

The Quantum Character Transform

Kagwe A. Muchane^{✉*}

Department of Computer Science, North Carolina State University, Raleigh, NC 27606, USA
Department of Mathematics, North Carolina State University, Raleigh, NC 27606, USA

Bastidas et al. (2025) highlighted the importance of exploiting group symmetries to reduce computational overhead and unify simulation methods across both condensed matter physics and ab initio quantum chemistry. I discuss the discrete character transform, which is the group Fourier transform restricted to $Z(\mathbb{C}[\mathcal{G}])$, the center of the group algebra, where class functions form a commutative convolution algebra and irreducible characters are its spectral basis. The normalized character table of a non-abelian group can be obtained through a restriction operator, which means an efficient quantum character transform for any finite group with an efficient quantum Fourier transform. I give complementary quantum algorithms for class convolution based on direct class-sum block encoding and spectral implementation in the character basis, explaining their applications to symmetry-adapted many-body dynamics, molecular and spectroscopic symmetry analysis, and quantum walks, as well as characterizing the central information used in hidden-subgroup problems, including its limitations for graph isomorphism.

I. INTRODUCTION

Quantum chemistry is widely considered one of the most promising and impactful commercial applications for early quantum computers, because molecular electronic structure is in essence quantum mechanical and its accurate classical description can require resources that grow prohibitively with system size [18, 48]. Yet, representing a molecular Hamiltonian on a quantum computer does not by itself make a calculation efficient. Practical algorithms must also identify and exploit the symmetries that constrain the physically relevant state space. Point-group, permutation, spin, and other finite symmetries decompose molecular Hilbert spaces into invariant sectors, where states and observables can be organized by irreducible representations. Character tables describe exactly this decomposition, supporting the classification of molecular orbitals and normal modes, the derivation of infrared and Raman selection rules, and the block diagonalization of symmetry-preserving Hamiltonians [9, 14, 23, 37]. Their pure mathematical foundations reach back to Frobenius’s analysis of the group determinant [29, 30], Burnside’s systematization of finite-group representation theory [16, 17], and Schur’s development of linear and projective representation theory [53, 54]. As chemistry becomes a major target for quantum hardware, computing and manipulating character data has become more than a classificatory concern.

The abelian quantum Fourier transform and Kitaev’s phase-estimation algorithm showed that spectral information can be extracted efficiently under a change of basis. Beyond the abelian case, efficient transforms over specific nonabelian families, most prominently Beals’s transform for the symmetric group [6] and the Schur transform of Bacon, Chuang, and Harrow [4], proved that noncommutative harmonic analysis can also be implemented efficiently when representation bases and branching rules allow recursive quantum circuits. In the literature, the term “character transform” has described two distinct operations. Greg Kuperberg’s early use of the term [46] should not be confused with the class-condensed unitary operator used today. While related, in his 2003 work on the

Hidden Subgroup Problem, Kuperberg uses the term to refer to a non-abelian Fourier transform on the full group representation space. Over a decade prior, a more compact, class-based change of basis was already implicit in character-theoretic literature used throughout classical chemistry to analyze molecular systems by symmetry. In Cotton’s book *Chemical Applications of Group Theory*, the Great Orthogonality Theorem gives the class-weighted orthogonality relations underlying molecular character tables [23]. With normalized conjugacy-class coordinates, these relations make the matrix with entries

$$U_{\Gamma,C} = \sqrt{\frac{|\mathcal{C}|}{|\mathcal{G}|}} \chi_{\Gamma}(C)$$

unitary, with rows indexed by irreducible characters and columns by conjugacy classes. In 2015, Bermejo-Vega and Zatloukal placed this class-to-character change of basis within quantum Fourier analysis over abelian hypergroups, showing how the group quantum Fourier transform can access the character basis of conjugation-invariant states [8]. Most recently, Bastidas, Liu, Chuang, and colleagues elevated this class-based unitary to a definitive quantum operation [5]. Recognizing that element-level tracking can become resource-intensive on near-term quantum hardware, Bastidas et al. deployed this class-based unitary directly inside quantum circuits as an active symmetry transform. By shifting the state vector into conjugacy-class and irreducible-representation coordinates, they enable symmetry-adapted projection and block diagonalization, such that physically relevant symmetry sectors can be isolated and manipulated coherently; they further demonstrated this symmetry-adapted subroutine for molecular hydrogen, H_2 , experimentally on IBM quantum hardware in the presence of noise.

That transform is the *Quantum Character Transform*. For a finite group $\mathcal{G}$, the class sums form a basis of the center $Z(\mathbb{C}[\mathcal{G}])$, while the irreducible group characters provide its spectral basis. The class algebra is the commutative convolution algebra of conjugacy-invariant functions on $\mathcal{G}$, and the normalized character table gives the unitary change of basis between nor-

malized conjugacy-class and irreducible-character coordinates. The discrete character transform is non-abelian Fourier analysis restricted to the center of the group algebra, and because commutativity is often easier to work with, it enables efficient spectral calculus for $Z(\mathbb{C}[\mathcal{G}])$. This has an important algorithmic consequence, as once the QCT is available, one can pass in both directions between central operators and their spectra. Given a class function f , the transform exposes the spectrum of a class convolution operator T_f ; conversely, a prescribed eigenvalue profile over irreducible sectors creates, through the inverse QCT, an associated class operator. This operator synthesis approach situates the quantum character transform in the same conceptual family as modern quantum singular-value and polynomial transformation methods, where diagonal manipulations in the character basis become operator transformations in the original basis.

Seen through its history, the expected applications of the discrete character transform are wide. If f is supported on a conjugacy class or union of classes, then T_f is the adjacency operator of a normal Cayley graph, so the character transform allows direct access to its spectrum and to the quantum-walk Hamiltonians generated by such graphs. If T_f is used inside a phase-estimation-style routine, irreducible sectors with distinguishable eigenvalues can be discriminated, which connects directly to symmetry projection and to the generalized symmetry-adapted constructions advocated by Bastidas et al. If one instead applies polynomial or rational functions of T_f , the same mechanism becomes a central spectral filter, a Green's-function/resolvent primitive, or a tool for isolating selected symmetry sectors before simulation or measurement. The character transform is also most compelling from the standpoint of compression. The full nonabelian QFT acts on $\mathbb{C}[\mathcal{G}]$, whose dimension is $|\mathcal{G}|$, whereas the QCT acts only on the class algebra, whose dimension is $d = k(\mathcal{G})$, the number of conjugacy classes. The natural state-space dimension of the transform is $k(\mathcal{G}) \leq |\mathcal{G}|$, and for binary encoding, $\lceil \log_2 k(\mathcal{G}) \rceil$ qubits. For noncommutative groups such as S_n , this becomes a dramatic asymptotic compression of the representation space. That compression is capitalizable whenever the relevant data is already conjugacy-invariant, and it has a useful connection to symmetry-aware machine learning. Class-algebra features are the compressed harmonic representatives of data defined only up to relabeling or conjugation, much as invariant and equivariant architectures exploit group structure to reduce model complexity and improve generalization.

I develop an algorithmic theory of the center $Z(\mathbb{C}[\mathcal{G}])$. We first establish the discrete character transform's correct relation to the full non-abelian quantum Fourier transform, showing that the normalized character table of $\mathcal{G}$ is procured by restricting the group Fourier transform to the conjugation-invariant subspace and coherently removing the redundant matrix-index degrees of freedom. This general algorithm is efficient for all known finite groups

admitting an efficient quantum Fourier transform, and unites earlier character-basis works in a common and comprehensive operator-theoretic framework. I then use modern group convolution methods to build quantum algorithms specifically for class convolution, and show how they support central-operator synthesis and spectral manipulation. I also apply this transform to symmetry-adapted many-body and molecular problems, quantum walks, and hidden-subgroup settings, where the information retained by restriction to the center and the limit of its usefulness for graph-isomorphism-type problems is clarified. An open-source implementation and benchmarking framework for the QCT and the class-convolution algorithms is under active development. Preliminary quantum-hardware experiments have been conducted for the QCT over D_n and class convolution over S_4 , with scaling results and empirical comparisons against classical GPU implementations to be reported in a subsequent version of this manuscript.

II. GROUP ALGEBRAS

The group algebra $\mathbb{C}[\mathcal{G}]$ is the complex vector space spanned by the elements of $\mathcal{G}$, consisting of all formal sums

$$x = \sum_{g \in \mathcal{G}} f(g) g, \quad f : \mathcal{G} \rightarrow \mathbb{C}.$$

Multiplication is defined by bilinearly extending the group law, making $\mathbb{C}[\mathcal{G}]$ a generally noncommutative group ring of $\mathcal{G}$ over $\mathbb{C}$. The function algebra $\mathbb{C}(\mathcal{G})$ is the space of complex-valued functions on $\mathcal{G}$, equipped with pointwise multiplication, $(f_1 f_2)(g) = f_1(g) f_2(g)$. Every group algebra element x is uniquely determined by its coefficient function $f : \mathcal{G} \rightarrow \mathbb{C}$, so there is a canonical vector space isomorphism $\mathbb{C}[\mathcal{G}] \cong \mathbb{C}(\mathcal{G})$. Although isomorphic as vector spaces, they differ as algebras, $\mathbb{C}(\mathcal{G})$ uses pointwise multiplication, whereas under this identification, the product of two group algebra elements $x_1 = \sum_g f_1(g) g$ and $x_2 = \sum_g f_2(g) g$ is represented by the convolution

$$(f_1 * f_2)(h) = \sum_{g \in \mathcal{G}} f_1(g) f_2(g^{-1}h), \quad (1)$$

which forms the algebraic foundation for Fourier analysis on $\mathcal{G}$. The Fourier transform diagonalizes convolution, reducing convolution to pointwise multiplication in the Fourier basis.

The center of the group algebra, $Z(\mathbb{C}[\mathcal{G}])$, consists of the elements x commuting with every other element. Since the elements $g \in \mathcal{G}$ form a basis of $\mathbb{C}[\mathcal{G}]$, it suffices to verify commutation with each basis element. Under the canonical vector space isomorphism $\mathbb{C}[\mathcal{G}] \cong \mathbb{C}(\mathcal{G})$, this condition is the same as the coefficient function being constant on conjugacy classes, that is, to being a class function. The class algebra, $\mathbb{C}(\mathcal{G})$, is the subalgebra of $\mathbb{C}[\mathcal{G}]$ consisting of all class functions,

$$\mathbb{C}(\mathcal{G}) := \{f \in \mathbb{C}(\mathcal{G}) \mid f(hgh^{-1}) = f(g), \forall g, h \in \mathcal{G}\}. \quad (2)$$

As a vector space, $\mathbb{C}(\mathcal{G})$ is indexed by the conjugacy classes of $\mathcal{G}$. For each conjugacy class $C \in \text{Conj}(\mathcal{G})$, define its indicator function by $\mathbb{1}_C(g)$. The collection $\{\mathbb{1}_C\}$ forms a basis for $\mathbb{C}(\mathcal{G})$. To obtain an orthonormal basis, these indicator functions must be normalized. The normalized counting inner product on $\mathbb{C}(\mathcal{G})$ is $\langle f_1, f_2 \rangle = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} f_1(g) \overline{f_2(g)}$, and following, $\langle \mathbb{1}_C, \mathbb{1}_C \rangle = \frac{|C|}{|\mathcal{G}|}$. The size of a conjugacy class is related to the size of its centralizer $Z_g = \{h \in \mathcal{G} \mid hg = gh\}$ by the orbit-stabilizer theorem, $|C| = \frac{|\mathcal{G}|}{|Z_g|}$. Hence the normalized class basis is

$$\{\hat{\mathbb{1}}_C(g)\} = \begin{cases} \sqrt{|Z_g|}, & g \in C, \\ 0, & \text{otherwise,} \end{cases} \quad (3)$$

where g is any representative of the conjugacy class C . The class algebra inherits the inner product $\langle f_1, f_2 \rangle = \sum_C \frac{|C|}{|\mathcal{G}|} f_1(C) \overline{f_2(C)}$, under which the irreducible characters form an orthonormal basis.

A. The Quantum Character Transform

The irreducible characters are a special family of class functions, namely, the Fourier basis for the center of the group algebra, satisfying

$$\chi_\Gamma(g) = \chi_\Gamma(hgh^{-1}), \quad \forall g, h \in \mathcal{G}.$$

Together with the normalized conjugacy-class indicator functions, they form two orthonormal bases of the class algebra.

Definition. The *Discrete Character Transform* is the change of basis between the conjugacy-class basis of $Z(\mathbb{C}[\mathcal{G}])$ and the irreducible character basis,

$$U_{\text{CT}} : Z(\mathbb{C}[\mathcal{G}]) \longrightarrow \mathbb{C}[\hat{\mathcal{G}}].$$

For each conjugacy class $C \in \text{Conj}(\mathcal{G})$, define its class sum $K = \sum_{g \in C} g \in Z(\mathbb{C}[\mathcal{G}])$, these class sums form a basis of the center $Z(\mathbb{C}[\mathcal{G}])$. Under amplitude encoding, the normalized indicator functions become the class states $|C\rangle = |C|^{-1/2} \sum_{g \in C} |g\rangle$. Likewise, each irreducible character defines the normalized quantum state $|\chi_\Gamma\rangle = \sum_C \sqrt{|C|/|\mathcal{G}|} \chi_\Gamma(C) |C\rangle$. From this unitary representation, comes the Quantum Character Transform (QCT),

$$\sum_C a_C |C\rangle \mapsto \sum_{\Gamma \in \hat{\mathcal{G}}} \left(\sum_C a_C \chi_\Gamma(C)^* \sqrt{\frac{|C|}{|\mathcal{G}|}} \right) |\Gamma\rangle. \quad (4)$$

Bastidas et al. introduced the Quantum Character Transform as the normalized character table appearing in a symmetry-adapted projection circuit. Their framework indexes the input register by conjugacy classes and the output register by irreducible representations, thereby obtaining the normalized character matrix, whose orientation as a forward or inverse transform determines whether its entries are $\sqrt{|C|/|\mathcal{G}|} \chi_\Gamma(C)^*$ or their complex conjugates. This

operational definition established the transform as a useful ingredient for coherent projection onto symmetry sectors, but left implicit the algebraic space on which it acts and its precise relation to the non-abelian group Fourier transform. Additionally, they extract the character transform from the full group Fourier through a probabilistic postselection of the matrix-index registers, which introduces a success-probability overhead. The following theorem shows that this projection is unnecessary; the redundant matrix-index degrees of freedom can instead be removed coherently and deterministically.

