



Quantum isomorphism is equivalent to equality of homomorphism counts from planar graphs

Mancinska, Laura; Roberson, David E.

Published in:

Proceedings of IEEE 61st Annual Symposium on Foundations of Computer Science

Link to article, DOI:

[10.1109/FOCS46700.2020.00067](https://doi.org/10.1109/FOCS46700.2020.00067)

Publication date:

2020

Document Version

Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):

Mancinska, L., & Roberson, D. E. (2020). Quantum isomorphism is equivalent to equality of homomorphism counts from planar graphs. In *Proceedings of IEEE 61st Annual Symposium on Foundations of Computer Science* (pp. 661-672). [9317964] IEEE Computer Society Press. Proceedings - Annual IEEE Symposium on Foundations of Computer Science, FOCS Vol. 2020-November <https://doi.org/10.1109/FOCS46700.2020.00067>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Quantum isomorphism is equivalent to equality of homomorphism counts from planar graphs

Laura Mančinska

*QMATH, Department of Mathematical Sciences
University of Copenhagen
Copenhagen, Denmark
mancinska@math.ku.dk*

David E. Roberson

*Department of Applied Mathematics and Computer Science
Technical University of Denmark
Lyngby, Denmark
dero@dtu.dk*

Abstract—Over 50 years ago, Lovász proved that two graphs are isomorphic if and only if they admit the same number of homomorphisms from any graph. Other equivalence relations on graphs, such as cospectrality or fractional isomorphism, can be characterized by equality of homomorphism counts from an appropriately chosen class of graphs. Dvořák [J. Graph Theory 2010] showed that taking this class to be the graphs of treewidth at most k yields a tractable relaxation of graph isomorphism known as k -dimensional Weisfeiler-Leman equivalence. Together with a famous result of Cai, Fürer, and Immerman [FOCS 1989], this shows that homomorphism counts from graphs of bounded treewidth do not determine a graph up to isomorphism. Dell, Grohe, and Rattan [ICALP 2018] raised the questions of whether homomorphism counts from planar graphs determine a graph up to isomorphism, and what is the complexity of the resulting relation. We answer the former in the negative by showing that the resulting relation is equivalent to the so-called quantum isomorphism [Mančinska et al, ICALP 2017]. Using this equivalence, we further resolve the latter question, showing that testing whether two graphs have the same number of homomorphisms from any planar graph is, surprisingly, an undecidable problem, and moreover is complete for the class **coRE** (the complement of recursively enumerable problems). Quantum isomorphism is defined in terms of a one-round, two-prover interactive proof system in which quantum provers, who are allowed to share entanglement, attempt to convince the verifier that the graphs are isomorphic. Our combinatorial proof leverages the quantum automorphism group of a graph, a notion from noncommutative mathematics.

Keywords—graphs; quantum information; quantum groups; homomorphism counting.

I. INTRODUCTION

This paper is about a surprising connection between counting complexity and the expressive power of quantum (commuting) multi-prover interactive proof systems.

A. Multi-prover interactive proof systems with entangled provers

Interactive proof systems are a central topic in complexity theory. Shortly after their introduction, Babai, Fortnow, and Lund characterized the power of multi-prover interactive proof systems by showing that $\text{MIP} = \text{NEXP}$ [4] and paving the way for the celebrated PCP theorem [2], [1].

In fact, already $\text{MIP}(2,1) = \text{NEXP}$; so only two-provers and a single round of interaction suffice to reach NEXP [15]. Such single-round proof systems can be viewed as families of (nonlocal) games \mathcal{G}_x , where the provers attempt to convince the verifier that the word x belongs to a language L . For a given game \mathcal{G} , we are interested to find the largest probability, $\omega(\mathcal{G})$, with which provers can make the verifier accept. This is called the value of \mathcal{G} .

Other than the requirement of no-communication, the provers are allowed unlimited computation time and are often described as all-powerful. From this perspective, it is natural to ask what happens if they are allowed to make use of quantum-mechanical strategies and most notably shared entanglement. The investigation of entanglement from the perspective of interactive proof systems was first initiated in [8]. Among other results, the authors define the complexity class MIP^* , the quantum value of a nonlocal game, $\omega_q(\mathcal{G})$, and point out the connection between nonlocal games and Bell inequalities [6] from the foundations of quantum physics.

There are two different models for quantum strategies available to quantum provers. Let us refer to them as the *tensor-product* model and the *commuting* model. Any tensor-product model quantum strategy is also a valid strategy in the commuting model but not the other way around. The definition of MIP^* and the quantum value, $\omega_q(\mathcal{G})$, from [8] follow the tensor-product model. The analogous notions for the commuting model are the complexity class MIP^{co} and the quantum commuting value, $\omega_{qc}(\mathcal{G})$ of a game. For any game \mathcal{G} , we have $\omega(\mathcal{G}) \leq \omega_q(\mathcal{G}) \leq \omega_{qc}(\mathcal{G})$ simply because any classical strategy is a tensor-product quantum strategy and any tensor-product quantum strategy is a commuting quantum strategy. It is worth noting that the same argument does not apply to complexity classes, since giving more power to the provers can change the soundness of proof systems and thus potentially decrease the expressive power of the corresponding complexity class.

In comparison to MIP , its quantum counterparts, MIP^* and MIP^{co} , have proved much harder to understand. Over the years, MIP^* was shown to contain increasingly larger

complexity classes ranging from NEXP [16] to NEEEXP [25]. A breakthrough result of Slofstra showed that determining if $\omega_{qc}(\mathcal{G}) = 1$ is an undecidable problem [29] and that the same holds for the quantum value ω_q [28]. In combination with the lack of upper bounds on MIP*, these undecidability results could make one speculate if MIP* could also contain undecidable problems. To the surprise of multiple communities, including complexity theory, operator algebras, and quantum computing, just earlier this year this question was settled in the affirmative [17]. To reach the groundbreaking MIP* = RE result, [17] reduces the halting problem to the problem of deciding if a nonlocal game has quantum value 1 or at most $\frac{1}{2}$.

