathsf{A}\,\mathsf{R}} \cdot C_{\mathsf{A}} \cdot A_{i\,\mathsf{AB}}\right) && (5.55) \\
&= \prod_{i=1}^{t} \left((\Omega_{\mathcal{A}_i})_{\mathsf{R}_{\leq i}} \cdot \mathrm{App}_{\mathsf{A}\,\mathsf{R_i}} \cdot C_{\mathsf{A}} \cdot A_{i\,\mathsf{AB}}\right) && \text{(Definition of } \widetilde{V})
\end{aligned}$$

[45] For convenience of not having to write the last t'th factor separately, we take $\lambda_{t+1} := \lambda_t^+$.

$$= \prod_{i=1}^{t} \left((\Omega_{\mathcal{A}_i})_{\mathsf{R}_{\leq i}} \cdot C_{\mathsf{R}_{\mathsf{X},\mathsf{i}}} \cdot \mathrm{App}_{\mathsf{A}\,\mathsf{R}_\mathsf{i}} \cdot A_{i\,\mathsf{AB}} \right) \qquad \text{(Lemma 5.15)}$$

Next, notice that for all $j < i$, $\mathrm{App}_{\mathsf{A}\,\mathsf{R}_\mathsf{i}} \cdot A_{i\,\mathsf{AB}}$ commutes with both $C_{\mathsf{R}_{\mathsf{X},\mathsf{j}}}$ and $\left(\Omega_{\mathcal{A}_j}\right)_{\mathsf{R}_{\leq \mathsf{j}}}$, since the former acts only on the A, B, and R_i registers, while both of the latter operators act only on $(\mathsf{R}_1 \ldots \mathsf{R}_\mathsf{j})$. Therefore, we can commute them all past:

$$= \prod_{i=1}^{t} \left((\Omega_{\mathcal{A}_i})_{\mathsf{R}_{\leq i}} \cdot C_{\mathsf{R}_{\mathsf{X},\mathsf{i}}} \right) \cdot \prod_{i=1}^{t} \left(\mathrm{App}_{\mathsf{A}\,\mathsf{R}_\mathsf{i}} \cdot A_{i\,\mathsf{AB}} \right) \qquad (5.56)$$

$$= \left(\prod_{i=1}^{t} (\Omega_{\mathcal{A}_i})_{\mathsf{R}_{\leq i}} \right) \cdot \left(\prod_{i=1}^{t} C_{\mathsf{R}_{\mathsf{X},\mathsf{i}}} \right) \cdot \prod_{i=1}^{t} \left(\mathrm{App}_{\mathsf{A}\,\mathsf{R}_\mathsf{i}} \cdot A_{i\,\mathsf{AB}} \right)$$

($C_{\mathsf{R}_{\mathsf{X},\mathsf{i}}}$ similarly commutes with $\left(\Omega_{\mathcal{A}_j}\right)_{\mathsf{R}_{\leq \mathsf{j}}}$, for all $j < i$)

$$= \left(\prod_{i=1}^{t} (\Omega_{\mathcal{A}_i})_{\mathsf{R}_{\leq i}} \right) \cdot \left(C^{\otimes t} \right)_{\mathsf{R}_\mathsf{X}} \cdot \mathsf{Adv}_t^{\mathrm{App}} \qquad \square$$

Lemma 5.17.

$$\prod_{k=1}^{t} (\Omega_{\mathcal{A}_k})_{\mathsf{R}_{\leq k}} = \Omega_{\mathcal{A}_t} \qquad (5.57)$$

Proof. For each k, write $S_k : \mathcal{A}_k \to \mathrm{End}(V^{\otimes k})$ for the Schur representation action of $\mathcal{A}_k$ on the first k path registers. When $k \leq t$, we view $\Omega_{\mathcal{A}_k}$ as an operator on $\mathsf{R}_{\leq t}$ by tensoring with $\mathbf{I}$ on the later registers $\mathsf{R}_{k+1}, \ldots, \mathsf{R}_t$.

We first recall the subspace onto which $\Omega_{\mathcal{A}_k}$ projects. Namely, $\Omega_{\mathcal{A}_k}$ is the orthogonal projector onto the subspace of states $|\psi\rangle$ satisfying

$$\left(S_k(a)\right)_{\mathsf{X}_{\leq k}} |\psi\rangle = \left(S_k(a)^T\right)_{\mathsf{Y}_{\leq k}} |\psi\rangle \qquad \forall a \in \mathcal{A}_k.$$

Here the $\mathsf{X}_{\leq k}$ registers are identified with $V^{\otimes k}$, whereas the $\mathsf{Y}_{\leq k}$ registers are identified with $\overline{V}^{\otimes k}$, using the conjugate basis labeled by the same basis elements.

To see the characterization, pass to the double Schur basis for $V^{\otimes k} \otimes \overline{V}^{\otimes k}$. In this basis,

$$V^{\otimes k} \otimes \overline{V}^{\otimes k} \cong \bigoplus_{\lambda,\mu} \left(V_G^\lambda \otimes V_{\mathcal{A}_k}^\lambda \right) \otimes \left(\overline{V_G^\mu} \otimes \overline{V_{\mathcal{A}_k}^\mu} \right).$$

The algebra $\mathcal{A}_k$ acts on the X half by

$$S_k(a) = \bigoplus_\lambda \mathbf{I}_{V_G^\lambda} \otimes \rho_\lambda(a),$$

and the transposed action on the Y half is

$$S_k(a)^T = \bigoplus_\mu \mathbf{I}_{\overline{V_G^\mu}} \otimes \rho_\mu(a)^T.$$

Taking a to be central idempotents of $\mathcal{A}_k$ first forces any state satisfying the displayed relations to have support only on the (λ, λ) blocks. On such a block, after treating the two G-registers as spectators, the relations become

$$\left(\rho_\lambda(a) \otimes \mathbf{I}\right) |\psi_{\lambda,\lambda}\rangle = \left(\mathbf{I} \otimes \rho_\lambda(a)^T\right) |\psi_{\lambda,\lambda}\rangle \qquad \forall a \in \mathcal{A}_k.$$

Since $\rho_\lambda(\mathcal{A}_k) = \mathrm{End}(V^\lambda_{\mathcal{A}_k})$, this is equivalent to

$$(M \otimes \mathbf{I}) \left|\psi_{\lambda,\lambda}\right\rangle = (\mathbf{I} \otimes M^T) \left|\psi_{\lambda,\lambda}\right\rangle \qquad \forall M \in \mathrm{End}(V^\lambda_{\mathcal{A}_k}).$$

The solutions are precisely arbitrary states on the two G-registers tensored with the EPR state on the two commutant registers,

$$\left|\mathsf{EPR}^\lambda_{\mathcal{A}_k}\right\rangle := \frac{1}{\sqrt{\dim(V^\lambda_{\mathcal{A}_k})}} \sum_{T \in \mathcal{B}(V^\lambda_{\mathcal{A}_k})} |T\rangle \otimes |T\rangle ,$$

where the second $|T\rangle$ denotes the corresponding conjugate-basis vector in $\overline{V^\lambda_{\mathcal{A}_k}}$. This is exactly the image of $\Omega_{\mathcal{A}_k}$ by definition.

The lemma now follows from the fact that the $\Omega_{\mathcal{A}_k}$ are all orthogonal projectors onto subspaces that form a linear chain. This is because the commutant algebras satisfy $S(\mathcal{A}_k) \subset S(\mathcal{A}_t)$ for all $k \leq t$, so the constraints

$$\left(S_k(a)\right)_{\mathsf{X}_{\leq k}} |\psi\rangle = \left(S_k(a)^T\right)_{\mathsf{Y}_{\leq k}} |\psi\rangle \qquad \forall a \in \mathcal{A}_k$$

are contained within the corresponding set of constraints for $\mathcal{A}_t$. This implies that

$$\Omega_{\mathcal{A}_t} \Omega_{\mathcal{A}_k} = \Omega_{\mathcal{A}_t}$$

for all $k \leq t$, which proves the lemma. □

Combining Lemmas 5.14, 5.16 and 5.17, we obtain an equivalent expanded form for an adversary querying $V \cdot C$:

Corollary 5.18.

$$\mathsf{Adv}^{V\cdot C}_t = \underbrace{\left(\Lambda_{\mathcal{A}_t}\right)_{\mathsf{R}_{\leq t}}}_{\text{ReweightEnd}} \cdot \underbrace{\left(\Omega_{\mathcal{A}_t}\right)_{\mathsf{R}_{\leq t}}}_{\text{Symmetrize}} \cdot \underbrace{\left(C^{\otimes t}\right)_{\mathsf{R}_\mathsf{X}}}_{C^{\otimes t}} \cdot \underbrace{\mathsf{Adv}^{\mathrm{App}}_t}_{\text{Append}} \cdot \underbrace{\left(\Lambda^+_{\mathcal{A}_0}\right)_{\mathsf{R}_0}}_{\text{ReweightStart}} \tag{5.58}$$

Both reweighting steps, as well as the Symmetrize step, depend on the group G (in particular, they depend on the commutant algebra of its tensor power representation), while $C^{\otimes t}$ and Append are independent of the group and its representation.

When we start with R_0 containing the empty database, this simplifies further, since initial reweighting step disappears:

Remark 5.19. When R contains the empty database, $|\emptyset\rangle$, then $\Lambda^+_{\mathcal{A}_0} |\emptyset\rangle_\mathsf{R} = |\emptyset\rangle_\mathsf{R}$.

Proof. Since the database register R is empty, it corresponds to the 1-dimensional space $(\mathbb{C}^N)^{\otimes 0} = \mathbb{C}$, on which G acts via the trivial irrep ($\rho^{\otimes 0} = 1$), and its commutant $\mathcal{A}_0 = \mathbb{C}$ is a simple module (its only irrep is the one that acts via scalar multiplication). Since the irreps are both 1-dimensional, the ratio β_λ is 1, and the rescaling has no effect. □

6 Approximation and Specialization to Specific Groups

Our general update rule can recover an interpretable recording oracle for any group representation. To demonstrate its generality, in this section, we show how it recovers existing recording oracles, including Zhandry's original recording oracle for functions [Zha19] and Ma and Huang's path-recording oracle [MH25] for unitaries as special cases. For the latter, since the path recording oracle of [MH25] is only approximate, we describe two explicit approximations that relate our path-recording oracle to theirs.

We also analyze the *unitary Haar cipher* model, in which unitaries $U_1, \ldots, U_K$ are each chosen i.i.d. from the Haar measure on $U(N)$ and queried in a controlled manner as $\sum_{k\in[K]} |k\rangle\langle k| \otimes U_k$.

6.1 The Unitary Path-Recording Oracle of [MH25]

We will use the following lemma from [MH25]:

Lemma 6.1 ([MH25, Lemma 2.2])**.** *Let ρ_{ST} be a density matrix on registers S, T and let Π_{T} be a projector that acts on register T. Then*

$$\|\mathrm{Tr}_{\mathsf{T}}(\rho_{\mathsf{ST}}) - \mathrm{Tr}_{\mathsf{T}}(\Pi_{\mathsf{T}}\rho_{\mathsf{ST}}\Pi_{\mathsf{T}})\|_1 = 1 - \mathrm{Tr}(\Pi_{\mathsf{T}}\rho_{\mathsf{ST}}). \tag{6.1}$$

6.1.1 The Distinct Subspace

Definition 6.2. For $1 \le t \le N$, define

$$[N]^t_{\mathrm{dist}} := \{(x_1, \ldots, x_t) \in [N]^t : \text{all } x_i\text{'s distinct}\} \tag{6.2}$$

as the set of all distinct ordered t-tuples over $[N]$.

Definition 6.3. The *distinct subspace* $\mathsf{Dist}_{N,t} \subseteq (\mathbb{C}^N)^{\otimes t}$ is defined as the span of all distinct strings, i.e.

$$\mathsf{Dist}_{N,t} := \mathrm{span}\{|x\rangle : x \in [N]^t_{\mathrm{dist}}\} \tag{6.3}$$

6.1.2 The [MH25] oracle

The [MH25] path-recording oracle acts on an adversary input register A and recording register R_{MH} whose basis consists of *sets* of input-output pairs $\{(x_1, y_1), \ldots, (x_t, y_t)\}$ with the property that $(y_1, \ldots, y_t)$ are distinct.

To compare this with our compressed oracle, we consider the following simple isometry mapping R_{MH} to our recording space R:

$$|\{(x_1, y_1), \ldots, (x_t, y_t)\}\rangle \mapsto \frac{1}{\sqrt{t!}} \sum_{\pi\in S_t} \left|x_{\pi^{-1}(1)}, \ldots, x_{\pi^{-1}(t)}, y_{\pi^{-1}(1)}, \ldots, y_{\pi^{-1}(t)}\right\rangle = \sqrt{t!}\cdot\Omega_{S_t} |x_1, \ldots, x_t, y_1, \ldots, y_t\rangle .$$

Via this isometry, we identify the Ma-Huang path recording update with the following update rule

$$|x_t\rangle_{\mathsf{A}} \otimes |x_1, \ldots, x_{t-1}, y_1, \ldots, y_{t-1}\rangle \mapsto \Omega_{S_t} \cdot \frac{\sqrt{t}}{\sqrt{N-t+1}} \sum_{y_t \notin \{y_1,\ldots,y_{t-1}\}} |y_t\rangle_{\mathsf{A}} \otimes |x_1, \ldots, x_t, y_1, \ldots, y_t\rangle \tag{6.4}$$

$$= \frac{\sqrt{t}}{\sqrt{N-t+1}} \cdot \left(\Omega_{S_t}\right)_{\mathsf{R}} \cdot \left(\Pi_{\mathsf{Dist}}\right)_{\mathsf{R}_Y} \mathsf{App}_{\mathsf{AR}}. \tag{6.5}$$

Thus, the statement that the Ma-Huang oracle is indistinguishable from our oracle amounts to two informal statements:

- The application of Π_{Dist} is approximately the identity map, and
- The overall effect of inverse subspace reweighting $\Lambda^{+}_{S_{t-1}}$ and reweighting Λ_{S_t} is approximately constant, with a scalar factor of $\sqrt{\frac{t}{N-t+1}}$.

Next, we will give a formal justification for the approximate equivalence of V_{MH} and $V_{U(N)}$ for t-query quantum algorithms.

6.1.3 Proof of approximate equivalence

Using the fact that $(\Pi_{\mathsf{Dist}})_{\mathsf{R}_{Y,\leq t}}$ commutes with Ω_{S_t}, we can then express the isometry implemented by a t-query adaptive adversary $\mathsf{Adv}^{V_{MH}}$ as

$$\mathsf{Adv}_t^{V_{MH}} = \sqrt{t!}\prod_{i=1}^{t}\frac{1}{\sqrt{N-i+1}}\cdot\left(\Pi_{\mathsf{Dist}}\right)_{\mathsf{R}_Y}\cdot\left(\Omega_{S_t}\right)_{\mathsf{R}}\cdot\mathsf{Adv}_t^{\mathrm{App}} \tag{6.6}$$

$$= \frac{1}{\sqrt{\binom{N}{t}}}\cdot\left(\Pi_{\mathsf{Dist}}\right)_{\mathsf{R}_Y}\cdot\left(\Omega_{S_t}\right)_{\mathsf{R}}\cdot\mathsf{Adv}_t^{\mathrm{App}}. \tag{6.7}$$

On the other hand, our own path-recording oracle implements

$$\mathsf{Adv}_t^{V_{U(N)}} = \left(\Lambda_{S_t}\right)_{\mathsf{R}}\cdot\left(\Omega_{S_t}\right)_{\mathsf{R}}\cdot\mathsf{Adv}_t^{\mathrm{App}}\cdot\left(\Lambda^{+}_{S_0}\right)_{\mathsf{R}_0} \tag{6.8}$$

$$= \left(\Lambda_{S_t}\right)_{\mathsf{R}}\cdot\left(\Omega_{S_t}\right)_{\mathsf{R}}\cdot\mathsf{Adv}_t^{\mathrm{App}}. \tag{6.9}$$

To compare these, define the quantum states

$$\left|\mathsf{Adv}_t^{V_{MH}}\right\rangle = \mathsf{Adv}_t^{V_{MH}}\left|0\right\rangle_{\mathsf{AB}}\left|\emptyset\right\rangle_{\mathsf{R}}, \qquad \left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle = \mathsf{Adv}_t^{V_{U(N)}}\left|0\right\rangle_{\mathsf{AB}}\left|\emptyset\right\rangle_{\mathsf{R}}. \tag{6.10}$$

We wish to derive the fact that

$$\left\|\mathrm{Tr}_{\mathsf{R}}\left[\left|\mathsf{Adv}_t^{V_{MH}}\right\rangle\left\langle\mathsf{Adv}_t^{V_{MH}}\right|\right] - \mathrm{Tr}_{\mathsf{R}}\left[\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\left\langle\mathsf{Adv}_t^{V_{U(N)}}\right|\right]\right\|_1 \leq O\left(\frac{t^2}{N}\right) \tag{6.11}$$

We prove this via the following sequence of steps. First, we show that

$$\left\|(\Pi_{\mathsf{Dist}})_{\mathsf{R}_Y}\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\right\|^2 \geq 1 - O\left(\frac{t^2}{N}\right). \tag{6.12}$$

Then, by Lemma 6.1, we have that

$$\left\|\mathrm{Tr}_{\mathsf{R}}\left[(\Pi_{\mathsf{Dist}})_{\mathsf{R}_Y}\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\left\langle\mathsf{Adv}_t^{V_{U(N)}}\right|(\Pi_{\mathsf{Dist}})_{\mathsf{R}_Y}\right] - \mathrm{Tr}_{\mathsf{R}}\left[\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\left\langle\mathsf{Adv}_t^{V_{U(N)}}\right|\right]\right\|_1 \leq O\left(\frac{t^2}{N}\right) \tag{6.13}$$

Finally, we prove that

$$\left\|(\Pi_{\mathsf{Dist}})_{\mathsf{R}_Y}\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\left\langle\mathsf{Adv}_t^{V_{U(N)}}\right|(\Pi_{\mathsf{Dist}})_{\mathsf{R}_Y} - \left|\mathsf{Adv}_t^{V_{MH}}\right\rangle\left\langle\mathsf{Adv}_t^{V_{MH}}\right|\right\|_1 \leq O\left(\frac{t^2}{N}\right) \tag{6.14}$$

For the remainder of this subsection, write

$$\Pi^Y_{\mathsf{Dist}} := \left(\Pi_{\mathsf{Dist}_{N,t}}\right)_{\mathsf{R}_Y}, \qquad \Pi^Y_{\mathsf{Dist},t-1} := \left(\Pi_{\mathsf{Dist}_{N,t-1}}\right)_{\mathsf{R}_{Y,<t}}, \qquad \alpha_t := \frac{1}{\sqrt{\binom{N}{t}}}.$$

We also abbreviate the append-only state by

$$|\psi_t\rangle := \mathsf{Adv}_t^{\mathrm{App}} |0\rangle_{\mathsf{AB}} |\emptyset\rangle_{\mathsf{R}}.$$

With this notation, Eq. (6.6) says

$$\left|\mathsf{Adv}_t^{V_{MH}}\right\rangle = \alpha_t \cdot \Pi^Y_{\mathsf{Dist}} \cdot \Omega_{S_t} |\psi_t\rangle.$$

while Eq. (6.9) says

$$\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle = \Lambda_{S_t} \cdot \Omega_{S_t} |\psi_t\rangle.$$

Lemma 6.4 (Analyzing $\Pi^Y_{\mathsf{Dist}} |\psi_t\rangle$)**.** *For every non-identity $\pi \in S_t$,*

$$\langle\psi_t| \Pi^Y_{\mathsf{Dist}} \Big(S(\pi) \otimes S(\pi)\Big) \Pi^Y_{\mathsf{Dist}} |\psi_t\rangle = 0.$$

Moreover,

$$\langle\psi_t| \Pi^Y_{\mathsf{Dist}} |\psi_t\rangle = t! \binom{N}{t}.$$

Consequently,

$$\left\|\Omega_{S_t} \Pi^Y_{\mathsf{Dist}} |\psi_t\rangle\right\|^2 = \binom{N}{t} = \frac{1}{\alpha_t^2}.$$

Proof. We prove the first two claims by induction on t. The case $t = 1$ is trivial (there are no non-identity permutations and $\Pi^Y_{\mathsf{Dist}} = \mathbf{I}$).

For the inductive step, we write

$$|\psi_t\rangle = \mathsf{App}_{\mathsf{AR}_t} \cdot (A_t)_{\mathsf{AB}} \cdot |\psi\rangle_{t-1} = \mathsf{App}_{\mathsf{AR}_t} |\psi_t^{\mathrm{pre}}\rangle.$$

Moreover, since $\Pi^Y_{\mathsf{Dist}} = \Pi^Y_{\mathsf{Dist}} \cdot \Pi^Y_{\mathsf{Dist},t-1}$ and $\Pi^Y_{\mathsf{Dist},t-1}$ commutes with $\mathsf{App}_{\mathsf{AR}_t}$, we have that

$$\Pi^Y_{\mathsf{Dist}} |\psi_t\rangle = \Pi^Y_{\mathsf{Dist}} \cdot \mathsf{App}_{\mathsf{AR}_t} \cdot \Pi^Y_{\mathsf{Dist},t-1} |\psi_t^{\mathrm{pre}}\rangle.$$

We now write explicitly in the standard basis

$$\Pi^Y_{\mathsf{Dist},t-1} |\psi_t^{\mathrm{pre}}\rangle = \sum_{\substack{x_1,\ldots,x_t \\ \text{distinct } y_1,\ldots,y_{t-1}}} |x_t\rangle_{\mathsf{A}} \left|\phi_{x_{1:t},y_{1:t-1}}\right\rangle_{\mathsf{B}} |x_1,\ldots x_{t-1},y_1,\ldots,y_{t-1}\rangle_{\mathsf{R}_{<t}}.$$

On this state, the effect of $\Pi^Y_{\mathsf{Dist}} \cdot \mathsf{App}_{\mathsf{AR}_t}$ is appending with respect to an (unnormalized) EPR state over $y_t \notin \{y_1,\ldots,y_{t-1}\}$ (controlled on $\{y_1,\ldots,y_{t-1}\}$), which tells us that

$$\Pi^Y_{\mathsf{Dist}} |\psi_t\rangle = \sum_{\substack{x_1,\ldots,x_t \\ \text{distinct } y_1,\ldots,y_t}} |y_t\rangle_{\mathsf{A}} \left|\phi_{x_{1:t},y_{1:t-1}}\right\rangle_{\mathsf{B}} |x_1,\ldots x_t,y_1,\ldots,y_t\rangle_{\mathsf{R}_{\le t}}.$$

Given that $\Pi^Y_{\mathsf{Dist},t-1}$ also commutes with $(A_t)_{\mathsf{AB}}$, this immediately tells us that

$$\left|\left|\Pi^Y_{\mathsf{Dist}} |\psi_t\rangle\right|\right|^2 = (N-t+1)\left|\left|\Pi^Y_{\mathsf{Dist},t-1} |\psi_{t-1}\rangle\right|\right|^2,$$

completing the second part of the induction. For the first part, we have that

$$\Big(S(\pi)\otimes S(\pi)\Big)\Pi^Y_{\mathsf{Dist}}\left|\psi_t\right\rangle = \sum_{\substack{x_1,\ldots,x_t\\ \text{distinct } y_1,\ldots,y_t}} \left|y_t\right\rangle_{\mathsf{A}}\left|\phi_{x_{1:t},y_{1:t-1}}\right\rangle_{\mathsf{B}}\left|x_{\pi^{-1}(1)},\ldots x_{\pi^{-1}(t)},y_{\pi^{-1}(1)},\ldots,y_{\pi^{-1}(t)}\right\rangle_{\mathsf{R}_{\le t}}.$$

Now, suppose that $\pi(t)\neq t$. Then, the inner product of these two states is zero because a nonzero term in the expansion of this inner product would require terms $y_t = y'_t$ on the A register but also $y_t = y'_{\pi^{-1}(t)}$ on the $\mathsf{R}_{Y,t}$ register.

On the other hand, suppose that $\pi(t)=t$ and identify π as a non-identity element of S_{t-1}; in particular, $S(\pi)$ commutes with $\mathsf{App}_{\mathsf{AR}_t}\cdot(A_t)_{\mathsf{AB}}$. Then, we write

$$\left\langle\psi_t\right|\Pi^Y_{\mathsf{Dist}}\Big(S(\pi)\otimes S(\pi)\Big)\Pi^Y_{\mathsf{Dist}}\left|\psi_t\right\rangle = \left\langle\psi^{\mathrm{pre}}_t\right|\Pi^Y_{\mathsf{Dist},t-1}\mathsf{App}^\dagger_{\mathsf{A},\mathsf{R}_t}\Pi^Y_{\mathsf{Dist}}\mathsf{App}_{\mathsf{A},\mathsf{R}_t}\Pi^Y_{\mathsf{Dist},t-1}\Big(S(\pi)\otimes S(\pi)\Big)\left|\psi^{\mathrm{pre}}_t\right\rangle$$

Next, we make use of the identity

$$\Pi^Y_{\mathsf{Dist},t-1}\mathsf{App}^\dagger_{\mathsf{A},\mathsf{R}_t}\Pi^Y_{\mathsf{Dist}}\mathsf{App}_{\mathsf{A},\mathsf{R}_t}\Pi^Y_{\mathsf{Dist},t-1} = (N-t+1)\Pi^Y_{\mathsf{Dist},t-1},$$

which again holds because the effect of $\Pi^Y_{\mathsf{Dist}}\mathsf{App}_{\mathsf{A},\mathsf{R}_t}$ on the image of $\Pi^Y_{\mathsf{Dist},t-1}$ is to append an un-normalized EPR state over $y_t\notin\{y_1,\ldots,y_{t-1}\}$:

$$\Pi^Y_{\mathsf{Dist}}\mathsf{App}_{\mathsf{A},\mathsf{R}_t}\Pi^Y_{\mathsf{Dist},t-1} = \sum_{\substack{x_1,\ldots x_t\\ \text{distinct } y_1,\ldots,y_t}} \left|y_t\right\rangle\!\left\langle x_t\right|_{\mathsf{A}}\otimes\left|x_1,\ldots,x_t,y_1,\ldots,y_t\right\rangle\!\left\langle x_1,\ldots,x_{t-1},y_1,\ldots,y_{t-1}\right|.$$

The claimed identity follows from this by a direct computation.