Theorem II.1 (Center restriction). *The Quantum Character Transform is the restriction of the group Fourier transform to the center,*

$$\mathcal{F}_\mathcal{G}|_{Z(\mathbb{C}[\mathcal{G}])} \cong U_{\text{QCT}}.$$

Proof. Let $\mathcal{G}$ be a finite group, $\hat{\mathcal{G}}$ a complete set of inequivalent irreducible unitary representations, and $\mathcal{F}_\mathcal{G} : \mathbb{C}[\mathcal{G}] \longrightarrow \bigoplus_{\Gamma \in \hat{\mathcal{G}}} V_\Gamma \otimes V_\Gamma^*$ the nonabelian Fourier transform. Applying the Fourier transform to a normalized conjugacy-class state,

$$\begin{aligned} \mathcal{F}_\mathcal{G}|C\rangle &= \frac{1}{\sqrt{|C|}} \sum_{g \in C} \mathcal{F}_\mathcal{G}|g\rangle \\ &= \sum_{\Gamma \in \hat{\mathcal{G}}} \sqrt{\frac{d_\Gamma}{|\mathcal{G}||C|}} \sum_{i,j} \left(\sum_{g \in C} \rho_\Gamma(g^{-1}) \right)_{ij} |\Gamma, i, j\rangle. \end{aligned}$$

Since C^{-1} is a conjugacy class, the operator $\sum_{g \in C} \rho_\Gamma(g^{-1})$ commutes with every $\rho_\Gamma(h)$. By Schur's lemma, $\sum_{g \in C} \rho_\Gamma(g^{-1}) = \alpha_{\Gamma, C} I_{d_\Gamma}$ for some scalar $\alpha_{\Gamma, C}$. Taking traces, $\alpha_{\Gamma, C} = \frac{|C| \chi_\Gamma(C^{-1})}{d_\Gamma} = \frac{|C| \chi_\Gamma(C)^*}{d_\Gamma}$, and therefore

$$\sum_{g \in C} \rho_\Gamma(g^{-1}) = \frac{|C| \chi_\Gamma(C)^*}{d_\Gamma} I_{d_\Gamma}.$$

The identity matrix spans the unique invariant line in $V_\Gamma \otimes V_\Gamma^*$, namely $\text{vec}(I_{V_\Gamma}) = \sum_i |i, i\rangle$. Let $|I_\Gamma\rangle = \frac{1}{\sqrt{d_\Gamma}} \text{vec}(I_{V_\Gamma})$. Substituting,

$$\begin{aligned} \mathcal{F}_\mathcal{G}|C\rangle &= \sum_{\Gamma \in \hat{\mathcal{G}}} \sqrt{\frac{|C|}{|\mathcal{G}|d_\Gamma}} \chi_\Gamma(C)^* \sum_{i=1}^{d_\Gamma} |\Gamma, i, i\rangle \\ &= \sum_{\Gamma \in \hat{\mathcal{G}}} \sqrt{\frac{|C|}{|\mathcal{G}|}} \chi_\Gamma(C)^* |\Gamma\rangle |I_\Gamma\rangle. \end{aligned}$$

Thus every conjugacy-class state is mapped into the invariant line of each irreducible block, and $\text{Im}(\mathcal{F}_\mathcal{G}|_{Z(\mathbb{C}[\mathcal{G}])}) = \bigoplus_{\Gamma \in \hat{\mathcal{G}}} \mathbb{C} |\Gamma\rangle |I_\Gamma\rangle$. Under the canonical unitary identification $|\Gamma\rangle |I_\Gamma\rangle \mapsto |\Gamma\rangle$, the restriction acts by

$$|C\rangle \mapsto \sum_{\Gamma \in \hat{\mathcal{G}}} \sqrt{\frac{|C|}{|\mathcal{G}|}} \chi_\Gamma(C)^* |\Gamma\rangle,$$

whose matrix is the normalized character table. $\square$

Remark. The state $|I_\Gamma\rangle = d_\Gamma^{-1/2} \text{vec}(I_{V_\Gamma})$ is the unique (up to phase) normalized $(\rho_\Gamma \otimes \rho_\Gamma^*)$ -invariant vector in $V_\Gamma \otimes V_\Gamma^*$. Under the Choi–Jamiołkowski correspondence [20, 42], it represents the identity intertwiner on V_Γ , and is the canonical maximally entangled state associated with the irrep V_Γ . Restricting the Fourier transform to the center removes every noncentral matrix direction, leaving this canonical invariant vector in each irreducible block. The entries of the normalized character table are the expansion coefficients of conjugacy-class states in the basis of these canonical invariant vectors.

This character-basis reduction appears in Ref. [8, Appendix D.1]. The hypergroup algebra of the abelian conjugacy-class hypergroup $\widehat{\mathcal{G}}$ is isomorphic to $Z(\mathbb{C}[\mathcal{G}])$, so our formulation positions the character transform inside the non-abelian Fourier transform as a canonical restriction to the central subspace.

III. CLASS CONVOLUTION

Let $f_1, f_2 \in \mathbb{C}(\mathcal{G})$ be class functions. Left convolution by f_1 defines the action of a convolution operator T_{f_1} by $T_{f_1}f_2 = f_1 * f_2$. When $\mathcal{G}$ is abelian, $Z(\mathbb{C}[\mathcal{G}]) = \mathbb{C}[\mathcal{G}]$, and T_f is the usual translation-invariant group-circulant convolution operator. For cyclic groups, this is the familiar circulant Toeplitz operator. For a general finite group, T_f is its central, conjugacy-invariant counterpart. Since the convolution of two class functions is again a class function, T_f restricts to an operator on $Z(\mathbb{C}[\mathcal{G}])$. Let $C_1, \dots, C_d$ be the conjugacy classes of $\mathcal{G}$. In the normalized class basis, the matrix elements of T_f are

$$(T_f)_{ij} = \langle C_i | T_f | C_j \rangle = \frac{1}{\sqrt{|C_i||C_j|}} \sum_{g \in C_i} \sum_{h \in C_j} f(gh^{-1}). \quad (5)$$

Choosing the alternative convolution convention replaces $f(gh^{-1})$ by $f(g^{-1}h)$; the two are related by the involution $f^\vee(g) = f(g^{-1})$. Class convolution operators commute because $Z(\mathbb{C}[\mathcal{G}])$ is commutative, $T_{f_1}T_{f_2} = T_{f_1*f_2} = T_{f_2*f_1} = T_{f_2}T_{f_1}$. Under the involution $f^*(g) = \overline{f(g^{-1})}$, one has $T_f^\dagger = T_{f^*}$. Also, $T_f T_f^\dagger = T_{f*f^*} = T_{f^*f} = T_f^\dagger T_f$, so every class convolution operator is *normal*. It is Hermitian when $f(g^{-1}) = \overline{f(g)}$ for all $g \in \mathcal{G}$, and real symmetric when f is real-valued and invariant under inversion.

The normalized irreducible characters form the common eigenbasis of these operators. To determine the eigenvalue, associate to f the central group-algebra element $z_f^\vee = \sum_{g \in \mathcal{G}} f(g)g^{-1}$.

$$\begin{aligned} \rho_\Gamma(z_f^\vee) &= \widehat{f}(\Gamma)I_{V_\Gamma}, \quad \widehat{f}(\Gamma) = \frac{1}{d_\Gamma} \sum_{g \in \mathcal{G}} f(g)\chi_\Gamma(g^{-1}) \\ &= \frac{1}{d_\Gamma} \sum_{i=1}^d |C_i| f(C_i) \overline{\chi_\Gamma(C_i)}. \end{aligned}$$

It follows that $T_f|\chi_\Gamma\rangle = \widehat{f}(\Gamma)|\chi_\Gamma\rangle$. Writing $D_{\widehat{f}} = \text{diag}(\widehat{f}(\Gamma))_{\Gamma \in \widehat{\mathcal{G}}}$, we obtain the following.

Corollary III.1. *The Discrete Character Transform simultaneously diagonalizes the entire algebra of class convolution operators,*

$$U_{\text{CT}} T_f U_{\text{CT}}^{-1} = D_{\widehat{f}}. \quad (6)$$

Rearranging, $T_f = U_{\text{CT}}^{-1} D_{\widehat{f}} U_{\text{CT}}$. Thus the regular representation of the class convolution algebra is equivalent to the algebra of diagonal operators on the irreducible-character labels,

$$\{T_f : f \in Z(\mathbb{C}[\mathcal{G}])\} \cong \bigoplus_{\Gamma \in \widehat{\mathcal{G}}} \mathbb{C}.$$

Class convolution over S_4 . As a running example, we focus on S_4 . Let the conjugacy classes C_i be ordered as $C_1 = (1^4)$, $C_2 = (2, 1^2)$, $C_3 = (2^2)$, $C_4 = (3, 1)$, $C_5 = (4)$, with respective sizes $|C_1| = 1$, $|C_2| = 6$, $|C_3| = 3$, $|C_4| = 8$, $|C_5| = 6$. In this ordering, the normalized character table defines the S_4 quantum character transform below.

TABLE I. Normalized character table of S_4 .

QCT $_{S_4}$	(1^4)	$(2, 1^2)$	(2^2)	$(3, 1)$	(4)
$[4]$	$\frac{1}{\sqrt{24}}$	$\frac{1}{2}$	$\frac{1}{2\sqrt{2}}$	$\frac{1}{\sqrt{3}}$	$\frac{1}{2}$
$[1^4]$	$\frac{1}{\sqrt{24}}$	$-\frac{1}{2}$	$\frac{1}{2\sqrt{2}}$	$\frac{1}{\sqrt{3}}$	$-\frac{1}{2}$
$[3, 1]$	$\frac{1}{\sqrt{24}}$	$\frac{1}{2}$	$-\frac{1}{2\sqrt{2}}$	0	$-\frac{1}{2}$
$[2, 1, 1]$	$\frac{1}{\sqrt{24}}$	$-\frac{1}{2}$	$-\frac{1}{2\sqrt{2}}$	0	$\frac{1}{2}$
$[2, 2]$	$\frac{1}{\sqrt{24}}$	0	$\frac{1}{\sqrt{2}}$	$-\frac{1}{\sqrt{3}}$	0

The rows are indexed by the irreducible representations of S_4 , labeled by the partitions $[4]$, $[1^4]$, $[3, 1]$, $[2, 1, 1]$, and $[2, 2]$. For a class function f , write $f = \sum_{i=1}^d f(C_i)\mathbb{1}_{C_i}$. Ergo,

$$T_f = \sum_{i=1}^d f(C_i)T_{C_i}, \quad T_{C_i} := T_{\mathbb{1}_{C_i}}. \quad (7)$$

The operators T_{C_i} form a basis for the algebra of class convolution operators. In the normalized conjugacy-class basis, for S_4 these operators are

$$\begin{aligned} T_{(1^4)} &= I_5, \\ T_{(2, 1^2)} &= \begin{pmatrix} 0 & \sqrt{6} & 0 & 0 & 0 \\ \sqrt{6} & 0 & \sqrt{2} & 2\sqrt{3} & 0 \\ 0 & \sqrt{2} & 0 & 0 & 2\sqrt{2} \\ 0 & 2\sqrt{3} & 0 & 0 & 2\sqrt{3} \\ 0 & 0 & 2\sqrt{2} & 2\sqrt{3} & 0 \end{pmatrix}, \\ T_{(2^2)} &= \begin{pmatrix} 0 & 0 & \sqrt{3} & 0 & 0 \\ 0 & 1 & 0 & 0 & 2 \\ \sqrt{3} & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 3 & 0 \\ 0 & 2 & 0 & 0 & 1 \end{pmatrix}, \\ T_{(3, 1)} &= \begin{pmatrix} 0 & 0 & 0 & 2\sqrt{2} & 0 \\ 0 & 4 & 0 & 0 & 4 \\ 0 & 0 & 0 & 2\sqrt{6} & 0 \\ 2\sqrt{2} & 0 & 2\sqrt{6} & 4 & 0 \\ 0 & 4 & 0 & 0 & 4 \end{pmatrix}, \\ T_{(4)} &= \begin{pmatrix} 0 & 0 & 0 & 0 & \sqrt{6} \\ 0 & 0 & 2\sqrt{2} & 2\sqrt{3} & 0 \\ 0 & 2\sqrt{2} & 0 & 0 & \sqrt{2} \\ 0 & 2\sqrt{3} & 0 & 0 & 2\sqrt{3} \\ \sqrt{6} & 0 & \sqrt{2} & 2\sqrt{3} & 0 \end{pmatrix}. \end{aligned} \quad (8)$$

These matrices commute and form the regular matrix representation of $Z(\mathbb{C}[S_4])$. By (6), the Character Transform simultaneously diagonalizes the basis operators, $U_{\text{CT}}^{(S_4)} T_{C_i} \left(U_{\text{CT}}^{(S_4)} \right)^\dagger = D_{C_i}$, and by the eigenvalue formula (III),

$$\begin{aligned} D_{(1^4)} &= \text{diag}(1, 1, 1, 1, 1), \\ D_{(2,1^2)} &= \text{diag}(6, -6, 2, -2, 0), \\ D_{(2^2)} &= \text{diag}(3, 3, -1, -1, 3), \\ D_{(3,1)} &= \text{diag}(8, 8, 0, 0, -4), \\ D_{(4)} &= \text{diag}(6, -6, -2, 2, 0). \end{aligned}$$

Every class convolution operator T_f is diagonalized by the same transform where

$$\begin{aligned} D_{\hat{f}} &= \text{diag}(\hat{f}([4]), \hat{f}([1^4]), \hat{f}([3, 1]), \hat{f}([2, 1, 1]), \hat{f}([2, 2])) \\ &= \text{diag} \begin{pmatrix} f(C_1)+6f(C_2)+3f(C_3)+8f(C_4)+6f(C_5) \\ f(C_1)-6f(C_2)+3f(C_3)+8f(C_4)-6f(C_5) \\ f(C_1)+2f(C_2)-f(C_3)-2f(C_5) \\ f(C_1)-2f(C_2)-f(C_3)+2f(C_5) \\ f(C_1)+3f(C_3)-4f(C_4) \end{pmatrix}. \end{aligned}$$

Conversely, reconstructing from the character-domain spectrum,

$$\begin{aligned} T_f &= \left(U_{\text{QCT}}^{(S_4)} \right)^\dagger D_{\hat{f}} U_{\text{QCT}}^{(S_4)} \\ &= \begin{pmatrix} a & b & c & d & e \\ b & r & g & h & i \\ c & g & j & k & \ell \\ d & h & k & m & h \\ e & i & \ell & h & r \end{pmatrix}, \end{aligned}$$

$$\begin{aligned} a &= f(C_1), & r &= f(C_1) + f(C_3) + 4f(C_4), \\ b &= \sqrt{6} f(C_2), & g &= \sqrt{2}(f(C_2) + 2f(C_5)), \\ c &= \sqrt{3} f(C_3), & h &= 2\sqrt{3}(f(C_2) + f(C_5)), \\ d &= 2\sqrt{2} f(C_4), & i &= 2f(C_3) + 4f(C_4), \\ e &= \sqrt{6} f(C_5), & j &= f(C_1) + 2f(C_3), \\ k &= 2\sqrt{6} f(C_4), & \ell &= \sqrt{2}(2f(C_2) + f(C_5)), \\ & & m &= f(C_1) + 3f(C_3) + 4f(C_4). \end{aligned}$$