The complexity class that corresponds to testing if $\omega_{qc}(\mathcal{G}) = 1$ is obtained by considering multi-prover interactive proof systems with the completeness and soundness parameters equal to one. The resulting complexity class, MIP_0^{co} , is sometimes referred to as the zero gap variant of MIP^{co} . We can similarly define the zero gap variant, MIP_0^* , for the class MIP*. Just recently MIP_0^* was shown to equal the class Π_2^0 from the second level of the arithmetical hierarchy [24]. Perhaps surprisingly, the more powerful commuting model yields a less powerful class. Indeed, by combining existing results, we can get that $\text{MIP}_0^{co} = \text{coRE}$ where coRE is the complement of the class of recursively enumerable languages, RE. To see the inclusion $\text{MIP}_0^{co} \subseteq \text{coRE}$, we can use the semidefinite programming hierarchies from [26], [13] which give converging upper-bounds on the commuting value, ω_{qc} . The equality follows from the previously mentioned work of Slofstra [29] and [11]. He shows that the complement of the word problem, which is complete for coRE , can be reduced to deciding if $\omega_{qc}(\mathcal{G}) = 1$ where \mathcal{G} is a linear binary constraint system (LCS) game [10], [9]. This reduction also demonstrates that testing if $\omega_{qc}(\mathcal{G}) = 1$ for the restricted class of LCS games is a complete problem for MIP_0^{co} .

Quantum isomorphism: In our previous work [3] we introduce the graph isomorphism game, where two provers attempt to convince the verifier that they know an isomorphism between graphs G and H :

The (G, H) -isomorphism game [3], [21]

Each player (prover) is given a vertex of either G or H , and must respond with a vertex of the *other* graph. Thus Alice receives or sends a vertex of each graph, which we denote by g_A and h_A . We define Bob's vertices g_B and h_B similarly. The players win (*i.e.* the verifier accepts) if $\text{rel}(g_A, g_B) = \text{rel}(h_A, h_B)$, where rel is a function indicating whether two vertices are equal, adjacent, or distinct and non-adjacent.

It is not difficult to see that the (G, H) -isomorphism game has value one if and only if the two graphs are isomorphic, denoted $G \cong H$. Motivated by this correspondence, [3] defines graphs G and H to be *quantum isomorphic*, de-

noted $G \cong_{qc} H$, if the commuting value of the (G, H) -isomorphism game is equal to one. In other words, $G \cong_{qc} H$ if and only if there is a quantum commuting strategy that wins the (G, H) -isomorphism game with probability one¹.

Given a linear constraint system (LCS) game \mathcal{F} (see [9]), [3] constructs graphs $G(\mathcal{F})$ and $G_0(\mathcal{F})$ such that $G(\mathcal{F}) \cong_{qc} G_0(\mathcal{F})$ if and only if $\omega_{qc}(\mathcal{F}) = 1$. Hence, the problem of testing if $\omega_{qc}(\mathcal{G}) = 1$ for an LCS game \mathcal{G} can be reduced to quantum isomorphism. It follows that the problem of testing if graphs G and H are quantum isomorphic is complete for the class MIP_0^{co} . The main result of the current work is that quantum isomorphism is equivalent to a counting problem:

Main Theorem. *Let G and H be graphs. Then $G \cong_{qc} H$ if and only if G and H have the same number of homomorphisms from every planar graph.*

If graphs G and H are quantum isomorphic then there is a natural, albeit not necessarily finite, certificate for it, namely, the (potentially infinite dimensional) quantum commuting strategy that wins the (G, H) -isomorphism game with probability one. In contrast, it is not *a priori* clear how to certify that G and H are not quantum isomorphic. Yet with the above theorem in hand, we can use a planar graph K which admits a different number of homomorphisms to G as opposed to H to certify that $G \not\cong_{qc} H$.

B. Graph isomorphism and homomorphism counting

Over 50 years ago, Lovász proved that graphs G and H are isomorphic if and only if they have the same number of homomorphisms from any graph K [18]. Here, a homomorphism from a graph K to G is an adjacency-preserving map that takes vertices of K to those of G . According to Lovász' result, we can specify a graph G by its *homomorphism vector* $\text{HOM}(G) := (\text{hom}(K, G))_{K \in \text{Graphs}}$, where $\text{hom}(K, G)$ is the number of homomorphisms from K to G . Even though in general the entries of the homomorphism vectors are NP-hard to compute, many efficiently computable relations on graphs can be expressed by restricting the homomorphism vector to entries which correspond to a specific family of graphs \mathcal{F} . We refer to such restricted homomorphism vectors by $\text{HOM}_{\mathcal{F}}(G)$. Trivial examples include counting homomorphisms from just the single vertex graph or the two vertex graph with a single edge, which simply test whether G and H have the same number of vertices or edges respectively. Less trivially, counting homomorphisms from all star graphs or from all cycles determines a graph's degree sequence or spectrum respectively, the latter being a classical result of algebraic graph theory. Very recently, a surprising result of this form was proven by Dvořák [14] and recently rediscovered by Dell, Grohe, and Rattan [12]: graphs G and

¹Even though it might be more accurate to refer to the relation \cong_{qc} as "quantum commuting isomorphism", for brevity we use the term *quantum isomorphism*.