We conclude that

$$\begin{aligned}
\left\langle\psi_t\right|\Pi^Y_{\mathsf{Dist}}\Big(S(\pi)\otimes S(\pi)\Big)\Pi^Y_{\mathsf{Dist}}\left|\psi_t\right\rangle &= (N-t+1)\left\langle\psi^{\mathrm{pre}}_t\right|\Pi^Y_{\mathsf{Dist},t-1}\Big(S(\pi)\otimes S(\pi)\Big)\left|\psi^{\mathrm{pre}}_t\right\rangle\\
&= (N-t+1)\left\langle\psi_{t-1}\right|\Pi^Y_{\mathsf{Dist},t-1}(A^\dagger_t)_{\mathsf{AB}}\Big(S(\pi)\otimes S(\pi)\Big)(A_t)_{\mathsf{AB}}\Pi^Y_{\mathsf{Dist},t-1}\left|\psi_{t-1}\right\rangle\\
&= (N-t+1)\left\langle\psi_{t-1}\right|\Pi^Y_{\mathsf{Dist},t-1}\Big(S(\pi)\otimes S(\pi)\Big)\Pi^Y_{\mathsf{Dist},t-1}\left|\psi_{t-1}\right\rangle.
\end{aligned}$$

This completes the induction.

Finally, using

$$\Omega_{S_t} = \frac{1}{t!}\sum_{\pi\in S_t} S(\pi)\otimes S(\pi),$$

the cross terms below vanish and so

$$\left\|\Omega_{S_t}\Pi^Y_{\mathsf{Dist}}\left|\psi_t\right\rangle\right\|^2 = \frac{1}{t!}\left\|\Pi^Y_{\mathsf{Dist}}\left|\psi_t\right\rangle\right\|^2 = \binom{N}{t}. \qquad \square$$

Lemma 6.5 (Reweighting is approximately scalar)**.** *Assume $t^2 = O(N)$. Then*

$$\|\Lambda_{S_t} - \alpha_t\mathbf{I}\| \le \alpha_t\cdot O\left(\frac{t^2}{N}\right).$$

Equivalently,

$$\alpha_t^2\left(1-O\left(\frac{t^2}{N}\right)\right)\mathbf{I} \preceq \Lambda^\dagger_{S_t}\Lambda_{S_t} \preceq \alpha_t^2\left(1+O\left(\frac{t^2}{N}\right)\right)\mathbf{I}. \tag{6.15}$$

Proof. We equivalently calculate

$$\left\| \frac{N^t}{t!} \Lambda_{S_t}^\dagger \Lambda_{S_t} - \mathbf{I} \right\| = \left\| \frac{N^t}{t!} \left(\sum_{\lambda \in \widehat{S}_t} \beta_{\lambda,t}^2 \Pi_{\lambda,t} \right) - \mathbf{I} \right\| \tag{6.16}$$

$$= \left\| \sum_{\lambda \in \widehat{S}_t} \left(\frac{N^t}{t!} \beta_{\lambda,t}^2 - 1 \right) \Pi_{\lambda,t} \right\| \tag{6.17}$$

$$\leq \max_{\lambda \in \widehat{S}_t} \left| \frac{N^t}{t!} \beta_{\lambda,t}^2 - 1 \right| \tag{6.18}$$

$$\leq O\left(\frac{t^2}{N} \right). \qquad \text{(Lemma 3.32)}$$

□

Lemma 6.6 (The exact unitary recording is almost distinct on R_Y)**.**

$$\left\| \Pi_{\mathsf{Dist}}^Y \left| \mathsf{Adv}_t^{V_{U(N)}} \right\rangle \right\|^2 \geq 1 - O\left(\frac{t^2}{N} \right).$$

Proof. We have that

$$\Pi_{\mathsf{Dist}}^Y \left| \mathsf{Adv}_t^{V_{U(N)}} \right\rangle = \Lambda_{S_t} \Omega_{S_t} \Pi_{\mathsf{Dist}}^Y |\psi_t\rangle .$$

By Lemma 6.5,

$$\Lambda_{S_t}^\dagger \Lambda_{S_t} \succeq \alpha_t^2 \left(1 - O\left(\frac{t^2}{N} \right) \right) \mathbf{I},$$

and by Lemma 6.4,

$$\left\| \Omega_{S_t} \Pi_{\mathsf{Dist}}^Y |\psi_t\rangle \right\|^2 = \frac{1}{\alpha_t^2}.$$

Together, these prove the lemma. □

Theorem 6.7. *For any t-query adversary with $t^2 = O(N)$,*

$$\left\| \left| \mathsf{Adv}_t^{V_{MH}} \right\rangle \left\langle \mathsf{Adv}_t^{V_{MH}} \right| - \Pi_{\mathsf{Dist}}^Y \left| \mathsf{Adv}_t^{V_{U(N)}} \right\rangle \left\langle \mathsf{Adv}_t^{V_{U(N)}} \right| \Pi_{\mathsf{Dist}}^Y \right\|_1 \leq O\left(\frac{t^2}{N} \right).$$

Proof. We write

$$\Pi_{\mathsf{Dist}}^Y \left| \mathsf{Adv}_t^{V_{U(N)}} \right\rangle - \left| \mathsf{Adv}_t^{V_{MH}} \right\rangle = \left(\Lambda_{S_t} - \alpha_t \mathbf{I} \right) \Omega_{S_t} \Pi_{\mathsf{Dist}}^Y |\psi_t\rangle .$$

By Lemmas 6.4 and 6.5,

$$\left\| \Pi_{\mathsf{Dist}}^Y \left| \mathsf{Adv}_t^{V_{U(N)}} \right\rangle - \left| \mathsf{Adv}_t^{V_{MH}} \right\rangle \right\| \leq \alpha_t \cdot O\left(\frac{t^2}{N} \right) \cdot \frac{1}{\alpha_t} = O\left(\frac{t^2}{N} \right).$$

Therefore,

$$\left\| \left| \mathsf{Adv}_t^{V_{MH}} \right\rangle \left\langle \mathsf{Adv}_t^{V_{MH}} \right| - \Pi_{\mathsf{Dist}}^Y \left| \mathsf{Adv}_t^{V_{U(N)}} \right\rangle \left\langle \mathsf{Adv}_t^{V_{U(N)}} \right| \Pi_{\mathsf{Dist}}^Y \right\|_1 \tag{6.19}$$

$$\leq 2 \cdot \left\| \Pi_{\mathsf{Dist}}^Y \left| \mathsf{Adv}_t^{V_{U(N)}} \right\rangle - \left| \mathsf{Adv}_t^{V_{MH}} \right\rangle \right\| \tag{6.20}$$

$$\leq O\left(\frac{t^2}{N} \right). \tag{6.21}$$

□

This completes the full comparison between the [MH25] path-recording oracle and our exact unitary path recording oracle.

6.2 The Unitary Haar Cipher

We now consider the case of the unitary Haar cipher, which consists of a collection of keyed unitaries $U_1, \ldots, U_K$, each sampled independently from the Haar measure on $U(N)$. The group is therefore $G = U(N)^K$, and the representation is its block-diagonal embedding:

$$\rho(U_1, \ldots, U_K) = \sum_{k\in[K]} |k\rangle\langle k| \otimes U_k \,. \tag{6.22}$$

This is the idealization of a pseudorandom unitary, in which the unitary for each key is chosen independently from the Haar measure.

As the reader might rightfully guess, we can get a recording oracle for the unitary Haar cipher by simply maintaining K different recordings of each of the K underlying unitaries. And in fact, we can compress it to size that grows only with the number of queries by only explicitly recording the keys whose individual recording has a non-zero number of entries. This is in fact the correct intuition, and for practical purposes, it might suffice to stop there. But the beauty of our framework is that it eliminates the guesswork. We can in fact just blindly take the representation above and plug it into our framework to recover this outcome.

For this, we must first consider the commutant algebra of the tensor power of the representation in Equation (6.22). It turns out to be the Schur representation of the diagram algebra that comes out of the wreath product $\mathbb{Z}_K \wr S_t$ of the cyclic group mod K and the symmetric group.

Definition 6.8. The *K-colored permutation group* is the wreath product $\mathbb{Z}_K \wr S_t$. That is, it is the group whose group elements $(f, \pi) \in \mathbb{Z}_K \wr S_t$ are parameterized by

- a "function" $f \in \mathbb{Z}_K^t$ (we often write group elements of $\mathbb{Z}_K^t$ as functions $f : [t] \to \mathbb{Z}_K$, or as vectors $\vec{f} = (f(1), \ldots, f(t))$), and
- a permutation $\pi \in S_t$,

with a group operation that is given by entry-wise addition in $\mathbb{Z}_K^t$, composing permutations in S_t, and permuting a function $f \in \mathbb{Z}_K^t$ by a permutation $\pi \in S_t$ whenever they commute past each other. Specifically,

$$\mathbb{Z}_K \wr S_t = \left\langle (f, \pi) \;\middle|\; f \in \mathbb{Z}_K^t,\, \pi \in S_t,\, (f_1, \pi_1) \cdot (f_2, \pi_2) = \big(f_1 + \pi_1(f_2),\, \pi_1 \circ \pi_2\big) \right\rangle \tag{6.23}$$

This group is also often called

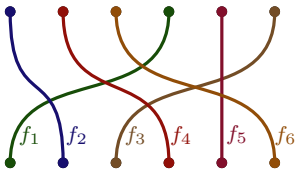
- the K-phased permutation group, since the $\mathbb{Z}_K$ can be viewed as phases that come from K'th roots of unity,
- the monomial unitary group with modulus K, or
- the complex reflection group $G(K, 1, t)$.

When specializing to $K = 2$, it is also known as the hyperoctahedral group.

Note that when it is clear from context we will often write the group elements (f, e_{S_t}) and $(e_{\mathbb{Z}_K^t}, \pi)$ more simply as f and π, respectively,[46] so that we can write any group element (f, π) as a product $f\,\pi$, or equivalently, as $\pi\, f'$, for $f' = \pi^{-1}(f)$.

[46] e_{S_t} and $e_{\mathbb{Z}_K^t}$ are the identity elements of the respective groups.

The defining representation of the colored permutation group is often taken to be as permutation matrices with nonzero entries taken from the K'th roots of unity. We can also view group elements as colored permutation diagrams, where each strand has an attached "color" in $\mathbb{Z}_K$:



(6.24)

Figure 16: Colored permutation diagram on 6 elements with colors $f_1, \ldots, f_6 \in \mathbb{Z}_K$ attached to the strands.

Diagram multiplication occurs by concatenation of the underlying permutations and summing (mod K) the colors along each merged strand.

We will be interested here in the Schur representation of $\mathbb{Z}_K \wr S_t$ on $(\mathbb{C}^{KN})^{\otimes t}$, which generalizes the Schur representation of S_t (Definition 3.16) by adding a phase of $\omega_K^{f_i k_i}$ whenever a register containing $|k_i, x_i\rangle$ passes through a strand with color f_i.

For example,

$$S\left(\begin{array}{c}\end{array}\right) = \sum_{\substack{\vec{k}\in[K]^6 \\ \vec{x}\in[N]^6}} \omega_K^{f_1k_1+\cdots+f_6k_6} \left|{}^{(k_2,x_2),(k_4,x_4),(k_6,x_6),(k_1,x_1),(k_5,x_5),(k_3,x_3)}\right\rangle \left\langle{}_{(k_1,x_1),(k_2,x_2),(k_3,x_3),(k_4,x_4),(k_5,x_5),(k_6,x_6)}\right| . \tag{6.25}$$

Formally,

$$S(\pi\, f) := \sum_{\substack{\vec{k}\in[K]^t \\ \vec{x}\in[N]^t}} \omega_K^{\vec{f}\cdot\vec{k}} \left|(k_{\pi^{-1}(1)}, x_{\pi^{-1}(1)}), \ldots, (k_{\pi^{-1}(t)}, x_{\pi^{-1}(t)})\right\rangle\!\left\langle(k_1,x_1),\ldots,(k_t,x_t)\right| \tag{6.26}$$

The *colored permutation algebra* is the group algebra $\mathbb{C}[\mathbb{Z}_K \wr S_t]$, which is the diagram algebra spanned by colored permutation diagrams.

Lemma 6.9. *The commutant of $\rho(\vec{U})^{\otimes t}$ (see Equation (6.22)) is the Schur representation of the colored permutation algebra $\mathbb{C}[\mathbb{Z}_K \wr S_t]$.*

Proof. First, we show that $S(\pi\, f)$ commutes with $\rho(\vec{U})^{\otimes t}$.

$$S(\pi\, f)\ \rho(\vec{U})^{\otimes t} = S(\pi)\ S(f)\ \rho(\vec{U})^{\otimes t} \tag{6.27}$$

$$= S(\pi) \left(\sum_{\vec{k}\in[K]^t} \omega_K^{\vec{f}\cdot\vec{k}} |\vec{k}\rangle\langle\vec{k}| \otimes \mathbf{I}\right) \left(\sum_{\vec{k}\in[K]^t} |\vec{k}\rangle\langle\vec{k}| \otimes U_{k_1} \otimes \cdots \otimes U_{k_t}\right) \tag{6.28}$$

$$= S(\pi) \left(\sum_{\vec{k}\in[K]^t} \omega_K^{\vec{f}\cdot\vec{k}} |\vec{k}\rangle\langle\vec{k}| \otimes U_{k_1} \otimes \cdots \otimes U_{k_t}\right) \tag{6.29}$$

$$= S(\pi) \left(\sum_{\vec{k}\in[K]^t} |\vec{k}\rangle\langle\vec{k}| \otimes U_{k_1} \otimes \cdots \otimes U_{k_t}\right) \left(\sum_{\vec{k}\in[K]^t} \omega_K^{\vec{f}\cdot\vec{k}} |\vec{k}\rangle\langle\vec{k}| \otimes \mathbf{I}\right) \tag{6.30}$$

$$= S(\pi)\ \rho(\vec{U})^{\otimes t}\ S(f) \tag{6.31}$$

$$= \rho(\vec{U})^{\otimes t}\ S(\pi)\ S(f) \tag{6.32}$$

This shows that $\mathbb{C}[\mathbb{Z}_K \wr S_t]$ is contained in the commutant.

It now remains to prove that every operator that commutes with $\rho(\vec{U})^{\otimes t}$ must have the form $S(a)$ for some $a \in \mathbb{C}[\mathbb{Z}_K \wr S_t]$. Let $T \in \mathrm{End}((\mathbb{C}^{KN})^{\otimes t})$ be a generic operator on this space, and let's say that it commutes with $\rho(\vec{U})^{\otimes t}$ for all tuples of unitaries $\vec{U} \in U(N)^K$. We first write T in terms of its blocks as

$$T = \sum_{\vec{p},\,\vec{q} \in [K]^t} |\vec{p}\rangle\langle\vec{q}| \otimes T_{\vec{p},\vec{q}} \tag{6.33}$$

where the $|\vec{p}\rangle\langle\vec{q}|$ acts on all the control (or key) registers, and each $T_{\vec{p},\vec{q}} \in \mathrm{End}((\mathbb{C}^N)^{\otimes t})$ acts on the rest of the registers (where each of the unitaries U_{k_i} acts).

We know that for all $\vec{U}$, we have $[\rho(\vec{U})^{\otimes t}, T] = \rho(\vec{U})^{\otimes t}\, T - T\, \rho(\vec{U})^{\otimes t} = 0$. For brevity, we write $U_{\vec{k}} \coloneqq U_{k_1} \otimes \cdots \otimes U_{k_t}$. The commutator is then

$$\rho(\vec{U})^{\otimes t}\, T - T\, \rho(\vec{U})^{\otimes t} \tag{6.34}$$

$$= \left(\sum_{\vec{k} \in [K]^t} |\vec{k}\rangle\langle\vec{k}| \otimes U_{\vec{k}}\right) \left(\sum_{\vec{p},\,\vec{q} \in [K]^t} |\vec{p}\rangle\langle\vec{q}| \otimes T_{\vec{p},\vec{q}}\right) - \left(\sum_{\vec{p},\,\vec{q} \in [K]^t} |\vec{p}\rangle\langle\vec{q}| \otimes T_{\vec{p},\vec{q}}\right) \left(\sum_{\vec{k} \in [K]^t} |\vec{k}\rangle\langle\vec{k}| \otimes U_{\vec{k}}\right) \tag{6.35}$$

$$= \left(\sum_{\vec{k},\,\vec{p},\,\vec{q} \in [K]^t} |\vec{k}\rangle\langle\vec{k}| \cdot |\vec{p}\rangle\langle\vec{q}| \otimes U_{\vec{k}} \cdot T_{\vec{p},\vec{q}}\right) - \left(\sum_{\vec{k},\,\vec{p},\,\vec{q} \in [K]^t} |\vec{p}\rangle\langle\vec{q}| \cdot\ |\vec{k}\rangle\langle\vec{k}| \otimes T_{\vec{p},\vec{q}} \cdot U_{\vec{k}}\right) \tag{6.36}$$

$$= \left(\sum_{\vec{p},\,\vec{q} \in [K]^t} |\vec{p}\rangle\langle\vec{q}| \otimes U_{\vec{p}} \cdot T_{\vec{p},\vec{q}}\right) - \left(\sum_{\vec{p},\,\vec{q} \in [K]^t} |\vec{p}\rangle\langle\vec{q}| \otimes T_{\vec{p},\vec{q}} \cdot U_{\vec{q}}\right) \tag{6.37}$$

$$= \sum_{\vec{p},\,\vec{q} \in [K]^t} |\vec{p}\rangle\langle\vec{q}| \otimes \left(U_{\vec{p}} \cdot T_{\vec{p},\vec{q}} - T_{\vec{p},\vec{q}} \cdot U_{\vec{q}}\right)\ . \tag{6.38}$$

Now, if this operator is 0, then it must be 0 on each block. That is, for every $\vec{p}, \vec{q} \in [K]^t$ (and every $\vec{U}$), we have that $U_{\vec{p}} \cdot T_{\vec{p},\vec{q}} - T_{\vec{p},\vec{q}} \cdot U_{\vec{q}} = 0$.

Since this holds for all $\vec{U} \in U(N)^K$, it must also hold for the case where each unitary is simply $U_k = e^{i\,\theta_k}\,\mathbf{I}$. In this case, we have that

$$U_{\vec{p}} \cdot T_{\vec{p},\vec{q}} - T_{\vec{p},\vec{q}} \cdot U_{\vec{q}} = e^{i\sum_{j\in[t]}\theta_{p_j}} T_{\vec{p},\vec{q}} - e^{i\sum_{j\in[t]}\theta_{q_j}} T_{\vec{p},\vec{q}} \tag{6.39}$$

$$= \left(e^{i\sum_{j\in[t]}\theta_{p_j}} - e^{i\sum_{j\in[t]}\theta_{q_j}}\right) T_{\vec{p},\vec{q}}. \tag{6.40}$$

Therefore, $T_{\vec{p},\,\vec{q}}$ must be 0 unless $\vec{p}$ is a permutation of $\vec{q}$ (that is, they have the same entries in a different order).

It must also hold when all the unitaries are identical (that is $U_k = U$ for all k). In this case, we have $U_{\vec{p}} \cdot T_{\vec{p},\vec{q}} - T_{\vec{p},\vec{q}} \cdot U_{\vec{q}} = U^{\otimes t} \cdot T_{\vec{p},\vec{q}} - T_{\vec{p},\vec{q}} \cdot U^{\otimes t}$, and thus from classical Schur-Weyl duality, we know that the blocks of T must be in the span of the permutations: $T_{\vec{p},\vec{q}} = \sum_{\pi\in S_t} c_\pi^{\vec{p},\vec{q}}\, S(\pi)$, for some coefficients $c_\pi^{\vec{p},\vec{q}}$.

Now for the full group $U(N)^K$, we have that

$$U_{\vec{p}} \cdot T_{\vec{p},\vec{q}} - T_{\vec{p},\vec{q}} \cdot U_{\vec{q}} = \sum_{\pi \in S_t} c_\pi^{\vec{p},\vec{q}} \left(U_{\vec{p}} \cdot S(\pi) - S(\pi) \cdot U_{\vec{q}} \right) \tag{6.41}$$

$$= \sum_{\pi \in S_t} c_\pi^{\vec{p},\vec{q}} \left(S(\pi) \cdot U_{\pi^{-1}(\vec{p})} - S(\pi) \cdot U_{\vec{q}} \right) \tag{6.42}$$

$$= \sum_{\pi \in S_t} c_\pi^{\vec{p},\vec{q}} S(\pi) \left(U_{\pi^{-1}(\vec{p})} - U_{\vec{q}} \right) . \tag{6.43}$$

since the $S(\pi)$ operators are linearly independent, it must hold that for all $\pi \in S_t$, $c_\pi^{\vec{p},\vec{q}} \left(U_{\pi^{-1}(\vec{p})} - U_{\vec{q}} \right) = 0$. So $c_\pi^{\vec{p},\vec{q}} = 0$ unless $\vec{p} = \pi(\vec{q})$.

We thus have that

$$T = \sum_{\substack{\vec{q} \in [K]^t \\ \pi \in S_t}} c_\pi^{\pi(\vec{q}),\vec{q}} \, |\pi(\vec{q})\rangle\langle\vec{q}| \otimes S(\pi) \tag{6.44}$$

$$= \sum_{\substack{\vec{q} \in [K]^t \\ \pi \in S_t}} c_\pi^{\pi(\vec{q}),\vec{q}} \, S(\pi) \left(|\vec{q}\rangle\langle\vec{q}| \otimes \mathbf{I} \right) , \tag{6.45}$$

where in the first line, $S(\pi)$ acts on only the second register, while in the second line $S(\pi)$ acts on both registers together. Thus, T is in the span of operators of the form $M_{\pi,\vec{q}} \coloneqq S(\pi) \left(|\vec{q}\rangle\langle\vec{q}| \otimes \mathbf{I} \right)$. It thus remains to show that $M_{\pi,\vec{q}}$ is in the span of $S(\pi f)$ for $(\pi, f) \in \mathbb{Z}_K \wr S_t$. To this end, observe that

$$|\vec{q}\rangle\langle\vec{q}| \otimes \mathbf{I} = \frac{1}{K^t} \sum_{\vec{f} \in \mathbb{Z}_K^t} \omega_K^{-\vec{f}\cdot\vec{q}} \left(\sum_{\vec{r} \in \mathbb{Z}_K^t} \omega_K^{\vec{f}\cdot\vec{r}} \; |\vec{r}\rangle\langle\vec{r}| \otimes \mathbf{I} \right) \tag{6.46}$$

$$= \frac{1}{K^t} \sum_{\vec{f} \in \mathbb{Z}_K^t} \omega_K^{-\vec{f}\cdot\vec{q}} \, S(f) . \tag{6.47}$$

Thus T is in the span of $S(\pi f)$, which completes the proof. □

Lemma 6.10. *The irreps of both the group $G = U(N)^K$ and the commutant algebra $\mathbb{C}[\mathbb{Z}_K \wr S_t]$ are parameterized by K-tuples of Young diagrams, $\left(\lambda^{(1)}, \lambda^{(2)}, \ldots, \lambda^{(K)}\right)$, where each $\lambda^{(i)} \in \widehat{S}_{t_i} \cap \widehat{U(N)}$, and $\sum_{i \in [K]} t_i = t$.*

Furthermore, the dimensions of the irreps are as follows:

Since the group $U(N)^K$ is a direct product, its irreps are also direct products of irreps of $U(N)$, and so the dimension of an irrep $\vec{\lambda} \coloneqq \left(\lambda^{(1)}, \lambda^{(2)}, \ldots, \lambda^{(K)}\right)$ is the product of the individual dimensions

$$\dim\left(V^{\vec{\lambda}}_{U(N)^K}\right) = \prod_{i \in [K]} \dim\left(V^{\lambda^{(i)}}_{U(N)}\right) . \tag{6.48}$$

Similarly, the dimension of the corresponding irrep of $\mathbb{Z}_K \wr S_t$ is [Str24, Mac98]

$$\dim\left(V^{\vec{\lambda}}_{\mathbb{Z}_K \wr S_t}\right) = \frac{t!}{\prod_{i \in [K]}(t_i!)} \prod_{i \in [K]} \dim\left(V^{\lambda^{(i)}}_{S_{t_i}}\right) = \frac{t!}{\prod_{i \in [K]}(t_i!)} \prod_{i \in [K]} f^{\lambda^{(i)}} . \tag{6.49}$$

Therefore, the irrep ratios are

$$\begin{aligned}\frac{\dim\left(V^{\vec{\lambda}}_{\mathbb{Z}_K \wr S_t}\right)}{\dim\left(V^{\vec{\lambda}}_{U(N)^K}\right)} &= \frac{t!}{\prod_{i\in[K]}(t_i!)} \prod_{i\in[K]} \frac{\dim\left(V^{\lambda^{(i)}}_{S_{t_i}}\right)}{\dim\left(V^{\lambda^{(i)}}_{U(N)}\right)} && (6.50)\\ &= \frac{t!}{\prod_{i\in[K]}(t_i!)} \prod_{i\in[K]} \frac{t_i!}{N^{t_i}} \left(1 \pm O\left(\frac{t_i^2}{N}\right)\right) && \text{(Lemma 3.32)}\\ &= \frac{t!}{N^t}\left(1 \pm O\left(\frac{t^2}{N}\right)\right) && \left(\textstyle\sum_i t_i^2 \le \left(\sum_i t_i\right)^2 = t^2\right)\end{aligned}$$

Note that since the colored permutation algebra is a group algebra, its dual basis (Example 3.18) is proportional to the inverse of each group element:

$$[\pi\, f]^* = \frac{1}{t!\, K^t}[f^{-1}\pi^{-1}]\,. \tag{6.51}$$

6.2.1 Recording Oracle for the Unitary Haar Cipher

We will use the following notation for the registers: Let the adversary's query register be split as $\mathsf{A} \coloneqq (\mathsf{A_k}, \mathsf{A_u})$, corresponding respectively to the key/control register, and the register on which the unitary U_k acts. That is, the query representation acts as $\rho(\vec{U}) = \sum_{k\in[K]} |k\rangle\langle k|_{\mathsf{A_k}} \otimes (U_k)_{\mathsf{A_u}}$.