The reconstructed class operator is symmetric and exhibits a pronounced Hankel-like repeated-entry structure. In fact, the apparent 5×5 density is highly redundant. Since the first row already determines the class-function coefficients,

$$\begin{aligned} f(C_1) &= a, & f(C_2) &= \frac{b}{\sqrt{6}}, & f(C_3) &= \frac{c}{\sqrt{3}}, \\ f(C_4) &= \frac{d}{2\sqrt{2}}, & f(C_5) &= \frac{e}{\sqrt{6}}, \end{aligned}$$

every remaining entry is fixed by a, b, c, d, e .

$$\begin{aligned} r &= a + \frac{c}{\sqrt{3}} + \sqrt{2} d, & g &= \frac{b+2e}{\sqrt{3}}, & h &= \sqrt{2}(b+e), \\ i &= \frac{2c}{\sqrt{3}} + \sqrt{2} d, & k &= \sqrt{3} d, & \ell &= \frac{2b+e}{\sqrt{3}}. \end{aligned}$$

The entire 25-entry operator can be written in terms of its five-entry first row. The class-convolution operators form a five-dimensional complex subspace of the 15-dimensional complex vector space of symmetric 5×5 matrices. Indeed, eliminating the class-function coefficients in favor of a, b, c, d, e shows

$$T_f = \begin{pmatrix} a & b & c & d & e \\ b & a + \frac{c}{\sqrt{3}} + \sqrt{2} d & \frac{b+2e}{\sqrt{3}} & \sqrt{2}(b+e) & \frac{2c}{\sqrt{3}} + \sqrt{2} d \\ c & \frac{b+2e}{\sqrt{3}} & a + \frac{2c}{\sqrt{3}} & \sqrt{3} d & \frac{2b+e}{\sqrt{3}} \\ d & \sqrt{2}(b+e) & \sqrt{3} d & a + \sqrt{3} c + \sqrt{2} d & \sqrt{2}(b+e) \\ e & \frac{2c}{\sqrt{3}} + \sqrt{2} d & \frac{2b+e}{\sqrt{3}} & \sqrt{2}(b+e) & a + \frac{c}{\sqrt{3}} + \sqrt{2} d \end{pmatrix}.$$

The entry relations and Hankel-like patterns visible in this S_4 example are from the constrained class algebra, as opposed to the translation-invariant abelian group algebras in which convolution operators take Toeplitz form. While this specific reflection pattern depends on the group, and the ordering of conjugacy classes, the strong global redundancy that it exhibits persists for general class convolution operators.

A. Implementation Methods

Quantum class convolution is now of interest as a symmetry-reduced form of group convolution, but the underlying class algebra also has a long history in classical representation theory and physical applications. For symmetric groups, class-sum products have been studied in connection with systems of identical particles [43], while Harrison and Scott used the standard representation of the S_3 class algebra to construct a charged-lepton mass matrix whose spectral structure is determined by S_3 character data [39]. Earlier this decade, quantum algorithms for the more general problem of group convolution and cross-correlation were developed by Castelazo et al. [19]. Naively, convolution is computed through the full nonabelian Fourier transform, where it becomes multiplication within irreducible Fourier blocks. For a class function, Schur's lemma makes the Fourier blocks scalar, so the internal matrix indices are unnecessary for the convolution spectrum. The diagonalization above provides two complementary implementations specialized to this central case. One may construct T_f directly as a linear combination of class-sum operators, or materialize it spectrally by conjugating a diagonal operator with the character transform. For Hermitian T_f , the same spectral representation computes the evolution e^{-itT_f} .

Direct Class-Sum Block Encoding

A direct implementation is natural when the class function is specified explicitly in the conjugacy-class basis, particularly when only a small number of conjugacy classes have nonzero coefficients, as occurs in sparse class functions and in many weighted Cayley-graph, combinatorial-Laplacian, and symmetry-adapted averaging operators. One may di-

rectly realize the convolution through a coherent superposition of left-regular group actions, yielding a block encoding whose complexity depends on the required state preparation and controlled group actions. This does not require an implementation of the QCT.

Let L_g denote the left-regular action $(L_g f_2)(h) = f_2(g^{-1}h)$. For the convolution convention adopted above, $T_{C_i} = \sum_{g \in C_i} L_g$, and

$$T_f = \sum_{i=1}^d f(C_i) \sum_{g \in C_i} L_g = \sum_{g \in \mathcal{G}} f(g) L_g.$$

Thus, a class convolution operator is a linear combination of unitary left-regular group actions with coefficients constant on conjugacy classes. Define the normalization

$$\alpha_f = \sum_{i=1}^d |C_i| |f(C_i)| = \sum_{g \in \mathcal{G}} |f(g)|.$$

Assuming coherent access to the class weights and phases of f , efficient preparation of normalized conjugacy-class states, and efficient controlled group multiplication, the following operations can be implemented coherently:

$$\text{PREP}_f |0\rangle = \frac{1}{\sqrt{\alpha_f}} \sum_{i=1}^d \sqrt{|C_i| |f(C_i)|} |i\rangle \frac{1}{\sqrt{|C_i|}} \sum_{g \in C_i} |g\rangle,$$

$$\text{SELECT}_f : |i\rangle |g\rangle |\psi\rangle \mapsto \frac{f(C_i)}{|f(C_i)|} |i\rangle |g\rangle L_g |\psi\rangle. \quad (9)$$

The phase may be chosen arbitrarily whenever $f(C_i) = 0$. Then $U_f = (\text{PREP}_f^\dagger \otimes I) \text{SELECT}_f (\text{PREP}_f \otimes I)$ satisfies $(\langle 0| \otimes I) U_f (|0\rangle \otimes I) = \frac{T_f}{\alpha_f}$, so U_f is an $(\alpha_f, a, 0)$ -block encoding of T_f , where a is the number of ancilla qubits needed by the class and group-element registers [33].



FIG. 1. Direct class-sum block encoding of a class convolution operator. The preparation oracle PREP_f prepares a weighted superposition over conjugacy classes and their elements, while SELECT_f coherently applies the corresponding left-regular action with the class-dependent phase. Uncomputing the ancilla yields a block encoding of T_f/α_f .

Proposition III.2 (Direct class-sum block encoding). *Let $0 \neq f \in \mathcal{C}(\mathcal{G})$. T_f/α_f has an exact block encoding with complexity*

$$2 \text{cost}(\text{PREP}_f) + \text{cost}(\text{SELECT}_f).$$

Because each L_g acts on the full group algebra and T_f preserves its center, the same block encoding restricts to $Z(\mathbb{C}[\mathcal{G}])$. This is the central specialization of

the linear combination of unitaries method for general group convolution in Ref. [19]. The block encoding is efficient whenever the required class-state preparation and controlled left-regular action are also of low gate complexity.

Consider the conjugacy class $C_{(3,1)}$ in S_4 ,

$$\{(123), (132), (124), (142), (134), (143), (234), (243)\}.$$

The associated class-convolution operator on the group-element basis is $T_{(3,1)} = \sum_{g \in C_{(3,1)}} L_g$ where $L_g |x\rangle = |gx\rangle$. $T_{(3,1)}$ is an equal-weight linear combination of eight permutation unitaries, and because $|C_{(3,1)}| = 8$, a three-qubit selector register can be prepared using only Hadamard gates: $H^{\otimes 3} |000\rangle = \frac{1}{\sqrt{8}} \sum_{r=0}^7 |r\rangle$.




FIG. 2. Direct class-sum block encoding of $T_{(3,1)}$. The selector register is prepared in a uniform superposition over the eight elements of $C_{(3,1)}$, and $\text{SELECT}_{(3,1)}$ coherently applies the corresponding left-regular actions. Uncomputing the selector yields $T_{(3,1)}/8$ in the $|000\rangle$ block.

Choose an ordering $C_{(3,1)} = \{g_0, \dots, g_7\}$ and define

$$\text{SELECT}_{(3,1)} = \sum_{r=0}^7 |r\rangle \langle r| \otimes L_{g_r}.$$

The LCU circuit has upper-left block $(\langle 000| H^{\otimes 3} \otimes I) \text{SELECT}_{(3,1)} (H^{\otimes 3} |000\rangle \otimes I) = \frac{1}{8} T_{(3,1)}$. Hence the circuit is an exact block encoding of the class sum with normalization factor $\alpha_{(3,1)} = 8$. On the normalized conjugacy-class states $|C_i\rangle$, the same operator has matrix entries

$$(T_{(3,1)})_{kj} = p_{(3,1),j}^k \sqrt{\frac{|C_k|}{|C_j|}},$$

where $p_{(3,1),j}^k$ are the class-algebra structure constants (also intersection numbers) defined by

$$K_{(3,1)} K_j = \sum_k p_{(3,1),j}^k K_k, \quad K_j = \sum_{x \in C_j} x.$$

For instance,

$$T_{(3,1)} |C_1\rangle = \sum_{g \in C_{(3,1)}} |g\rangle = \sqrt{8} |C_{(3,1)}\rangle = 2\sqrt{2} |C_{(3,1)}\rangle,$$

which is the entry $2\sqrt{2}$ in the first column of the compressed class-basis matrix of $T_{(3,1)}$.

Under the isomorphism $S_4 \cong T_d$, the class $C_{(3,1)}$ is the eight proper threefold rotations of tetrahedral symmetry [23]. Under a molecular representation ρ , the class sum operator $\sum_{g \in C_{(3,1)}} \rho(g)$ represents the

equal-weight sum of the $C_{3,1}$ symmetry operations, and applied to a molecular state, its normalized form acts as a symmetry-adapted filter according to the state's response under threefold rotational symmetry. In the character basis, $\frac{1}{8}U_{\text{QCT}}^{(S_4)}T_{(3,1)}\left(U_{\text{QCT}}^{(S_4)}\right)^\dagger = \text{diag}\left(1, 1, 0, 0, -\frac{1}{2}\right)$, so the sectors with eigenvalue 1 are preserved, those with eigenvalue 0 are annihilated, and the remaining sector is multiplied by $-1/2$.

Spectral Implementation through the QCT

Access to the quantum character transform provides a second implementation strategy in which a class convolution operator is specified by its eigenvalues instead of its expansion in group actions. In the irreducible-character basis, T_f is diagonal, with entries given by the spectral coefficients $\hat{f}(\Gamma)$, which is useful for operations defined spectrally. The spectral representation can also improve the block-encoding normalization, since it depends on $\max_{\Gamma} |\hat{f}(\Gamma)|$ instead of the ℓ_1 weight of the class-sum decomposition.

Let

$$\beta_f \geq \max_{\Gamma \in \hat{\mathcal{G}}} |\hat{f}(\Gamma)| = \|T_f\|.$$

Let $\hat{f}_\epsilon(\Gamma)$ be a finite-precision encoding of $\hat{f}(\Gamma)$, chosen such that $\max_{\Gamma \in \hat{\mathcal{G}}} |\hat{f}_\epsilon(\Gamma) - \hat{f}(\Gamma)| \leq \epsilon\beta_f$, $|\hat{f}_\epsilon(\Gamma)| \leq \beta_f$. Define the reversible spectral oracle $O_{\hat{f}}: |\Gamma\rangle|y\rangle \mapsto |\Gamma\rangle|y \oplus \hat{f}_\epsilon(\Gamma)\rangle$, so that $O_{\hat{f}}|\Gamma\rangle|0\rangle = |\Gamma\rangle|\hat{f}_\epsilon(\Gamma)\rangle$. Using reversible arithmetic, condition an ancilla amplitude-and-phase operation on the encoded value so that $(|0\rangle_{\text{anc}} \otimes I)U_{D_{\hat{f}_\epsilon}}(|0\rangle_{\text{anc}} \otimes I) = \frac{D_{\hat{f}_\epsilon}}{\beta_f}$, where the ancilla contains the spectral-value and signal registers. Conjugating this block encoding by the QCT, $U_f^{\text{spec}} = \left(I \otimes \left(U_{\text{QCT}}^{(g)}\right)^\dagger\right) U_{D_{\hat{f}_\epsilon}} \left(I \otimes U_{\text{QCT}}^{(g)}\right)$, which satisfies $\left\|(|0\rangle_{\text{anc}} \otimes I)U_f^{\text{spec}}(|0\rangle_{\text{anc}} \otimes I) - \frac{T_f}{\beta_f}\right\| \leq \epsilon$. U_f^{spec} is an ϵ -approximate block encoding of T_f with spectral normalization β_f .




FIG. 3. Spectral implementation of a class convolution operator. The QCT maps the input into the irreducible-character basis, where the oracle $O_{\hat{f}}$ computes the spectral coefficients $\hat{f}(\Gamma)$. The gate Amp performs a controlled amplitude encoding using the normalized spectral values $\hat{f}(\Gamma)/\beta_f$. After uncomputing the spectral register and applying QCT^{-1} , the circuit is an ϵ -approximate block encoding of T_f/β_f on the input register, relative to the $|0\rangle_{\text{val}}|0\rangle_{\text{sig}}$ ancilla state.

Proposition III.3 (Spectral block encoding through the QCT). *Let $f \in Z(\mathbb{C}[\mathcal{G}])$, and $\beta_f \geq \max_{\Gamma \in \hat{\mathcal{G}}} |\hat{f}(\Gamma)|$. Then T_f/β_f has an ϵ -approximate block encoding with complexity*

$$2 \text{cost}(\text{QCT}) + 2 \text{cost}(O_{\hat{f}}) + \text{polylog}(1/\epsilon),$$

where ϵ is the precision of the diagonal block encoding.

Choosing the optimal spectral normalization, $\beta_f = \max_{\Gamma \in \hat{\mathcal{G}}} |\hat{f}(\Gamma)| = \|T_f\|$,

$$\beta_f \leq \sum_i |C_i| |f(C_i)| = \alpha_f.$$

The inequality may be strict, so the spectral implementation can have a smaller normalization than the direct LCU construction.

The same spectral construction applies to functions of T_f . For a computable scalar function ϕ ,

$$\phi(T_f) = U_{\text{QCT}}^\dagger \text{diag}\left(\phi(\hat{f}(\Gamma))\right) U_{\text{QCT}}.$$

If $\phi(x) = e^{-itx}$,

$$e^{-itT_f} = U_{\text{QCT}}^\dagger \text{diag}\left(e^{-it\hat{f}(\Gamma)}\right) U_{\text{QCT}}.$$

Hamiltonian simulation can be performed by replacing the amplitude-encoding operation in Fig. 3 with the diagonal phase operation $U_{\text{phase}}(t)|\Gamma\rangle = e^{-it\hat{f}(\Gamma)}|\Gamma\rangle$.