H are not distinguished by the k -dimensional Weisfeiler-Leman algorithm² if and only if they admit the same number of homomorphisms from all graphs with treewidth at most k . Since for fixed k , the k -dimensional Weisfeiler-Leman algorithm runs in polynomial-time, we can efficiently test if $\text{HOM}_{\mathcal{T}_k}(G) = \text{HOM}_{\mathcal{T}_k}(H)$, where \mathcal{T}_k is the class of graphs of treewidth at most k . In the same work Dell, Grohe, and Rattan ask whether homomorphism counts from planar graphs determine a graph up to isomorphism and what is the complexity of testing if $\text{HOM}_{\mathcal{P}}(G) = \text{HOM}_{\mathcal{P}}(H)$, where \mathcal{P} is the class of planar graphs. Our main theorem answers both of these questions. Firstly, since there are examples of pairs of non-isomorphic graphs that are quantum isomorphic, we get that homomorphism counts from planar graphs *do not* determine graph up to isomorphism. Secondly, a surprising complexity-theoretic consequence of our main theorem is that testing whether $\text{HOM}_{\mathcal{P}}(G) = \text{HOM}_{\mathcal{P}}(H)$ is undecidable and even complete for coRE:

Corollary. Given graphs G and H , the problem of testing if there is a planar graph K such that $\text{hom}(K, G) \neq \text{hom}(K, H)$ is complete for the complexity class RE.

The above also implies that there is no computable function of two graphs G and H which gives an upper bound on the size of planar graphs that must be checked to determine whether G and H are quantum isomorphic. In the classical case, it always suffices to count homomorphisms from graphs on $|V(G)|$ or fewer vertices.

II. PRELIMINARIES

Here we will give a brief introduction to the main notions used in this work, namely bi-labeled graphs and homomorphism matrices, quantum automorphism groups of graphs, and quantum isomorphisms.

A. Bi-labeled graphs and homomorphism matrices

A *bi-labeled graph* \mathbf{K} is a triple $(K, \mathbf{a}, \mathbf{b})$ where K is a graph and $\mathbf{a} = (a_1, \dots, a_\ell) \in V(K)^\ell$, $\mathbf{b} = (b_1, \dots, b_k) \in V(K)^k$ are tuples/vectors of vertices of K , where ℓ, k are nonnegative integers. Note that a vertex can appear multiple times in \mathbf{a} and/or \mathbf{b} , or not appear in them at all. We refer to \mathbf{a} and \mathbf{b} as the *output* and *input* tuples/vectors respectively, and to their entries as the output/input vertices. We use \emptyset to denote the empty tuple. The term “bi-labeled graph” comes from the work of Lovász on graph limits [19], where he defined them in an equivalent but slightly different manner. However, according to [23] essentially the same notion was introduced in unpublished notes of Neumaier in 1989.

²The k -dimensional Weisfeiler-Leman algorithm is a well studied graph isomorphism heuristic based on partitioning k -tuples of vertices. In the case of $k = 1$, it is known to be equivalent to another well studied graph relation: fractional isomorphism, which can be viewed as a linear relaxation of isomorphism.

Drawing a bi-labeled graph: Graphs are ubiquitously depicted and thought of diagrammatically, which of course helps to provide intuition not supplied by formal definitions. The same approach aids in understanding bi-labeled graphs. We can depict the graph part of a bi-labeled graph in the usual way, as points in the plane connected by curves representing edges, but we must add something to this to represent the additional data of input/output vertices. We do this by including input/output “wires” that are attached to the input and output vertices of the bi-labeled graph (see Figure 1). Specifically, to draw a bi-labeled graph $\mathbf{K} = (K, \mathbf{a}, \mathbf{b})$ we draw the underlying graph K , and we attach the i^{th} output wire to a_i and the j^{th} input wire to b_j . The input and output wires extend to the far right and far left of the picture respectively. We indicate which input/output wire is which by drawing them so that they occur in numerical order (first at the top) at the edges of the picture. The wires differ from the edges in that they only have a vertex at one end. We distinguish them by drawing the wires thinner. These drawings of bi-labeled graphs may be somewhat reminiscent of pictures of circuits (though with inputs on the right instead of left), and in fact some of the operations we will perform on bi-labeled graphs are analogous to combining circuits in series or parallel.

Homomorphism matrices: Given a graph G and a bi-labeled graph $\mathbf{K} = (K, \mathbf{a}, \mathbf{b})$ with ℓ outputs and k inputs, the G -homomorphism matrix of \mathbf{K} , denoted $T^{\mathbf{K} \rightarrow G}$, is the $V(G)^\ell \times V(G)^k$ matrix defined entrywise as

$$\left(T^{\mathbf{K} \rightarrow G}\right)_{\mathbf{u}, \mathbf{v}} = |\{\varphi \in \text{Hom}(K, G) : \varphi(\mathbf{a}) = \mathbf{u}, \varphi(\mathbf{b}) = \mathbf{v}\}|$$

where $\text{Hom}(K, G)$ is the set of all homomorphisms from K to G and $\varphi(\mathbf{a}) = \mathbf{u}$ indicates that $\varphi(a_i) = u_i$ for all i . In other words, the entries of the matrix $T^{\mathbf{K} \rightarrow G}$ count the number of homomorphisms from K to G , but partitioned according to the images of the input/output vertices. In particular, this means that the sum of the entries of $T^{\mathbf{K} \rightarrow G}$ is equal to the total number of homomorphisms from K to G , denoted $\text{hom}(K, G)$.