Similarly, let the purification register R of the recording be split into t registers $\mathsf{R}_1, \ldots, \mathsf{R}_t$, and let each R_i contain a 4-tuple $(\mathsf{R_{X_k},}i, \mathsf{R_{X_u},}i, \mathsf{R_{Y_k},}i, \mathsf{R_{Y_u},}i)$, with $\mathsf{R_{X},}i \coloneqq (\mathsf{R_{X_k},}i, \mathsf{R_{X_u},}i)$, and $\mathsf{R_{Y},}i \coloneqq (\mathsf{R_{Y_k},}i, \mathsf{R_{Y_u},}i)$. Let $\mathsf{R_{X_k}} \coloneqq (\mathsf{R_{X_k},}1, \ldots, \mathsf{R_{X_k},}t)$ $\mathsf{R_{X_u}} \coloneqq (\mathsf{R_{X_u},}1, \ldots, \mathsf{R_{X_u},}t)$, and so forth.

We have from Equation (5.40) that the path recording update rule for the unitary Haar cipher, $G = U(N)^K$, has the form

$$V_{U(N)^K} \coloneqq \left(\Lambda_{\mathbb{Z}_K \wr S_i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathbb{Z}_K \wr S_i}\right)_{\mathsf{R}_{\le i}} \left(\mathrm{App}\right)_{\mathsf{A}\,\mathsf{R}_i} \left(\Lambda^+_{\mathbb{Z}_K \wr S_{i-1}}\right)_{\mathsf{R}_{\le i-1}}, \tag{6.52}$$

and from Corollary 5.18 that a t-query adaptive algorithm Adv_t querying $V_{U(N)^K}$ has the form[47]

$$\mathsf{Adv}_t^{V_{U(N)^K}} = \left(\Lambda_{\mathbb{Z}_K \wr S_t}\right)_{\mathsf{R}} \cdot \left(\Omega_{\mathbb{Z}_K \wr S_t}\right)_{\mathsf{R}} \cdot \mathsf{Adv}_t^{\mathrm{App}}\,. \tag{6.53}$$

To understand the path recording oracle for the unitary Haar cipher, we therefore only need to look at the form of $\Lambda_{\mathbb{Z}_K \wr S_t}$ and $\Omega_{\mathbb{Z}_K \wr S_t}$.

Lemma 6.11 (Factoring Ω along the wreath product)**.**

$$\Omega_{\mathbb{Z}_K \wr S_t} = \Omega_{S_t}\, \Omega_{\mathbb{Z}_K^t} \tag{6.54}$$

Proof. We have that

$$S(\pi\, f) = S(\pi)\, S(f)\,, \tag{6.55}$$

[47] when starting with an empty recording. Otherwise, we of course also have an additional rescaling of $\left(\Lambda^+_{\mathbb{Z}_K \wr S_{|\mathsf{R}_0|}}\right)_{\mathsf{R}_0}$ at the beginning.

and

$$
\begin{aligned}
S\big((\pi f)^*\big)^T &= \frac{1}{t!\,K^t} S\left(f^{-1}\pi^{-1}\right)^T && \text{(by Equation (6.51))} \\
&= \frac{1}{t!\,K^t} \left(S\left(\pi^{-1}\right)\, S\left(f^{-1}\right)\right)^T && (6.56) \\
&= \frac{1}{t!\,K^t} S\left(\pi^{-1}\right)^T\, S\left(f^{-1}\right)^T && (6.57) \\
&= \frac{1}{t!\,K^t} S\left(\pi\right)\, \overline{S\left(f\right)}, && \text{(since } S(\pi) \text{ is real, and } S(f) \text{ is diagonal)}
\end{aligned}
$$

and therefore

$$
\begin{aligned}
\Omega_{\mathbb{Z}_K \wr S_t} &= \sum_{\pi f \in \mathbb{Z}_K \wr S_t} S(\pi f) \otimes S((\pi f)^*)^T && \text{(Lemma 5.11)} \\
&= \frac{1}{t!\,K^t} \sum_{\pi f \in \mathbb{Z}_K \wr S_t} S\big(\pi\big)\, S\big(f\big) \otimes S\left(\pi\right)\, \overline{S\left(f\right)} && (6.58) \\
&= \left(\frac{1}{t!} \sum_{\pi \in S_t} S\big(\pi\big) \otimes S\left(\pi\right)\right) \left(\frac{1}{K^t} \sum_{f \in \mathbb{Z}_K^t} S\big(f\big) \otimes \overline{S\left(f\right)}\right) && (6.59) \\
&= \Omega_{S_t}\, \Omega_{\mathbb{Z}_K^t} && \square
\end{aligned}
$$

Note that the same argument also shows that $\Omega_{\mathbb{Z}_K \wr S_t} = \Omega_{\mathbb{Z}_K^t}\, \Omega_{S_t}$ (just write the elements of $\mathbb{Z}_K \wr S_t$ as $f'\pi$ instead of πf), implying that $\Omega_{\mathbb{Z}_K^t}$ and Ω_{S_t} commute.

Lemma 6.12.

$$
\left(\Omega_{\mathbb{Z}_K^t}\right)_{\mathsf{R}_{\mathsf{X}_k}, \mathsf{R}_{\mathsf{X}_u}, \mathsf{R}_{\mathsf{Y}_k}, \mathsf{R}_{\mathsf{Y}_u}} = \left(\Pi_{\mathsf{X}_k = \mathsf{Y}_k}\right)_{\mathsf{R}_{\mathsf{X}_k}, \mathsf{R}_{\mathsf{Y}_k}} \tag{6.60}
$$

where $\Pi_{\mathsf{X}_k = \mathsf{Y}_k}$ is the projector onto the registers $\mathsf{R}_{\mathsf{X}_k}$ and $\mathsf{R}_{\mathsf{Y}_k}$ being equal in the standard basis.

Proof.

$$
\begin{aligned}
\Omega_{\mathbb{Z}_K^t} &= \frac{1}{K^t} \sum_{f \in \mathbb{Z}_K^t} S\big(f\big) \otimes \overline{S\left(f\right)} && (6.61) \\
&= \frac{1}{K^t} \sum_{f \in \mathbb{Z}_K^t} \sum_{\substack{\vec{k} \in [K]^t \\ \vec{x} \in [N]^t}} \omega_K^{\vec{f} \cdot \vec{k}} \left|\vec{k}, \vec{x}\right\rangle\!\left\langle \vec{k}, \vec{x}\right|_{\mathsf{R}_{\mathsf{X}_k}, \mathsf{R}_{\mathsf{X}_u}} \otimes \sum_{\substack{\vec{k'} \in [K]^t \\ \vec{y} \in [N]^t}} \omega_K^{-\vec{f} \cdot \vec{k'}} \left|\vec{k'}, \vec{y}\right\rangle\!\left\langle \vec{k'}, \vec{y}\right|_{\mathsf{R}_{\mathsf{Y}_k}, \mathsf{R}_{\mathsf{Y}_u}} && (6.62) \\
&= \sum_{\substack{\vec{k}, \vec{k'} \in [K]^t \\ \vec{x}, \vec{y} \in [N]^t}} \underbrace{\left(\frac{1}{K^t} \sum_{f \in \mathbb{Z}_K^t} \omega_K^{\vec{f} \cdot (\vec{k} - \vec{k'})}\right)}_{\delta_{\vec{k}, \vec{k'}}} \left|\vec{k}, \vec{x}\right\rangle\!\left\langle \vec{k}, \vec{x}\right|_{\mathsf{R}_{\mathsf{X}_k}, \mathsf{R}_{\mathsf{X}_u}} \otimes \left|\vec{k'}, \vec{y}\right\rangle\!\left\langle \vec{k'}, \vec{y}\right|_{\mathsf{R}_{\mathsf{Y}_k}, \mathsf{R}_{\mathsf{Y}_u}} && (6.63) \\
&= \sum_{\substack{\vec{k} \in [K]^t \\ \vec{x}, \vec{y} \in [N]^t}} \left|\vec{k}, \vec{x}\right\rangle\!\left\langle \vec{k}, \vec{x}\right|_{\mathsf{R}_{\mathsf{X}_k}, \mathsf{R}_{\mathsf{X}_u}} \otimes \left|\vec{k}, \vec{y}\right\rangle\!\left\langle \vec{k}, \vec{y}\right|_{\mathsf{R}_{\mathsf{Y}_k}, \mathsf{R}_{\mathsf{Y}_u}} && (6.64) \\
&= \sum_{\vec{k} \in [K]^t} \left|\vec{k}, \vec{k}, \right\rangle\!\left\langle \vec{k}, \vec{k}, \right|_{\mathsf{R}_{\mathsf{X}_k}, \mathsf{R}_{\mathsf{Y}_k}} \otimes \mathbf{I}_{\mathsf{R}_{\mathsf{X}_u}, \mathsf{R}_{\mathsf{Y}_u}} && \square
\end{aligned}
$$

We see, therefore, that $\Omega_{\mathbb{Z}_K^t}$ plays the role of ensuring that the control register remains the same before and after the query, and moreover, that the same key is recorded in both the R_X and the R_Y halves of the recording. That is, the recording will always have pairs of the form (k_i, x_i, k_i, y_i). Specifically, by the following lemma, we have that when the R register is already correctly formatted in a state of this form, the update rule without the rescaling proceeds by appending a new tuple of the form (k_i, x_i, k_i, y_i), and symmetrizing over all permutations with the existing tuples.

Lemma 6.13.

$$\Big(\Omega_{\mathbb{Z}_K \wr S_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\operatorname{App}\Big)_{\mathsf{A}\,\mathsf{R}_i} = \Big(\Omega_{S_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\operatorname{App}_{(\mathsf{X}_k=\mathsf{Y}_k)}\Big)_{\mathsf{A}\,\mathsf{R}_i} \Big(\Pi_{\mathsf{X}_k=\mathsf{Y}_k}\Big)_{\mathsf{R}_{\leq i-1}} \tag{6.65}$$

where

$$\operatorname{App}_{(\mathsf{X}_k=\mathsf{Y}_k)} \coloneqq \sum_{k\in[K],\, x,y\in[N]} |k,y\rangle\langle k,x|_\mathsf{A} \otimes |k,x,k,y\rangle_{\mathsf{R}_i} \tag{6.66}$$

Proof.

$$\Big(\Omega_{\mathbb{Z}_K \wr S_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\operatorname{App}\Big)_{\mathsf{A}\,\mathsf{R}_i} = \Big(\Omega_{S_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\Pi_{\mathsf{X}_k=\mathsf{Y}_k}\Big)_{\mathsf{R}_{\leq i}} \Big(\operatorname{App}\Big)_{\mathsf{A}\,\mathsf{R}_i} \tag{6.67}$$

$$= \Big(\Omega_{S_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\Pi_{\mathsf{X}_k=\mathsf{Y}_k}\Big)_{\mathsf{R}_i} \Big(\Pi_{\mathsf{X}_k=\mathsf{Y}_k}\Big)_{\mathsf{R}_{\leq i-1}} \Big(\operatorname{App}\Big)_{\mathsf{A}\,\mathsf{R}_i} \tag{6.68}$$

$$= \Big(\Omega_{S_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\Pi_{\mathsf{X}_k=\mathsf{Y}_k}\Big)_{\mathsf{R}_i} \Big(\operatorname{App}\Big)_{\mathsf{A}\,\mathsf{R}_i} \Big(\Pi_{\mathsf{X}_k=\mathsf{Y}_k}\Big)_{\mathsf{R}_{\leq i-1}} \tag{6.69}$$

$$= \Big(\Omega_{S_i}\Big)_{\mathsf{R}_{\leq i}} \left(\sum_{\substack{k,k'\in[K]\\ x,y\in[N]}} |k',y\rangle\langle k,x|_\mathsf{A} \otimes \Pi_{\mathsf{X}_k=\mathsf{Y}_k} |k,x,k',y\rangle_{\mathsf{R}_i} \right) \Big(\Pi_{\mathsf{X}_k=\mathsf{Y}_k}\Big)_{\mathsf{R}_{\leq i-1}} \tag{6.70}$$

$$= \Big(\Omega_{S_i}\Big)_{\mathsf{R}_{\leq i}} \left(\sum_{\substack{k\in[K]\\ x,y\in[N]}} |k,y\rangle\langle k,x|_\mathsf{A} \otimes |k,x,k,y\rangle_{\mathsf{R}_i} \right) \Big(\Pi_{\mathsf{X}_k=\mathsf{Y}_k}\Big)_{\mathsf{R}_{\leq i-1}} \tag{6.71}$$

$$= \Big(\Omega_{S_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\operatorname{App}_{(\mathsf{X}_k=\mathsf{Y}_k)}\Big)_{\mathsf{A}\,\mathsf{R}_i} \Big(\Pi_{\mathsf{X}_k=\mathsf{Y}_k}\Big)_{\mathsf{R}_{\leq i-1}} \qquad \square$$

Corollary 6.14. *Assuming we start with a validly formatted recording register, we can write the path recording update rule for the unitary Haar cipher as*

$$V_{U(N)^K} = \Big(\Lambda_{\mathbb{Z}_K \wr S_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\Omega_{S_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\operatorname{App}_{(\mathsf{X}_k=\mathsf{Y}_k)}\Big)_{\mathsf{A}\,\mathsf{R}_i} \Big(\Lambda^+_{\mathbb{Z}_K \wr S_{i-1}}\Big)_{\mathsf{R}_{\leq i-1}}, \tag{6.72}$$

and we can write a t-query adaptive algorithm Adv_t querying $V_{U(N)^K}$ starting with an empty recording as

$$\mathsf{Adv}_t^{V_{U(N)^K}} = \Big(\Lambda_{\mathbb{Z}_K \wr S_t}\Big)_{\mathsf{R}} \cdot \Big(\Omega_{S_t}\Big)_{\mathsf{R}} \cdot \mathsf{Adv}_t^{\operatorname{App}_{(\mathsf{X}_k=\mathsf{Y}_k)}}. \tag{6.73}$$

Thus, we have that up to the reweighting operator $\Lambda_{\mathbb{Z}_K \wr S_t}$ (which for the sake of simplicity, we can safely set aside, as it is $O\left(\frac{t^2}{N}\right)$-close to a scalar multiple of the identity, see Lemma 6.10[48]), our

[48] A more careful treatment can show that even the $\Lambda_{\mathbb{Z}_K \wr S_t}$ operator can be properly decomposed, and it is therefore not necessary to incur any error in transforming this path recording to a K-tuple of individual path recordings. Since this is slightly more involved and offers little additional clarity or intuition at this point, we will not do so here. However, we do work this out explicitly in the special case of diagonal unitaries ($N = 1$) below.

recording oracle is encoding multiset states of the form

$$|\{(k_1,x_1,y_1),(k_2,x_2,y_2),\ldots,(k_t,x_t,y_t)\}\rangle \coloneqq \Omega_{S_t}\,|(k_1,x_1,k_1,y_1),(k_2,x_2,k_2,y_2),\ldots,(k_t,x_t,k_t,y_t)\rangle\ , \tag{6.74}$$

and each i'th query adds (k_i,x_i,y_i) to the set (by appending (App) and then reshuffling the order (Ω_{S_t})). We can, of course, alternatively associate each such multiset state with a K-tuple $(R_1,\ldots,R_K)$ of multisets of the form $R_k=\{(x_i,y_i)\mid k_i=k\}$, where the k'th multiset in the K-tuple contains the (x,y) pairs corresponding to querying the k'th unitary.

6.2.2 Diagonal Unitaries

Specializing further, we can take $N=1$, to get the group $U(1)^K$ of K-dimensional diagonal unitaries (the K-torus). In this case, each U_k in Equation (6.22) is simply a unit-norm complex phase, and x_i and y_i are 0-qubit registers, and therefore disappear from the update rule, dropping out of Equation (6.74) to give recording states of the form

$$|\{k_1,k_2,\ldots,k_t\}\rangle :\propto \Omega_{S_t}\,|(k_1,k_1),(k_2,k_2),\ldots,(k_t,k_t)\rangle\ . \tag{6.75}$$

Furthermore, the commutant of $\rho(\vec{U})^{\otimes t}$ is still a representation of $\mathbb{C}[\mathbb{Z}_K\wr S_t]$, but it is no longer a faithful representation: as noted in Lemma 6.10 since $t>N=1$, the irreps $\vec{\lambda}$ that survive are those where each $\lambda^{(i)}$ indexes an irrep of $U(1)$. That is, we have that each $\lambda^{(i)}$ is a Young diagram with a single row. Let $\mathcal{A}_t$ be this quotient algebra, that is $\mathcal{A}_t\coloneqq\mathbb{C}[\mathbb{Z}_K\wr S_t]/\ker(S)$ (where $S(\cdot)$ is the Schur representation of $\mathbb{Z}_K\wr S_t$, see Equation (6.26)).

Lemma 6.15. *Let $G=U(1)^K$ be the group of K dimensional diagonal unitaries, with representation $\rho(\theta_1,\ldots,\theta_K)=\sum_{k\in[K]}e^{-i\theta_k}\,|k\rangle\langle k|$. Then the path recording update rule for this group simplifies to a multiset insertion*

$$V_{U(1)^K}=\sum_{\substack{k\in[K]\\ R\text{ multiset from }[K]}}|k\rangle\langle k|_{\mathsf{A}}\otimes|R\uplus\{k\}\rangle\langle R|_{\mathsf{R}} \tag{6.76}$$

on multiset states of the form $|\{k_1,k_2,\ldots,k_t\}\rangle\coloneqq\frac{1}{\sqrt{t!\prod_{i\in[K]}(t_i!)}}\sum_{\pi\in S_t}|\pi(k_1,k_2,\ldots,k_t)\rangle$, where t_i is the number of times i appears in the multiset.[49]

Proof. We start with the recording oracle for $U(N)^K$ from Corollary 6.14, specialized to the case where $N=1$:

$$V_{U(1)^K}=\left(\Lambda_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}}\left(\Omega_{S_i}\right)_{\mathsf{R}_{\le i}}\left(\mathrm{App}_{(\mathsf{X_k}=\mathsf{Y_k})}\right)_{\mathsf{A}\,\mathsf{R}_i}\left(\Lambda^{+}_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\le i-1}}\ , \tag{6.77}$$

Note that the purification register here contains both an $\mathsf{R_X}=(\mathsf{R_{X1}},\ldots,\mathsf{R_{X}}_t)$ and a corresponding $\mathsf{R_Y}=(\mathsf{R_{Y1}},\ldots,\mathsf{R_{Y}}_t)$. But we know by Corollary 6.14 that $\mathsf{R_X}$ and $\mathsf{R_Y}$ will contain identical strings (recall that we have thrown out the part of R that was not identical when we took $N=1$). So of course, there is redundancy. We can therefore write the update rule here as if it were acting on just a single purification register $\mathsf{R}=(\mathsf{R}_1,\ldots,\mathsf{R}_t)$ containing a string of t k-values, These are off by a simple relabeling on the purification register, so this distinction will not matter, but it will simplify notation, so we will write it here as if it were a single register $\mathsf{R}=(\mathsf{R}_1,\ldots,\mathsf{R}_t)$. We thus

[49] Note that when the multiset contains collisions/repeats, different permutations can give the same basis string, and so the prefactor provides the proper normalization to make this a normalized state.

replace $\Big(\operatorname{App}_{(\mathsf{X_k}=\mathsf{Y_k})}\Big)_{\mathsf{A}\,\mathsf{R}_i} = \sum_{k\in[K]} |k\rangle\langle k|_{\mathsf{A}} \otimes |k,k\rangle_{\mathsf{R_{X}}_i\,\mathsf{R_{Y}}_i}$ with $\Big(\operatorname{App}\Big)_{\mathsf{A}\,\mathsf{R}_i} = \sum_{k\in[K]} |k\rangle\langle k|_{\mathsf{A}} \otimes |k\rangle_{\mathsf{R}_i}$ and write Ω_{S_t} as $\frac{1}{t!}\sum_{\pi\in S_t} S(\pi)$.

Now, we argue about the reweighting operators $\Lambda_{\mathcal{A}_i}$ and $\Lambda^{+}_{\mathcal{A}_{i-1}}$: Consider the irrep ratios in Lemma 6.10. The group $U(1)$ is Abelian, which implies that all the irreps in the denominator of Equation (6.50) have dimension 1, and furthermore, the irreps of $U(1)$ are parameterized by Young diagrams of a single-row, which means that each $\lambda^{(i)}$ in the numerator is the trivial irrep of the corresponding symmetric group S_{t_i}, all of which have dimension 1. We thus have that the irrep ratios simplify to $\frac{t!}{\prod_{i\in[K]}(t_i!)}$, where t_i is the size of $\lambda^{(i)}$, or the number of times that i appears in the recording register. Therefore, the reweighting operator (for this case when $N = 1$) produces a type-dependent scalar:

$$\Lambda_{\mathcal{A}_t} = \sum_{(t_1,\ldots,t_K)} \sqrt{\frac{t!}{\prod_{i\in[K]}(t_i!)}}\,\Pi_{(t_1,\ldots,t_K)}\,, \quad \text{and} \quad \Lambda^{+}_{\mathcal{A}_{t-1}} = \sum_{(t_1,\ldots,t_K)} \sqrt{\frac{\prod_{i\in[K]}(t_i!)}{t!}}\,\Pi_{(t_1,\ldots,t_K)}\,, \tag{6.78}$$

where $\Pi_{(t_1,\ldots,t_K)}$ is the projector onto strings with a type of the form $(t_1,\ldots,t_K)$.

We thus have that

$$\begin{aligned}
\Lambda_{\mathcal{A}_t}\Omega_{S_t} &= \sum_{(t_1,\ldots,t_K)} \sqrt{\frac{t!}{\prod_{i\in[K]}(t_i!)}}\,\Pi_{(t_1,\ldots,t_K)} \cdot \frac{1}{t!}\sum_{\pi\in S_t} S(\pi) && (6.79)\\
&= \sum_{(t_1,\ldots,t_K)} \sqrt{\frac{t!}{\prod_{i\in[K]}(t_i!)}}\cdot \frac{1}{t!}\sum_{\pi\in S_t} S(\pi)\,\Pi_{(t_1,\ldots,t_K)} \\
&\qquad\qquad (\Pi_{(t_1,\ldots,t_K)} \text{ is invariant under permutations}) \\
&= \sum_{(t_1,\ldots,t_K)} \frac{1}{\sqrt{t!\prod_{i\in[K]}(t_i!)}} \sum_{\pi\in S_t} S(\pi)\,\Pi_{(t_1,\ldots,t_K)} && (6.80)\\
&= \sum_{k_1,k_2,\ldots,k_t\in[K]} |\{k_1,k_2,\ldots,k_t\}\rangle\langle k_1,k_2,\ldots,k_t| && (6.81)\\
&=: \sum_{\vec{k}\in[K]^t} |\{\vec{k}\}\rangle\langle\vec{k}|\,. && (6.82)
\end{aligned}$$

So this is the normalized symmetrization operator. Similarly, we have that

$$\begin{aligned}
\Omega_{S_{t-1}}\Lambda^{+}_{\mathcal{A}_{t-1}} &= \frac{1}{(t-1)!}\sum_{\pi\in S_{t-1}} S(\pi) \sum_{(t_1,\ldots,t_K)} \sqrt{\frac{\prod_{i\in[K]}(t_i!)}{(t-1)!}}\,\Pi_{(t_1,\ldots,t_K)} && (6.83)\\
&= \sum_{(t_1,\ldots,t_K)} \sqrt{\frac{\prod_{i\in[K]}(t_i!)}{(t-1)!}}\cdot \frac{1}{(t-1)!}\sum_{\pi\in S_{t-1}} \Pi_{(t_1,\ldots,t_K)}\,S(\pi) \\
&\qquad\qquad (\Pi_{(t_1,\ldots,t_K)} \text{ is invariant under permutations}) \\
&= \sum_{(t_1,\ldots,t_K)} \frac{\prod_{i\in[K]}(t_i!)}{(t-1)!}\sqrt{\frac{1}{(t-1)!\prod_{i\in[K]}(t_i!)}} \sum_{\pi\in S_{t-1}} \Pi_{(t_1,\ldots,t_K)}\,S(\pi) && (6.84)\\
&= \sum_{k_1,k_2,\ldots,k_t\in[K]} \frac{\prod_{i\in[K]}\left(t_i^{(\vec{k})}!\right)}{(t-1)!}\,|k_1,k_2,\ldots,k_t\rangle\langle\{k_1,k_2,\ldots,k_t\}| && (6.85)
\end{aligned}$$

$$=: \sum_{\vec{k}\in[K]^t} \frac{\prod_{i\in[K]}\left(t_i^{(\vec{k})}!\right)}{(t-1)!} |\vec{k}\rangle\langle\{\vec{k}\}| \,. \tag{6.86}$$

where we take $t_i^{(\vec{k})}!$ to be the number of times that a value $i \in [K]$ appears inside $\vec{k}$.