Proposition III.4 (Hamiltonian simulation through the QCT). *Let T_f be Hermitian. Given coherent access to the spectral coefficients through $O_{\hat{f}}$, the evolution e^{-itT_f} can be implemented to precision ϵ with complexity*

$$2 \text{cost}(\text{QCT}) + 2 \text{cost}(O_{\hat{f}}) + \text{cost}(U_{\text{rot}}(t, \epsilon)).$$

The final term accounts for the reversible arithmetic and phase-rotation circuit applying $U_{\text{phase}}(t)$ to precision ϵ once $\hat{f}(\Gamma)$ has been computed. Quantum phase estimation is not required; the QCT supplies the eigenbasis, while $O_{\hat{f}}$ supplies the eigenvalues. Other efficiently computable unit-modulus spectral functions can be implemented by replacing the phase e^{-itx} , while bounded spectral functions that are not unit modulus require an appropriate diagonal block encoding or a polynomial-transformation procedure such as quantum singular value transformation [33].

IV. ALGORITHM BY RESTRICTION

The Restriction Theorem facilitates a universal implementation strategy for the QCT. By Theorem II.1, the non-abelian Fourier transform maps the center into the invariant lines

$$\mathcal{F}_{\mathcal{G}}|\psi\rangle = \sum_{\Gamma \in \hat{\mathcal{G}}} \hat{\psi}(\Gamma) |\Gamma\rangle |I_{\Gamma}\rangle,$$

where $|I_\Gamma\rangle = \frac{1}{\sqrt{d_\Gamma}} \text{vec}(I_{V_\Gamma}) = \frac{1}{\sqrt{d_\Gamma}} \sum_{i=1}^{d_\Gamma} |i, i\rangle$ is the canonical maximally entangled state in Theorem II.1. Thus, on the Fourier image of the center, the matrix-index registers contain no independent information about $|\psi\rangle$; conditioned on Γ , they occupy the known state $|I_\Gamma\rangle$. A restriction algorithm coherently unprepares these states, leaving only the irreducible-character amplitudes.

Definition (Restriction operator). Define the *restriction operator*

$$\mathcal{R}_\mathcal{G} := \sum_{\Gamma \in \widehat{\mathcal{G}}} |\Gamma\rangle\langle\Gamma| \otimes \langle I_\Gamma|.$$

In the irreducible-block decomposition of the Fourier space, $\mathcal{R}_\mathcal{G} = \text{diag}(\langle I_{\Gamma_1}|, \dots, \langle I_{\Gamma_d}|)$, and its action on the invariant subspace is

$$\mathcal{R}_\mathcal{G}(|\Gamma\rangle |I_\Gamma\rangle) = |\Gamma\rangle.$$

The restriction operator is a partial isometry from the full Fourier space onto the character space. In particular, $\mathcal{R}_\mathcal{G} \mathcal{R}_\mathcal{G}^\dagger = I_{|\widehat{\mathcal{G}}|}$, while $\mathcal{R}_\mathcal{G}^\dagger \mathcal{R}_\mathcal{G} = \sum_{\Gamma \in \widehat{\mathcal{G}}} |\Gamma\rangle\langle\Gamma| \otimes |I_\Gamma\rangle\langle I_\Gamma|$, which is the orthogonal projector onto $\text{Im}(\mathcal{F}_\mathcal{G}|_{Z(\mathbb{C}[\mathcal{G}]})}) = \bigoplus_{\Gamma \in \widehat{\mathcal{G}}} \mathbb{C} |\Gamma\rangle |I_\Gamma\rangle$. On this subspace, $\mathcal{R}_\mathcal{G}$ is the canonical unitary identification with the character space established by the Restriction Theorem. Consequently,

$$U_{\text{QCT}} = \mathcal{R}_\mathcal{G} \mathcal{F}_\mathcal{G}|_{Z(\mathbb{C}[\mathcal{G}])}.$$

To implement this identification coherently, for each positive integer d we need a maximally-entangled-state preparation unitary W_d satisfying

$$W_d |0, 0\rangle = |I_d\rangle := \frac{1}{\sqrt{d}} \sum_{i=0}^{d-1} |i, i\rangle. \quad (10)$$

In a logical d -dimensional encoding, such a unitary can be compiled by applying the Fourier transform F_d to the first matrix-index register, followed by the controlled-addition operation $\text{SUM}_d |i, j\rangle = |i, j + i \bmod d\rangle$, so that

$$W_d = \text{SUM}_d(F_d \otimes I). \quad (11)$$

Then, $W_d^\dagger |I_d\rangle = (F_d^\dagger \otimes I) \text{SUM}_d^\dagger |I_d\rangle = |0, 0\rangle$. For $d = 2$, $F_2 = H$ and $\text{SUM}_2 = \text{CNOT}$, and one can use $W_2 = \text{CNOT}(H \otimes I)$, the standard Bell-state preparation circuit. Taking the matrix-index registers to have a common ambient dimension, extend each W_{d_Γ} arbitrarily to a unitary on these registers. The resulting controlled unitary is

$$\mathcal{R}_\mathcal{G}^W := \sum_{\Gamma \in \widehat{\mathcal{G}}} |\Gamma\rangle\langle\Gamma| \otimes W_{d_\Gamma}^\dagger. \quad (12)$$

It satisfies $\mathcal{R}_\mathcal{G}^W(|\Gamma\rangle |I_\Gamma\rangle) = |\Gamma\rangle |0, 0\rangle$, and hence, for every $|\psi\rangle \in Z(\mathbb{C}[\mathcal{G}])$, $\mathcal{R}_\mathcal{G}^W \mathcal{F}_\mathcal{G} |\psi\rangle = (U_{\text{QCT}} |\psi\rangle) |0, 0\rangle$. The matrix-index registers are returned to the fixed state $|0, 0\rangle$ and can be discarded without loss of coherence.

1 Restriction Algorithm for $\text{QCT}_\mathcal{G}$

- 1: **Input:** State $|\psi\rangle \in Z(\mathbb{C}[\mathcal{G}])$, encoded in the full group register as a superposition of normalized conjugacy-class states.
- 2: Apply the group QFT: $|\phi\rangle = \mathcal{F}_\mathcal{G} |\psi\rangle$.
- 3: Apply the controlled restriction:

$$\mathcal{R}_\mathcal{G}^W |\phi\rangle = (U_{\text{QCT}} |\psi\rangle) |0, 0\rangle.$$

- 4: Discard the matrix-index registers $|0, 0\rangle$.
 - 5: **Output:** $U_{\text{QCT}} |\psi\rangle$ in the irreducible-character basis.
-

The additional cost of the restriction step is that of the controlled inverses $W_{d_\Gamma}^\dagger$, so

$$\text{cost}(\text{QCT}_\mathcal{G}) = \text{cost}(\text{QFT}_\mathcal{G}) + \text{cost}(\mathcal{R}_\mathcal{G}^W).$$

The restriction overhead is the cost of implementing $W_{d_\Gamma}^\dagger$ coherently from the irrep label Γ . Because $d_\Gamma^2 \leq |\mathcal{G}|$, the joint matrix-index space has dimension at most $|\mathcal{G}|$, so the restriction step uses $\mathcal{O}(\log |\mathcal{G}|)$ matrix-index qubits under a binary encoding. For families of finite groups in which d_Γ can be efficiently computed from Γ , the restriction step reduces to coherently controlled generalized Bell-state unpreparation and can be implemented with complexity polynomial in $\log |\mathcal{G}|$ and $\log(1/\epsilon)$. This holds for finite-group families with an efficient QFT and access to representation data, such as D_n , whose irreducible representations have dimension one or two, and S_n , where irreducible representations Γ are indexed by partitions $\lambda \vdash n$ and d_Γ is efficiently computable from the hook-length formula.

While the restriction algorithm operates on the full group space, its output has dimension $k(\mathcal{G}) = |\widehat{\mathcal{G}}|$, with one label per irreducible representation. A full group register uses $\lceil \log_2 |\mathcal{G}| \rceil$ qubits, while the character register requires only $\lceil \log_2 k(\mathcal{G}) \rceil$. This reduction provides a natural space-complexity parameter for comparing a *class-native QCT* with the restriction of a full group Fourier transform. When $k(\mathcal{G})$ grows substantially more slowly than $|\mathcal{G}|$, a QCT algorithm that works completely in the class basis acts on an asymptotically smaller state space, meaning smaller group-specific transform circuits may be possible.

Remark (Central compression). For a finite group $\mathcal{G}$, the *commuting probability* is

$$\text{cp}(\mathcal{G}) = \frac{|\{(g, h) \in \mathcal{G}^2 : gh = hg\}|}{|\mathcal{G}|^2} = \frac{k(\mathcal{G})}{|\mathcal{G}|},$$

where $k(\mathcal{G}) = |\text{Conj}(\mathcal{G})|$. Since $\dim Z(\mathbb{C}[\mathcal{G}]) = k(\mathcal{G})$, $\dim \mathbb{C}[\mathcal{G}] = |\mathcal{G}|$, it follows that

$$\text{cp}(\mathcal{G}) = \frac{\dim Z(\mathbb{C}[\mathcal{G}])}{\dim \mathbb{C}[\mathcal{G}]}.$$

The ratio between the QCT and QFT state-space dimensions is the commuting probability of the group, with $\frac{|\mathcal{G}|}{k(\mathcal{G})} = \text{cp}(\mathcal{G})^{-1}$ being the state-space compression factor for non-abelian groups.

For metacyclic groups such as D_n , however, this dimensional reduction does not produce an asymptotic reduction in register size. The dihedral group D_n has order $2n$, while

$$k(D_n) = \begin{cases} (n+3)/2, & n \text{ odd}, \\ n/2 + 3, & n \text{ even}, \end{cases}$$

so $k(D_n) = n/2 + \mathcal{O}(1)$. $\log_2 |D_n| - \log_2 k(D_n) \rightarrow 2$, and the class representation saves asymptotically only two qubits. Efficient quantum Fourier transforms for metacyclic groups are already known [40] with polylogarithmic gate complexity for D_n , so a constant reduction in register size shows no immediate indication of an asymptotically faster class-native character transform, although the existence of such a quantum circuit has not been formally ruled out. The classical complexity follows similarly. Individual irreducible characters of D_n have explicit closed forms and can be evaluated in $\mathcal{O}(1)$ arithmetic operations, while materializing the normalized character matrix requires writing $\Theta(n^2)$ entries. A classical restriction algorithm instead embeds the class representation into the group representation, applies a fast Fourier transform, and compresses the result back to the character representation. The embedding sends a class coefficient a_i to the constant value $a_i/\sqrt{|C_i|}$ on every element of C_i , where expanding the class representation over D_n requires $\Theta(|D_n|) = \Theta(n)$ operations, and compression is likewise $\Theta(n)$. Fast Fourier transforms for dihedral groups run in time $\mathcal{O}(|D_n| \log |D_n|) = \mathcal{O}(n \log n)$ [21], so the embedding and compression costs do not change the $\mathcal{O}(n \log n)$ complexity of the classical restriction algorithm.

The situation is different for the symmetric group S_n . Its order is $n!$, but $k(S_n) = p(n)$, where $p(n) \sim \frac{1}{4n\sqrt{3}} \exp(\pi\sqrt{2n/3})$ is the partition number. The group register is of size $\log |S_n| = \Theta(n \log n)$, whereas $\log k(S_n) = \Theta(\sqrt{n})$, a nontrivial asymptotic gap. A class-native algorithm can provide a polynomial compression in register size. In 2016, Kawano and Sekigawa reported an $\mathcal{O}(n^3 \log n)$ algorithm for the S_n quantum Fourier transform [44], and a recent reanalysis by Bruinsma, Grinko, and Ozols improves the gate bound to $\tilde{\mathcal{O}}(n^3)$ [13]. Shortly after the initial dissemination of this manuscript, Rathore and Zhou shared the first known algorithm for the QCT over S_n without passing through the full group QFT, using $\tilde{\mathcal{O}}(n^{5/2})$ gates and $\mathcal{O}(n)$ qubits [52]. While the restriction algorithm inherits the complexity of the QFT and the new Hamiltonian-based method achieves a square-root-factor improvement in gate complexity while using $\mathcal{O}(n)$ qubits, a native QCT acting entirely on this compressed class space could provide a further polynomial reduction in register size. Whether that space reduction can also yield a lower asymptotic gate complexity remains the more important question, with potential consequences for class convolution, symmetry-adapted spectral algorithms, and conjugacy-invariant formulations of group isomorphism.

V. HIDDEN SUBGROUPS AND CENTRAL SHADOWS

One of the primary motivations discussed for this transform was its application to the hidden subgroup problem over non-abelian groups, which turns out to be quite limited. This result is not surprising, as the restriction algorithm provides a rigorous way to understand the information retained by the QCT in hidden-subgroup problems. In the most general hidden-subgroup problem, an oracle is constant on left cosets of an unknown subgroup $H \leq \mathcal{G}$ and factors through the left-coset space $\mathcal{G}/H$ [27, 41], while the data relevant to the QCT is conjugacy invariant and factors through the set of conjugacy classes $\mathcal{G} // \mathcal{G}$. It is obvious that the character transform cannot reproduce the full information available under nonabelian Fourier sampling, but which information about H survives restriction to the center $Z(\mathbb{C}[\mathcal{G}])$? We call the invariant the *central shadow* of H , and it presents both a positive and a negative finding. It identifies the subgroup information accessible to class-convolution operators, while showing that distinct nonconjugate subgroups with the same central shadow cannot be distinguished by any protocol whose dependence on the hidden subgroup factors entirely through the class algebra.

For each irreducible representation $\rho_\Gamma : \mathcal{G} \rightarrow U(V_\Gamma)$, the normalized subgroup average determines the operator

$$\Pi_H^{(\Gamma)} = \frac{1}{|H|} \sum_{h \in H} \rho_\Gamma(h). \quad (13)$$

Since averaging over H is idempotent and self-adjoint in a unitary representation, $\Pi_H^{(\Gamma)}$ is the orthogonal projector onto the H -invariant subspace [55]

$$V_\Gamma^H = \{v \in V_\Gamma : \rho_\Gamma(h)v = v \text{ for all } h \in H\}.$$

The full nonabelian Fourier description of the subgroup average retains, within each irreducible sector Γ , the matrix-valued projector $\Pi_H^{(\Gamma)}$. Its rank is

$$\text{rank } \Pi_H^{(\Gamma)} = \text{Tr } \Pi_H^{(\Gamma)} = \dim V_\Gamma^H.$$

Projecting the subgroup average onto $Z(\mathbb{C}[\mathcal{G}])$ removes this internal matrix structure while retaining its conjugacy-invariant content.