As a simple example, consider the bi-labeled graph $\mathbf{M}^{\ell, k} := (K, (a^\ell), (a^k))$, where K is the graph with a single vertex a , and (a^ℓ) indicates a tuple of length ℓ with every entry equal to a . It is straightforward to see that, for $\ell + k > 0$, the following holds:

$$\left(T^{\mathbf{M}^{\ell, k} \rightarrow G}\right)_{\mathbf{u}, \mathbf{v}} = \begin{cases} 1 & \text{if } u_1 = \dots = u_\ell = v_1 = \dots = v_k \\ 0 & \text{o.w.} \end{cases}$$

The bi-labeled graphs $\mathbf{M}^{1,0}$ and $\mathbf{M}^{1,2}$ (see Figure 2) will be of particular importance.

Another important example is the bi-labeled graph $\mathbf{A} := (K, (a), (b))$ (shown in Figure 2) where K is the complete graph on vertex set $\{a, b\}$. Then $T^{\mathbf{A} \rightarrow G}$ is a $V(G) \times V(G)$ matrix whose uv -entry is 1 if $u \sim v$ (we use \sim to denote adjacency), and is 0 otherwise. In other words, $T^{\mathbf{A} \rightarrow G}$ is the *adjacency matrix* of the graph G , denoted A_G .



Figure 1. How to draw bi-labeled graphs.

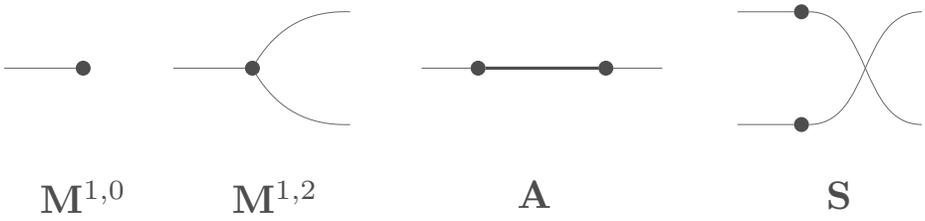


Figure 2. Some important bi-labeled graphs.

Operations on bi-labeled graphs: There are several possibilities for how one might act on or combine some given bi-labeled graphs in order to construct new ones. Here we will describe a few such possibilities. For a completely rigorous description of these operations, please see the full verison of the paper [22, Section 3.1]. Here we only aim to give an intuitive understanding.

The first operation we present is the *composition* of two bi-labeled graphs. Given $\mathbf{K}_1 = (K_1, \mathbf{a}, \mathbf{b})$ and $\mathbf{K}_2 = (K_2, \mathbf{c}, \mathbf{d})$ such that \mathbf{b} and \mathbf{c} have the same length, the composition $\mathbf{K}_1 \circ \mathbf{K}_2$ is the bi-labeled graph \mathbf{K} whose underlying graph K is obtained by taking the disjoint union of K_1 and K_2 and then merging the vertices b_i and c_i for all i . The output vector of \mathbf{K} is \mathbf{a} and its input vector is \mathbf{d} . An illustration is given in Figure 3.

Given bi-labeled graphs $\mathbf{K}_1 = (K_1, \mathbf{a}, \mathbf{b})$ and $\mathbf{K}_2 = (K_2, \mathbf{c}, \mathbf{d})$, their *tensor product*, denoted $\mathbf{K}_1 \otimes \mathbf{K}_2$, is the bi-labeled graph $\mathbf{K} = (K_1 \cup K_2, \mathbf{ac}, \mathbf{bd})$, where \mathbf{ac} denotes the concatenation of the tuples \mathbf{a} and \mathbf{c} . Diagrammatically, $\mathbf{K}_1 \otimes \mathbf{K}_2$ is obtained by simply drawing \mathbf{K}_1 above \mathbf{K}_2 .

The *transpose* of a bi-labeled graph $\mathbf{K} = (K, \mathbf{a}, \mathbf{b})$ is $\mathbf{K}^* := (K, \mathbf{b}, \mathbf{a})$, i.e., we simply swap the input/output vectors. Diagrammatically, we reflect \mathbf{K} about the vertical axis.

Our last operation is the *Schur product* of bi-labeled graphs. Given $\mathbf{K}_1 = (K_1, \mathbf{a}, \mathbf{b})$ and $\mathbf{K}_2 = (K_2, \mathbf{c}, \mathbf{d})$ such that \mathbf{a}, \mathbf{c} have the same length, and \mathbf{b}, \mathbf{d} have the same length, their Schur product, denoted $\mathbf{K}_1 \bullet \mathbf{K}_2$, is the bi-

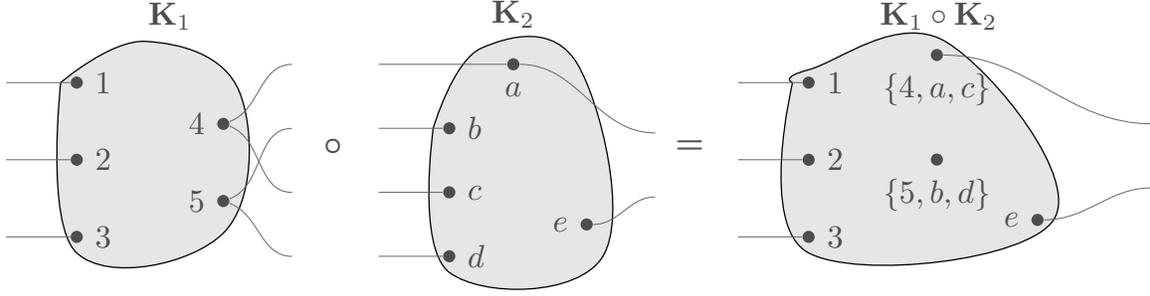
labeled graph $\mathbf{K} = (K, \mathbf{a}', \mathbf{b}')$ where K is obtained from the $K_1 \cup K_2$ by merging a_i with c_i and b_j with d_j for all i, j . These merged vertices are the elements of the input/output vectors of \mathbf{K} , i.e., a'_i is the vertex formed by merging a_i with c_i .