We now combine everything to get the result:

$$V_{U(1)^K} = \left(\Lambda_{\mathcal{A}_i}\right)_{\mathsf{R}_{\leq i}} \left(\Omega_{S_i}\right)_{\mathsf{R}_{\leq i}} \left(\text{App}\right)_{\mathsf{A}\,\mathsf{R}_i} \left(\Lambda^+_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\leq i-1}} \tag{6.87}$$

$$= \left(\Lambda_{\mathcal{A}_i}\right)_{\mathsf{R}_{\leq i}} \left(\Omega_{S_i}\right)_{\mathsf{R}_{\leq i}} \left(\Omega_{S_{i-1}}\right)_{\mathsf{R}_{\leq i-1}} \left(\text{App}\right)_{\mathsf{A}\,\mathsf{R}_i} \left(\Lambda^+_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\leq i-1}} \qquad \text{(by Lemma 5.17)}$$

$$= \left(\Lambda_{\mathcal{A}_i}\right)_{\mathsf{R}_{\leq i}} \left(\Omega_{S_i}\right)_{\mathsf{R}_{\leq i}} \left(\text{App}\right)_{\mathsf{A}\,\mathsf{R}_i} \left(\Omega_{S_{i-1}}\right)_{\mathsf{R}_{\leq i-1}} \left(\Lambda^+_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\leq i-1}}$$

($\Omega_{S_{i-1}}$ and App act on different registers)

$$= \left(\Lambda_{\mathcal{A}_i}\right)_{\mathsf{R}_{\leq i}} \left(\Omega_{S_i}\right)_{\mathsf{R}_{\leq i}} \left(\sum_{\substack{k\in[K]\\ \vec{k}\in[K]^{i-1}}} |k\rangle\langle k|_{\mathsf{A}} \otimes |\vec{k},k\rangle_{\mathsf{R}_{\leq i}}\langle \vec{k}\,|_{\mathsf{R}_{\leq i-1}} \right)_{\mathsf{A}\,\mathsf{R}_i} \left(\Omega_{S_{i-1}}\right)_{\mathsf{R}_{\leq i-1}} \left(\Lambda^+_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\leq i-1}}$$

(definition of App)

$$= \left(\sum_{\vec{k}'\in[K]^{i-1}} |\{\vec{k}'\}\rangle\langle\vec{k}'| \right)_{\mathsf{R}_{\leq i}} \left(\sum_{\substack{k\in[K]\\ \vec{k}\in[K]^{i-1}}} |k\rangle\langle k|_{\mathsf{A}} \otimes |\vec{k},k\rangle_{\mathsf{R}_{\leq i}}\langle \vec{k}\,|_{\mathsf{R}_{\leq i-1}} \right)_{\mathsf{A}\,\mathsf{R}_i} \left(\sum_{\vec{k}''\in[K]^{i-1}} \frac{\prod_{i\in[K]}\left(t_i^{(\vec{k}'')}!\right)}{(t-1)!} |\vec{k}''\rangle\langle\{\vec{k}''\}| \right)_{\mathsf{R}_{\leq i-1}}$$

(by Equations (6.82) and (6.86))

$$= \sum_{\substack{k\in[K]\\ \vec{k}\in[K]^{i-1}}} \frac{\prod_{i\in[K]}\left(t_i^{(\vec{k})}!\right)}{(t-1)!} |k\rangle\langle k|_{\mathsf{A}} \otimes |\{\vec{k}\} \uplus \{k\}\rangle_{\mathsf{R}_{\leq i}}\langle\{\vec{k}\}|_{\mathsf{R}_{\leq i-1}} \tag{6.88}$$

$$= \sum_{\substack{k\in[K]\\ R \text{ multiset from } [K]}} |k\rangle\langle k|_{\mathsf{A}} \otimes |R \uplus \{k\}\rangle_{\mathsf{R}_{\leq i}} \langle R\,|_{\mathsf{R}_{\leq i-1}} \tag{6.89}$$

Where the last equality follows since each multiset $R = \{\vec{k}\}$ corresponds to $\frac{(t-1)!}{\prod_{i\in[K]}\left(t_i^{(\vec{k})}!\right)}$ distinct strings. This completes the proof. □

We can see, therefore, that the path recording oracle for $U(1)^K$ is simply keeping a running total that keeps track of how many times each key k has been queried. This is, of course, the recording oracle of Zhandry [Zha19] for the case of queries to a Haar random diagonal unitary (that is, a function $f : [K] \to U(1)$).

6.3 Zhandry's Compressed Phase Oracle

We can, of course, specialize this further to random Boolean functions, $f : \{0,1\}^n \to \{0,1\}$. Without loss of generality, we can view these functions as being applied in the phase:

$$\rho(f) = \sum_{x\in\{0,1\}^n} (-1)^{f(x)} |x\rangle\langle x| \,. \tag{6.90}$$

The group of such Boolean functions is thus $G = \mathbb{Z}_2^N$ (again, taking $N = 2^n$), with the group operation given by an XOR on the truth tables, and a representation as diagonal unitaries with ± 1 values.

Lemma 6.16. *Let $G = \mathbb{Z}_2^N$ be the group of Boolean functions, with the representation in Equation (6.90) above. Then the path recording update rule for this group simplifies to a set insertion mod 2 (that is, add to the set if it does not yet appear, and remove it if it does):*

$$V_{\mathbb{Z}_2^N} = \sum_{\substack{x \in [N] \\ R \subseteq [N]}} |x\rangle\langle x|_{\mathsf{A}} \otimes |R \oplus \{x\}\rangle\langle R|_{\mathsf{R}} \,. \tag{6.91}$$

This is the compressed oracle of [Zha19] for Boolean functions applied in the phase.

We remark that Zhandry's compressed oracle can also be derived from the tableau-recording oracle: since all irreducible representations of a finite Abelian group are one-dimensional (i.e., characters of the group), the relevant Clebsch-Gordan transforms all simplify to scalars. Nevertheless, we prove Lemma 6.16 to demonstrate how the path-recording oracle can be used to derive a wide variety of existing compressed oracles, and to improve our general-purpose understanding of the path-recording oracle.

Proof. We can easily compute the commutant of $\rho(f)^{\otimes t}$ as follows. Let $\vec{f}$ be the bit string encoding of f as $(f(1), f(2), \ldots, f(N)) \in \{0,1\}^N$, and similarly let $m_2(x) \in \{0,1\}^N$ be the one-hot encoding of $x \in [N]$. Then $f(x) = \vec{f} \cdot m_2(x)$. We write

$$\begin{aligned}
\rho(f)^{\otimes t} &= \sum_{x_1,\ldots,x_t \in [N]} (-1)^{f(x_1) \oplus \cdots \oplus f(x_t)} \, |x_1, \ldots, x_t\rangle\langle x_1, \ldots, x_t| && (6.92)\\
&= \sum_{x_1,\ldots,x_t \in [N]} (-1)^{\vec{f} \cdot (m_2(x_1) \oplus \cdots \oplus m_2(x_t))} \, |x_1, \ldots, x_t\rangle\langle x_1, \ldots, x_t| && (6.93)\\
&= \sum_{\boldsymbol{x} \in [N]^t} (-1)^{\vec{f} \cdot (m_2(\boldsymbol{x}))} \, |\boldsymbol{x}\rangle\langle \boldsymbol{x}| \,. && (6.94)
\end{aligned}$$

where $\boldsymbol{x} \coloneqq (x_1, \ldots, x_t)$, and $m_2(\boldsymbol{x}) \coloneqq m_2(x_1) \oplus \cdots \oplus m_2(x_t)$ is the count mod 2 of $\boldsymbol{x}$.

By extending it linearly, we can view the tensor power representation $\rho(f)^{\otimes t}$ as a representation of the corresponding group algebra $\mathbb{C}[\mathbb{Z}_2^N]$: For each algebra element $a \in \mathbb{C}[\mathbb{Z}_2^N]$, where $a = \sum_{f \in \mathbb{Z}_2^N} a_f \, f$, we have that

$$R(a) \coloneqq \sum_{f \in \mathbb{Z}_2^N} a_f \, \rho(f)^{\otimes t} \,. \tag{6.95}$$

Now consider the Fourier basis of the group algebra:

$$\left\{ \hat{\mu} \coloneqq \frac{1}{2^N} \sum_{f \in \mathbb{Z}_2^N} (-1)^{\vec{f} \cdot \vec{\mu}} \, f \right\}_{\mu \in \mathbb{Z}_2^N} . \tag{6.96}$$

The representations $R(\hat{\mu})$ of the Fourier basis form another basis for the span of the tensor power

operators. We can compute that

$$R(\hat{\mu}) = \frac{1}{2^N} \sum_{f \in \mathbb{Z}_2^N} (-1)^{\vec{f} \cdot \vec{\mu}} \, \rho(f)^{\otimes t} \tag{6.97}$$

$$= \frac{1}{2^N} \sum_{\substack{f \in \mathbb{Z}_2^N \\ \boldsymbol{x} \in [N]^t}} (-1)^{\vec{f} \cdot (\vec{\mu} \oplus m_2(\boldsymbol{x}))} \, |\boldsymbol{x}\rangle\langle\boldsymbol{x}| \tag{6.98}$$

$$= \sum_{\substack{\boldsymbol{x} \in [N]^t \\ m_2(\boldsymbol{x}) = \vec{\mu}}} |\boldsymbol{x}\rangle\langle\boldsymbol{x}| \tag{6.99}$$

Note that $m_2(\boldsymbol{x})$ always has Hamming weight at most t, so this is 0 whenever μ has Hamming weight larger than t. Otherwise, if $|\mu| \leq t$ we interpret μ as a subset of $[N]$ and observe that $R(\hat{\mu})$ is a projector onto strings $x_1, \ldots, x_t$ such that each $x \in \mu$ appears an odd number of times in the string and each $x \notin \mu$ appears an even number (potentially zero times). Note further that the subspaces corresponding to these projectors are mutually orthogonal: $R(\hat{\mu})R(\hat{\nu}) = \delta_{\mu\nu} R(\hat{\mu})$.

Their commutant $\mathcal{A}_t$ is therefore the algebra of operators that preserve these subspaces, and is spanned by operators of the form

$$\left\{ E^{\mu}_{\boldsymbol{a},\boldsymbol{b}} \coloneqq |\boldsymbol{a}\rangle\langle\boldsymbol{b}| \;\middle|\; m_2(\boldsymbol{a}) = m_2(\boldsymbol{b}) = \mu \right\}_{\boldsymbol{a},\boldsymbol{b} \in [N]^t,\, \mu \subseteq [N],} \tag{6.100}$$

This basis of the commutant algebra already presents it as a direct sum of matrix algebras. That is, this is the Fourier basis of the algebra. We can thus directly see that the irreps of this algebra are parameterized by $\mu \subseteq [N]$ such that $|\mu| \leq t$, and furthermore, that the dimension of irrep μ is the total number of ways to choose $\boldsymbol{a} \in [N]^t$ such that $m_2(a) = \mu$:

$$\dim(V^{\mu}_{\mathcal{A}_t}) = \sum_{\substack{t_1, \ldots, t_N \geq 0 \\ t_i \text{ odd iff } i \in \mu \\ \sum_i t_i = t}} \frac{t!}{t_1! \ldots t_N!} \tag{6.101}$$

$$= \sum_{\substack{\vec{t} \in \mathbb{Z}^N_{\geq 0},\, |\vec{t}| = t \\ \vec{\mu} = \vec{t} \mod 2}} \frac{t!}{t_1! \ldots t_N!} \,. \tag{6.102}$$

By Lemma 3.31, Item 1, we have that $S\left(\left(E^{\mu}_{\boldsymbol{a},\boldsymbol{b}}\right)^*\right)^T = \frac{1}{\dim(V^{\mu}_{\mathcal{A}_t})} S\left(E^{\mu}_{\boldsymbol{a},\boldsymbol{b}}\right)$.

We thus have that the reweighting and symmetrization operators have the following form:

$$\Omega_{\mathcal{A}_t} = \sum_{\substack{\mu \subseteq [N],\, \boldsymbol{a},\boldsymbol{b} \in [N]^t \\ m_2(\boldsymbol{a}) = m_2(\boldsymbol{b}) = \mu}} \frac{1}{\dim(V^{\mu}_{\mathcal{A}_t})} \, S\left(E^{\mu}_{\boldsymbol{a},\boldsymbol{b}}\right) \otimes S\left(E^{\mu}_{\boldsymbol{a},\boldsymbol{b}}\right) \tag{6.103}$$

$$= \sum_{\substack{\boldsymbol{a},\boldsymbol{b} \in [N]^t \\ m_2(\boldsymbol{a}) = m_2(\boldsymbol{b})}} \frac{1}{\dim\left(V^{m_2(\boldsymbol{a})}_{\mathcal{A}_t}\right)} \, |\boldsymbol{a}\rangle\langle\boldsymbol{b}| \otimes |\boldsymbol{a}\rangle\langle\boldsymbol{b}| \,, \tag{6.104}$$

and as before, since we now have that the two registers are always equal in the standard basis, we will abuse notation and write it as if it were a single register

$$\Omega_{\mathcal{A}_t} = \sum_{\substack{\boldsymbol{a},\boldsymbol{b}\in[N]^t \\ m_2(\boldsymbol{a})=m_2(\boldsymbol{b})}} \frac{1}{\dim\left(V_{\mathcal{A}_t}^{m_2(\boldsymbol{a})}\right)} |\boldsymbol{a}\rangle\langle\boldsymbol{b}| \,. \tag{6.105}$$

$$\Lambda_{\mathcal{A}_t} = \sum_{\mu\subseteq[N]} \sqrt{\frac{\dim(V_{\mathcal{A}_t}^{\mu})}{\dim(V_{\mathbb{Z}_2^N}^{\mu})}}\,\Pi_\mu \qquad \Lambda^+_{\mathcal{A}_t} = \sum_{\mu\subseteq[N]} \sqrt{\frac{\dim(V_{\mathbb{Z}_2^N}^{\mu})}{\dim(V_{\mathcal{A}_t}^{\mu})}}\,\Pi_\mu \tag{6.106}$$

$$= \sum_{\mu\subseteq[N]} \sqrt{\dim(V_{\mathcal{A}_t}^{\mu})}\,\Pi_\mu \qquad = \sum_{\mu\subseteq[N]} \frac{1}{\sqrt{\dim(V_{\mathcal{A}_t}^{\mu})}}\Pi_\mu \tag{6.107}$$

where $\Pi_\mu \coloneqq \sum_{\boldsymbol{x}\in[N]^t,\, m_2(\boldsymbol{x})=\vec{\mu}} |\boldsymbol{x}\rangle\langle\boldsymbol{x}|$ are the projectors from Equation (6.99), and the second line follows since $\mathbb{Z}_2^N$ is an Abelian group, so its irreps are all 1-dimensional.

For $\boldsymbol{x}\in[N]^t$, let $\{\boldsymbol{x}\}_{\text{mod }2}\subseteq[N]$ be the set corresponding to $m_2(\boldsymbol{x})$. That is, it is the set of values appearing an odd number of times in $\boldsymbol{x}$. Let

$$|\{\boldsymbol{x}\}_{\text{mod }2}\rangle \coloneqq \frac{1}{\sqrt{\dim\left(V_{\mathcal{A}_t}^{m_2(\boldsymbol{a})}\right)}} \sum_{\substack{\boldsymbol{y}\in[N]^t \\ m_2(\boldsymbol{y})=m_2(\boldsymbol{x})}} |\boldsymbol{y}\rangle \tag{6.108}$$

be the uniform superposition over strings of the same parity.

We then have that

$$\Lambda_{\mathcal{A}_t}\Omega_{\mathcal{A}_t} = \sum_{\substack{\boldsymbol{a},\boldsymbol{b}\in[N]^t \\ m_2(\boldsymbol{a})=m_2(\boldsymbol{b})}} \sqrt{\dim\left(V_{\mathcal{A}_t}^{m_2(\boldsymbol{a})}\right)}\,\frac{1}{\dim\left(V_{\mathcal{A}_t}^{m_2(\boldsymbol{a})}\right)} |\boldsymbol{a}\rangle\langle\boldsymbol{b}| \tag{6.109}$$

$$= \sum_{\substack{\boldsymbol{a},\boldsymbol{b}\in[N]^t \\ m_2(\boldsymbol{a})=m_2(\boldsymbol{b})}} \frac{1}{\sqrt{\dim\left(V_{\mathcal{A}_t}^{m_2(\boldsymbol{a})}\right)}} |\boldsymbol{a}\rangle\langle\boldsymbol{b}| \tag{6.110}$$

$$= \sum_{\boldsymbol{x}\in[N]^t} |\{\boldsymbol{x}\}_{\text{mod }2}\rangle\langle\boldsymbol{x}| \,, \tag{6.111}$$

and similarly,

$$\Omega_{\mathcal{A}_{t-1}}\Lambda^+_{\mathcal{A}_{t-1}} = \sum_{\substack{\boldsymbol{a},\boldsymbol{b}\in[N]^{t-1} \\ m_2(\boldsymbol{a})=m_2(\boldsymbol{b})}} \frac{1}{\sqrt{\dim\left(V_{\mathcal{A}_{t-1}}^{m_2(\boldsymbol{a})}\right)}}\,\frac{1}{\dim\left(V_{\mathcal{A}_{t-1}}^{m_2(\boldsymbol{a})}\right)} |\boldsymbol{a}\rangle\langle\boldsymbol{b}| \tag{6.112}$$

$$= \sum_{\boldsymbol{x}\in[N]^{t-1}} \frac{1}{\dim\left(V_{\mathcal{A}_{t-1}}^{m_2(\boldsymbol{x})}\right)} |\boldsymbol{x}\rangle\langle\{\boldsymbol{x}\}_{\text{mod }2}| \,. \tag{6.113}$$

Let $d_{\boldsymbol{x}} \coloneqq \frac{1}{\dim\left(V_{\mathcal{A}_{t-1}}^{m_2(\boldsymbol{x})}\right)}$ for brevity. As before, we now combine these to conclude:

$$V_{\mathbb{Z}_2^N} = \Big(\Lambda_{\mathcal{A}_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\Omega_{\mathcal{A}_i}\Big)_{\mathsf{R}_{\leq i}} \Big(\text{App}\Big)_{\mathsf{A}\,\mathsf{R}_i} \Big(\Lambda^+_{\mathcal{A}_{i-1}}\Big)_{\mathsf{R}_{\leq i-1}} \tag{6.114}$$

$$= \left(\Lambda_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\le i-1}} \left(\text{App}\right)_{\mathsf{A}\,\mathsf{R}_i} \left(\Lambda^+_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\le i-1}} \qquad \text{(by Lemma 5.17)}$$

$$= \left(\Lambda_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\text{App}\right)_{\mathsf{A}\,\mathsf{R}_i} \left(\Omega_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\le i-1}} \left(\Lambda^+_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\le i-1}}$$

($\Omega_{\mathcal{A}_{i-1}}$ and App act on different registers)

$$= \left(\Lambda_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\Omega_{\mathcal{A}_i}\right)_{\mathsf{R}_{\le i}} \left(\sum_{\substack{x\in[N]\\ \boldsymbol{x}\in[N]^{i-1}}} |x\rangle\langle x|_{\mathsf{A}} \otimes |\boldsymbol{x}, x\rangle_{\mathsf{R}_{\le i}} \langle \boldsymbol{x}\,|_{\mathsf{R}_{\le i-1}} \right)_{\mathsf{A}\,\mathsf{R}_i} \left(\Omega_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\le i-1}} \left(\Lambda^+_{\mathcal{A}_{i-1}}\right)_{\mathsf{R}_{\le i-1}}$$

(definition of App)

$$= \left(\sum_{\boldsymbol{x}'\in[N]^{i-1}} |\{\boldsymbol{x}'\}_{\text{mod }2}\rangle\langle \boldsymbol{x}'| \right)_{\mathsf{R}_{\le i}} \left(\sum_{\substack{x\in[N]\\ \boldsymbol{x}\in[N]^{i-1}}} |x\rangle\langle x|_{\mathsf{A}} \otimes |\boldsymbol{x}, x\rangle_{\mathsf{R}_{\le i}} \langle \boldsymbol{x}\,|_{\mathsf{R}_{\le i-1}} \right)_{\mathsf{A}\,\mathsf{R}_i} \left(\sum_{\boldsymbol{x}''\in[N]^{i-1}} \frac{1}{d_{\boldsymbol{x}''}} |\boldsymbol{x}''\rangle\langle\{\boldsymbol{x}''\}_{\text{mod }2}| \right)_{\mathsf{R}_{\le i-1}}$$

(by Equations (6.111) and (6.113))

$$= \sum_{\substack{x\in[N]\\ \boldsymbol{x}\in[N]^{i-1}}} \frac{1}{d_{\boldsymbol{x}}} |x\rangle\langle x|_{\mathsf{A}} \otimes |\{\boldsymbol{x}\}_{\text{mod }2} \oplus \{x\}\rangle_{\mathsf{R}_{\le i}} \langle\{\boldsymbol{x}\}|_{\mathsf{R}_{\le i-1}} \tag{6.115}$$

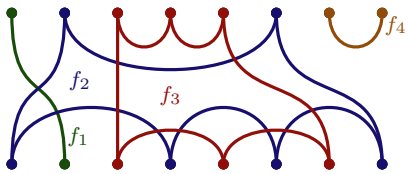
$$= \sum_{\substack{x\in[N]\\ R\subseteq[N]}} |x\rangle\langle x|_{\mathsf{A}} \otimes |R \oplus \{x\}\rangle_{\mathsf{R}_{\le i}} \langle R\,|_{\mathsf{R}_{\le i-1}} \tag{6.116}$$

Where the last equality follows since each multiset $R = \{\boldsymbol{x}\}_{\text{mod }2}$ corresponds to $d_{\boldsymbol{x}}$ distinct strings. This completes the proof. □

It turns out that there is another characterization of this commutant algebra in terms of colored even-partition diagrams.

Definition 6.17 (Colored Even-Partition Algebra)**.** The *even-partition algebra* is the algebra of partition diagrams in which every component has even number of vertices. This is equivalent to the Tanabe algebra $T_t(N, 2)$ [Tan97].

The $\mathbb{Z}_N$*-colored even-partition algebra* is the algebra that has a basis of even-partition diagrams with each component assigned a color in $\mathbb{Z}_N$ (see Figure 17 for an example of such a colored diagram).[50] Diagram multiplication accumulates the colors of each of the components that get merged together.



(6.117)

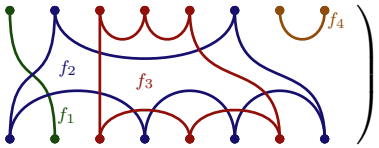
Figure 17: Colored even-partition diagram on 8 elements with colors $f_1, f_2, f_3, f_4 \in \mathbb{Z}_K$ attached to the four components.

[50] This diagram algebra is closely related to the G-colored partition algebras studied by [Blo03].

The Schur representation on $(\mathbb{C}^N)^{\otimes t}$ differs from that of the partition algebra by attaching a phase of $\omega_N^{f_i x_i}$ to the sum over x_i of a component of color f_i (similar to Schur representation of $\mathbb{Z}_K \wr S_t$ on $(\mathbb{C}^{KN})^{\otimes t}$).

For example,

$$S\left(\quad\right) = \sum_{x_1,x_2,x_3,x_4 \in [N]} \omega_K^{f_1 x_1 + \cdots + f_4 x_4} \left|x_1,x_2,x_3,x_3,x_3,x_2,x_4,x_4\right\rangle \left\langle x_2,x_1,x_3,x_2,x_3,x_2,x_3,x_2\right| . \quad (6.118)$$



Lemma 6.18. *The commutant algebra of $\rho(f)^{\otimes t}$ is isomorphic to the* colored even-partition algebra.

Proof sketch. This comes from simply performing a change of basis on the basis of Equation (6.100) in much the same way as in Equation (6.46). □

7 PC is a Pseudorandom Unitary

In this section, we prove that any algorithm making t queries cannot distinguish a Haar-random unitary from a random Clifford circuit (or arbitrary 2-design) composed with a random permutation with advantage better than $O(t^2/N)$. Throughout this section, we assume $t \leq c\sqrt{N}$ for a sufficiently small absolute constant $c > 0$. In particular, this guarantees that the partition algebra $P_t(N)$ is semisimple and all $O(t^2/N)$ error terms below are smaller than an absolute constant.

In order to prove this, we carefully analyze the path-recording oracles $Q^{\text{Path}}_{U(N)}$, $Q^{\text{Path}}_{S_N}$ for the defining representation of $U(N)$ and the in-place representation of S_N. In particular, we prove that they are close in trace distance on the distinct, nonplussed subspace DNP (Definition 7.5), and prove that a random 2-design C causes almost all of the recording state to land in DNP. Proving the first of these statements, in turn, requires two main steps: an analysis of the dual elements of permutations $[\pi]^* \in P_t(N)$ within the partition algebra, and an analysis of λ-irrep dimensions in $(U(N), S_t)$ and $(S_N, P_t(N))$ Schur-Weyl duality.

7.1 The Distinct, Nonplussed Subspace

7.1.1 Preliminaries

Lemma 7.1 ([MPSY24, MH25])**.** *Let C be any unitary 2-design on $\mathbb{C}^N$. Then,*

$$\mathbb{E}_C\left[\operatorname{Tr}\left(\Pi_{\mathsf{Dist}_{N,t}} C^{\otimes t} \rho(C^\dagger)^{\otimes t}\right)\right] \geq 1 - \frac{t(t-1)}{N+1} \quad (7.1)$$

Definition 7.2. Let $|+\rangle$ denote the uniform superposition on the standard basis of $\mathbb{C}^N$, and define $W = \mathbf{I} - |+\rangle\langle+|$, i.e. the projector away from $|+\rangle$. For a system of t registers, the *nonplussed subspace*[51] $\mathsf{NoPlus}_{N,t} = W^{\otimes t}$ is obtained by projecting away the uniform superposition on each register.

An orthonormal basis for $\mathsf{NoPlus}_{N,t}$ is given by

$$\left\{ \frac{1}{N^{t/2}} \sum_{\vec{x} \in [N]^t} \omega_N^{\langle \vec{x}, \vec{y} \rangle} |\vec{x}\rangle \right\}_{(y_1,\ldots,y_t) \in (\mathbb{Z}_N \setminus \{0\})^t} \quad (7.2)$$

[51] In other words, we are completely unfazed by states in this subspace.

Another basis for $\mathsf{NoPlus}_{N,t}$ that will often be convenient for us (but not orthonormal) is given by

$$\left\{ \bigotimes_{i=1}^{t} \left(|1\rangle - |x_i\rangle \right) \right\}_{(x_1,\ldots,x_t)\in([N]\setminus\{1\})^t} \tag{7.3}$$

Note that, while not orthonormal, these basis states are linearly independent, since only the state corresponding to $\vec{x}$ has support on $|\vec{x}\rangle$.