Definition (Central shadow of a subgroup). Let $H \leq \mathcal{G}$. The *central shadow* of H is the conjugacy average of its normalized subgroup sum,

$$\begin{aligned} z_H &:= \frac{1}{|\mathcal{G}|} \sum_{x \in \mathcal{G}} x \left(\frac{1}{|H|} \sum_{h \in H} h \right) x^{-1} \\ &= \sum_{C \in \text{Conj}(\mathcal{G})} \frac{|H \cap C|}{|H||C|} K_C, \quad K_C := \sum_{g \in C} g. \end{aligned}$$

In the class basis, z_H records the distribution of H across the conjugacy classes of $\mathcal{G}$. It forgets which individual elements occur within a class and how those elements are jointly embedded as a subgroup.

Theorem V.1 (Character spectrum of a subgroup central shadow). *For every $H \leq \mathcal{G}$ and $\Gamma \in \widehat{\mathcal{G}}$, define $m_\Gamma(H) := \frac{1}{|\widehat{H}|} \sum_{h \in H} \chi_\Gamma(h)$. Then*

$$m_\Gamma(H) = \dim V_\Gamma^H \\ = \langle \text{Res}_H^\mathcal{G} \chi_\Gamma, 1_H \rangle_H = \left\langle \chi_\Gamma, \text{Ind}_H^\mathcal{G} 1_H \right\rangle_{\mathcal{G}}.$$

where 1_H denotes the trivial character of H . Moreover,

$$\rho_\Gamma(z_H) = \frac{m_\Gamma(H)}{d_\Gamma} I_{d_\Gamma},$$

and hence, under the spectral convention for class convolution,

$$\widehat{z}_H(\Gamma) = \frac{m_\Gamma(H)}{d_\Gamma} = \frac{\dim V_\Gamma^H}{d_\Gamma}.$$

The multiplicity profile $(m_\Gamma(H))_{\Gamma \in \widehat{\mathcal{G}}}$ is the irreducible-multiplicity vector of $\mathbb{C}[\mathcal{G}/H] \cong \text{Ind}_H^\mathcal{G} 1_H$.

This also makes clear what is lost under the restriction $\mathcal{R}_\mathcal{G}$. Before restriction, the Fourier block associated with Γ contains the projector $\Pi_H^{(\Gamma)}$, whereas its central component is captured only by the scalar $\text{Tr} \Pi_H^{(\Gamma)}$. At the level of the Γ -block, centralization acts as

$$\Pi_H^{(\Gamma)} \mapsto \frac{\text{Tr} \Pi_H^{(\Gamma)}}{d_\Gamma} I_{d_\Gamma} = \frac{m_\Gamma(H)}{d_\Gamma} I_{d_\Gamma}.$$

The character transform retains the dimension, but not the orientation of the invariant subspace V_Γ^H within V_Γ . That missing information is encoded in the internal matrix indices present in $\mathcal{F}_\mathcal{G}$ and removed by restriction to the center.

A. Weak Fourier Sampling

The same invariant appears in weak Fourier sampling. For a left coset state $|gH\rangle = \frac{1}{\sqrt{|\widehat{H}|}} \sum_{h \in H} |gh\rangle$, application of the nonabelian Fourier transform $\mathcal{F}_\mathcal{G}$, followed by measurement of only the irreducible-representation label Γ , gives

$$\Pr_H(\Gamma) = \frac{|H|d_\Gamma}{|\mathcal{G}|} \text{rank} \Pi_H^{(\Gamma)} = \frac{|H|d_\Gamma}{|\mathcal{G}|} \dim V_\Gamma^H.$$

Thus weak Fourier sampling depends on H only through the multiplicity profile

$$m_H = (m_\Gamma(H))_{\Gamma \in \widehat{\mathcal{G}}} = (\dim V_\Gamma^H)_{\Gamma \in \widehat{\mathcal{G}}}.$$

Equivalently, $\Pr_H(\Gamma) = \frac{|H|d_\Gamma^2}{|\mathcal{G}|} \widehat{z}_H(\Gamma)$. The distinction from strong Fourier sampling is that weak sampling retains only the irreducible-sector weights, whereas measurements of the internal Fourier indices may additionally probe the matrix structure of $\Pi_H^{(\Gamma)}$ removed by $\mathcal{R}_\mathcal{G}$.

The class-intersection profile, central shadow, and multiplicity profile are three descriptions of the same

conjugacy-invariant information. For two subgroups $H_1, H_2 \leq \mathcal{G}$,

$$z_{H_1} = z_{H_2} \iff \frac{|H_1 \cap C|}{|H_1|} = \frac{|H_2 \cap C|}{|H_2|} \quad \forall C \in \text{Conj}(\mathcal{G}),$$

and, since the QCT is invertible on the center, equivalently

$$\frac{1}{|H_1|} \sum_{h \in H_1} \chi_\Gamma(h) = \frac{1}{|H_2|} \sum_{g \in H_2} \chi_\Gamma(g) \quad \forall \Gamma \in \widehat{\mathcal{G}}.$$

Taking $C = \{e\}$ in the class-profile equality gives $|H_1| = |H_2|$, so equality of central shadows is equivalent to

$$|H_1 \cap C| = |H_2 \cap C| \quad \forall C \in \text{Conj}(\mathcal{G}).$$

This is the classical relation of *almost conjugacy*, or Gassmann equivalence [32, 51, 57]. Equivalently, $\text{Ind}_{H_1}^\mathcal{G} 1_{H_1} \cong \text{Ind}_{H_2}^\mathcal{G} 1_{H_2}$, and hence $\dim V_\Gamma^{H_1} = \dim V_\Gamma^{H_2}$ for every $\Gamma \in \widehat{\mathcal{G}}$ by Frobenius reciprocity [51, 55, 57]. Thus

$$z_{H_1} = z_{H_2} \iff H_1 \text{ and } H_2 \text{ are Gassmann equivalent.}$$

Corollary V.2 (Central-shadow indistinguishability). *Let $H_1, H_2 \leq \mathcal{G}$ satisfy $z_{H_1} = z_{H_2}$. Then their character spectra coincide, and every class-convolution operator has identical spectral response on the two subgroup shadows. More generally, any procedure whose dependence on the hidden subgroup factors entirely through $Z(\mathbb{C}[\mathcal{G}])$ cannot distinguish H_1 from H_2 .*

This limitation cannot be overcome by changing the class function, combining multiple class-convolution operators, or selecting such operators adaptively. Optimizing a central probe can change how the surviving information is accessed, but cannot recover information removed by projection onto $Z(\mathbb{C}[\mathcal{G}])$. Nonconjugate Gassmann-equivalent subgroups give intrinsic central-shadow collisions, not artifacts of a particular circuit or measurement prescription.

B. Hidden Subgroup Conjugacy Problem

A problem especially well matched to the conjugacy invariance of the central shadow is the *hidden subgroup conjugacy problem* (HSCP). In the HSCP, an oracle hides a subgroup $H \leq \mathcal{G}$ as in the ordinary hidden subgroup problem, but the objective is only to know its conjugacy class

$$[H]_\mathcal{G} = \{gHg^{-1} : g \in \mathcal{G}\},$$

instead of the particular hidden subgroup [3]. This is exactly the symmetry respected by the central shadow, since $z_{xHx^{-1}} = z_H$ for every $x \in \mathcal{G}$.

The central shadow therefore defines a conjugacy invariant of the hidden subgroup. The relevant question for the HSCP is whether this invariant is complete on the family of subgroups under consideration.

Proposition V.3 (Central-shadow criterion for the HSCP). *Let $\mathcal{H}$ be a promised family of subgroups of $\mathcal{G}$. The QCT central shadow is a complete invariant for the HSCP on $\mathcal{H}$ if and only if*

$$H_1, H_2 \in \mathcal{H}, \quad z_{H_1} = z_{H_2} \implies H_1 \sim_{\mathcal{G}} H_2.$$

Equivalently, $\mathcal{H}$ contains no pair of nonconjugate Gassmann-equivalent subgroups.

Proof. Conjugate subgroups have identical central shadows. Since the QCT is invertible on $Z(\mathbb{C}[\mathcal{G}])$, it preserves equality and inequality of central shadows. Thus the QCT central shadow reveals $[H]_{\mathcal{G}}$ on $\mathcal{H}$ exactly when $[H]_{\mathcal{G}} \mapsto z_H$ is injective on the promised family, or equivalently when $\mathcal{H}$ contains no pair of nonconjugate Gassmann-equivalent subgroups. $\square$

The criterion admits both positive and negative examples even among small symmetric groups. A simple positive example occurs for the family of Klein four subgroups of S_4 . Consider the Klein four subgroups

$$\begin{aligned} H_1 &= \{e, (12)(34), (13)(24), (14)(23)\} \\ H_2 &= \{e, (12), (34), (12)(34)\}. \end{aligned}$$

Their class-intersection profiles are $(1, 0, 3, 0, 0)$ and $(1, 2, 1, 0, 0)$, respectively. They are nonconjugate, and $z_{H_1} \neq z_{H_2}$, so their subgroup conjugacy classes are distinguishable by the central shadow.

However, this separation fails in general. In S_6 , consider

$$\begin{aligned} H_1 &= \{e, (34)(56), (35)(46), (36)(45)\}, \\ H_2 &= \{e, (34)(56), (12)(56), (12)(34)\}. \end{aligned}$$

Each subgroup intersects the identity class once and the class of cycle type $2^2 1^2$ three times, with empty intersection with every other conjugacy class. So their central shadows coincide: $z_{H_1} = z_{H_2}$. The subgroups are Gassmann equivalent but nevertheless nonconjugate. Their orbit-size decompositions in the natural action on six points are respectively $1 + 1 + 4$ and $2 + 2 + 2$, which are invariant under relabeling.

This also places the QCT relative to existing HSCP and character-sampling algorithms. Writing $\rho_H = \frac{1}{|\mathcal{G}|} \sum_{x \in \mathcal{G}} |xH\rangle\langle xH|$ for the single-copy hidden-subgroup state, Bacon showed that the single-copy hidden-subgroup conjugacy state $\rho_{[H]} = \frac{1}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \rho_{gHg^{-1}}$ decomposes after Fourier transformation into scalar irreducible blocks whose coefficients contain the same character averages $\sum_{h \in H} \chi_{\Gamma}(h)$ [3]. At the single-copy level, the central shadow captures the conjugacy-invariant multiplicity data. The multicopy problem is richer: the conjugacy average of several copies retains correlations between copies, and Bacon’s Heisenberg-group algorithm uses a Clebsch–Gordan transform and its multiplicity register to solve the HSCP with two hidden-subgroup states.

For S_n , Bravyi, Gosset, Havlicek, and Schatzki construct a polynomial-size quantum circuit that prepares the uniform state over the conjugacy class $\mathcal{C}_g$ of a specified permutation g and, after the S_n QFT, obtains character amplitudes proportional to $\chi_{\Gamma}(g)$ [10].

Their circuit exploits the same scalarization of central class sums underlying the QCT, but addresses a different task: it takes a specified conjugacy class as input, whereas the HSCP begins with an unknown hidden subgroup. The circuit alone does not solve the HSCP.

The proposition above is information-theoretic; by itself, it makes no claim about efficiency. Even when distinct promised subgroup conjugacy classes have distinct central shadows, an efficient HSCP algorithm further requires efficient access to the shadow, an efficient QCT, and sufficient separation of the resulting spectral data to permit discrimination with polynomial resources. Conversely, a Gassmann collision is absolute for central methods. No improvement in spectral probing or measurement can recover information already removed by projection onto $Z(\mathbb{C}[\mathcal{G}])$. Multicopy transforms such as the Clebsch–Gordan construction lie beyond this one-copy central restriction and may access additional representation-theoretic correlations.

C. Graph Isomorphism

Graph Isomorphism provides an important test case for the limitations of central information. The standard reduction of Graph Isomorphism to a hidden-subgroup problem produces a nonabelian HSP over a symmetric-group-related permutation group. This problem was one of the principal motivations for the study of nonabelian Fourier sampling, but it is also a setting in which some of its strongest limitations are known [36, 49].

For two rigid graphs on n vertices, the reduction is formulated over the wreath product

$$\mathcal{G} = S_n \wr S_2 = (S_n \times S_n) \rtimes S_2.$$

In the nonisomorphic case the hidden subgroup is trivial, whereas in the isomorphic case it is generated by an involution exchanging the two graph copies according to an isomorphism between their vertex sets. Under the standard embedding into S_{2n} , these generators are fixed-point-free involutions and hence have a definite conjugacy-class signature. This makes the reduction superficially well matched to a character-based analysis. The difficulty, however, is that this conjugacy-class signature does not make the two HSP hypotheses efficiently distinguishable by single-register Fourier information, much less identify the underlying graph isomorphism.

Indeed, Moore, Russell, and Schulman showed that the Fourier-sampling framework does not provide sufficient information for the Graph-Isomorphism-related hidden-subgroup problem over symmetric groups [49]. Their result rules out even strong Fourier sampling in this setting; the obstruction is information-theoretic and cannot be overcome by choosing a different basis within an irreducible block. Since the QCT retains only the conjugacy-invariant component of this representation-theoretic information, the central restriction cannot bypass the Fourier-sampling barrier, and the limitation extends to joint measurements on a bounded number of coset states. Hallgren, Moore,

Rottler, Russell, and Sen showed that obtaining non-negligible information for the Graph Isomorphism HSP requires measurements entangled across $\Omega(n \log n)$ coset states [36], so the missing information can reside in correlations among many hidden-subgroup states and remain inaccessible to any single-copy central statistic. Nor does the need for multicopy information automatically make representation-theoretic sieving sufficient. Moore, Russell, and Śniady showed that a broad class of Kuperberg-style quantum sieve algorithms, which iteratively combine representations according to tensor-product decompositions, cannot yield a polynomial-time algorithm for the Graph Isomorphism HSP; within their model the required running time is $e^{\Omega(\sqrt{n})}$ [50].

These results clarify the scope of the QCT rather than diminishing its role. Full subgroup recovery and central property testing are different computational tasks, and Graph Isomorphism demonstrates that the multiplicity data retained by the central shadow does not suffice for the full hidden-subgroup problem. For families with distinct central shadows, class-convolution operators may nevertheless be used as discriminating central probes by weighting the irreducible sectors differently. Such spectral processing can increase the separation between distinct shadows but, by Corollary V.2, cannot distinguish subgroups whose central shadows coincide.