It should come as no surprise that the above operations on bi-labeled graphs correspond to algebraic operations on the corresponding homomorphism matrices. The proofs of these correspondences are given in the full version of the paper [22, Section 3.2]. Summarizing the results there, we have the following:

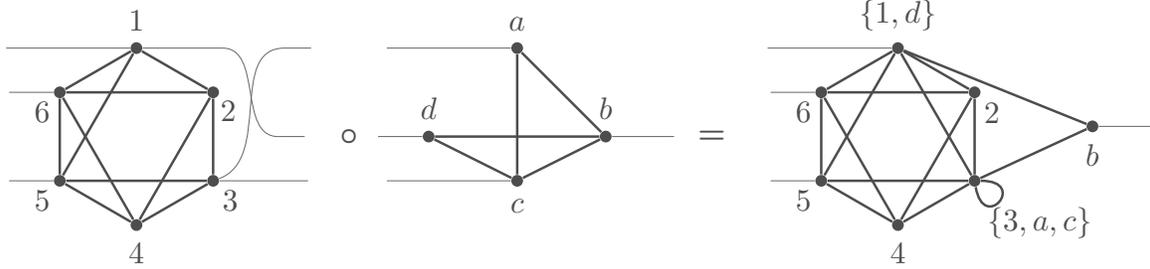
$$\begin{aligned}
 T^{\mathbf{K}_1 \rightarrow G} T^{\mathbf{K}_2 \rightarrow G} &= T^{\mathbf{K}_1 \circ \mathbf{K}_2 \rightarrow G}; \\
 T^{\mathbf{K}_1 \rightarrow G} \otimes T^{\mathbf{K}_2 \rightarrow G} &= T^{\mathbf{K}_1 \otimes \mathbf{K}_2 \rightarrow G}; \\
 (T^{\mathbf{K} \rightarrow G})^* &= T^{\mathbf{K}^* \rightarrow G}; \\
 T^{\mathbf{K}_1 \rightarrow G} \bullet T^{\mathbf{K}_2 \rightarrow G} &= T^{\mathbf{K}_1 \bullet \mathbf{K}_2 \rightarrow G}.
 \end{aligned}$$

Note that the Schur product of matrices is simply the entrywise product, also sometimes called the Hadamard product.

Planar bi-labeled graphs: Her we introduce a class of bi-labeled graphs based on planarity, which will allow us to make a connection to quantum automorphism groups of graphs. To define this class, we must first define the following: Given an ℓ -output, k -input bi-labeled graph $\mathbf{K} = (K, \mathbf{a}, \mathbf{b})$, define the graph $K^\circ := K^\circ(\mathbf{a}, \mathbf{b})$ as the graph obtained from K by adding the cycle $C =$



(a) Generic example of composition.



(b) Concrete example of composition.

Figure 3. Illustration of the composition of bi-labeled graphs.

$(\alpha_1, \dots, \alpha_\ell, \beta_k, \dots, \beta_1)$ of new vertices, and edges $a_i \alpha_i$, $b_j \beta_j$ for all $i \in [\ell], j \in [k]$. We refer to the cycle C as the *enveloping cycle* of K° . We further define $K^\circ := K^\circ(\mathbf{a}, \mathbf{b})$ as the graph obtained from K° by adding an additional vertex adjacent to every vertex of the enveloping cycle.

Although $K^\circ(\mathbf{a}, \mathbf{b})$ and $K^\circ(\mathbf{a}, \mathbf{b})$ depend on the input and output vectors of $\mathbf{K} = (K, \mathbf{a}, \mathbf{b})$, we will typically refer to them as simply K° and K° when there should be no confusion.

We now define our class \mathcal{P} of *planar bi-labeled graphs* as those bi-labeled graphs \mathbf{K} such that K° is planar, or equivalently such that K° has a planar embedding in which its enveloping cycle is the boundary of a face (typically chosen to be the outer face). We use $\mathcal{P}(\ell, k)$ to denote the elements of \mathcal{P} with ℓ outputs and k inputs. Note that this definition implies that if $(K, \mathbf{a}, \mathbf{b}) \in \mathcal{P}(\ell, k)$, then K is planar. However, this is not sufficient. For example, if $\mathbf{K} = (K, (a, b), (b, a))$ where K is the edgeless graph on $\{a, b\}$, then K is clearly planar, but K° is a subdivision of a complete graph on five vertices and thus non-planar. However, we remark that if $\ell + k \leq 1$, then the planarity of K is both necessary and sufficient for membership in $\mathcal{P}(\ell, k)$ [22, Lemma 5.5].

One reason not to define \mathcal{P} to be the bi-labeled graphs $(K, \mathbf{a}, \mathbf{b})$ such that K is planar, is that this condition is not stable under composition of bi-labeled graphs [22, Example 5.8]. As we will see in Section III, the closure of \mathcal{P} under the operations of composition, tensor product, and transpose is crucial to the proof of our main result.

Finally, let us remark that though the class $\mathcal{P}(\ell, k)$ is not closed under Schur product for $\ell + k \geq 3$, it is for $\ell + k \leq 2$, and this is particularly easy to see for $\ell + k \leq 1$ which is all we will need here.

B. Quantum automorphism groups of graphs

A complete introduction to the aspects of quantum groups required for this work is given in the full version [22, Section 2.2]. Here we only aim to briefly introduce the quantum automorphism group of a graph and the associated notion of intertwiners.