Lemma 7.3. *Let C be sampled from any unitary 1-design on $\mathbb{C}^N$. Then,*

$$\mathbb{E}_C \left[\mathrm{Tr} \left(\Pi_{\mathsf{NoPlus}_{N,t}} C^{\otimes t} \rho (C^\dagger)^{\otimes t} \right) \right] \geq 1 - \frac{t}{N} \tag{7.4}$$

Proof. Since $\{|+\rangle\langle+|_i\}_{i=1}^t$ are commuting projectors, we can use the operator form of the union bound:

Lemma 7.4 ([KOMW19])**.** *Let ρ be a density operator, and let $\Pi_1, \ldots, \Pi_m$ be commuting orthogonal projectors. Then*

$$1 - \mathrm{Tr}\left(\Pi_m \cdots \Pi_1 \rho \Pi_1 \cdots \Pi_m\right) \leq \sum_{i=1}^{m} \mathrm{Tr}\left((\mathbf{I} - \Pi_i)\rho\right). \tag{7.5}$$

which implies that

$$W^{\otimes t} = (\mathbf{I} - |+\rangle\langle+|)^{\otimes t} \succeq \mathbf{I} - \sum_{i=1}^{t} |+\rangle\langle+| \tag{7.6}$$

Therefore,

$$\mathrm{Tr}\left(\Pi_{\mathsf{NoPlus}_{N,t}} C^{\otimes t} \rho (C^\dagger)^{\otimes t} \right) \geq 1 - \sum_{i=1}^{t} \mathrm{Tr}\left(|+\rangle\langle+|_i \, C^{\otimes t} \rho (C^\dagger)^{\otimes t} \right) \tag{7.7}$$

$$\implies \mathbb{E}_C \left[\mathrm{Tr} \left(\Pi_{\mathsf{NoPlus}_{N,t}} C^{\otimes t} \rho (C^\dagger)^{\otimes t} \right) \right] \geq 1 - \sum_{i=1}^{t} \mathbb{E}_C \left[\mathrm{Tr}\left(|+\rangle\langle+|_i \, C^{\otimes t} \rho (C^\dagger)^{\otimes t} \right) \right] \tag{7.8}$$

Now, consider the summand on the right hand side. By the cyclic property of trace and linearity of expectation, the summand on the right hand side is equivalent to

$$\mathrm{Tr}\left(\mathbb{E}_C [(C^\dagger)^{\otimes t} |+\rangle\langle+|_i \, C^{\otimes t} \rho] \right) = \mathrm{Tr}\left(\mathbb{E}_C [(C^\dagger |+\rangle\langle+| \, C)_i] \rho \right) \tag{7.9}$$

For any 1-design, $\mathbb{E}_C[C^\dagger \sigma C] = \frac{\mathrm{Tr}(\sigma)}{N}\mathbf{I}$. Since $\sigma = |+\rangle\langle+|$ has trace 1, Eq. (7.9) simplifies to $\mathrm{Tr}(\rho)/N = 1/N$. Plugging this into the summand in Eq. (7.8) completes the proof. □

Definition 7.5. The *distinct nonplussed subspace* $\mathsf{DNP}_{N,t} := \mathsf{Dist}_{N,t} \cap \mathsf{NoPlus}_{N,t}$ is the intersection of the distinct and nonplussed subspaces.

Lemma 7.6. *The distinct nonplussed subspace $\mathsf{DNP}_{N,t}$ has dimension $N^t \left(1 - O\left(\frac{t^2}{N}\right)\right)$. In particular, this means that the fully mixed state has probability at least $1 - O\left(\frac{t^2}{N}\right)$ of being in the distinct nonplussed subspace.*

Proof. We can take the following to be a partial and non-orthogonal basis for the distinct nonplussed subspace:

$$S = \left\{ \bigotimes_{i=1}^{t} \left(|i\rangle - |x_i\rangle \right) \right\}_{(x_1,\ldots,x_t)\in([N]\setminus\{1,\ldots,t\})^t_{\text{dist}}} \tag{7.10}$$

As in Equation (7.3), these basis states are linearly independent (though not orthogonal) because only the state corresponding to $\vec{x}$ has support on $|\vec{x}\rangle$. Furthermore, S is contained in the distinct nonplussed subspace: Fix $|\psi\rangle \in S$. Clearly, $|\psi\rangle \in \mathsf{Dist}_{N,t}$, since all t registers are supported on pairwise disjoint sets of computational basis vectors. Moreover,

$$\begin{aligned} W(|i\rangle - |x_i\rangle) &= |i\rangle - |x_i\rangle + (\langle +|i\rangle - \langle +|x_i\rangle)\,|+\rangle && (7.11)\\ &= |i\rangle - |x_i\rangle + \left(\frac{1}{\sqrt{N}} - \frac{1}{\sqrt{N}}\right)|+\rangle && (7.12)\\ &= |i\rangle - |x_i\rangle && (7.13)\\ \implies W^{\otimes t}\,|\psi\rangle &= |\psi\rangle && (7.14) \end{aligned}$$

and so $|\psi\rangle \in \mathsf{NoPlus}_{N,t}$ as well. So the size of S is a lower bound on the dimension of the distinct nonplussed subspace. This gives us that

$$\begin{aligned} \dim(\mathsf{DNP}_{N,t}) \geq |S| &= (N-t)(N-t-1)\ldots(N-2t+1) && (7.15)\\ &\geq N^t\left(1 - O\left(\frac{t^2}{N}\right)\right). && \square \end{aligned}$$

7.1.2 Unitary 2-designs and the Distinct Nonplussed Subspace

In this section, we will show that applying a unitary 2-design with high probability maps an arbitrary state $\rho \in (\mathbb{C}^N)^{\otimes t}$ to a state that has high overlap with the distinct nonplussed subspace. More precisely, we will prove the following theorem:

Theorem 7.7. *Let $\Pi_{\mathsf{DNP}_{N,t}}$ be the projection onto the distinct nonplussed subspace, and C be any unitary 2-design on $\mathbb{C}^N$. Then,*

$$\mathbb{E}_C\left[\mathrm{Tr}\left(\Pi_{\mathsf{DNP}_{N,t}} C^{\otimes t}\rho(C^\dagger)^{\otimes t}\right)\right] \geq 1 - O\left(\frac{t^2}{N}\right) \tag{7.16}$$

for any state $\rho \in (\mathbb{C}^N)^{\otimes t}$.

To prove Theorem 7.7, we first collect some facts about projectors:

Lemma 7.8 ([Gao15])**.** *Let ρ be a density operator, and let $\Pi_1, \ldots, \Pi_m$ be orthogonal projectors. Then*

$$1 - \mathrm{Tr}\left(\Pi_m \cdots \Pi_1 \rho \Pi_1 \cdots \Pi_m\right) \leq 4\sum_{i=1}^{m} \mathrm{Tr}\left((\mathbf{I} - \Pi_i)\rho\right). \tag{7.17}$$

Equivalently,

$$\mathrm{Tr}\left(\Pi_m \cdots \Pi_1 \rho \Pi_1 \cdots \Pi_m\right) \geq 1 - 4\sum_{i=1}^{m} \mathrm{Tr}\left((\mathbf{I} - \Pi_i)\rho\right). \tag{7.18}$$

Definition 7.9. Assume that A and B are subspaces of $\mathbb{C}^N$, and let $C \coloneqq A \cap B$ be the intersection. The *Friedrichs angle* $c(A, B)$ is defined as

$$c(A,B) \coloneqq \sup_{\substack{|x\rangle \in A \cap C^\perp \\ |y\rangle \in B \cap C^\perp \\ \||x\rangle\| = \||y\rangle\| = 1}} |\langle x|y\rangle| \tag{7.19}$$

Informally, the Friedrichs angle is the highest overlap between two unit vectors in A and B, excluding the intersection. For example, when Π_A and Π_B are commuting projectors, $c(A, B) = 0$. Notice that since $|y\rangle \in B \cap C^\perp$,

$$\|\Pi_A |y\rangle\| = \sup_{\substack{|a\rangle \in A \\ \||a\rangle\| = 1}} |\langle a|y\rangle| = \sup_{\substack{|a\rangle \in A \cap C^\perp \\ \||a\rangle\| = 1}} |\langle a|y\rangle| \tag{7.20}$$

$$\implies \|\Pi_A |_{B \cap C^\perp}\|^2 = \sup_{\substack{|y\rangle \in B \cap C^\perp \\ \||y\rangle\| = 1}} |\langle y|\Pi_A|y\rangle| \tag{7.21}$$

$$= \sup_{\substack{|x\rangle \in A \cap C^\perp \\ |y\rangle \in B \cap C^\perp \\ \||x\rangle\| = \||y\rangle\| = 1}} |\langle y|x\rangle\langle x|y\rangle| = c(A,B)^2 \tag{7.22}$$

We make use of two more characterizations of $c(A, B)$. First,

Claim 7.10. *$c(A, B)$ is the largest singular value of $\Pi_A \Pi_B$ aside from 1.*

Proof. Since A decomposes into $C \oplus (A \cap C^\perp)$, we have that

$$\Pi_A \Pi_B = \Pi_C \Pi_B + \Pi_{A \cap C^\perp} \Pi_B = \Pi_C + \Pi_{A \cap C^\perp} \Pi_B.$$

Moreover, decomposing $B = C \oplus (B \cap C^\perp)$, we can write

$$\Pi_A \Pi_B = \Pi_C + \Pi_{A \cap C^\perp} \Pi_C + \Pi_{A \cap C^\perp} \Pi_{B \cap C^\perp} = \Pi_C + \Pi_{A \cap C^\perp} \Pi_{B \cap C^\perp}.$$

Thus, $\Pi_A \Pi_B$ block-diagonalizes according to the $(C, A \cap C^\perp)$, $(C, B \cap C^\perp)$ vector space decompositions. $c(A, B)$ is defined to be the top singular value of the $(A \cap C^\perp, B \cap C^\perp)$ block, so this proves the claim. □

Kayalar and Weinert proved that the Friedrichs angle captures the difference between first projecting onto A, then projecting onto B, as opposed to a single projection onto $A \cap B$:

Theorem 7.11 ([KW89])**.**

$$||\Pi_A \Pi_B - \Pi_{A \cap B}|| = c(A,B). \tag{7.23}$$

From Theorem 7.11, we can derive a related bound which will be useful in our setting:

Lemma 7.12.

$$\|\Pi_B \Pi_A \Pi_B - \Pi_{A \cap B}\| \leq c(A,B)^2 \tag{7.24}$$

Proof.

$$\begin{aligned}
\|\Pi_B\Pi_A\Pi_B - \Pi_{A\cap B}\| &= \|\Pi_B\Pi_A\Pi_A\Pi_B - \Pi_{A\cap B}\| && (\text{since } \Pi_A^2 = \Pi_A)\\
&= \|\Pi_B\Pi_A\Pi_A\Pi_B - \Pi_{A\cap B} - \Pi_{A\cap B} + \Pi_{A\cap B}\| && (7.25)\\
&= \|\Pi_B\Pi_A\Pi_A\Pi_B - \Pi_B\Pi_A\Pi_{A\cap B} - \Pi_{A\cap B}\Pi_A\Pi_B + \Pi_{A\cap B}\| \\
& && (\text{since } \Pi_A\Pi_{A\cap B} = \Pi_{A\cap B})\\
&= \|(\Pi_B\Pi_A - \Pi_{A\cap B})(\Pi_A\Pi_B - \Pi_{A\cap B})\| && (7.26)\\
&\leq \|\Pi_B\Pi_A - \Pi_{A\cap B}\|\ \|\Pi_A\Pi_B - \Pi_{A\cap B}\| && (\text{submultiplicativity})\\
&= \|\Pi_A\Pi_B - \Pi_{A\cap B}\|^2 && (7.27)\\
&\leq c(A,B)^2 && (\text{Theorem 7.11})
\end{aligned}$$

□

Corollary 7.13.

$$|\mathrm{Tr}(\Pi_A\Pi_B\rho\Pi_B\Pi_A) - \mathrm{Tr}(\Pi_{A\cap B}\rho\Pi_{A\cap B})| \leq c(A,B)^2. \tag{7.28}$$

Proof.

$$\begin{aligned}
|\mathrm{Tr}(\Pi_A\Pi_B\rho\Pi_B\Pi_A) - \mathrm{Tr}(\Pi_{A\cap B}\rho\Pi_{A\cap B})| &= |\mathrm{Tr}(\Pi_B\Pi_A\Pi_A\Pi_B\rho) - \mathrm{Tr}(\Pi_{A\cap B}\Pi_{A\cap B}\rho)| && (7.29)\\
&= |\mathrm{Tr}(\Pi_B\Pi_A\Pi_B\rho) - \mathrm{Tr}(\Pi_{A\cap B}\rho)| && (7.30)\\
&= \left|\mathrm{Tr}\left((\Pi_B\Pi_A\Pi_B - \Pi_{A\cap B})\rho\right)\right| && (7.31)\\
&\leq ||\Pi_B\Pi_A\Pi_B - \Pi_{A\cap B}||\ \|\rho\|_1 && (\text{Hölder's inequality})\\
&= ||\Pi_B\Pi_A\Pi_B - \Pi_{A\cap B}|| && (\|\rho\|_1 = 1)\\
&= c(A,B)^2 && (\text{Lemma 7.12})
\end{aligned}$$

□

Finally, we will make use of one more fact about the Friedrichs angle:

Lemma 7.14. $\Pi_A\Pi_B + \Pi_B\Pi_A \succeq -c(A,B)\Big(\Pi_A + \Pi_B\Big)$

Proof. By Jordan's lemma, Π_A and Π_B are simultaneously block-diagonalizable into one- and two-dimensional subspaces where on the two dimensional subspaces we have

$$\Pi_A = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \qquad \Pi_B = \begin{pmatrix} \cos^2\theta & \cos\theta\sin\theta \\ \cos\theta\sin\theta & \sin^2\theta \end{pmatrix}$$

for some angle $0 < \theta < \frac{\pi}{2}$. The one-dimensional subspaces together span $A \cap B$, $A^\perp \cap B$, $A \cap B^\perp$, and $A^\perp \cap B^\perp$. Thus, in this language, the Friedrichs angle is given by

$$c(A,B) = \max\cos(\theta),$$

where the maximum is taken over two-dimensional Jordan blocks. But on any fixed Jordan block, one can see that

$$\Pi_A\Pi_B + \Pi_B\Pi_A + \cos(\theta)(\Pi_A + \Pi_B) = \cos(\theta)\cdot\begin{pmatrix} (1+\cos(\theta))^2 & \sin\theta(1+\cos(\theta)) \\ \sin\theta(1+\cos(\theta)) & \sin^2\theta \end{pmatrix} \succeq 0,$$

which proves the claim. □

Our first goal will be to show that $c(\mathsf{Dist}_{N,t}, \mathsf{NoPlus}_{N,t})^2$ is small when $t \ll N$:

Lemma 7.15.

$$c(\mathsf{Dist}_{N,t}, \mathsf{NoPlus}_{N,t})^2 \leq \frac{t(t-1)}{N-t+1}. \tag{7.32}$$

Proof. Define the averaging operator A_i as follows:

$$A_i |\vec{x}\rangle = \begin{cases} \dfrac{1}{N-t+1} \displaystyle\sum_{a \notin \{x_j : j \neq i\}} |x_1, \ldots, x_{i-1}, a, x_{i+1}, \ldots, x_t\rangle & \text{if } \vec{x} \in [N]^t_{\text{dist}}, \\ |\vec{x}\rangle\,, & \text{otherwise.} \end{cases} \tag{7.33}$$

From the definition, it is clear that A_i is an orthogonal projector, whose normalized $+1$ eigenvectors within the distinct subspace are given by

$$\frac{1}{\sqrt{N-t+1}} \sum_{a \notin \{z_j : j \neq i\}} |z_1, \ldots, z_{i-1}, a, z_{i+1}, \ldots, z_t\rangle\,, \quad \vec{z} = (z_1, \ldots, z_{i-1}, z_{i+1}, \ldots, z_t) \in [N]^{t-1}_{\text{dist}} \tag{7.34}$$

and a (overcomplete) basis for $\ker A_i$ (which is within the distinct subspace) is given by

$$\frac{1}{\sqrt{2}}(|z_1, \ldots, z_{i-1}, a, z_{i+1} \ldots, z_t\rangle - |z_1, \ldots, z_{i-1}, b, z_{i+1} \ldots, z_t\rangle), \quad a \neq b \tag{7.35}$$

We will be interested in the (Hermitian, PSD) operator

$$A = \sum_i A_i.$$

We first show that the kernel of A is exactly the intersection $\mathsf{Dist}_{N,t} \cap \mathsf{NoPlus}_{N,t}$:

Claim 7.16.

$$\mathsf{Dist}_{N,t} \cap \mathsf{NoPlus}_{N,t} = \bigcap_{i=1}^{t} \ker A_i = \ker A \tag{7.36}$$

Proof. We know that $\bigcap_{i=1}^t \ker A_i = \ker A$ because the A_i are all PSD.

For the main claim, by definition of $\mathsf{NoPlus}_{N,t}$, $\mathsf{NoPlus}_{N,t} = \bigcap_{i=1}^t \operatorname{Im} W^{(i)}$. So it is sufficient for us to show that for each i,

$$\ker A_i = \mathsf{Dist} \cap \operatorname{Im} W^{(i)} \tag{7.37}$$

For the forward inclusion, it is clear by definition of A_i that $\ker A_i \subseteq \mathsf{Dist}$. Moreover, for any state $|\phi\rangle$ in Eq. (7.35), $\operatorname{Tr}_{[t]\setminus i}(|\phi\rangle\langle\phi|) = \frac{1}{\sqrt{2}}(|a\rangle - |b\rangle)$ is orthogonal to $|+\rangle$, and therefore in $\operatorname{Im} W$. Thus

$$\ker A_i \subseteq \mathsf{Dist} \cap \operatorname{Im} W^{(i)}. \tag{7.38}$$

Conversely, if $|\psi\rangle \notin \ker A_i$, then either $|\psi\rangle$ has support on a non-distinct t-tuple (in which case $|\psi\rangle \notin \mathsf{Dist}$), or has support on a basis state $|\phi\rangle$ in Eq. (7.34), which satisfies

$$\operatorname{Tr}_{[t]\setminus i}(|\phi\rangle\langle\phi|) = \frac{1}{\sqrt{N-t+1}} \sum_{a \notin \{z_j : j \neq i\}} |a\rangle \tag{7.39}$$

which has nonzero overlap with $|+\rangle \in \ker W$. Hence, $|\phi\rangle \notin \operatorname{Im} W^{(i)}$, and so

$$|\psi\rangle \notin \ker A_i \implies |\psi\rangle \notin \mathsf{Dist} \cap \operatorname{Im} W^{(i)}. \tag{7.40}$$

□

Next, we calculate the pairwise Friedrichs angles $c(\mathrm{Im}(A_i), \mathrm{Im}(A_j))$ for $i \neq j$.

Claim 7.17. *$c(\mathrm{Im}(A_i), \mathrm{Im}(A_j)) \leq \frac{1}{N-t+1}$ for all $i < j$.*

Proof. The projections A_i and A_j are simultaneously block-diagonal in the coordinates $\ell \notin \{i, j\}$. Thus, it suffices to prove the claim restricting to states of the

$$\bigotimes_{\ell \notin \{i,j\}} |x_\ell\rangle \otimes \left|\psi_{x_{\neq i,j}}\right\rangle$$

where $\left|\psi_{x_{\neq i,j}}\right\rangle$ is on the ith and jth registers. Restricted to these states and looking only at registers i and j, $\mathrm{Im}(A_i)$ is spanned by states of the form

$$\left|\psi^{(i)}_{x_j}\right\rangle = \frac{1}{\sqrt{N-t+1}} \sum_{x'_i \notin \{x_\ell\}_{\ell \neq i}} \left|x'_i\right\rangle |x_j\rangle$$

for all $x_j \notin \{x_\ell\}_{\ell \notin \{i,j\}}$, while $\mathrm{Im}(A_j)$ is similarly spanned by

$$\left|\psi^{(j)}_{x_i}\right\rangle = \frac{1}{\sqrt{N-t+1}} \sum_{x'_j \notin \{x_\ell\}_{\ell \neq j}} |x_i\rangle \left|x'_j\right\rangle$$

for all $x_i \notin \{x_\ell\}_{\ell \notin \{i,j\}}$. The corresponding Gram matrix of all such inner products $\left\langle \psi^{(j)}_{x_i} \middle| \psi^{(i)}_{x_j} \right\rangle$ is given by

$$\frac{1}{N-t+1}\Big(J_{N-t+2} - \mathbf{I}_{N-t+2}\Big),$$

where J denotes an all ones matrix, as the diagonal terms are indeed zero while the off-diagonal terms are $\frac{1}{N-t+1}$. By Claim 7.10, the $c(\mathrm{Im}(A_i), \mathrm{Im}(A_j))$ is the largest singular value of this matrix aside from 1. Because J has rank 1, all remaining singular values of the Gram matrix are $\frac{1}{N-t+1}$, so the claim follows. □

In order to finish proving Lemma 7.15, we will need to establish one more key claim.

Claim 7.18. *$A \succeq (1 - \frac{t-1}{N-t+1})(\mathbf{I} - \Pi_{\ker(A)})$.*

Proof. By Claim 7.17 and Lemma 7.14, we have that

$$A_i A_j + A_j A_i \succeq -\frac{1}{N-t+1}\Big(A_i + A_j\Big)$$

for all $i \neq j$. Thus, we can calculate

$$A^2 = A + \sum_{i<j} (A_i A_j + A_j A_i) \tag{7.41}$$

$$\succeq A - \frac{1}{N-t+1} \sum_{i<j} (A_i + A_j) \tag{7.42}$$

$$= (1 - \frac{t-1}{N-t+1})A. \tag{7.43}$$

Since A is Hermitian PSD, this means that all nonzero eigenvalues of A are at least $(1 - \frac{t-1}{N-t+1})$ and so

$$A \succeq (1 - \frac{t-1}{N-t+1})(\mathbf{I} - \Pi_{\ker A}).$$ □

We are finally ready to prove Lemma 7.15. Fix any $|\psi\rangle \in \mathsf{NoPlus}_{N,t} \cap (\mathsf{DNP}_{N,t})^{\perp}$, and define $|\phi\rangle := \Pi_{\mathsf{Dist}_{N,t}} |\psi\rangle$ and $|\eta\rangle := (\mathbf{I} - \Pi_{\mathsf{Dist}_{N,t}}) |\psi\rangle$. By Eq. (7.22), bounding $c(\mathsf{Dist}_{N,t}, \mathsf{NoPlus}_{N,t})^2$ reduces to bounding

$$\langle\psi|\Pi_{\mathsf{Dist}_{N,t}}|\psi\rangle = \langle\phi|\phi\rangle = \langle\phi|\mathbf{I} - \Pi_{\mathsf{DNP}_{N,t}}|\phi\rangle \leq \frac{N-t+1}{N-2t+2}\sum_{i=1}^{t}\langle\phi|A_i|\phi\rangle = \frac{N-t+1}{N-2t+2}\sum_{i=1}^{t}\|A_i|\phi\rangle\|^2 \tag{7.44}$$

where the second equality holds because $|\phi\rangle$ is orthogonal to $\mathsf{DNP}_{N,t}$: for any $|\gamma\rangle \in \mathsf{DNP}_{N,t}$,

$$\langle\gamma|\phi\rangle = \langle\gamma|\Pi_{\mathsf{Dist}_{N,t}}|\psi\rangle = \langle\gamma|\psi\rangle = 0 \tag{7.45}$$

Now, we turn to bounding $\|A_i|\phi\rangle\|^2$. For any $\vec{z} \in [N]^{t-1}$ and $a \in [N]$, let $\alpha_{\vec{z},a}$ be the coefficient of $|z_1, \ldots, z_{i-1}, a, z_{i+1}, \ldots, z_t\rangle$ when expanding $|\psi\rangle$ in the computational basis. Since $|\psi\rangle \in \mathsf{NoPlus}_{N,t}$,

$$\sum_{a\in[N]} \alpha_{\vec{z},a} = 0 \implies -\sum_{a\in\{z_j : j\neq i\}} \alpha_{\vec{z},a} = \sum_{a\notin\{z_j : j\neq i\}} \alpha_{\vec{z},a} \tag{7.46}$$

But then,

$$\begin{aligned}
A_i|\phi\rangle &= A_i\left(\Pi_{\mathsf{Dist}_{N,t}}|\psi\rangle\right) && (7.47)\\
&= A_i\left(\sum_{\vec{z}\in[N]^{t-1}_{\mathrm{dist}}}\sum_{a\notin\{z_j : j\neq i\}} \alpha_{\vec{z},a}|z_1,\ldots,z_{i-1},a,z_{i+1},\ldots,z_t\rangle\right) && (7.48)\\
&= \frac{1}{N-t+1}\sum_{\vec{z}\in[N]^{t-1}_{\mathrm{dist}}}\sum_{b\notin\{z_j : j\neq i\}}\sum_{a\notin\{z_j : j\neq i\}} \alpha_{\vec{z},a}|z_1,\ldots,z_{i-1},b,z_{i+1},\ldots,z_t\rangle && (7.49)\\
&= -\frac{1}{N-t+1}\sum_{\vec{z}\in[N]^{t-1}_{\mathrm{dist}}}\sum_{b\notin\{z_j : j\neq i\}}\sum_{a\in\{z_j : j\neq i\}} \alpha_{\vec{z},a}|z_1,\ldots,z_{i-1},b,z_{i+1},\ldots,z_t\rangle && (7.50)\\
\implies \|A_i|\phi\rangle\|^2 &= \frac{1}{(N-t+1)^2}\sum_{\vec{z}\in[N]^{t-1}_{\mathrm{dist}}}\sum_{b\notin\{z_j : j\neq i\}}\left|\sum_{a\in\{z_j : j\neq i\}} \alpha_{\vec{z},a}\right|^2 && (7.51)\\
&= \frac{1}{N-t+1}\sum_{\vec{z}\in[N]^{t-1}_{\mathrm{dist}}}\left|\sum_{a\in\{z_j : j\neq i\}} \alpha_{\vec{z},a}\right|^2 && (7.52)\\
&\leq \frac{t-1}{N-t+1}\sum_{\vec{z}\in[N]^{t-1}_{\mathrm{dist}}}\sum_{a\in\{z_j : j\neq i\}} \left|\alpha_{\vec{z},a}\right|^2 && (7.53)\\
&\leq \frac{t-1}{N-t+1}\||\eta\rangle\|^2 = \frac{t-1}{N-t+1}\left(1 - \||\phi\rangle\|^2\right) && (7.54)
\end{aligned}$$

On the first line we expanded out the definition of $|\phi\rangle$, and on the second line we expanded out the application of A_i. On the third line, we plug in Eq. (7.46), the fourth is by the definition of the Euclidean norm, and the fifth uses that $|\{z_j : j \neq i\}| = t-1$. The sixth line follows by an application of Cauchy Schwarz, and on the last line, we use that the remaining sum is over non-distinct basis states and is therefore bounded by $\||\eta\rangle\|^2 = 1 - \||\phi\rangle\|^2$.