The Graph Isomorphism obstruction also indicates where additional algebraic structure must enter. Passing beyond $Z(\mathbb{C}[\mathcal{G}])$ to noncentral operator algebras can retain internal representation and relational information erased by centralization. For permutation actions, orbital and coherent algebras provide nice examples of such extensions, connecting representation-theoretic operators with coherent configurations and Weisfeiler–Leman-type refinement. Developing transforms for these larger algebras is beyond the scope here, but the central case establishes the conjugacy-invariant endpoint of what survives centralization and when noncommutativity may become necessary.

VI. APPLICATIONS: GRAPHS, WALKS, AND SYMMETRIC HAMILTONIANS

The connection between Fourier transforms and graph spectra is already present in the abelian setting. A circulant graph on n vertices is a Cayley graph of the cyclic group $\mathbb{Z}_n$, and its adjacency matrix is circulant and diagonalized by the discrete Fourier transform [58], or on a quantum computer, the quantum Fourier transform. Generally, for a finite abelian group $\mathcal{G}$, translation-invariant graph operators are convolution operators on $\mathbb{C}[\mathcal{G}]$ and are diagonalized by the group Fourier transform; the eigenvalues of an abelian Cayley graph are the character sums over its connection set [2, 58]. For a non-abelian group, ordinary convolution is no longer diagonalized by scalar characters, and the non-abelian Fourier transform decomposes into non-commutative matrix-valued blocks. However, a distinguished commutative sector survives.

If the connection weights are invariant under conjugation, the convolution operator belongs to the center $Z(\mathbb{C}[\mathcal{G}])$. The passage from QFT on circulant graphs to QCT on conjugacy-invariant Cayley graphs extends the familiar Fourier description of cyclic graph spectra from abelian translation symmetry to non-abelian conjugacy symmetry.




FIG. 4. The permutohedral Cayley graph and its conjugacy completion for S_4 . The solid edges are the adjacent transpositions (12), (23), and (34) and form the permutohedral 1-skeleton $\text{Cay}(S_4, \{(12), (23), (34)\})$. The dashed edges are (13), (14), and (24). Together these six transpositions form the conjugacy class $C_{(2,1^2)}$ and the normal Cayley graph $\text{Cay}(S_4, C_{(2,1^2)})$. The permutohedral 1-skeleton is a spanning subgraph of the normal Cayley graph. The vertical embedding is graded by inversion number, or Coxeter length, from the identity 0123 at the bottom to the longest element 3210 at the top; orienting the solid edges upward is the weak Bruhat order on S_4 .

For each $C \in \text{Conj}(\mathcal{G})$, define the relation

$$R_C = \{(g, h) \in \mathcal{G} \times \mathcal{G} : gh^{-1} \in C\}$$

and its adjacency matrix $A_C \in \mathbb{C}^{\mathcal{G} \times \mathcal{G}}$ by

$$(A_C)_{g,h} = \mathbb{1}_{\{gh^{-1} \in C\}}. \quad (14)$$

A_C is the adjacency matrix of the relation graph associated with C . Moreover, $A_C^\top = A_{C^{-1}}$, so the relation graph is undirected exactly when $C = C^{-1}$. After ordering the conjugacy classes as $C_1, \dots, C_d$, let $A_i := A_{C_i}$, and $K_i := K_{C_i}$. The collection $\{A_1, \dots, A_d\}$ forms the *conjugacy-class association scheme* of $\mathcal{G}$, and its span is a commutative Bose–Mesner algebra [1, 12, 34]. If $\rho : \mathcal{G} \rightarrow GL(\mathbb{C}[\mathcal{G}])$ is the left-regular

representation, then $\rho(K_i) = \sum_{g \in C_i} \rho(g)$ is the left-convolution operator associated with the class sum K_i . In the group-element basis $\{|g\rangle\}_{g \in \mathcal{G}}$, this operator has matrix entries

$$(\rho(K_i))_{h,h'} = \mathbb{1}_{\{hh'^{-1} \in C_i\}} = (A_i)_{h,h'}.$$

Thus the adjacency algebra of the conjugacy-class scheme is isomorphic to the center of the group algebra, $\text{span}\{A_i\}_{i=1}^d \cong Z(\mathbb{C}[\mathcal{G}])$. In fact,

$$A_i A_j = \sum_{k=1}^d p_{ij}^k A_k, \quad (15)$$

where p_{ij}^k are the *intersection numbers* of the scheme, counting the number of ways a path of types i, j whose endpoints have relation k . The sums of the classes K_i are in $Z(\mathbb{C}[\mathcal{G}])$, so the matrices A_i can also be diagonalized simultaneously.

The matrices A_i act on the full $|\mathcal{G}|$ -dimensional group-element space $\mathbb{C}[\mathcal{G}]$, and yet preserve $Z(\mathbb{C}[\mathcal{G}])$. Its restriction to the center, expressed in the normalized conjugacy-class basis, is the class-convolution matrix T_C defined in Section III.

$$T_C = \left[A_C|_{Z(\mathbb{C}[\mathcal{G}])} \right]_{\{|C_i\rangle\}}, \quad (T_C)_{ij} = \langle C_i | A_C | C_j \rangle$$

For S_4 , the Cayley-graph adjacency operators A_i are 24×24 , while their central restrictions T_C are the 5×5 class-convolution matrices displayed in (8).




FIG. 5. Bipartite class graph for the transposition class of S_4 . The vertices are the normalized conjugacy-class states $|C_i\rangle$, and edge weights are the nonzero matrix elements of the central restriction $T_{(2,1^2)}$ of $A_{(2,1^2)} = \sum_{\tau \in C_{(2,1^2)}} L_\tau$. Since multiplication by a transposition reverses permutation parity, $T_{(2,1^2)}$ couples only even and odd classes.

A *normal Cayley graph* is obtained by choosing a connection set $\mathcal{S} \subseteq \mathcal{G} \setminus \{e\}$ that is both inverse closed and conjugacy invariant. Equivalently, $\mathcal{S}$ is an inversion-closed union of conjugacy classes,

$$\mathcal{S} = \bigcup_{i \in \mathcal{I}} C_i, \quad C_i \subseteq \mathcal{S} \implies C_i^{-1} \subseteq \mathcal{S}.$$

The graph $\text{Cay}(\mathcal{G}, \mathcal{S})$ has general adjacency operator $A_{\mathcal{S}} = \sum_{i \in \mathcal{I}} A_i$. Assigning a weight w_i to each conju-

gacy class produces the weighted central adjacency

$$A_w = \sum_{i=1}^d w_i A_i. \quad (16)$$

Related work treats general weighted Cayley graphs through irreducible matrix coefficients [7]. A normal Cayley graph specifies one operator in the larger commuting adjacency algebra of the conjugacy-class scheme. The QCT provides a common spectral basis for this entire algebra, simultaneously covering every choice of connection set and class weights.

Association schemes have two canonical bases. The matrices A_i form the adjacency, or relation, basis, while the primitive idempotent projectors P_Γ form their common spectral basis.

Remark (Peirce interpretation). Define the primitive central idempotents of $\mathbb{C}[\mathcal{G}]$

$$P_\Gamma = \frac{d_\Gamma}{|\mathcal{G}|} \sum_{g \in \mathcal{G}} \chi_\Gamma(g) \rho(g). \quad (17)$$

P_Γ projects onto the dual Γ^* -isotypic component, and satisfies $P_\Gamma^\dagger = P_\Gamma$,

$$P_\Gamma P_\mu = \delta_{\Gamma\mu} P_\Gamma, \quad \sum_{\Gamma \in \widehat{\mathcal{G}}} P_\Gamma = I_{|\mathcal{G}|},$$

This is a *central Peirce decomposition*

$$\text{span}\{A_i\}_{i=1}^d = \bigoplus_{\Gamma \in \widehat{\mathcal{G}}} \mathbb{C} P_\Gamma.$$

Writing $A_i = \sum_\Gamma p_i(\Gamma) P_\Gamma$, the coefficients $p_i(\Gamma)$ are the eigenvalues of A_i on the range of P_Γ . By Schur's lemma,

$$p_i(\Gamma) = \frac{|C_i| \chi_\Gamma(C_i^{-1})}{d_\Gamma} = \frac{|C_i| \overline{\chi_\Gamma(C_i)}}{d_\Gamma}. \quad (18)$$

In the regular representation, each projector has rank d_Γ^2 . For an individual operator, distinct sectors may share the same eigenvalue.

The matrix $P_{\text{eig}} = (p_i(\Gamma))_{\Gamma, i}$ is the first eigenmatrix of the conjugacy-class association scheme. The standard eigenmatrix orthogonality relations supply the class-size and multiplicity weights that render the character transform unitary [1, 12]. After normalization, the change between the adjacency and primitive idempotent bases is expressed by the QCT.

Recent work on association schemes also provides useful context for future direction. Translation association schemes make the connection between convolution, group characters, and graph eigenmatrices, while quantum association schemes extend the formalism to noncommutative finite quantum spaces and quantum graphs [35]. These coherent algebras carry both composition and a generalized Schur product. Under composition, each is a finite-dimensional C^* -algebra whose center admits a QCT-type transform: an orthonormal central basis and the primitive central projections, normalized using a faithful trace, are related by a unitary

transform that diagonalizes multiplication by central operators. The center need not be closed under the generalized Schur product and therefore need not itself define a quantum association scheme.

Returning to the conjugacy-class association scheme, its algebra contains more than adjacency matrices. Polynomials and functions of central adjacency operators remain diagonal in the common character basis. Their restrictions to class space can therefore be represented and processed in the smaller d -dimensional basis. These operators include graph Laplacians, diffusion operators, random-walk kernels, and Hamiltonians generating continuous-time quantum walks. The following subsections develop these applications.

A. Laplacians, Diffusion, and Graph Filtering

Let $A_w = \sum_{i \in I} w_i A_i$ be the weighted adjacency operator of a normal Cayley graph of $\mathcal{G}$, where the weights are real and nonnegative and satisfy $w_i = w_{i^*}$ whenever $C_{i^*} = C_i^{-1}$. The graph is weighted regular with degree $\kappa = \sum_{i \in I} w_i |C_i|$. For $\kappa > 0$, its combinatorial and normalized Laplacians are

$$\mathcal{L} = D - A_w = \kappa I - A_w, \quad \mathcal{L}_{\text{norm}} = I - \frac{1}{\kappa} A_w. \quad (19)$$

Writing θ_Γ for the adjacency eigenvalues obtained above, the Laplacian eigenvalues in sector Γ are

$$\mu_\Gamma = \kappa - \theta_\Gamma, \quad \nu_\Gamma = 1 - \frac{\theta_\Gamma}{\kappa},$$

for the combinatorial and normalized Laplacians, respectively. The trivial irrep, for which $\chi_{\text{triv}}(g) = 1$ for every $g \in \mathcal{G}$, gives $\theta_{\text{triv}} = \kappa$ and $\mu_{\text{triv}} = 0$, with eigenvector the constant vector. This zero eigenvalue is unique when the graph is connected.

On the central subspace, let $T_w = A_w|_{Z(\mathbb{C}[\mathcal{G}])}$ and $\mathcal{L}_{\text{cl}} = \kappa I - T_w$. The QCT diagonalizes these compressed operators:

$$U_{\text{QCT}} \mathcal{L}_{\text{cl}} U_{\text{QCT}}^\dagger = \text{diag}(\mu_\Gamma), \\ U_{\text{QCT}} \left(I - \frac{1}{\kappa} T_w \right) U_{\text{QCT}}^\dagger = \text{diag}(\nu_\Gamma).$$

For a general undirected graph $\mathcal{X}$ with Laplacian $\mathcal{L}$, graph Fourier analysis chooses an orthonormal eigenbasis $\{u_j\}$ satisfying $\mathcal{L}u_j = \mu_j u_j$ and defines the graph Fourier coefficients of a signal f by $\hat{f}_j = \langle u_j, f \rangle$. In matrix form, if $\mathcal{L} = U\Gamma U^\dagger$, the graph Fourier transform is $\mathcal{F}_{\mathcal{X}} = U^\dagger$. Unlike the ordinary Fourier transform, U depends on the particular graph and must be obtained from its spectral decomposition [56]. For conjugacy-invariant Cayley graphs, the QCT basis is fixed by the representation theory of $\mathcal{G}$ and does not depend on the class weights.

Functional calculus preserves the same spectral projectors. On the full group-element space,

$$e^{-t\mathcal{L}} = \sum_{\Gamma \in \hat{\mathcal{G}}} e^{-t\mu_\Gamma} P_\Gamma, \quad e^{-it\mathcal{L}} = \sum_{\Gamma \in \hat{\mathcal{G}}} e^{-it\mu_\Gamma} P_\Gamma.$$

The first operator is the heat semigroup, or continuous-time random-walk evolution, while the second is continuous-time quantum-walk evolution with the Laplacian as its Hamiltonian [28]. Restricting these operators to the central subspace,

$$e^{-t\mathcal{L}_{\text{cl}}} = U_{\text{QCT}}^\dagger \text{diag}(e^{-t\mu_\Gamma}) U_{\text{QCT}}, \\ e^{-it\mathcal{L}_{\text{cl}}} = U_{\text{QCT}}^\dagger \text{diag}(e^{-it\mu_\Gamma}) U_{\text{QCT}}.$$

A graph filter $\varphi(\mathcal{L}_{\text{cl}})$, such as the resolvent low-pass filter $(I + \alpha\mathcal{L}_{\text{cl}})^{-1}$ or a band-pass filter, also acts diagonally in the QCT basis,

$$\varphi(\mathcal{L}_{\text{cl}}) = U_{\text{QCT}}^\dagger \text{diag}(\varphi(\mu_\Gamma)) U_{\text{QCT}}.$$

The irreducible-character sectors Γ define the graph-frequency decomposition, ordered by the Laplacian eigenvalues μ_Γ , with larger values indicating higher graph frequencies [38, 56].

B. Random and Quantum Walks

Let $p : \mathcal{G} \rightarrow [0, 1]$ be a conjugacy-invariant probability distribution, and write $p_i = p(s)$ for $s \in C_i$. Its normalization is $\sum_{i=1}^d |C_i| p_i = 1$. Using the class-sum adjacency operators defined above, the transition operator of the discrete-time random walk is

$$P = \sum_{i=1}^d p_i A_i. \quad (20)$$

Viewing p as a class function, define $\theta_\Gamma := \hat{p}(\Gamma)$, where $\hat{p}(\Gamma)$ is given by Eq. (18). Let $P_{\text{cl}} = P|_{\mathbb{C}(\mathcal{G})}$ denote the induced action on the class-function space. The discrete character transform exposes the scalar eigenvalues:

$$U_{\text{QCT}} P_{\text{cl}} U_{\text{QCT}}^\dagger = \text{diag}(\theta_\Gamma)_{\Gamma \in \hat{\mathcal{G}}}. \quad (21)$$

A central random walk on the full group algebra is compressed by the QCT to one spectral coordinate for each irreducible character.