An isomorphism from a graph G to itself is known as an *automorphism*, and these form a group under composition known as the *automorphism group* of G , denoted $\text{Aut}(G)$. An element of $\text{Aut}(G)$ is necessarily a permutation of the elements of $V(G)$, and thus they can be encoded as permutation matrices whose rows/columns are indexed by $V(G)$. Then $P \in \text{Aut}(G)$ if and only if $PA_G = A_G P$, i.e., P commutes with the adjacency matrix of G . In order to define the quantum automorphism group of G , denoted $\text{Qut}(G)$, we will introduce a quantum analog of permutation matrices.

A typical definition of a permutation matrix might be that it is a 01-matrix with precisely one 1 in every row and column. However, we can define it somewhat more abstractly: A matrix $P = (p_{ij}) \in \mathbb{C}^{n \times n}$ is a permutation matrix if $p_{ij} = p_{ij}^2$, and $\sum_\ell p_{i\ell} = 1 = \sum_k p_{kj}$ for all $i, j \in [n]$. A *quantum permutation matrix* (or *magic unitary*) is an operator-valued generalization of this definition.

Precisely, an $n \times n$ matrix $U = (u_{ij})$ whose entries are elements of a C^* -algebra \mathcal{A} (equivalently they are bounded linear operators on a Hilbert space) is a quantum permutation matrix if $u_{ij} = u_{i,j}^* = u_{i,j}^2$ and $\sum_{\ell} u_{i\ell} = \mathbf{1} = \sum_k u_{k j}$ for all $i, j \in [n]$. Here, the element $\mathbf{1}$ is the identity in \mathcal{A} . It is not difficult to see that a permutation matrix is exactly a quantum permutation matrix with $\mathcal{A} = \mathbb{C}$. Note that the definition of a quantum permutation matrix implies that its entries along a row or column are pairwise orthogonal.

Now consider a graph G and its automorphism group $\text{Aut}(G)$ represented as permutation matrices. The \mathbb{C} -valued functions on $\text{Aut}(G)$ form a (commutative) C^* -algebra under pointwise multiplication denoted $C(\text{Aut}(G))$. For $i, j \in V(G)$, let $u_{ij} \in C(\text{Aut}(G))$ denote the function that takes $P \in \text{Aut}(G)$ to its ij -entry. In other words, u_{ij} is the characteristic function of automorphisms of G that map i to j . The functions u_{ij} are 01-valued, and thus $u_{ij} = u_{i,j}^* = u_{i,j}^2$. Moreover, since any permutation matrix has exactly one 1 in its i^{th} row, we have that $\sum_{\ell} u_{i\ell}(P) = 1$ for all $P \in \text{Aut}(G)$. In other words, $\sum_{\ell} u_{i\ell}$ is the constant 1 function, which is the identity in $C(\text{Aut}(G))$. Similarly, $\sum_k u_{kj} = \mathbf{1}$, and thus $U = (u_{ij})$ is a quantum permutation matrix. This matrix U has an additional nice property: since $PA_G = A_G P$ for all $P \in \text{Aut}(G)$, it follows that $U A_G = A_G U$, where $(A_G U)_{ij} := \sum_k (A_G)_{ik} u_{kj} = \sum_{k:k \sim i} u_{kj}$ and similarly for the entries of $U A_G$. It turns out that the properties of the matrix U above can be used to define $C(\text{Aut}(G))$ in a more abstract manner. The algebra $C(\text{Aut}(G))$ is isomorphic to the universal C^* -algebra³ with *commutative* generators u_{ij} for $i, j \in V(G)$ such that $U = (u_{ij})$ is a quantum permutation matrix and $U A_G = A_G U$. We now define the ‘‘algebra of functions’’ on the quantum automorphism group of G , denoted $C(\text{Qut}(G))$, in the same way but without the restriction that the entries of U commute. The quantum permutation matrix U is referred to as the *fundamental representation* of $\text{Qut}(G)$.

One might notice that we have not defined the quantum automorphism group of G , which we denote $\text{Qut}(G)$, but only the algebra $C(\text{Qut}(G))$. In the classical case, the (abstractly defined) algebra $C(\text{Aut}(G))$ along with the matrix U completely determine $\text{Aut}(G)$, and a similar statement holds more generally for compact groups and certain types of commutative algebras. This motivates the *viewing* of noncommutative analogs of these algebras, such as $C(\text{Qut}(G))$, as being the function algebras of noncommutative (or quantum) groups. Note however that $C(\text{Qut}(G))$ is not *actually* the algebra of functions of *any* object, since any such algebra is necessarily commutative. Thus these quantum groups exist mostly as a useful analogy for studying such noncommutative algebras. Some authors state explicitly that these quantum groups do not exist as concrete mathematical

³Without going into details, the *universal C^* -algebra* construction used here is analogous to the construction of groups using generators and relations.

objects [5], whereas others say that $\text{Qut}(G)$ is the pair $(C(\text{Qut}(G)), U)$ [27].

The most important objects associated to $\text{Qut}(G)$ for us are its intertwiners. Let G be a graph and U be the fundamental representation of $\text{Qut}(G)$. An (ℓ, k) -intertwiner of $\text{Qut}(G)$ is a $V(G)^\ell \times V(G)^k$ complex valued matrix T satisfying $U^{\otimes \ell} T = T U^{\otimes k}$. Here $U^{\otimes r}$ is a $V(G)^r \times V(G)^r$ matrix whose $i_1 \dots i_r, j_1 \dots j_r$ -entry is $u_{i_1 j_1} u_{i_2 j_2} \dots u_{i_r j_r}$. We use $C_q^G(\ell, k)$ to denote the set of (ℓ, k) -intertwiners of $\text{Qut}(G)$, and C_q^G to denote the union of all these sets.