We therefore conclude that

$$\begin{aligned}
\||\phi\rangle\|^2 &\leq \frac{N-t+1}{N-2t+2}\sum_{i=1}^{t}\|A_i\,|\phi\rangle\|^2 && \text{(by Equation (7.44))}\\
&\leq \frac{t(t-1)}{N-2t+2}(1-\||\phi\rangle\|^2) && \text{(by Equation (7.54))}
\end{aligned}$$

Rearranging gives

$$\||\phi\rangle\|^2 \leq \frac{t(t-1)}{N-2t+2+t(t-1)} \leq \frac{t(t-1)}{N-t+1}.$$

Since this holds for every $|\psi\rangle$ in the optimization in Eq. (7.22), this gives the claimed bound, which completes the proof of Lemma 7.15. □

We can now prove Theorem 7.7:

Proof of Theorem 7.7. First,

$$\begin{aligned}
&\mathbb{E}_C\left[\mathrm{Tr}\left(\Pi_{\mathsf{DNP}_{N,t}}C^{\otimes t}\rho(C^\dagger)^{\otimes t}\right)\right] && (7.55)\\
&= \mathbb{E}_C\left[\mathrm{Tr}\left(\Pi_{\mathsf{DNP}_{N,t}}C^{\otimes t}\rho(C^\dagger)^{\otimes t}\Pi_{\mathsf{DNP}_{N,t}}\right)\right] && (7.56)\\
&\geq \mathbb{E}_C\left[\mathrm{Tr}\left(\Pi_{\mathsf{Dist}_{N,t}}\Pi_{\mathsf{NoPlus}_{N,t}}C^{\otimes t}\rho(C^\dagger)^{\otimes t}\Pi_{\mathsf{NoPlus}_{N,t}}\Pi_{\mathsf{Dist}_{N,t}}\right)\right] - c(\mathsf{Dist}_{N,t},\mathsf{NoPlus}_{N,t})^2 \\
& && \text{(by Corollary 7.13)}\\
&\geq \mathbb{E}_C\left[\mathrm{Tr}\left(\Pi_{\mathsf{Dist}_{N,t}}\Pi_{\mathsf{NoPlus}_{N,t}}C^{\otimes t}\rho(C^\dagger)^{\otimes t}\Pi_{\mathsf{NoPlus}_{N,t}}\Pi_{\mathsf{Dist}_{N,t}}\right)\right] - \frac{t(t-1)}{N-t+1} && \text{(by Lemma 7.15)}\\
&\geq 1 - 4\cdot\mathbb{E}_C\left[\mathrm{Tr}\left((\mathbf{I}-\Pi_{\mathsf{Dist}_{N,t}})C^{\otimes t}\rho(C^\dagger)^{\otimes t}\right)\right] && (7.57)\\
&\qquad - 4\cdot\mathbb{E}_C\left[\mathrm{Tr}\left((\mathbf{I}-\Pi_{\mathsf{NoPlus}_{N,t}})C^{\otimes t}\rho(C^\dagger)^{\otimes t}\right)\right] - \frac{t(t-1)}{N-t+1} && \text{(by Lemma 7.8)}\\
&= 1 - 4\left(1-\mathbb{E}_C\left[\mathrm{Tr}\left(\Pi_{\mathsf{Dist}_{N,t}}C^{\otimes t}\rho(C^\dagger)^{\otimes t}\right)\right]\right) && (7.58)\\
&\qquad - 4\left(1-\mathbb{E}_C\left[\mathrm{Tr}\left(\Pi_{\mathsf{NoPlus}_{N,t}}C^{\otimes t}\rho(C^\dagger)^{\otimes t}\right)\right]\right) - \frac{t(t-1)}{N-t+1} && \text{(since } \mathrm{Tr}\left(C^{\otimes t}\rho(C^\dagger)^{\otimes t}\right)=1\text{)}\\
&\geq 1 - \frac{4t(t-1)}{N+1} - \frac{4t}{N} - \frac{t(t-1)}{N-t+1} && \text{(by Lemmas 7.1 and 7.3)}\\
&\geq 1 - \frac{6t^2}{N-t+1} && (7.59)
\end{aligned}$$

which is $1-O(t^2/N)$. □

7.2 The Action of the Partition Algebra on the Distinct Nonplussed Subspace

Lemma 7.19. *Let $D \in \mathcal{B}(P_t(N))$ be a partition diagram, and let $S(\cdot)$ denote the Schur representation (Definition 3.16) of $P_t(N)$. Then $S(D)\,\Pi_{\mathsf{DNP}_{N,t}} = 0$ if and only if D is not a permutation diagram.*

Proof. The "only if" direction is straightforward, since $S(D)$ is unitary if D is a permutation. For the "if" direction, assuming that D is not a permutation diagram, there are one of two possible cases to consider:

a) D has a connected component with at least two vertices in the input row, and/or

b) D has a singleton in the input row.

This is because if no two input vertices are in the same connected component, and every input vertex is propagating (no singletons), then the input vertices must each connect to a distinct output vertex, producing a permutation.

First, assume that D has a connected component C_1 containing two input vertices i_{in} and j_{in}. By definition of the Schur representation, all summands in Eq. (3.10) will have $\langle y_i| = \langle y_j|$, which implies that $S(D)\,|\vec{z}\rangle = 0$ unless $z_i = z_j$. Therefore, $S(D)\,|\psi\rangle = 0$ for all $|\psi\rangle \in \mathsf{Dist}_{N,t}$.

Next, assume that D has a singleton i_{in} in the input row. Again using Eq. (3.10), this implies that $S(D)$ factors as $V \otimes \langle +|_i$, where $V : (\mathbb{C}^N)^{\otimes t-1} \to (\mathbb{C}^N)^{\otimes t}$. $V \otimes \langle +|_i$ annihilates all states that are orthogonal to $|+\rangle_i$, including all states of the form $|\psi\rangle \otimes (|1\rangle - |z_t\rangle)$ (Eq. (7.3)). Therefore, $S(D)\,|\psi\rangle = 0$ for all $|\psi\rangle \in \mathsf{NoPlus}_{N,t}$. Combining the two cases, it follows that

$$S(D)\,\Pi_{\mathsf{DNP}_{N,t}} = 0 \tag{7.60}$$

□

Lemma 7.20. $\mathsf{Schur}_{P_t(N)}$ *maps the distinct nonplussed subspace* $\mathsf{DNP}_{N,t}$ *to the subspace spanned by* $(S_N, P_t(N))$*-Schur basis states* $|\lambda, X, Y\rangle$ *with* $|\lambda| = t$.

In order to prove Lemma 7.20, we first show that an irrep of $P_t(N)$ which annihilates all non-permutation diagrams must correspond to a full-box irrep, i.e. $|\lambda| = t$:

Lemma 7.21. *Let* $J \trianglelefteq P_t(N)$ *be the two-sided ideal spanned by all non-permutation diagrams in* $\mathcal{B}(P_t(N))$.[52] *Then,* $\rho_\lambda(J) = 0$ *if and only if* $|\lambda| = t$.

Proof. The "if" direction follows from [FND26] Theorem 4.17, which proves that if $|\lambda| = t$, then $\rho_\lambda(D) = 0$ for all diagrams D with fewer than t propagating components.[53] The only diagrams in $\mathcal{B}(P_t(N))$ with t propagating components are permutation diagrams, which implies that $\rho_\lambda(j) = 0$ for all $j \in J$.

To prove the reverse direction, note that there is a bijection between the irreps of $P_t(N)/J$ and the irreps of $P_t(N)$ which annihilate J ([EH18], Lemma 2.37, Lemma 3.5):

$$\{\lambda : \lambda \in \widehat{P_t(N)},\quad \rho_\lambda(j) = 0 \text{ for all } j \in J\} \quad\Longleftrightarrow\quad \{\lambda' : \lambda' \in \widehat{P_t(N)/J}\} \tag{7.61}$$

Since $P_t(N)/J \simeq \mathbb{C}[S_t]$ ([HR05], Equation 2.11), the number of irreducible representations of $P_t(N)$ which annihilate J is equal to the number of irreducible representations of $\mathbb{C}[S_t]$. These are indexed by partitions of t. Since the t-box irreducible representations of $P_t(N)$ already annihilate J, it follows that no other irreducible representation of $P_t(N)$ (with $|\lambda| < t$) can annihilate J. □

Proof of Lemma 7.20. We begin with the forward direction, and let $j \in J$, where J is defined as in Lemma 7.21. For any $|\psi\rangle \in (\mathbb{C}^N)^{\otimes t}$, we can write the state $S(j)\,|\psi\rangle$ in the $(S_N, P_t(N))$-Schur basis as follows:

[52] That J is in fact a two-sided ideal is a consequence of [HR05], Equation 1.3.

[53] A connected component is *propagating* if it intersects both the input and output rows of the diagram.

$$|\psi\rangle = \sum_{\substack{\lambda \in \widehat{P_t(N)}, \\ X \in \mathcal{B}(V_{S_N}^{(N-|\lambda|,\lambda)}), \\ Y \in \mathcal{B}(V_{P_t(N)}^{\lambda})}} \alpha_{\lambda,X,Y} |\lambda, X, Y\rangle \tag{7.62}$$

$$\implies S(j)|\psi\rangle = \sum_{\substack{\lambda \in \widehat{P_t(N)}, \\ X \in \mathcal{B}(V_{S_N}^{(N-|\lambda|,\lambda)}), \\ Y \in \mathcal{B}(V_{P_t(N)}^{\lambda})}} \alpha_{\lambda,X,Y} |\lambda, X\rangle \otimes \rho_\lambda(j)|Y\rangle \tag{7.63}$$

Next, fix some λ with $|\lambda| < t$. By Lemma 7.21, there is some $k \in J$ such that $\rho_\lambda(k) \neq 0$. Therefore, $\rho_\lambda(J)$ is a non-zero ideal of $\rho_\lambda(A)$. But since λ is an irrep, $\rho_\lambda(A)$ is a simple algebra, which has no non-zero proper ideals. Therefore, $\rho_\lambda(J) = \rho_\lambda(A)$, and so there exists $j_\lambda \in J$ such that $\rho_\lambda(j_\lambda) = \rho_\lambda(\mathbf{1}) = \mathbf{I}_{V_{P_t(N)}^{\lambda}}$, where $\mathbf{1}$ is the identity diagram in A.

Setting $j = j_\lambda$ in Eq. (7.63) implies that the coefficient on $|\lambda, X, Y\rangle$ in $S(j_\lambda)|\psi\rangle$ equal to $\alpha_{\lambda,X,Y}$. But, if $|\psi\rangle \in \mathsf{DNP}_{N,t}$, then Lemma 7.19 implies that $S(j_\lambda)|\psi\rangle = 0$, so $\alpha_{\lambda,X,Y} = 0$. Since λ was chosen arbitrarily among all irreps with fewer than t boxes, it follows that in the $(S_N, P_t(N))$-Schur basis, $|\psi\rangle$ is supported entirely on full-box irreps.

For the reverse direction, assume that in the $(S_N, P_t(N))$-Schur basis, $|\psi\rangle$ is supported entirely on states $|\lambda, X, Y\rangle$ with $|\lambda| = t$. By Lemma 7.21, $\rho_\lambda(J) = 0$, and so the block-diagonal decomposition on the right hand side of Eq. (7.63) implies that $S(j)|\psi\rangle = 0$ for all $j \in J$. In particular, $S(B_{ij})|\psi\rangle = 0$, where B_{ij} is the diagram with a single block containing the input and output vertices in columns i and j, with all other blocks propagating vertically along a single column. From Definition 3.16, it follows $S(B_{ij})$ is the equality projector between registers i and j, and therefore

$$\forall i, j,\ S(B_{ij})|\psi\rangle = 0 \implies |\psi\rangle \in \mathsf{Dist}_{N,t}. \tag{7.64}$$

Similarly, we also have that $S(P_i)|\psi\rangle = 0$, where P_i is the diagram containing singleton in the ith column, again all other blocks propagating vertically along a single column. In this case, $S(P_i)$ is proportional to the projector onto $\mathbf{I} \otimes |+\rangle_i$, and so

$$\forall i,\ S(P_i)|\psi\rangle = 0 \implies |\psi\rangle \in \mathsf{NoPlus}_{N,t} \tag{7.65}$$

Combining the equations above, we conclude that $|\psi\rangle \in \mathsf{Dist}_{N,t} \cap \mathsf{NoPlus}_{N,t} = \mathsf{DNP}_{N,t}$. □

7.2.1 The Dual of a Permutation Element

Next, we prove a lemma which relates the dual of a permutation $\pi \in \mathbb{C}[S_t]$ with the dual of the same permutation, but in the partition algebra $P_t(N)$:

Lemma 7.22. *For any $\pi \in \mathcal{B}(\mathbb{C}[S_t])$, let $\pi^{*_{P_t(N)}}$ denote the dual element of π with respect to $P_t(N)$, and let $\pi^{*_{\mathbb{C}[S_t]}}$ denote the dual element of π with respect to $\mathbb{C}[S_t]$. For any $\sigma \in \mathcal{B}(\mathbb{C}[S_t])$,*

$$\langle \pi^{*_{\mathbb{C}[S_t]}}, \sigma \rangle_2 = \langle \pi^{*_{P_t(N)}}, \sigma \rangle_2 \tag{7.66}$$

Equivalently, the expansions of $\pi^{*_{\mathbb{C}[S_t]}}$ and $\pi^{*_{P_t(N)}}$ in the basis $\mathcal{B}(P_t(N))$ agree on all permutations $\sigma \in \mathcal{B}(\mathbb{C}[S_t])$. (Consequently, among all elements of S_t only π^{-1} appears in this expansion.)

Proof. Since $\mathbb{C}[S_t]$ is a group algebra, $\pi^{*_{\mathbb{C}[S_t]}} = \pi^{-1}/|S_t|$. So, it suffices to show that

$$\langle \pi^{*_{P_t(N)}}, \sigma \rangle_2 = \frac{\delta_{\pi^{-1},\sigma}}{|S_t|} = \frac{\delta_{\pi^{-1},\sigma}}{t!} \tag{7.67}$$

As in Lemma 7.20, we let $J \trianglelefteq P_t(N)$ be the two-sided ideal spanned by all non-permutation diagrams, so that we have the canonical quotient map

$$\begin{aligned} q : P_t(N) &\longrightarrow P_t(N)/J, \\ a &\longmapsto a + J. \end{aligned} \tag{7.68}$$

Moreover, since $P_t(N)$ is semisimple, there exists a central idempotent[54] e such that [Lam01]:

$$P_t(N) \simeq eP_t(N) \oplus J \tag{7.69}$$

Therefore, $eP_t(N) \simeq P_t(N)/J$, with one possible isomorphism given by restricting the quotient map in Eq. (7.68):

$$\begin{aligned} q_e : eP_t(N) &\xrightarrow{\sim} P_t(N)/J, \\ ea &\longmapsto ea + J. \end{aligned} \tag{7.70}$$

By definition, $q_e(e) = q_e(e \cdot \mathbf{1}) = e + J$. Moreover, e is a unit in $eP_t(N)$, so e also maps to $\mathbf{1} + J$, where $\mathbf{1}$ is the identity permutation. Therefore, $\mathbf{1} + J = e + J$, which implies that $e = \mathbf{1} + j$ for some $j \in J$.

Using that $e = \mathbf{1} + j$, we now show that $\pi^{*_{P_t(N)}} = \pi^{-1}e/t!$:

$$\begin{aligned}
\langle D, \pi^{*_{P_t(N)}} \rangle_{\mathrm{reg}_{P_t(N)}} &= \frac{1}{t!} \langle D, \pi^{-1} e \rangle_{\mathrm{reg}_{P_t(N)}} && (7.71) \\
&= \frac{1}{t!} \sum_{E \in \mathcal{B}(P_t(N))} \langle E, D\pi^{-1} e \cdot E \rangle_2 && (\text{Definition of } \langle \cdot, \cdot \rangle_{\mathrm{reg}_{P_t(N)}}) \\
&= \frac{1}{t!} \sum_{E \in \mathcal{B}(P_t(N))} \langle E, eD\pi^{-1} E \rangle_2 && (e \text{ is central in } P_t(N)) \\
&= \frac{1}{t!} \cdot \delta_{D \in \mathcal{B}(\mathbb{C}[S_t])} \sum_{\sigma \in \mathcal{B}(\mathbb{C}[S_t])} \langle \sigma, eD\pi^{-1}\sigma \rangle_2 && (eP_t(N) \cap J = \{0\}, \text{ Eq. (7.69)}) \\
&= \frac{1}{t!} \cdot \delta_{D \in \mathcal{B}(\mathbb{C}[S_t])} \sum_{\sigma \in \mathcal{B}(\mathbb{C}[S_t])} \langle \sigma, (\mathbf{1} + j)D\pi^{-1}\sigma \rangle_2 && (7.72) \\
&= \frac{1}{t!} \cdot \delta_{D \in \mathcal{B}(\mathbb{C}[S_t])} \sum_{\sigma \in \mathcal{B}(\mathbb{C}[S_t])} \langle \sigma, D\pi^{-1}\sigma \rangle_2 && (\langle \sigma, k \rangle = 0 \text{ for all } k \in J) \\
&= \frac{1}{t!} \cdot \delta_{D \in \mathcal{B}(\mathbb{C}[S_t])} \cdot (\delta_{D,\pi} \cdot t!) && (\text{Eq. (3.12)}) \\
&= \delta_{D,\pi} && (7.73)
\end{aligned}$$

[54] A central idempotent $e \in A$ is an element satisfying $ea = ae$ for all $a \in A$, and $e^2 = e$. If a semisimple algebra A decomposed into a direct sum $I \oplus J$ of two ideals, then both I and J must be of the form eA for some central idempotent e. See the discussion at the start of Chapter 22 in [Lam01].

Finally, since $\pi^{*_{P_t(N)}} = \pi^{-1} e / t!$,

$$\begin{aligned}
\langle \pi^{*_{P_t(N)}}, \sigma \rangle_2 &= \frac{1}{t!} \langle \pi^{-1} e, \sigma \rangle_2 && (7.74)\\
&= \frac{1}{t!} \langle \pi^{-1}(\mathbf{1} + j), \sigma \rangle_2 && (7.75)\\
&= \frac{1}{t!} \langle \pi^{-1}, \sigma \rangle_2 && (\langle \sigma, k \rangle = 0 \text{ for all } k \in J)\\
&= \frac{1}{t!} \cdot \delta_{\pi^{-1}, \sigma} && (7.76)
\end{aligned}$$

This matches Eq. (7.67), concluding the proof. □

7.3 Additional Properties of the Update Operators

We give some useful properties of the update operators in the path-recording basis of Section 5.

Lemma 7.23 (Λ_{S_t} commutes with permutation-respecting maps on R_X or R_Y)**.** *Let M_{R_X} be a linear transformation such that $M \cdot S(\pi) = S(\pi) \cdot M$ for every permutation $\pi \in S_t$. Then, $M_{\mathsf{R}_X} \otimes \mathbf{I}_{\mathsf{R}_Y}$ commutes with Λ_{S_t}. This applies to $M \in \{\Pi_{\mathsf{Dist}}, \Pi_{\mathsf{NoPlus}}, \Pi_{\mathsf{DNP}}, \Pi^{\mathsf{FDist}}_{\mathsf{NoPlus}}\}$ (the latter is defined below).*

The analogous statement holds for M_{R_Y}.

Proof. By Theorem 3.11, we know that M (as an operator on $(\mathbb{C}^N)^{\otimes t}$) is in the span of $\{U^{\otimes t}\}_{U \in U(N)}$. As a result, in the Schur basis, M is block-diagonal with respect to the irrep $V^\lambda_{U(N)} \otimes V^\lambda_{\mathcal{A}_t}$. Since Λ_{S_t} is block-diagonal on the $|\lambda\rangle \otimes |\lambda'\rangle$ irrep spaces and acts as a scalar on each of them, we conclude that Λ_{S_t} commutes with $M_{\mathsf{R}_X} \otimes \mathbf{I}$ and $\mathbf{I} \otimes M_{\mathsf{R}_Y}$. □

Lemma 7.24 ($\Lambda_{P_t(N)}$ commutes with Π^X_{DNP} and Π^Y_{DNP})**.** *$\Lambda_{P_t(N)}$ commutes with Π^X_{DNP} and Π^Y_{DNP}.*

Proof. For Π^X_{DNP}, this follows directly from the fact that Π_{DNP} projects, in the $(S_N, P_t(N))$-Schur basis, onto irrep spaces with $|\lambda| = t$ (Lemma 7.20). Π^Y_{DNP} and $\Lambda_{P_t(N)}$ act on disjoint registers. □

Lemma 7.25 ($\Lambda_{P_t(N)}$ is approximately $\sqrt{\frac{t!}{N^t}}\mathbf{I}$ on full-box irreps)**.**

$$\frac{t!}{N^t}\left(1 - O\left(\frac{t^2}{N}\right)\right) \cdot \Pi^{X,Y}_{\mathsf{DNP}} \preceq \Pi^{X,Y}_{\mathsf{DNP}} \Lambda^\dagger_{P_t(N)} \Lambda_{P_t(N)} \Pi^{X,Y}_{\mathsf{DNP}} \preceq \frac{t!}{N^t}\left(1 + O\left(\frac{t^2}{N}\right)\right) \cdot \Pi^{X,Y}_{\mathsf{DNP}} \tag{7.77}$$

Proof. We have

$$\begin{aligned}
\left\| \frac{N^t}{t!} \Lambda^\dagger_{P_t(N)} \Lambda_{P_t(N)} \Pi_{\mathsf{DNP}} - \Pi_{\mathsf{DNP}} \right\| &= \left\| \frac{N^t}{t!} \left(\sum_{\lambda \in \widehat{P}_t(N)} \beta^2_{\lambda,t} \Pi_{\lambda,t} \right) \Pi_{\mathsf{DNP}} - \Pi_{\mathsf{DNP}} \right\| && (7.78)\\
&= \left\| \frac{N^t}{t!} \left(\sum_{\substack{\lambda \in \widehat{P}_t(N) \\ |\lambda| = t}} \beta^2_{\lambda,t} \Pi_{\lambda,t} \right) - \left(\sum_{\substack{\lambda \in \widehat{P}_t(N) \\ |\lambda| = t}} \Pi_{\lambda,t} \right) \right\| && \text{(Lemma 7.20)}\\
&= \left\| \sum_{\substack{\lambda \in \widehat{P}_t(N) \\ |\lambda| = t}} \left(\frac{N^t}{t!} \beta^2_{\lambda,t} - 1 \right) \Pi_{\lambda,t} \right\| && (7.79)
\end{aligned}$$

$$\leq \max_{\substack{\lambda \in \widehat{P_t(N)} \\ |\lambda|=t}} \left| \frac{N^t}{t!} \beta_{\lambda,t}^2 - 1 \right| \tag{7.80}$$

$$\leq O\left(\frac{t^2}{N}\right). \qquad \text{(Lemma 3.32)} \quad \square$$

Lemma 7.26 ($\Omega_{\mathcal{A}_t}$ on the Distinct Nonplussed Subspace)**.** *Let*

$$\Pi_{\mathsf{DNP}}^X := \left(\Pi_{\mathsf{DNP}_{N,t}}\right)_{R_X}, \qquad \Pi_{\mathsf{DNP}}^Y := \left(\Pi_{\mathsf{DNP}_{N,t}}\right)_{R_Y}, \qquad \Pi_{\mathsf{DNP}}^{X,Y} := \Pi_{\mathsf{DNP}}^X \Pi_{\mathsf{DNP}}^Y.$$

Let $\Omega_{P_t(N)}$ denote the commutant EPR projector for the partition algebra $P_t(N)$, and let

$$\Omega_{S_t} := \Omega_{\mathbb{C}[S_t]}$$

denote the commutant EPR projector for the symmetric group algebra. Then:

$$[\Pi_{\mathsf{DNP}}^X, \Omega_{P_t(N)}] = 0, \qquad [\Pi_{\mathsf{DNP}}^Y, \Omega_{P_t(N)}] = 0, \tag{7.81}$$

$$[\Pi_{\mathsf{DNP}}^X, \Omega_{S_t}] = 0, \qquad [\Pi_{\mathsf{DNP}}^Y, \Omega_{S_t}] = 0, \tag{7.82}$$

and

$$\Omega_{S_t} \Pi_{\mathsf{DNP}}^{X,Y} = \Omega_{P_t(N)} \Pi_{\mathsf{DNP}}^{X,Y}. \tag{7.83}$$

Proof. We prove the three identities in order.

First, let

$$\Pi_{|\lambda|=t} := \sum_{\substack{\lambda \in \widehat{P_t(N)} \\ |\lambda|=t}} \Pi_\lambda$$

be the projector onto the full-box irreps in the $(S_N, P_t(N))$-Schur decomposition of $(\mathbb{C}^N)^{\otimes t}$. By Lemma 7.20,

$$\Pi_{|\lambda|=t} = \Pi_{\mathsf{DNP}_{N,t}}.$$

Thus

$$\Pi_{\mathsf{DNP}}^X = (\Pi_{|\lambda|=t})_{\mathsf{R}_X}, \qquad \Pi_{\mathsf{DNP}}^Y = (\Pi_{|\lambda|=t})_{\mathsf{R}_Y}.$$

In the double Schur basis, $\Omega_{P_t(N)}$ is block-diagonal with respect to the irrep labels (λ, λ'). Since $(\Pi_{|\lambda|=t})_{\mathsf{R}_X}$ and $(\Pi_{|\lambda|=t})_{\mathsf{R}_Y}$ are just projections onto a subset of these blocks, we immediately obtain that they commute with $\Omega_{P_t(N)}$.