For a symmetric ergodic walk, the trivial representation satisfies $\theta_{\text{triv}} = 1$, while the remaining eigenvalues determine convergence to stationarity. The absolute spectral gap is $\gamma = 1 - \max_{\Gamma \neq \text{triv}} |\theta_\Gamma|$. The mixing behavior of a conjugacy-invariant walk can be expressed in terms of these character ratios.

A standard example is the random-transposition walk on S_n . Let C_{tr} denote the conjugacy class of transpositions, with $|C_{\text{tr}}| = \binom{n}{2}$, and choose each transposition $s \in C_{\text{tr}}$ with probability $p(s) = \binom{n}{2}^{-1}$. Equation (18) becomes

$$\theta_\Gamma = \frac{\chi_\Gamma((12))}{d_\Gamma}.$$

The spectrum of the walk is encoded by the normalized characters of the transposition class, as in the classical representation-theoretic analysis of random walks on finite groups [25]. The non-lazy random-transposition

walk is periodic because every step changes permutation parity; in particular, the sign representation has eigenvalue -1 . A lazy version of the walk is required for convergence to the uniform distribution in total variation.

Repeated evolution remains diagonal in character coordinates, which is the character-theoretic version of diagonalizing a translation-invariant walk by the ordinary Fourier transform. Convolution becomes multiplication, and the t -step evolution is obtained by raising each character-sector eigenvalue to the t th power. The QCT records one sector eigenvalue θ_Γ for each $\Gamma \in \widehat{\mathcal{G}}$, whereas on the regular representation $\mathbb{C}[\mathcal{G}]$, this eigenvalue has multiplicity d_Γ^2 . Therefore,

$$\text{Tr}(\mathbf{P}^t) = \sum_{\Gamma \in \widehat{\mathcal{G}}} d_\Gamma^2 \theta_\Gamma^t. \quad (22)$$

Translation invariance makes the diagonal entries of $\mathbf{P}^t$ equal, so the t -step return probability to any $g \in \mathcal{G}$ is

$$(\mathbf{P}^t)_{g,g} = \frac{1}{|\mathcal{G}|} \sum_{\Gamma \in \widehat{\mathcal{G}}} d_\Gamma^2 \theta_\Gamma^t. \quad (23)$$

The same spectral description applies to continuous-time quantum walks. Let $H = \sum_{i=1}^d w_i A_i$ be a central Hermitian Hamiltonian, with class weights satisfying $w_{i^*} = \overline{w_i}$ whenever $C_{i^*} = C_i^{-1}$. Let $w \in \mathbb{C}(\mathcal{G})$ be the class function defined by $w(g) = w_i$ for $g \in C_i$. The Hamiltonian eigenvalues are $\eta_\Gamma := \widehat{w}(\Gamma)$, where $\widehat{w}(\Gamma)$ is given by Eq. (18). Its compressed action $H_{\text{cl}} = H|_{\mathbb{C}(\mathcal{G})}$ has the evolution

$$e^{-itH_{\text{cl}}} = U_{\text{QCT}}^\dagger \text{diag}(e^{-it\eta_\Gamma})_{\Gamma \in \widehat{\mathcal{G}}} U_{\text{QCT}}. \quad (24)$$

Evolution under a central Hamiltonian reduces in QCT coordinates to independent phase evolution of the irreducible-character sectors [28].

The diagonal form also allows a direct measurement interface for spectral information. For a normalized class-function state $|\psi\rangle = \sum_{i=1}^d a_i |C_i\rangle$,

$$U_{\text{QCT}} |\psi\rangle = \sum_{\Gamma \in \widehat{\mathcal{G}}} \widehat{a}_\Gamma |\Gamma\rangle. \quad (25)$$

Measurement in the character basis samples Γ according to $\text{Pr}(\Gamma | \psi) = |\widehat{a}_\Gamma|^2$. The character-sector weights can be sampled without reconstructing the full transformed state classically, and observables diagonal in the character basis can be estimated from repeated measurements.

For a real-valued function F on the spectrum of H_{cl} ,

$$F(H_{\text{cl}}) = U_{\text{QCT}}^\dagger \left(\sum_{\Gamma \in \widehat{\mathcal{G}}} F(\eta_\Gamma) |\Gamma\rangle \langle \Gamma| \right) U_{\text{QCT}}. \quad (26)$$

Its expectation in the state $|\psi\rangle$ is

$$\langle F(H_{\text{cl}}) \rangle_\psi = \sum_{\Gamma \in \widehat{\mathcal{G}}} |\widehat{a}_\Gamma|^2 F(\eta_\Gamma). \quad (27)$$

$F(x) = x^k$ gives the spectral moments $\langle H_{\text{cl}}^k \rangle_\psi = \sum_{\Gamma \in \widehat{\mathcal{G}}} |\widehat{a}_\Gamma|^2 \eta_\Gamma^k$.

For a subset $S \subseteq \widehat{\mathcal{G}}$, the projector onto the sectors indexed by S is

$$\Pi_S = U_{\text{QCT}}^\dagger \left(\sum_{\Gamma \in S} |\Gamma\rangle \langle \Gamma| \right) U_{\text{QCT}}, \quad (28)$$

with $\langle \psi | \Pi_S | \psi \rangle = \sum_{\Gamma \in S} |\widehat{a}_\Gamma|^2$. This quantity measures the total spectral weight carried by a chosen family of symmetry sectors.

For quantum walks, the first moments determine average spectral quantities, higher moments probe the eigenvalue distribution seen by the initial state, and sector projectors determine its overlap with prescribed symmetry components. The phases $e^{-it\eta_\Gamma}$ control periodicity, interference, time-averaged mixing, and conditions for state transfer. The QCT acts both as a diagonalizing transform for class-invariant walks and as an interface for sampling their symmetry-adapted spectral observables.

C. Class-Symmetric Hamiltonians and Exchange Dynamics

Let $\rho : \mathcal{G} \rightarrow U(\mathcal{H})$ be a unitary representation on a physical Hilbert space $\mathcal{H}$. A *class-symmetric* Hamiltonian is one of the form

$$H = \sum_{g \in \mathcal{G}} J_g \rho(g), \quad (29)$$

where $J_{ghg^{-1}} = J_g$ and $J_{g^{-1}} = \overline{J_g}$. The first condition makes the coupling constant on each conjugacy class, while the second ensures that $H = H^\dagger$. In terms of the class sums K_i , the Hamiltonian is $H = \sum_{i=1}^d J_i \rho(K_i)$, where $J_i = J_g$ for $g \in C_i$.

Each $\rho(K_i)$ commutes with $\rho(\mathcal{G})$. If the representation decomposes as

$$\mathcal{H} \cong \bigoplus_{\Gamma \in \widehat{\mathcal{G}}} \mathbb{C}^{m_\Gamma} \otimes V_\Gamma,$$

then by Schur's lemma

$$\rho(K_i)|_{\mathbb{C}^{m_\Gamma} \otimes V_\Gamma} = I_{m_\Gamma} \otimes \frac{|C_i| \chi_\Gamma(C_i)}{d_\Gamma} I_{V_\Gamma}.$$

The energy associated with the irreducible sector Γ is

$$E_\Gamma = \sum_{i=1}^d J_i \frac{|C_i| \chi_\Gamma(C_i)}{d_\Gamma}. \quad (30)$$

The QCT diagonalizes the class-algebra representation of these operators, while their action on $\mathcal{H}$ follows from the irreducible decomposition of ρ .

A paradigmatic example is the uniform permutation-exchange interaction of n identical spin- $\frac{1}{2}$ particles. Let $\pi_{ij} := \rho((ij))$ denote the operator that swaps spins i and j . The group S_n acts on the n -spin space by permuting tensor factors, and the Hamiltonian

$$H_{\text{ex}} = J \sum_{1 \leq i < j \leq n} \pi_{ij} = J \rho(K_{(2,1^{n-2})}) \quad (31)$$

is the class sum of all transpositions in S_n . Its spectrum on the n -spin Hilbert space is $E_\Gamma = J \binom{n}{2} \frac{\chi_\Gamma((12))}{d_\Gamma}$, where only partitions $\lambda \vdash n$ with at most two rows occur in $(\mathbb{C}^2)^{\otimes n}$ by Schur–Weyl duality [31]. This recovers the energy levels of the permutation-symmetric exchange model without directly diagonalizing its $2^n \times 2^n$ Hamiltonian. The trivial representation has energy $E_{\text{triv}} = J \binom{n}{2}$; whether this is the highest or lowest energy depends on the sign of J . For even n , the total-spin-zero sector is the sector labeled by the partition $(n/2, n/2)$, not the sign representation, except when $n = 2$. On the other hand, a nearest-neighbor chain $H_{\text{nn}} = J \sum_{i=1}^{n-1} \pi_{i,i+1}$ is not a class sum, since the adjacent transpositions form only part of the full transposition class. The set of terms is not invariant under conjugation by S_n , so H_{nn} does not commute with the full permutation action and the character values alone do not diagonalize it.

Returning to S_4 , the most general central class-sum Hamiltonian on the 24-dimensional regular space is

$$H = J_0 I + J_2 \rho(K_{(2,1^2)}) + J_{22} \rho(K_{(2^2)}) + J_3 \rho(K_{(3,1)}) + J_4 \rho(K_{(4)}) \quad (32)$$

where $J_0, J_2, J_{22}, J_3, J_4 \in \mathbb{R}$. Modulo the overall energy shift $J_0 I$, this is a four-parameter family. Restricting to the five-dimensional center and applying the QCT gives the diagonal matrix $\text{diag}(E_\Gamma)$, where

$$E_\Gamma = J_0 + J_2 \frac{6 \chi_\Gamma((12))}{d_\Gamma} + J_{22} \frac{3 \chi_\Gamma((12)(34))}{d_\Gamma} + J_3 \frac{8 \chi_\Gamma((123))}{d_\Gamma} + J_4 \frac{6 \chi_\Gamma((1234))}{d_\Gamma}. \quad (33)$$

For the trivial representation, $\chi_{\text{triv}}(g) = 1$ for all $g \in S_4$, giving $E_{\text{triv}} = J_0 + 6J_2 + 3J_{22} + 8J_3 + 6J_4$. The remaining four sector energies follow from the S_4 character table.

D. Molecular and Permutation–Inversion Dynamics

Class-sum Hamiltonians also appear in molecular physics. Suppose that a finite permutation–inversion group $\mathcal{G}$ acts regularly on a set of equivalent molecular configurations, which may then be labeled by $g \in \mathcal{G}$. If tunneling or rearrangement processes are represented by elements $s \in \mathcal{G}$, an effective Hamiltonian can take the form

$$H = E_0 I + \sum_{s \in \mathcal{G}} t_s \rho(s), \quad (34)$$

where $\rho(s)$ permutes the configurations and $t_{s^{-1}} = \overline{t_s}$. The configuration space need not be the group itself; it can be a homogeneous $\mathcal{G}$ -set.

If the amplitudes are constant on conjugacy classes, so that $t_{hsh^{-1}} = t_s$ for all $s, h \in \mathcal{G}$, then

$$H = E_0 I + \sum_{i=1}^d t_i \rho(K_i)$$

is a class-symmetric Hamiltonian. In the language of permutation–inversion groups used in molecular spectroscopy [15, 59], this condition assigns equal amplitudes to rearrangements related by conjugation. The energy in each irreducible sector is obtained from the character formula of the preceding subsection. Such class-invariant models provide a compact description of tunneling splittings associated with large-amplitude molecular motions.




FIG. 6. Conjugate symmetry operations on ideal tetrahedral methane. The tetrahedral symmetry group $T_d \cong S_4$ acts by permuting the four hydrogen vertices. Taking $H_1 = (1, 1, 1)$, $H_2 = (1, -1, -1)$, $H_3 = (-1, 1, -1)$, and $H_4 = (-1, -1, 1)$, the rotation $s = C_2(z)$ induces $(H_1 H_4)(H_2 H_3) \in C_{(2^2)}$, while $h = C_3([111])$ induces $(H_2 H_3 H_4) \in C_{(3,1)}$. Conjugation gives $hsh^{-1} = C_2(x)$, inducing $(H_1 H_2)(H_3 H_4) \in C_{(2^2)}$. The diagram commutes, with $h(s \cdot x) = (hsh^{-1}) \cdot (h \cdot x) = h \cdot s \cdot x$.

Consider six equivalent configurations arranged in a cycle and let r generate the cyclic shift between them. If hopping in either direction has amplitude t , the Hermitian tunneling Hamiltonian is

$$H_{\text{hop}} = E_0 I + t(\rho(r) + \rho(r^{-1})).$$

For the cyclic group, the QCT reduces to the ordinary finite Fourier transform and yields the energies

$$E_k = E_0 + 2t \cos\left(\frac{2\pi k}{6}\right), \quad k = 0, \dots, 5.$$

If the model is instead formulated using a dihedral permutation–inversion group, the pair $\{r, r^{-1}\}$ belongs to one conjugacy class, and the same Hamiltonian is a class sum.

E. Symmetry-Aware Learning

Symmetry-aware learning architectures seek to encode known invariances directly into the representation and admissible learned operators without relying solely on their recovery from augmented data [11, 22]. For data represented by class functions $f \in C(\mathcal{G})$, conjugation invariance reduces the ambient function

space on $\mathcal{G}$ to the d -dimensional class space, where $d = |\text{Conj}(\mathcal{G})| = |\hat{\mathcal{G}}|$. The QCT coordinates $\hat{a}_\Gamma$ in Eq. (25) give a fixed symmetry-adapted representation that can precede either a classical or quantum learning model.

The QCT also provides a spectral parameterization of the entire algebra of class-convolution layers. A trainable central convolutional layer can be parameterized by its response on the irreducible-character sectors. Such a layer has d spectral parameters, compared with d^2 entries for an unrestricted linear operator on the class space. When a class-convolution operator serves as a graph shift, repeated linear propagation multiplies each character-sector amplitude by successive powers of its spectral multiplier. This describes the linear propagation step in message-passing architectures [24, 45].

The same representation permits symmetry-adapted spectral filtering and dimensionality reduction. Writing f in normalized class coordinates, let $\Gamma \subseteq \hat{\mathcal{G}}$ and define the diagonal selector $D_\Gamma = \sum_{\Gamma \in \Gamma} |\Gamma\rangle\langle\Gamma|$. The map $f \mapsto D_\Gamma U_{\text{QCT}} f$ retains only selected irreducible-character sectors, giving a $|\Gamma|$ -dimensional representation after the discarded coordinates are removed. When the class-convolution operators describe adjacency or propagation, this restriction can be viewed as symmetry-adapted spectral pooling, where information is retained or discarded at the level of representation sectors, not individual group elements. Multiple learned class-convolution filters can be combined with sector selection, nonlinear processing, or measurement for classification and discrimination tasks.