In the classical case, an (ℓ, k) -intertwiner of $\text{Aut}(G)$ is a matrix T satisfying $P^{\otimes \ell} T = T P^{\otimes k}$ for all permutation matrices $P \in \text{Aut}(G)$. This is equivalent to T being constant on the orbits of the induced action of $\text{Aut}(G)$ on $V(G)^\ell \times V(G)^k$, i.e., $T_{i_1 \dots i_\ell, j_1 \dots j_k} = T_{\sigma(i_1) \dots \sigma(i_\ell), \sigma(j_1) \dots \sigma(j_k)}$ for all $\sigma \in \text{Aut}(G)$. In other words, the (ℓ, k) -intertwiners of $\text{Aut}(G)$ are the span of the characteristic matrices of the orbits of the action of $\text{Aut}(G)$ on $V(G)^\ell \times V(G)^k$. In the quantum case, this unfortunately breaks down somewhat. For $\ell + k \geq 3$, the space $C_q^G(\ell, k)$ is not necessarily the span of any set of 01-matrices. However, for $\ell + k \leq 2$, the notion of orbits of $\text{Qut}(G)$ is sound. In [20], it was shown if $U = (u_{ij})$ is the fundamental representation of $\text{Qut}(G)$, then the relation $i \sim_1 j$ defined as $u_{ij} \neq 0$ is an equivalence relation, and its equivalence classes are defined to be the orbits of $\text{Qut}(G)$. Moreover, they show that $C_q^G(1, 0)$ is the span of the characteristic vectors of these orbits. We will see below that the orbits of the quantum automorphism group of a graph can be used to characterize quantum isomorphism.

It is known and straightforward to see that C_q^G is closed under matrix products, tensor products, conjugate transposition, and linear combinations. Moreover, it is known that using these operations, C_q^G is generated by just three of its elements [7]:

$$C_q^G = \langle M^{1,0}, M^{1,2}, A_G \rangle_{+, \otimes, *}.$$

Here, $M^{\ell,k}$ is the homomorphism matrix of the bi-labeled graph $\mathbf{M}^{\ell,k}$ described in Section II-A. It thus follows from the correspondence between bi-labeled graph operations and matrix operations that $C_q^G(\ell, k)$ is equal to the span of the matrices $T^{\mathbf{K} \rightarrow G}$ such that $\mathbf{K} \in \langle M^{1,0}, M^{1,2}, A_G \rangle_{\otimes, *, *}$ and \mathbf{K} has ℓ outputs and k inputs. Therefore, we can characterize the intertwiners of $\text{Qut}(G)$ by characterizing the bi-labeled graphs that are generated by $\mathbf{M}^{1,0}, \mathbf{M}^{1,2}$, and \mathbf{A} using the operations of composition, tensor product, and transposition. Notably, the latter is a purely combinatorial problem.

C. Quantum isomorphisms

We give a detailed overview of quantum strategies for the isomorphism game in the full version [22, Section 2.3]. Here we simply introduce a concise mathematical reformulation of the operationally defined notion of quantum isomorphism.

A well known reformulation of graph isomorphism states that $G \cong H$ if and only if there exists a *permutation matrix*

P such that $PA_G P^T = A_H$. A similar characterization of quantum isomorphism makes use of the notion of quantum permutation matrices we described above. The following theorem summarizes results from both [3] and [20]:

Theorem II.1. *Given G and H we have*

- 1) $G \cong H$ if and only if $\mathcal{U}A_G = A_H\mathcal{U}$ for a quantum permutation \mathcal{U} with commutative entries.
- 2) $G \cong_{qc} H$ if and only if $\mathcal{U}A_G = A_H\mathcal{U}$ for a quantum permutation matrix \mathcal{U} .

III. PROOF OVERVIEW

Here we give an overview of the proof that graphs G and H are quantum isomorphic if and only if $\text{hom}(K, G) = \text{hom}(K, H)$ for all planar graphs K . We will refer to the latter relation as *planar isomorphism*. We first present some theorems that connect the notions of quantum automorphism groups, quantum isomorphisms, and homomorphism counting. We will then apply these to prove our main result in Sections III-A and III-B. We begin with the main result of [20], which characterizes quantum isomorphism in terms of quantum automorphism groups:

Theorem III.1. *Let G and H be connected graphs. Then $G \cong_{qc} H$ if and only if there exists $g \in V(G)$ and $h \in V(H)$ that are in the same orbit of $\text{Qut}(G \cup H)$.*

Note that this characterization is perfectly analogous to the classical case. We also remark that the connectedness condition above is not really a restriction. By taking complements (which preserves quantum isomorphism) if necessary we may assume that G is connected, and connected graphs can only be quantum isomorphic to connected graphs. With Theorem III.1 in mind, we see that a combinatorial characterization of the orbits of $\text{Qut}(G)$ would be a big step towards a combinatorial characterization of quantum isomorphism, and this is indeed the route we take.