Second, the symmetric-group commutant EPR projector (Example 5.12) has the form

$$\Omega_{S_t} = \frac{1}{t!} \sum_{\pi \in S_t} S(\pi)_{\mathsf{R}_X} \otimes S(\pi)_{\mathsf{R}_Y},$$

For every $\pi \in S_t$, the operator $S(\pi)$ only permutes the t tensor factors. The subspace $\mathsf{Dist}_{N,t}$ is invariant under tensor-factor permutations, and

$$\mathsf{NoPlus}_{N,t} = (\mathbf{I} - |+\rangle\langle+|)^{\otimes t}$$

is also invariant under tensor-factor permutations. Hence

$$\mathsf{DNP}_{N,t} = \mathsf{Dist}_{N,t} \cap \mathsf{NoPlus}_{N,t}$$

is invariant under every $S(\pi)$. Therefore

$$S(\pi)\Pi_{\mathsf{DNP}_{N,t}} = \Pi_{\mathsf{DNP}_{N,t}} S(\pi) \qquad \forall \pi \in S_t.$$

Commuting term by term in the displayed formula for Ω_{S_t}, we obtain

$$[\Pi^X_{\mathsf{DNP}}, \Omega_{S_t}] = 0, \qquad [\Pi^Y_{\mathsf{DNP}}, \Omega_{S_t}] = 0.$$

For the third claim, we have the following calculation

$$\begin{aligned}
&\left(\Omega^{(t)}_{P_t(N)}\right)_{\mathsf{R}} \cdot \Pi^{X,Y}_{\mathsf{DNP}} && (7.84)\\
&= \left(\sum_{D \in \mathcal{B}(P_t(N))} S(D) \otimes S(D^{*_{P_t(N)}})^T\right) \cdot \Pi^{X,Y}_{\mathsf{DNP}} && \text{(Lemma 5.11)}\\
&= \left(\sum_{\pi \in S_t} S(\pi) \otimes S([\pi]^{*_{P_t(N)}})^T\right) \cdot \Pi^{X,Y}_{\mathsf{DNP}} && \text{(Lemma 7.19)}
\end{aligned}$$

Next, we expand $[\pi]^{*_{P_t(N)}}$ in the basis of $P_t(N)$ as $\sum_{E \in \mathcal{B}(P_t(N))} \alpha_{\pi,E}\, E$:

$$\begin{aligned}
&\left(\sum_{\pi \in S_t} S(\pi) \otimes \sum_{E \in \mathcal{B}(P_t(N))} \alpha_{\pi,E}\, S(E)^T\right) \cdot \Pi^{X,Y}_{\mathsf{DNP}} && (7.85)\\
&= \left(\sum_{\pi \in S_t} S(\pi) \otimes \sum_{\pi' \in \mathbb{C}[S_t]} \alpha_{\pi,\pi'}\, S(\pi')^T\right) \cdot \Pi^{X,Y}_{\mathsf{DNP}} && \text{(Lemma 7.19)}\\
&= \left(\sum_{\pi \in S_t} S(\pi) \otimes S([\pi]^{*_{\mathbb{C}[S_t]}})^T\right) \cdot \Pi^{X,Y}_{\mathsf{DNP}} && \text{(Lemma 7.22)}\\
&= \left(\Omega^{(t)}_{\mathbb{C}[S_t]}\right)_{\mathsf{R}} \cdot \Pi^{X,Y}_{\mathsf{DNP}}. && \square
\end{aligned}$$

7.4 Updates in the Distinct Nonplussed Subspace

Let V_{S_N} and $V_{U(N)}$ be the update operators when $G = S_N$ or $U(N)$, respectively. In this case, the commutant algebra $\mathcal{A}_t$ is equal to respectively either $P_t(N)$ or $\mathbb{C}[S_t]$, so Corollary 5.18 becomes

$$\mathsf{Adv}_t^{V_{S_N} \cdot C} = \underbrace{\left(\Lambda_{P_t(N)}\right)_{\mathsf{R}}}_{\text{ReweightEnd}} \cdot \underbrace{\left(\Omega_{P_t(N)}\right)_{\mathsf{R}}}_{\text{Symmetrize}} \cdot \underbrace{\left(C^{\otimes t}\right)_{\mathsf{R_X}}}_{C^{\otimes t}} \cdot \underbrace{\mathsf{Adv}_t^{\text{App}}}_{\text{Append}} \cdot \underbrace{\left(\Lambda^+_{P_0(N)}\right)_{\mathsf{R}_0}}_{\text{ReweightStart}} \tag{7.86}$$

$$\mathsf{Adv}_t^{V_{U(N)} \cdot C} = \underbrace{\left(\Lambda_{\mathbb{C}[S_t]}\right)_{\mathsf{R}}}_{\text{ReweightEnd}} \cdot \underbrace{\left(\Omega_{\mathbb{C}[S_t]}\right)_{\mathsf{R}}}_{\text{Symmetrize}} \cdot \underbrace{\left(C^{\otimes t}\right)_{\mathsf{R_X}}}_{C^{\otimes t}} \cdot \underbrace{\mathsf{Adv}_t^{\text{App}}}_{\text{Append}} \cdot \underbrace{\left(\Lambda^+_{P_0(N)}\right)_{\mathsf{R}_0}}_{\text{ReweightStart}} \tag{7.87}$$

Let

$$\left|\mathsf{Adv}_t^{V_{S_N} \cdot C}\right\rangle := \mathsf{Adv}_t^{V_{S_N} \cdot C} \left(|0\rangle_{\mathsf{AB}} \otimes |\emptyset\rangle_{\mathsf{R}}\right) \tag{7.88}$$

$$\left|\mathsf{Adv}_t^{V_{U(N)} \cdot C}\right\rangle := \mathsf{Adv}_t^{V_{U(N)} \cdot C} \left(|0\rangle_{\mathsf{AB}} \otimes |\emptyset\rangle_{\mathsf{R}}\right) \tag{7.89}$$

Lemma 7.27. *We can rearrange the terms in* $\mathsf{Adv}_t^{V_{U(N)}\cdot C}$ *so that* $C^{\otimes t}$ *appears at the end, as*

$$\mathsf{Adv}_t^{V_{U(N)}\cdot C} = \left(C^{\otimes t}\right)_{\mathsf{R_X}} \cdot \left(\Lambda_{\mathbb{C}[S_t]}\right)_{\mathsf{R}} \cdot \left(\Omega_{\mathbb{C}[S_t]}\right)_{\mathsf{R}} \cdot \mathsf{Adv}_t^{\mathrm{App}} \cdot \left(\Lambda^{+}_{\mathbb{C}[S_0]}\right)_{\mathsf{R}_0} \tag{7.90}$$

$$= \left(C^{\otimes t}\right)_{\mathsf{R_X}} \mathsf{Adv}_t^{V_{U(N)}} \tag{7.91}$$

Proof. To see this, it suffices to show that $C^{\otimes t}$ commutes with both $\Omega_{\mathbb{C}[S_t]}$ and $\Lambda_{\mathbb{C}[S_t]}$. This is clear when viewing the three operators in the Schur basis. All three are controlled by the irrep label, and $C^{\otimes t}$ acts only on the unitary (group) register, while $\Omega_{\mathbb{C}[S_t]}$ acts only on the commutant register, and $\Lambda_{\mathbb{C}[S_t]}$ is proportional to the identity within each block. □

Lemma 7.28 (Analyzing $\Pi^Y_{\mathsf{Dist}}\mathsf{Adv}_t^{\mathrm{App}}|0\rangle$)**.** *Let*

$$|\psi_t\rangle_{\mathsf{ABR}} := \mathsf{Adv}_t^{\mathrm{App}}|0\rangle$$

be the unnormalized append-only state,[55] *and let*

$$D_t^Y := (\Pi_{\mathsf{Dist}_{N,t}})_{\mathsf{R}_Y}.$$

For every non-identity $\pi \in S_t$*,*

$$\langle\psi_t| D_t^Y \Big(S(\pi)\otimes S(\pi)\Big) D_t^Y |\psi_t\rangle = 0.$$

Also,

$$\langle\psi_t| D_t^Y |\psi_t\rangle = N^t \cdot \prod_{i=1}^{t}\left(1-\frac{i-1}{N}\right) = N^t\cdot\left(1-O\left(\frac{t^2}{N}\right)\right).$$

Proof. We prove the claim by induction on t. The case $t=1$ is trivial (there are no non-identity permutations and $D_1^Y = \mathbf{I}$).

For the inductive step, we write

$$|\psi_t\rangle = \mathsf{App}_{\mathsf{AR}_t}\cdot(A_t)_{\mathsf{AB}}\cdot|\psi\rangle_{t-1} = \mathsf{App}_{\mathsf{AR}_t}|\psi_t^{\mathrm{pre}}\rangle.$$

Moreover, since $D_t^Y = D_t^Y\cdot D_{t-1}^Y$ and D_{t-1}^Y commutes with $\mathsf{App}_{\mathsf{AR}_t}$, we have that

$$D_t^Y|\psi_t\rangle = D_t^Y\cdot\mathsf{App}_{\mathsf{AR}_t}\cdot D_{t-1}^Y|\psi_t^{\mathrm{pre}}\rangle.$$

We now write explicitly in the standard basis

$$D_{t-1}^Y|\psi_t^{\mathrm{pre}}\rangle = \sum_{\substack{x_1,\ldots,x_t\\ \text{distinct } y_1,\ldots,y_{t-1}}} |x_t\rangle_{\mathsf{A}}\left|\phi_{x_{1:t},y_{1:t-1}}\right\rangle_{\mathsf{B}}|x_1,\ldots x_{t-1},y_1,\ldots,y_{t-1}\rangle_{\mathsf{R}_{<t}}.$$

On this state, the effect of $D_t^Y\cdot\mathsf{App}_{\mathsf{AR}_t}$ is appending with respect to an (unnormalized) EPR state over $y_t\notin\{y_1,\ldots,y_{t-1}\}$ (controlled on $\{y_1,\ldots,y_{t-1}\}$), which tells us that

$$D_t^Y|\psi_t\rangle = \sum_{\substack{x_1,\ldots,x_t\\ \text{distinct } y_1,\ldots,y_t}} |y_t\rangle_{\mathsf{A}}\left|\phi_{x_{1:t},y_{1:t-1}}\right\rangle_{\mathsf{B}}|x_1,\ldots x_t,y_1,\ldots,y_t\rangle_{\mathsf{R}_{\le t}}.$$

[55] For brevity, we will write $|0\rangle$ to represent $|0\rangle_{\mathsf{AB}}|\emptyset\rangle_{\mathsf{R}}$.

Given that D^Y_{t-1} also commutes with $(A_t)_{\mathsf{AB}}$, this immediately tells us that

$$\Big|\Big| D^Y_t \left|\psi_t\right\rangle \Big|\Big|^2 = \Big(N-(t-1)\Big)\Big|\Big| D^Y_{t-1} \left|\psi_{t-1}\right\rangle \Big|\Big|^2,$$

completing the first part of the induction. For the second part, we have that

$$\Big(S(\pi)\otimes S(\pi)\Big) D^Y_t \left|\psi_t\right\rangle = \sum_{\substack{x_1,\ldots,x_t \\ \text{distinct } y_1,\ldots,y_t}} \left|y_t\right\rangle_{\mathsf{A}} \left|\phi_{x_{1:t},y_{1:t-1}}\right\rangle_{\mathsf{B}} \left|x_{\pi^{-1}(1)},\ldots x_{\pi^{-1}(t)}, y_{\pi^{-1}(1)},\ldots,y_{\pi^{-1}(t)}\right\rangle_{\mathsf{R}_{\leq t}}.$$

Now, suppose that $\pi(t) \neq t$. Then, the inner product of these two states is zero because a nonzero term in the expansion of this inner product would require terms $y_t = y'_t$ for the A register but also $y_{\pi^{-1}(t)} = y'_{\pi^{-1}(t)}$ for the $\mathsf{R}_{Y,t}$ register.

On the other hand, suppose that $\pi(t) = t$ and identify π as an element of S_{t-1}; in particular, $S(\pi)$ commutes with $\mathsf{App}_{\mathsf{AR}_t} \cdot (A_t)_{\mathsf{AB}}$. Then, we write

$$\left\langle\psi_t\right| D^Y_t \Big(S(\pi)\otimes S(\pi)\Big) D^Y_t \left|\psi_t\right\rangle = \left\langle\psi^{\text{pre}}_t\right| D^Y_{t-1} \mathsf{App}^\dagger_{\mathsf{A},\mathsf{R}_t} D^Y_t \mathsf{App}_{\mathsf{A},\mathsf{R}_t} D^Y_{t-1} \Big(S(\pi)\otimes S(\pi)\Big) \left|\psi^{\text{pre}}_t\right\rangle$$

Next, we make use of the identity

$$D^Y_{t-1} \mathsf{App}^\dagger_{\mathsf{A},\mathsf{R}_t} D^Y_t \mathsf{App}_{\mathsf{A},\mathsf{R}_t} D^Y_{t-1} = (N-t+1) D^Y_{t-1},$$

which again holds because the effect of $D^Y_t \mathsf{App}_{\mathsf{A},\mathsf{R}_t}$ on the image of D^Y_{t-1} is to append an unnormalized EPR state over $y_t \notin \{y_1,\ldots,y_{t-1}\}$:

$$D^Y_t \mathsf{App}_{\mathsf{A},\mathsf{R}_t} D^Y_{t-1} = \sum_{\substack{x_1,\ldots x_t \\ \text{distinct } y_1,\ldots,y_t}} \left|y_t\right\rangle\!\left\langle x_t\right|_{\mathsf{A}} \otimes \left|x_1,\ldots,x_t,y_1,\ldots,y_t\right\rangle\!\left\langle x_1,\ldots,x_{t-1},y_1,\ldots,y_{t-1}\right|.$$

The claimed identity follows from this by a direct computation.

We conclude that

$$\begin{aligned}
\left\langle\psi_t\right| D^Y_t \Big(S(\pi)\otimes S(\pi)\Big) D^Y_t \left|\psi_t\right\rangle &= (N-t+1) \left\langle\psi^{\text{pre}}_t\right| D^Y_{t-1} \Big(S(\pi)\otimes S(\pi)\Big) \left|\psi^{\text{pre}}_t\right\rangle \\
&= (N-t+1) \left\langle\psi_{t-1}\right| D^Y_{t-1} (A^\dagger_t)_{\mathsf{AB}} \Big(S(\pi)\otimes S(\pi)\Big) (A_t)_{\mathsf{AB}} D^Y_{t-1} \left|\psi_{t-1}\right\rangle \\
&= (N-t+1) \left\langle\psi_{t-1}\right| D^Y_{t-1} \Big(S(\pi)\otimes S(\pi)\Big) D^Y_{t-1} \left|\psi_{t-1}\right\rangle.
\end{aligned}$$

This completes the induction. □

In order to argue about Π_{NoPlus}, we pass to a "smaller" projection,

$$\begin{aligned}
\Big(\Pi^{\mathsf{FDist}}_{\mathsf{NoPlus}}\Big)_{\mathsf{R}_Y} &:= \sum_{\substack{\widehat{y}_1,\ldots,\widehat{y}_t \\ \text{distinct and } \neq 0}} \left|\widehat{y}_1,\ldots,\widehat{y}_t\right\rangle\!\left\langle\widehat{y}_1,\ldots,\widehat{y}_t\right| && (7.92) \\
&= F_N^{\otimes t} \cdot (\mathbf{I} - \left|0\right\rangle\!\left\langle 0\right|)^{\otimes t} \cdot \Pi^Y_{\mathsf{Dist}} \cdot (F^\dagger_N)^{\otimes t}, && (7.93)
\end{aligned}$$

where $\left|\widehat{y}\right\rangle := \frac{1}{\sqrt{N}} \sum_{z\in\mathbb{Z}_N} \omega_N^{-yz} \left|z\right\rangle$ denotes the yth $\mathbb{Z}_N$ Fourier basis element of $\mathbb{C}^N$ and F_N denotes the $\mathbb{Z}_N$ quantum Fourier transform. In other words, $\Pi^{\mathsf{FDist}}_{\mathsf{NoPlus}}$ projects, in the $\mathbb{Z}_N$ Fourier basis, onto distinct tuples of nonzero basis states.

Lemma 7.29 (Analyzing $\Pi^{\mathsf{FDist}}_{\mathsf{NoPlus}}\mathsf{Adv}^{\mathrm{App}}_t|0\rangle$)**.** *Let*

$$|\psi_t\rangle_{\mathsf{ABR}} := \mathsf{Adv}^{\mathrm{App}}_t|0\rangle$$

be the unnormalized append-only state, and let

$$D^Y_t := (\Pi^{\mathsf{FDist}}_{\mathsf{NoPlus}})_{\mathsf{R}_Y}.$$

For every non-identity $\pi \in S_t$,

$$\langle\psi_t| D^Y_t \Big(S(\pi)\otimes S(\pi)\Big) D^Y_t |\psi_t\rangle = 0.$$

Also,

$$\langle\psi_t| D^Y_t |\psi_t\rangle = N^t \cdot \prod_{i=1}^{t}\left(1-\frac{i}{N}\right) = N^t\cdot\left(1-O\left(\frac{t^2}{N}\right)\right).$$

Proof sketch. The proof is extremely similar to that of Lemma 7.28, again proceeding by induction on t. The case $t=1$ is again trivial. For the inductive step, we use the same decomposition

$$|\psi_t\rangle = \mathsf{App}_{\mathsf{AR}_t}\cdot(A_t)_{\mathsf{AB}}\cdot|\psi\rangle_{t-1} = \mathsf{App}_{\mathsf{AR}_t}|\psi^{\mathrm{pre}}_t\rangle$$

and equation

$$D^Y_t|\psi_t\rangle = D^Y_t\cdot\mathsf{App}_{\mathsf{AR}_t}\cdot D^Y_{t-1}|\psi^{\mathrm{pre}}_t\rangle.$$

Next, we write $D^Y_{t-1}|\psi^{\mathrm{pre}}_t\rangle$ with R_Y in the *Fourier* basis:

$$D^Y_{t-1}|\psi^{\mathrm{pre}}_t\rangle = \sum_{\substack{x_1,\ldots,x_t\\ \text{distinct } y_1,\ldots,y_{t-1}\neq 0}} |x_t\rangle_{\mathsf{A}}\left|\phi_{x_{1:t},y_{1:t-1}}\right\rangle_{\mathsf{B}}|x_1,\ldots x_{t-1},\widehat{y}_1,\ldots,\widehat{y}_{t-1}\rangle_{\mathsf{R}_{<t}}.$$

On this state, the effect of $D^Y_t\cdot\mathsf{App}_{\mathsf{AR}_t}$ is appending with respect to an (unnormalized) *Fourier basis* EPR state

$$\sum_{y_t\notin\{0,\,y_1,\,\ldots,\,y_{t-1}\}}\left|\widehat{-y_t}\right\rangle|\widehat{y}_t\rangle$$

over $y_t\notin\{0,\,y_1,\,\ldots,\,y_{t-1}\}$ (controlled on $\{y_1,\ldots,y_{t-1}\}$), which tells us that

$$D^Y_t|\psi_t\rangle = \sum_{\substack{x_1,\ldots,x_t\\ \text{distinct } y_1,\ldots,y_t\neq 0}} \left|\widehat{-y_t}\right\rangle_{\mathsf{A}}\left|\phi_{x_{1:t},y_{1:t-1}}\right\rangle_{\mathsf{B}}|x_1,\ldots x_t,\widehat{y}_1,\ldots,\widehat{y}_t\rangle_{\mathsf{R}_{\le t}}.$$

This again completes the inductive norm calculation. In addition, an identical argument now shows that $\langle\psi_t| D^Y_t\Big(S(\pi)\otimes S(\pi)\Big)D^Y_t|\psi_t\rangle = 0$ whenever $\pi(t)\neq t$.

On the other hand, if $\pi(t)=t$, then an entirely analogous argument shows that

$$\langle\psi_t| D^Y_t\Big(S(\pi)\otimes S(\pi)\Big)D^Y_t|\psi_t\rangle = (N-t)\langle\psi_{t-1}| D^Y_{t-1}\Big(S(\pi)\otimes S(\pi)\Big)D^Y_{t-1}|\psi_{t-1}\rangle$$

(the only difference is that the factor $(N-t)$ replaces $(N-t+1)$, from ruling out 0 in addition to $y_1,\ldots,y_{t-1}$), completing the induction. □

Lemma 7.30. *Let $\Pi_{\mathsf{DNP}_{N,t}}$ be the orthogonal projector onto the distinct nonplussed subspace (Definition 7.5). Then*

$$\left|\left|\Pi^Y_{\mathsf{DNP}}\left|\mathsf{Adv}^{V_{U(N)}}_t\right\rangle\right|\right|^2 \ge 1-O\left(\frac{t^2}{N}\right) \tag{7.94}$$

Proof. By Corollary 7.13 and Lemma 7.15, we have that

$$\left\|\Pi^Y_{\mathsf{DNP}}\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\right\|^2 \geq \left\|\Pi^Y_{\mathsf{Dist}}\cdot\Pi^Y_{\mathsf{NoPlus}}\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\right\|^2 - O\left(\frac{t^2}{N}\right)$$

Thus, by the quantum union bound (Lemma 7.8), it suffices to prove the lemma statement for Π^Y_{Dist} and Π^Y_{NoPlus} rather than Π^Y_{DNP}.

For Π^Y_{Dist}, we begin by writing

$$\begin{aligned}\Pi^Y_{\mathsf{Dist}}\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle &= \Pi^Y_{\mathsf{Dist}}\cdot\Lambda_{S_t}\cdot\Omega_{S_t}\cdot\mathsf{Adv}_t^{\mathrm{App}}|0\rangle &(7.95)\\ &= \Lambda_{S_t}\cdot\Omega_{S_t}\cdot\Pi^Y_{\mathsf{Dist}}\cdot\mathsf{Adv}_t^{\mathrm{App}}|0\rangle\ . &\text{(by Lemmas 7.23 and 7.26)}\end{aligned}$$

Now, by Lemma 7.28, we have that

$$\begin{aligned}\left\|\Omega_{S_t}\cdot\Pi^Y_{\mathsf{Dist}}\cdot\mathsf{Adv}_t^{\mathrm{App}}|0\rangle\right\|^2 &= \frac{1}{(t!)^2}\left\|\sum_{\pi\in S_t}(S(\pi)\otimes S(\pi))\cdot\Pi^Y_{\mathsf{Dist}}\cdot\mathsf{Adv}_t^{\mathrm{App}}|0\rangle\right\|^2 &\text{(definition of }\Omega_{S_t})\\ &= \frac{1}{t!}\left\|\Pi^Y_{\mathsf{Dist}}\cdot\mathsf{Adv}_t^{\mathrm{App}}|0\rangle\right\|^2 &\text{(by the first part of Lemma 7.28)}\\ &\geq \frac{N^t}{t!}\left(1-O\left(\frac{t^2}{N}\right)\right)\ . &\text{(by the second part of Lemma 7.28)}\end{aligned}$$

Finally, we know from Lemma 6.5 that $\Lambda_{S_t}^2 \succeq \frac{t!}{N^t}\left(1-O\left(\frac{t^2}{N}\right)\right)\cdot\mathbf{I}$, so we conclude that

$$\begin{aligned}\left\|\Pi^Y_{\mathsf{Dist}}\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\right\|^2 &= \left\|\Lambda_{S_t}\cdot\Omega_{S_t}\cdot\Pi^Y_{\mathsf{Dist}}\cdot\mathsf{Adv}_t^{\mathrm{App}}|0\rangle\right\|^2 &(7.96)\\ &\geq \frac{t!}{N^t}\left(1-O\left(\frac{t^2}{N}\right)\right)\left\|\Omega_{S_t}\cdot\Pi^Y_{\mathsf{Dist}}\cdot\mathsf{Adv}_t^{\mathrm{App}}|0\rangle\right\|^2 &(7.97)\\ &\geq 1-O\left(\frac{t^2}{N}\right)\ . &(7.98)\end{aligned}$$

Next, to handle Π_{NoPlus}, we use the fact that

$$\left\|\Pi_{\mathsf{NoPlus}}\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\right\|^2 \geq \left\|\Pi^{\mathsf{FDist}}_{\mathsf{NoPlus}}\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\right\|^2. \tag{7.99}$$

An identical argument to the above then proves that

$$\left\|\Pi^{\mathsf{FDist}}_{\mathsf{NoPlus}}\left|\mathsf{Adv}_t^{V_{U(N)}}\right\rangle\right\|^2 \geq 1-O\left(\frac{t^2}{N}\right). \tag{7.100}$$

This completes the proof. □

Lemma 7.31. *Let C be sampled from a unitary 2-design $\mathcal{D}$, and let $\Pi_{\mathsf{DNP}_{N,t}}$ be the orthogonal projector onto the distinct nonplussed subspace (Definition 7.5). Then*

$$\mathop{\mathbb{E}}_{C\leftarrow\mathcal{D}}\mathrm{Tr}\left[\Pi^{X,Y}_{\mathsf{DNP}}\left|\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right|\right] \geq 1-O\left(\frac{t^2}{N}\right) \tag{7.101}$$

Proof. By a union bound, it suffices to prove the same statement for Π^X_{DNP} and Π^Y_{DNP} separately.

For Π^X_{DNP} this follows directly from Lemma 7.27 and Theorem 7.7.