F. Conjugacy Orbits and Geometric Invariance

A geometric interpretation is found when $\mathcal{G}$ is represented in $\text{Pin}(V, Q)$ or $\text{Spin}(V, Q)$, with Q nondegenerate. Group conjugation is the Clifford sandwich action,

$$R \mapsto URU^{-1}, \quad U, R \in \mathcal{G} \subset \mathcal{Cl}(V, Q)^\times, \quad (35)$$

so conjugacy classes are orbits of Clifford elements under this action. For a rotor $R = \exp(-B/2)$ with bivector $B \in \bigwedge^2 V$, conjugation is

$$URU^{-1} = \exp\left(-\frac{UBU^{-1}}{2}\right).$$

Conjugation transports the generating bivector by the group action [26]. The stabilizer of R under conjugation is its centralizer, $Z_{\mathcal{G}}(R) = \{U \in \mathcal{G} : UR = RU\}$, so its conjugacy orbit has the homogeneous-space description

$$\mathcal{O}_R \simeq \mathcal{G}/Z_{\mathcal{G}}(R).$$

When $\mathcal{G}$ is a Lie subgroup of $\text{Spin}(V, Q)$, its tangent space at R is $T_R \mathcal{O}_R = \{[X, R] : X \in \mathfrak{g}\}$, while the elements satisfying $[X, R] = 0$ form the Lie algebra of the centralizer [47]. For finite $\mathcal{G}$, these orbits are the discrete conjugacy classes C_R . Passing from individual

group elements to conjugacy-class coordinates discards position within an orbit while retaining its class label. The discrete QCT converts these class coordinates into irreducible-character sectors,

$$|C_R\rangle \xrightarrow{U_{\text{QCT}}} \sum_{\Gamma \in \hat{\mathcal{G}}} \hat{a}_R(\Gamma) |\Gamma\rangle. \quad (36)$$

This orbit-space formulation also motivates a continuous extension of the character transform, replacing the discrete conjugacy-class basis by functions on the conjugacy-orbit space of a compact group and the finite character table by the continuous character transform.

VII. CONCLUSION

The quantum character transform is the discrete Fourier transform over the center $Z(\mathbb{C}[\mathcal{G}])$ of a group algebra, which confirms its operator-theoretic setting. Conjugacy-class states and irreducible characters form dual bases of the same commutative algebra, and the QCT simultaneously diagonalizes every class-convolution operator. Central operators share a common spectral calculus in which convolution, Hamiltonian evolution, spectral filtering, and other matrix functions are diagonal operations on irreducible-character labels.

I further corroborate that the complexity of the QCT is closely tied to that of the non-abelian group Fourier transform, through its canonical restriction to the center. For finite groups with an efficient QFT, this gives the asymptotic QCT bound

$$\mathcal{O}(\text{cost}(\text{QFT}_{\mathcal{G}}) + \text{poly}(\log |\mathcal{G}|, \log(1/\epsilon))).$$

Whenever the restriction overhead is asymptotically bounded by the QFT cost, this simplifies to $\text{cost}(\text{QCT}_{\mathcal{G}}) = \mathcal{O}(\text{cost}(\text{QFT}_{\mathcal{G}}))$. The transform applies whenever the relevant operator is conjugacy invariant. Normal Cayley graphs and their quantum walks, class-symmetric many-body Hamiltonians, exchange and molecular permutation dynamics, and symmetry-sector filters are all instances of the same central spectral problem. Still, the scope of the character transform should be stated precisely. Because class-convolution operators are central, they are block-scalar on irreducible sectors but do not retain the internal matrix-index information present in the full non-abelian Fourier transform. For this reason, the QCT should not be oversold as a universal replacement for non-abelian Fourier sampling in hidden subgroup problems. Over groups such as S_n , there are strong limitations on what Fourier-sampling approaches can learn from single coset states, and those limitations should temper any claim that central class-operator measurements alone can overcome the known barriers surrounding graph-isomorphism-type instances. The correct message is narrower and stronger; central operators provide the right tool for conjugacy-invariant spectral questions, while non-central operators remain

necessary for finer, less understood measurements beyond the class algebra.

The restriction method leaves a central complexity question unresolved. Most importantly, it establishes an upper bound inherited from the quantum Fourier transform but does not settle the optimal complexity of generating character tables on a quantum computer. Subsequent work demonstrates that QCT algorithms can improve upon this inherited bound without passing through the full group Fourier transform [52]. For groups such as S_n , the gap between the group and class register sizes means that character transform algorithms specialized to the group and operating entirely in the compressed class space could use

substantially lower resources than restriction-based algorithms. Discovering if more sub-Fourier implementations exist across different groups, how efficiently their character data can be computed, and whether these reductions translate into end-to-end advantages for symmetry-adapted simulation is the foreseeable future for the character transform.

Acknowledgements. This research was supported by the TRIO Ronald E. McNair Postbaccalaureate Achievement Program at North Carolina State University in summer 2025, which was funded by the U.S. Department of Education under grant number P217A220189-24.

-
- [1] Arnbjörg Soffia Árnadóttir and Chris Godsil. A note on eigenvalues of Cayley graphs. *Linear Algebra and its Applications*, 721:240–249, 2025.
 - [2] László Babai. Spectra of cayley graphs. *Journal of Combinatorial Theory, Series B*, 27(2):180–189, 1979.
 - [3] Dave Bacon. How a clebsch–gordan transform helps to solve the heisenberg hidden subgroup problem. *Quantum Information and Computation*, 8(5):438–467, 2008.
 - [4] Dave Bacon, Isaac L. Chuang, and Aram W. Harrow. Efficient quantum circuits for schur and clebsch-gordan transforms. *Physical Review Letters*, 97(17):170502, 2006.
 - [5] Victor Bastidas, Nathan Fitzpatrick, K. Joven, Zane Rossi, Shariful Islam, Troy Voorhis, Isaac Chuang, and Yuan Liu. Unification of finite symmetries in the simulation of many-body systems on quantum computers. *Physical Review A*, 111, 05 2025.
 - [6] Robert Beals. Quantum computation of fourier transforms over symmetric groups. In *Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, STOC ’97, pages 48–53. Association for Computing Machinery, 1997.
 - [7] Kathryn Beck, Mahya Ghandehari, Skyler Hudson, and Jenna Paltenstein. Frames for signal processing on cayley graphs. *Journal of Fourier Analysis and Applications*, 30(6):66, 2024.
 - [8] Juan Bermejo-Vega and Kevin C. Zatloukal. Abelian hypergroups and quantum computation. arXiv:1509.05806 [quant-ph], 2015.
 - [9] Joseph L. Birman. *Theory of Crystal Space Groups and Infra-Red and Raman Lattice Processes of Insulating Crystals*. Springer-Verlag, 1974.
 - [10] Sergey Bravyi, David Gosset, Vojtech Havlicek, and Louis Schatzki. Classical and quantum algorithms for characters of the symmetric group. *PRX Quantum*, 6:030323, Aug 2025.
 - [11] Michael M. Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
 - [12] Andries E. Brouwer, Arjeh M. Cohen, and Arnold Neumaier. *Distance-Regular Graphs*, volume 18 of *Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge*. Springer-Verlag, Berlin, Heidelberg, 1989.
 - [13] Carli Bruinsma, Dmitry Grinko, and Maris Ozols. Quantum fourier transform for the symmetric group. *arXiv preprint arXiv:2608.28569*, 2026.
 - [14] Philip R. Bunker. *Molecular Symmetry and Spectroscopy*. Academic Press, New York, 1979.
 - [15] Philip R. Bunker and Per Jensen. *Molecular Symmetry and Spectroscopy*. NRC Research Press, 2nd edition, 1998.
 - [16] William Burnside. *Theory of Groups of Finite Order*. Cambridge University Press, Cambridge, 1st edition, 1897.
 - [17] William Burnside. *Theory of Groups of Finite Order*. Cambridge University Press, Cambridge, 2nd edition, 1911.
 - [18] Yudong Cao, Jonathan Romero, Jonathan P. Olson, Matthias Degroote, Peter D. Johnson, Mária Kieferová, Ian D. Kivlichan, Tim Menke, Borja Peropadre, Nicolas P. D. Sawaya, Sukin Sim, Libor Veis, and Alán Aspuru-Guzik. Quantum chemistry in the age of quantum computing. *Chemical Reviews*, 119(19):10856–10915, 2019.
 - [19] Grecia Castelazo, Quynh T. Nguyen, Giacomo De Palma, Dirk Englund, Seth Lloyd, and Bobak T. Kiani. Quantum algorithms for group convolution, cross-correlation, and equivariant transformations. *Physical Review A*, 106(3):032402, September 2022.
 - [20] Man-Duen Choi. Completely positive linear maps on complex matrices. *Linear Algebra and its Applications*, 10(3):285–290, 1975.
 - [21] Michael Clausen. Fast Fourier transforms for metabelian groups. *SIAM Journal on Computing*, 18(3):584–593, 1989.
 - [22] Taco S. Cohen and Max Welling. Group equivariant convolutional networks. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2990–2999, 2016.
 - [23] F. Albert Cotton. *Chemical Applications of Group Theory*. Wiley-Interscience, New York, 3rd edition, 1990.
 - [24] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*, volume 29, 2016.
 - [25] Persi Diaconis. *Group Representations in Probability and Statistics*. Institute of Mathematical Statistics, 1988.
 - [26] Chris Doran and Anthony Lasenby. *Geometric Algebra for Physicists*. Cambridge University Press, 2003.

- [27] Mark Ettinger, Peter Høyer, and Emanuel Knill. The quantum query complexity of the hidden subgroup problem is polynomial. *Information Processing Letters*, 91(1):43–48, 2004.
- [28] Edward Farhi and Sam Gutmann. Quantum computation and decision trees. *Physical Review A*, 58(2):915–928, 1998.
- [29] Ferdinand Georg Frobenius. Über die gruppencharaktere. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin*, pages 985–1021, 1896.
- [30] Ferdinand Georg Frobenius. Über die darstellung der endlichen gruppen durch lineare substitutionen. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin*, pages 994–1015, 1897.
- [31] William Fulton and Joe Harris. *Representation Theory: A First Course*, volume 129 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1991.
- [32] Fritz Gassmann. Bemerkungen zur vorstehenden arbeit von hurwitz: Über beziehungen zwischen den primidealen eines algebraischen körpers und den substitutionen seiner gruppe. *Mathematische Zeitschrift*, 25:665–675, 1926.
- [33] András Gilyén, Yuan Su, Guang Hao Low, and Nathan Wiebe. Quantum singular value transformation and beyond: Exponential improvements for quantum matrix arithmetics. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 193–204, 2019.
- [34] Christopher D. Godsil. *Algebraic Combinatorics*. Chapman & Hall/CRC Mathematics Series. Chapman & Hall, New York, 1993.
- [35] Daniel Gromada. Quantum association schemes, 2024.
- [36] Sean Hallgren, Cristopher Moore, Martin Rötteler, Alexander Russell, and Pranab Sen. Limitations of quantum coset states for graph isomorphism. *Journal of the ACM*, 57(6), 2010.
- [37] Morton Hamermesh. *Group Theory and Its Application to Physical Problems*. Addison-Wesley, 1962.
- [38] David K. Hammond, Pierre Vandergheynst, and Rémi Gribonval. Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis*, 30(2):129–150, 2011.
- [39] P. F. Harrison and W. G. Scott. Permutation symmetry, tri-bimaximal neutrino mixing and the s_3 group characters. *Physics Letters B*, 557(1–2):76–86, 2003.
- [40] Peter Hoyer. Efficient quantum transforms, 1997.
- [41] Peter Høyer. The hidden subgroup problem and quantum computation using group representations. *SIAM Journal on Computing*, 31(6):1775–1787, 2002.
- [42] Andrzej Jamiolkowski. Linear transformations which preserve trace and positive semidefiniteness of operators. *Reports on Mathematical Physics*, 3(4):275–278, 1972.
- [43] Jacob Katriel. Class-sum products in the symmetric group: Combinatorial interpretation of the reduced class coefficients. *International Journal of Quantum Chemistry*, 68(2):103–118, 1998.
- [44] Yasuhito Kawano and Hiroshi Sekigawa. Quantum fourier transform over symmetric groups—improved result. *Journal of Symbolic Computation*, 75:219–243, 2016.
- [45] Thomas N. Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.
- [46] Greg Kuperberg. A subexponential-time quantum algorithm for the dihedral hidden subgroup problem. *SIAM Journal on Computing*, 35(1):170–188, 2005.
- [47] John M. Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, 2 edition, 2012.
- [48] Sam McArdle, Suguru Endo, Alán Aspuru-Guzik, Simon C. Benjamin, and Xiao Yuan. Quantum computational chemistry. *Reviews of Modern Physics*, 92(1):015003, 2020.
- [49] Cristopher Moore, Alexander Russell, and Leonard J. Schulman. The symmetric group defies strong fourier sampling. *SIAM Journal on Computing*, 37(6):1842–1864, 2008.
- [50] Cristopher Moore, Alexander Russell, and Piotr Śniady. On the impossibility of a quantum sieve algorithm for graph isomorphism. *SIAM Journal on Computing*, 39(6):2377–2396, 2010.
- [51] Robert Perlis. On the equation $\zeta_k(s) = \zeta_{k'}(s)$. *Journal of Number Theory*, 9(3):342–360, 1977.
- [52] Dikshant Rathore and Leo Zhou. An efficient hamiltonian-based quantum algorithm for characters of the symmetric group, 2026.
- [53] Issai Schur. Über die darstellung der endlichen gruppen durch abgeleitete lineare substitutionen. *Journal für die reine und angewandte Mathematik*, 127:20–50, 1904.
- [54] Issai Schur. Neue begründung der theorie der gruppencharaktere. *Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften zu Berlin*, pages 406–432, 1905.
- [55] Jean-Pierre Serre. *Linear Representations of Finite Groups*, volume 42 of *Graduate Texts in Mathematics*. Springer, 1977.
- [56] David I. Shuman, Sunil K. Narang, Pascal Frossard, Antonio Ortega, and Pierre Vandergheynst. The emerging field of signal processing on graphs: Extending high-dimensional data analysis to networks and other irregular domains. *IEEE Signal Processing Magazine*, 30(3):83–98, 2013.
- [57] Toshikazu Sunada. Riemannian coverings and isospectral manifolds. *Annals of Mathematics*, 121(1):169–186, 1985.
- [58] Audrey Terras. *Fourier Analysis on Finite Groups and Applications*, volume 43 of *London Mathematical Society Student Texts*. Cambridge University Press, 1999.
- [59] Michael Tinkham. *Group Theory and Quantum Mechanics*. Dover, 2003.