For Π^Y_{DNP}, we begin with Lemma 7.27, which tells us that

$$\mathop{\mathbb{E}}_C \left|\left| \Pi^Y_{\mathsf{DNP}} \left| \mathsf{Adv}_t^{V_{U(N)} \cdot C} \right\rangle \right|\right|^2 = \mathop{\mathbb{E}}_C \left|\left| \Pi^Y_{\mathsf{DNP}} C_{\mathsf{R}_X}^{\otimes t} \left| \mathsf{Adv}_t^{V_{U(N)}} \right\rangle \right|\right|^2 \tag{7.102}$$

$$= \mathop{\mathbb{E}}_C \left|\left| C_{\mathsf{R}_X}^{\otimes t} \Pi^Y_{\mathsf{DNP}} \left| \mathsf{Adv}_t^{V_{U(N)}} \right\rangle \right|\right|^2 \tag{7.103}$$

$$= \left|\left| \Pi^Y_{\mathsf{DNP}} \left| \mathsf{Adv}_t^{V_{U(N)}} \right\rangle \right|\right|^2, \tag{7.104}$$

as $C_{\mathsf{R}_X}^{\otimes t}$ is unitary and commutes with Π^Y_{DNP}. The conclusion now follows from Lemma 7.30. □

Lemma 7.32. *Let C be sampled from a unitary 2-design $\mathcal{D}$, and let $\Pi^{X,Y}_{\mathsf{DNP}}$ be the distinct nonplussed subspace projector. Then*

$$\mathop{\mathbb{E}}_{C \leftarrow \mathcal{D}} \mathrm{Tr}\left[\Pi^{X,Y}_{\mathsf{DNP}} \left| \mathsf{Adv}_t^{V_{S_N} \cdot C} \right\rangle\!\!\left\langle \mathsf{Adv}_t^{V_{S_N} \cdot C} \right| \right] \geq 1 - O\left(\frac{t^2}{N}\right) \tag{7.105}$$

Proof.

$$\mathop{\mathbb{E}}_C \left|\left| \Pi^{X,Y}_{\mathsf{DNP}} \left| \mathsf{Adv}_t^{V_{S_N} \cdot C} \right\rangle \right|\right|^2 = \mathop{\mathbb{E}}_C \left|\left| \Pi^{X,Y}_{\mathsf{DNP}} \cdot \Lambda_{P_t(N)} \cdot \Omega_{P_t(N)} \cdot C_{\mathsf{R_X}}^{\otimes t} \cdot \mathsf{Adv}_t^{\mathrm{App}} |0\rangle \right|\right|^2 \tag{7.106}$$

$$= \mathop{\mathbb{E}}_C \left|\left| \Lambda_{P_t(N)} \cdot \Pi^{X,Y}_{\mathsf{DNP}} \cdot \Omega_{P_t(N)} \cdot C_{\mathsf{R_X}}^{\otimes t} \cdot \mathsf{Adv}_t^{\mathrm{App}} |0\rangle \right|\right|^2 \quad \text{(Lemma 7.24)}$$

$$= \mathop{\mathbb{E}}_C \left|\left| \Lambda_{P_t(N)} \cdot \Pi^{X,Y}_{\mathsf{DNP}} \cdot \Omega_{S_t} \cdot C_{\mathsf{R_X}}^{\otimes t} \cdot \mathsf{Adv}_t^{\mathrm{App}} |0\rangle \right|\right|^2 \quad \text{(Lemma 7.26)}$$

$$\geq \mathop{\mathbb{E}}_C \left|\left| \Lambda_{S_t} \cdot \Pi^{X,Y}_{\mathsf{DNP}} \cdot \Omega_{S_t} \cdot C_{\mathsf{R_X}}^{\otimes t} \cdot \mathsf{Adv}_t^{\mathrm{App}} |0\rangle \right|\right|^2 \left(1 - O\left(\frac{t^2}{N}\right)\right)$$

(Lemmas 6.5 and 7.25)

$$= \mathop{\mathbb{E}}_C \left|\left| \Pi^{X,Y}_{\mathsf{DNP}} \cdot \Lambda_{S_t} \cdot \Omega_{S_t} \cdot C_{\mathsf{R_X}}^{\otimes t} \cdot \mathsf{Adv}_t^{\mathrm{App}} |0\rangle \right|\right|^2 \left(1 - O\left(\frac{t^2}{N}\right)\right) \quad \text{(Lemma 7.23)}$$

$$= \mathop{\mathbb{E}}_C \left|\left| \Pi^{X,Y}_{\mathsf{DNP}} \left| \mathsf{Adv}_t^{V_{U(N)} \cdot C} \right\rangle \right|\right|^2 \left(1 - O\left(\frac{t^2}{N}\right)\right) \tag{7.107}$$

$$\geq 1 - O\left(\frac{t^2}{N}\right) \quad \text{(Lemma 7.31)} \quad \square$$

Corollary 7.33. *Using Lemmas 6.1, 7.31 and 7.32 we get the following trace distance bounds:*

$$\left\| \mathrm{Tr}_{\mathsf{R}} \left[\mathop{\mathbb{E}}_{C \leftarrow \mathcal{D}} \left| \mathsf{Adv}_t^{V_{U(N)} \cdot C} \right\rangle\!\!\left\langle \mathsf{Adv}_t^{V_{U(N)} \cdot C} \right| \right] - \mathrm{Tr}_{\mathsf{R}} \left[\mathop{\mathbb{E}}_{C \leftarrow \mathcal{D}} \Pi^{X,Y}_{\mathsf{DNP}} \left| \mathsf{Adv}_t^{V_{U(N)} \cdot C} \right\rangle\!\!\left\langle \mathsf{Adv}_t^{V_{U(N)} \cdot C} \right| \Pi^{X,Y}_{\mathsf{DNP}} \right] \right\|_1 \leq O\left(\frac{t^2}{N}\right) \tag{7.108}$$

$$\left\| \mathrm{Tr}_{\mathsf{R}} \left[\mathop{\mathbb{E}}_{C \leftarrow \mathcal{D}} \left| \mathsf{Adv}_t^{V_{S_N} \cdot C} \right\rangle\!\!\left\langle \mathsf{Adv}_t^{V_{S_N} \cdot C} \right| \right] - \mathrm{Tr}_{\mathsf{R}} \left[\mathop{\mathbb{E}}_{C \leftarrow \mathcal{D}} \Pi^{X,Y}_{\mathsf{DNP}} \left| \mathsf{Adv}_t^{V_{S_N} \cdot C} \right\rangle\!\!\left\langle \mathsf{Adv}_t^{V_{S_N} \cdot C} \right| \Pi^{X,Y}_{\mathsf{DNP}} \right] \right\|_1 \leq O\left(\frac{t^2}{N}\right) \tag{7.109}$$

7.5 Putting it Together

Finally, we compare the cases $G = U(N)$ and $G = S_N$ of Corollary 5.18, and show that the G-dependent steps are close in expectation whenever C is sampled from a unitary 2-design. This implies that queries to PC for a random permutation matrix $P \in S_N$ are (nearly) indistinguishable from queries to $UC = U'$ for a Haar random unitary $U \in U(N)$, leading to the conclusion that PC is a pseudorandom unitary for any pseudorandom permutation P and any unitary 2-design C.

Theorem 7.34.

$$\left\| \mathop{\mathbb{E}}_{\substack{P\leftarrow S_N \\ C\leftarrow\mathcal{D}}} \left[\left|\mathsf{Adv}_t^{P\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{P\cdot C}\right| \right] - \mathop{\mathbb{E}}_{U\leftarrow U(N)} \left[\left|\mathsf{Adv}_t^{U}\right\rangle\!\left\langle\mathsf{Adv}_t^{U}\right| \right] \right\|_1 \leq O\left(\frac{t^2}{N}\right) \tag{7.110}$$

Proof. By the correctness of the path-recording oracles, we know that for every C,

$$\mathrm{Tr}_{\mathsf{R}}\left[\left|\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right| \right] = \mathop{\mathbb{E}}_{U\leftarrow U(N)} \left|\mathsf{Adv}_t^{U\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{U\cdot C}\right| = \mathop{\mathbb{E}}_{U\leftarrow U(N)} \left|\mathsf{Adv}_t^{U}\right\rangle\!\left\langle\mathsf{Adv}_t^{U}\right| \tag{7.111}$$

and

$$\mathrm{Tr}_{\mathsf{R}}\left[\left|\mathsf{Adv}_t^{V_{S_N}\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{V_{S_N}\cdot C}\right| \right] = \mathop{\mathbb{E}}_{P\leftarrow S_N} \left|\mathsf{Adv}_t^{P\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{P\cdot C}\right| . \tag{7.112}$$

Therefore, by Corollary 7.33, it suffices to prove that

$$\left\| \mathrm{Tr}_{\mathsf{R}}\left[\mathop{\mathbb{E}}_{C\leftarrow\mathcal{D}} \Pi_{\mathsf{DNP}}^{X,Y} \left|\mathsf{Adv}_t^{V_{S_N}\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{V_{S_N}\cdot C}\right| \Pi_{\mathsf{DNP}}^{X,Y} \right] - \mathrm{Tr}_{\mathsf{R}}\left[\mathop{\mathbb{E}}_{C'\leftarrow\mathcal{D}} \Pi_{\mathsf{DNP}}^{X,Y} \left|\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right| \Pi_{\mathsf{DNP}}^{X,Y} \right] \right\|_1 \leq O\left(\frac{t^2}{N}\right) \tag{7.113}$$

In fact, we prove the following stronger statement: for every fixed C, we have that

$$\left\| \Pi_{\mathsf{DNP}}^{X,Y} \left|\mathsf{Adv}_t^{V_{S_N}\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{V_{S_N}\cdot C}\right| \Pi_{\mathsf{DNP}}^{X,Y} - \Pi_{\mathsf{DNP}}^{X,Y} \left|\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right| \Pi_{\mathsf{DNP}}^{X,Y} \right\|_1 \leq O\left(\frac{t^2}{N}\right) \tag{7.114}$$

To see this, we make use of the standard manipulation

$$\begin{aligned} &\left\| \Pi_{\mathsf{DNP}}^{X,Y} \left|\mathsf{Adv}_t^{V_{S_N}\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{V_{S_N}\cdot C}\right| \Pi_{\mathsf{DNP}}^{X,Y} - \Pi_{\mathsf{DNP}}^{X,Y} \left|\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right\rangle\!\left\langle\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right| \Pi_{\mathsf{DNP}}^{X,Y} \right\|_1 \\ \leq& 2\cdot \left\| \Pi_{\mathsf{DNP}}^{X,Y} \left|\mathsf{Adv}_t^{V_{S_N}\cdot C}\right\rangle - \Pi_{\mathsf{DNP}}^{X,Y} \left|\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right\rangle \right\| \end{aligned} \tag{7.115}$$

Now, given that our starting state is empty, we can write (for both $V = V_{U(N)}$ and $V = V_{S_N}$)

$$\left|\mathsf{Adv}_t^{V\cdot C}\right\rangle = \left(\Lambda_{\mathcal{A}_t}\right)_{\mathsf{R}_{\leq t}} \cdot \left(\Omega_{\mathcal{A}_t}\right)_{\mathsf{R}_{\leq t}} \cdot \left(C^{\otimes t}\right)_{\mathsf{R}_{\mathsf{X}}} \cdot \mathsf{Adv}_t^{\mathrm{App}} \cdot |0\rangle . \tag{7.116}$$

Letting $|\psi\rangle = \left(C^{\otimes t}\right)_{\mathsf{R}_{\mathsf{X}}} \cdot \mathsf{Adv}_t^{\mathrm{App}} \cdot |0\rangle$, we have

$$\begin{aligned} \left\| \Pi_{\mathsf{DNP}}^{X,Y} \left|\mathsf{Adv}_t^{V_{S_N}\cdot C}\right\rangle - \Pi_{\mathsf{DNP}}^{X,Y} \left|\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right\rangle \right\| &= \left\| \Pi_{\mathsf{DNP}}^{X,Y} \left(\Lambda_{P_t(N)}\Omega_{P_t(N)} - \Lambda_{S_t}\Omega_{S_t}\right) |\psi\rangle \right\| && (7.117)\\ &= \left\| \left(\Lambda_{P_t(N)}\Pi_{\mathsf{DNP}}^{X,Y}\Omega_{P_t(N)} - \Lambda_{S_t}\Pi_{\mathsf{DNP}}^{X,Y}\Omega_{S_t}\right) |\psi\rangle \right\| && (7.118)\\ &= \left\| \left(\Lambda_{P_t(N)} - \Lambda_{S_t}\right)\Pi_{\mathsf{DNP}}^{X,Y}\Omega_{S_t} |\psi\rangle \right\| && (7.119)\\ &= \left\| \left(\Lambda_{P_t(N)} - \Lambda_{S_t}\right)\Pi_{\mathsf{DNP}}^{X,Y}\Lambda_{S_t}^{+} \left|\mathsf{Adv}_t^{V_{U(N)}\cdot C}\right\rangle \right\| && (7.120)\\ &\leq \left\| \left(\Lambda_{P_t(N)} - \Lambda_{S_t}\right)\Pi_{\mathsf{DNP}}^{X,Y} \right\| \cdot \left\| \Lambda_{S_t}^{+} \right\| . && (7.121) \end{aligned}$$

By Lemmas 6.5 and 7.25, on the full-box/DNP blocks the squared reweighting matrices $\Lambda^{\dagger}_{P_t(N)}\Lambda_{P_t(N)}$ and $\Lambda^{\dagger}_{\mathbb{C}[S_t]}\Lambda_{\mathbb{C}[S_t]}$ differ by a multiplicative $1 + O(t^2/N)$. Taking square roots preserves this error up to constants, so it follows that

$$\left\| \left(\Lambda_{P_t(N)} - \Lambda_{S_t}\right) \Pi^{X,Y}_{\mathsf{DNP}} \right\| \leq \sqrt{\frac{t!}{N^t}} \cdot O\left(\frac{t^2}{N}\right). \tag{7.122}$$

Similarly, by Lemma 6.5 we have that

$$\left\| \Lambda^{+}_{S_t} \right\| \leq \sqrt{\frac{N^t}{t!}} \cdot \left(1 + O\left(\frac{t^2}{N}\right)\right). \tag{7.123}$$

This completes the proof of the theorem. □

The following corollary formalizes Theorems 1.6 and 1.7 to conclude that $P_k \cdot C$ is either a pseudorandom unitary or a t-design, if P_k is instantiated with respectively either a pseudorandom permutation or a $2t$-wise independent permutation.

Corollary 7.35. *Suppose that P_k is a (t, ε)-secure in-place post-quantum PRP, meaning that for all polynomial-time t-query quantum algorithms Adv, we have that $\mathsf{Adv}^{P_k} \approx_{\varepsilon} \mathsf{Adv}^{P}$ for a uniformly random permutation unitary P. Moreover, let $C \leftarrow \mathcal{D}$ be a unitary sampled from an exact unitary 2-design. Then, $P_k \cdot C$ is a $(t, \varepsilon + O(t^2/N))$-secure PRU. Moreover, if P_k is secure against inefficient Adv, so is $P_k \cdot C$.*

Proof. This immediately follows from Theorem 7.34 and a hybrid argument switching P_k to a uniformly random P. □

Acknowledgements

The authors thank Henry Yuen, John Bostanci, Quynh Nguyen, Natalie Parham, Angelos Pelecanos, Mark Zhandry, Joe Carolan, and Takashi Yamakawa for helpful discussions.

This work was done in part while AL and BN were visiting the Simons Institute for the Theory of Computing, where BN was supported by NSF QLCI Grant No. 2016245. BN acknowledges the following funding sources: AFOSR award FA9550-23-1-0363, NSF awards CCF-2530159, CCF-2144219, and CCF-2329939, and the Sloan Foundation. AL was supported in part by NSF CAREER award CNS-2541300 and an E. Lawrence Keyes, Jr./Emerson Electric Co. Faculty Award. J.W. was supported by the NSF CAREER award CCF-233971 and a fellowship from the Sloan Foundation.

References

[AAB+19] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.

[ABGL25a] Prabhanjan Ananth, John Bostanci, Aditya Gulati, and Yao-Ting Lin. Pseudorandom unitaries in the haar random oracle model. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part II*, volume 16001 of *LNCS*, pages 301–333. Springer, Cham, August 2025.

[ABGL25b] Prabhanjan Ananth, John Bostanci, Aditya Gulati, and Yao-Ting Lin. Pseudorandomness in the (inverseless) haar random oracle model. In Serge Fehr and Pierre-Alain Fouque, editors, *EUROCRYPT 2025, Part VII*, volume 15607 of *LNCS*, pages 138–166. Springer, Cham, May 2025.

[AGL25] Prabhanjan Ananth, Aditya Gulati, and Yao-Ting Lin. On the limitations of pseudorandom unitaries - or: Cryptographic applications of LOCC indistinguishability of identical versus independent haar unitaries. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part III*, volume 16270 of *LNCS*, pages 69–103. Springer, Cham, December 2025.

[BBBV97] Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. Strengths and weaknesses of quantum computing. *SIAM journal on Computing*, 26(5):1510–1523, 1997.

[BCH05] Dave Bacon, Isaac Chuang, and Aram Harrow. The quantum Schur transform: I. efficient qudit circuits. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, 2005.

[BFG+25] Adam Burchardt, Jiani Fei, Dmitry Grinko, Martin Larocca, Maris Ozols, Sydney Timmerman, and Vladyslav Visnevskyi. High-dimensional quantum schur transforms, 2025.

[BG26] James Bartusek and Eli Goldin. Unclonable encryption in the haar random oracle model. *arXiv preprint arXiv:2603.11437*, 2026.

[Blo03] Matthew Bloss. G-colored partition algebras as centralizer algebras of wreath products. *Journal of algebra*, 265(2):690–710, 2003.

[BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93*, pages 62–73. ACM Press, November 1993.

[Bra37] Richard Brauer. On algebras which are connected with the semisimple continuous groups. *Annals of Mathematics*, 38(4):857–872, 1937.

[Car26] Joseph Carolan. Compressed permutation oracles. In *Proceedings of the 58th Annual ACM Symposium on Theory of Computing*, STOC '26, page 150–161, New York, NY, USA, 2026. Association for Computing Machinery.

[CFHL21] Kai-Min Chung, Serge Fehr, Yu-Hsuan Huang, and Tai-Ning Liao. On the compressed-oracle technique, and post-quantum security of proofs of sequential work. In Anne Canteaut and François-Xavier Standaert, editors, *EUROCRYPT 2021, Part II*, volume 12697 of *LNCS*, pages 598–629. Springer, Cham, October 2021.

[DR90] Persi Diaconis and Daniel Rockmore. Efficient computation of the fourier transform on finite groups. *Journal of the American Mathematical Society*, 3(2):297–332, 1990.

[EH18] Karin Erdmann and Thorsten Holm. *Algebras and Representation Theory*. Springer Undergraduate Mathematics Series. Springer, 2018.

[FND26] Ben Foxman, Barak Nehoran, and Yongshan Ding. Efficient quantum fourier transforms for semisimple algebras, 2026.

[FTH23] Jiani Fei, Sydney Timmerman, and Patrick Hayden. Efficient quantum algorithm for port-based teleportation, 2023.

[Gao15] Jingliang Gao. Quantum union bounds for sequential projective measurements. *Physical Review A*, 92(5):052331, 2015.

[GBO23] Dmitry Grinko, Adam Burchardt, and Maris Ozols. Gelfand-tsetlin basis for partially transposed permutations, with applications to quantum information. *arXiv preprint arXiv:2310.02252*, 2023.

[GW09] Roe Goodman and Nolan R Wallach. *Symmetry, representations, and invariants*, volume 255. Springer, 2009.

[GY25] Dmitry Grinko and Satoshi Yoshida. Quantum simulation of random unitaries from clebsch-gordan transforms. *arXiv preprint arXiv:2509.26623*, 2025.

[Hal13] Brian C Hall. Lie groups, lie algebras, and representations. In *Quantum Theory for Mathematicians*, pages 333–366. Springer, 2013.

[Har05] Aram Harrow. *Applications of coherent classical communication and the Schur transform to quantum information theory*. PhD thesis, Massachusetts Institute of Technology, 2005.

[HJ20] Tom Halverson and Theodore N Jacobson. Set-partition tableaux and representations of diagram algebras. *Algebraic Combinatorics*, 3(2):509–538, 2020.

[HLS24] Yassine Hamoudi, Qipeng Liu, and Makrand Sinha. The NISQ complexity of collision finding. In Marc Joye and Gregor Leander, editors, *EUROCRYPT 2024, Part IV*, volume 14654 of *LNCS*, pages 3–32. Springer, Cham, May 2024.

[HR05] Tom Halverson and Arun Ram. Partition algebras. *European Journal of Combinatorics*, 26(6):869–921, 2005.

[HY25] Minki Hhan and Shogo Yamada. Pseudorandom function-like states from common haar unitary. In Benny Applebaum and Huijia (Rachel) Lin, editors, *TCC 2025, Part III*, volume 16270 of *LNCS*, pages 134–165. Springer, Cham, December 2025.

[JZ25] Stacey Jeffery and Sebastian Zur. The compressed oracle is a worthy (multiplicative) adversary. *arXiv preprint arXiv:2509.07876*, 2025.

[Kin70] Ronald C King. Generalized young tableaux and the general linear group. *Journal of Mathematical Physics*, 11(1):280–293, 1970.

[KOMW19] Samad Khabbazi Oskouei, Stefano Mancini, and Mark M. Wilde. Union bound for quantum information processing. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 475(2221):20180612, January 2019.

[Kro19] Hari Krovi. An efficient high dimensional quantum schur transform. *Quantum*, 3:122, 2019.

[KS08] Tali Kaufman and Madhu Sudan. Algebraic property testing: the role of invariance. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing*, pages 403–412, 2008.

[KW89] Selahattin Kayalar and Howard L. Weinert. Oblique projections: formulas, algorithms, and error bounds. *Math. Control Signals Systems*, 2(1):33–45, 1989.

[Kwo08] Jae-Hoon Kwon. Rational semistandard tableaux and character formula for the lie superalgebra glË†âˆž—âˆž. *Advances in Mathematics*, 217(2):713–739, 2008.

[Lam01] T. Y. Lam. *A First Course in Noncommutative Rings*, volume 131 of *Graduate Texts in Mathematics*. Springer, New York, 2 edition, 2001.

[LR88] Michael Luby and Charles Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM journal on computing*, 17(2):373–386, 1988.

[LZ19a] Qipeng Liu and Mark Zhandry. On finding quantum multi-collisions. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 189–218. Springer, Cham, May 2019.

[LZ19b] Qipeng Liu and Mark Zhandry. Revisiting post-quantum Fiat-Shamir. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 326–355. Springer, Cham, August 2019.

[Mac98] Ian Grant Macdonald. *Symmetric functions and Hall polynomials*. Oxford university press, 1998.

[Mar96] Paul Martin. The structure of the partition algebras. *Journal of Algebra*, 183(2):319–358, 1996.

[MH25] Fermi Ma and Hsin-Yuan Huang. How to construct random unitaries. In Michal Koucký and Nikhil Bansal, editors, *57th ACM STOC*, pages 806–809. ACM Press, June 2025.

[MMW25] Christian Majenz, Giulio Malavolta, and Michael Walter. Permutation superposition oracles for quantum query lower bounds. In Michal Koucký and Nikhil Bansal, editors, *57th ACM STOC*, pages 1508–1519. ACM Press, June 2025.

[MPSY24] Tony Metger, Alexander Poremba, Makrand Sinha, and Henry Yuen. Simple constructions of linear-depth t-designs and pseudorandom unitaries. In *65th FOCS*, pages 485–492. IEEE Computer Society Press, October 2024.

[MRR03] Cristopher Moore, Daniel Rockmore, and Alexander Russell. Generic quantum fourier transforms, 2003.

[MRW16] David Maslen, Daniel N. Rockmore, and Sarah Wolff. The efficient computation of fourier transforms on semisimple algebras, 2016.

[MU17] Michael Mitzenmacher and Eli Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, 2017.

[Ngu23] Quynh T Nguyen. The mixed schur transform: efficient quantum circuit and applications. *arXiv preprint arXiv:2310.01613*, 2023.

[PCX25] Tianci Peng, Shujiao Cao, and Rui Xue. On quantum query complexities of collision-finding in non-uniform random functions. In Goichiro Hanaoka and Bo-Yin Yang, editors, *ASIACRYPT 2025, Part VIII*, volume 16252 of *LNCS*, pages 193–223. Springer, Singapore, December 2025.

[Ros21] Ansis Rosmanis. Tight bounds for inverting permutations via compressed oracle arguments. *arXiv preprint arXiv:2103.08975*, 2021.

[Sag01] Bruce E Sagan. *The symmetric group: representations, combinatorial algorithms, and symmetric functions*. Springer, 2001.

[Sch27] Issai Schur. über die rationalen darstellungen der allgemeinen linearen gruppe. *Sitzungsberichte der Preussischen Akademie der Wissenschaften zu Berlin, Physikalisch-Mathematische Klasse*, pages 58–75, 1927.

[SML+25] Thomas Schuster, Fermi Ma, Alex Lombardi, Fernando Brandao, and Hsin-Yuan Huang. Strong random unitaries and fast scrambling. *arXiv preprint arXiv:2509.26310*, 2025.

[Ste87] John R Stembridge. Rational tableaux and the tensor algebra of gln. *Journal of Combinatorial Theory, Series A*, 46(1):79–120, 1987.

[Str24] Eugene Strahov. Generalized regular representations of big wreath products, 2024.

[Tan97] Kenichiro Tanabe. On the centralizer algebra of the unitary reflection group g (m, p, n). *Nagoya mathematical journal*, 148:113–126, 1997.

[TW25] Ewin Tang and John Wright. Amplitude amplification and estimation require inverses. *arXiv preprint arXiv:2507.23787*, 2025.

[Unr23] Dominique Unruh. Towards compressed permutation oracles. In Jian Guo and Ron Steinfeld, editors, *ASIACRYPT 2023, Part IV*, volume 14441 of *LNCS*, pages 369–400. Springer, Singapore, December 2023.

[VK92] Naum Vilenkin and Anatoli Klimyk. *Representation of Lie Groups and Special Functions. Volume 3: Classical and Quantum Groups and Special Functions*. Springer, 1992.

[Wey39] Hermann Weyl. *The Classical Groups: Their Invariants and Representations*. Princeton Mathematical Series. Princeton University Press, Princeton, NJ, 1939.

[Zha12] Mark Zhandry. How to construct quantum random functions. In *53rd FOCS*, pages 679–687. IEEE Computer Society Press, October 2012.

[Zha19] Mark Zhandry. How to record quantum queries, and applications to quantum indifferentiability. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part II*, volume 11693 of *LNCS*, pages 239–268. Springer, Cham, August 2019.

[Zha25] Mark Zhandry. How to model unitary oracles. In Yael Tauman Kalai and Seny F. Kamara, editors, *CRYPTO 2025, Part II*, volume 16001 of *LNCS*, pages 237–268. Springer, Cham, August 2025.