Recall from Section II-B, that the orbits of $\text{Qut}(G)$ are closely related to the $(1, 0)$ -intertwiners of $\text{Qut}(G)$. Precisely, $C_q^G(1, 0)$ is the span of the characteristic vectors of the orbits of $\text{Qut}(G)$. Thus to obtain our homomorphism-counting characterization of orbits, we in fact prove such a characterization for the intertwiners of $\text{Qut}(G)$ in general. We show that the intertwiners of $\text{Qut}(G)$ are the (span of) the G -homomorphism matrices of the planar bi-labeled graphs introduced in Section II-A:

Theorem III.2. *For any graph G , we have that $C_q^G(\ell, k) = \text{span}\{T^{\mathbf{K} \rightarrow G} : \mathbf{K} \in \mathcal{P}(\ell, k)\}$, $\forall \ell, k \in \mathbb{N}$.*

As discussed in Section II-B, it is previously known that $C_q^G = \langle M^{1,0}, M^{1,2}, A_G \rangle_{+, \circ, \otimes, *}$, and these three generators are the G -homomorphism matrices of $\mathbf{M}^{1,0}, \mathbf{M}^{1,2}, \mathbf{A}$ respectively. Thus the above theorem is immediate from the correspondence between matrix and bi-labeled graph operations described in Section II-A and the following:

Theorem III.3. $\mathcal{P} = \langle \mathbf{M}^{1,0}, \mathbf{M}^{1,2}, \mathbf{A} \rangle_{+, \circ, \otimes, *}$.

The proof of Theorem III.3 in fact comprises the majority of the work of this paper. Notably, the claim is a purely combinatorial one, as is its proof. We will leave the description of its proof for Section III-C, and proceed onwards for now.

Specializing to the case $\ell = 1, k = 0$, Theorem III.2 tells us that $C_q^G(1, 0)$, i.e., the span of the characteristic vectors of the orbits of $\text{Qut}(G)$, is equal to the span of the G -homomorphism matrices of $\mathbf{K} \in \mathcal{P}(1, 0)$. Recall that a bi-labeled graph $\mathbf{K} = (K, (a), \emptyset)$ is an element of $\mathcal{P}(1, 0)$ if and only if K is planar (there is no condition on the vertex a). We thus obtain the following lemma, where $\text{hom}((K, a), (G, u))$ denotes the number of homomorphisms from K to G that map a to u :

Corollary III.4. *Let G be a graph. Vertices $u, v \in V(G)$ are in the same orbit of $\text{Qut}(G)$ if and only if*

$$\text{hom}((K, a), (G, u)) = \text{hom}((K, a), (G, v))$$

for all connected planar graphs K and $a \in V(K)$.

We are able to restrict to connected planar graphs K in the above because additional components only contribute a scalar factor to the G -homomorphism matrix of \mathbf{K} .

A. Quantum isomorphism implies planar isomorphism

Suppose that graphs G and H are quantum isomorphic. By Theorem II.1, there exists a quantum permutation matrix \mathcal{U} such that $\mathcal{U}A_G = A_H\mathcal{U}$. We also have that $\mathcal{U}M^{1,0} = M^{1,0}\mathcal{U}^{\otimes 0}$ and $\mathcal{U}M^{1,2} = M^{1,2}\mathcal{U}^{\otimes 2}$, since these hold for any quantum permutation matrix. From this we are able to prove that if $T \in \langle M^{1,0}, M^{1,2}, A_G \rangle_{+, \circ, \otimes, *}$ is given by an expression involving $M^{1,0}, M^{1,2}$, and A_G , and $T' \in \langle M^{1,0}, M^{1,2}, A_H \rangle_{+, \circ, \otimes, *}$ is given by the same expression but with each occurrence of A_G replaced by A_H , then $\mathcal{U}^{\otimes \ell} T = T' \mathcal{U}^{\otimes k}$ for appropriate ℓ, k . By Theorem III.3 and the correspondence between matrix and bi-labeled graph operations, we obtain that $\mathcal{U}^{\otimes \ell} T^{\mathbf{K} \rightarrow G} = T'^{\mathbf{K} \rightarrow H} \mathcal{U}^{\otimes k}$ for all $\mathbf{K} \in \mathcal{P}(\ell, k)$. Furthermore, we are able to show that this correspondence is sum-preserving. As the sum of the entries of $T^{\mathbf{K} \rightarrow G}$ for $\mathbf{K} = (K, \mathbf{a}, \mathbf{b})$ is simply $\text{hom}(K, G)$, it immediately follows that $\text{hom}(K, G) = \text{hom}(K, H)$ for all planar graphs K .

B. Planar isomorphism implies quantum isomorphism

The return route is a bit longer, the difficulty being that in a certain sense we only have access to the sum of the entries of homomorphism matrices in this case. However, a simple trick using Schur products of bi-labeled graphs/homomorphism matrices will solve our problems. So suppose that $\text{hom}(K, G) = \text{hom}(K, H)$ for all planar graphs K . It follows that the sum of the entries of $T^{\mathbf{K} \rightarrow G}$ is equal to that of $T^{\mathbf{K} \rightarrow H}$ for any $\mathbf{K} \in \mathcal{P}$, and this extends also to linear combinations of homomorphism matrices. Recall that, unlike the general case, $\mathcal{P}(1, 0)$ is closed under Schur

products. Therefore, $\mathbf{K}^{\bullet m} \in \mathcal{P}(1, 0)$ for $\mathbf{K} \in \mathcal{P}(1, 0)$, and thus the sum of the entries of $(T^{\mathbf{K} \rightarrow G})^{\bullet m} = T^{\mathbf{K}^{\bullet m} \rightarrow G}$ is equal to that of $(T^{\mathbf{K} \rightarrow H})^{\bullet m} = T^{\mathbf{K}^{\bullet m} \rightarrow H}$. It then follows from Newton's relations that $(T^{\mathbf{K} \rightarrow G})$ and $(T^{\mathbf{K} \rightarrow H})$ have the same multiset of entries for $\mathbf{K} \in \mathcal{P}(1, 0)$, and this also extends to linear combinations. With a bit more work one can show that the expressions for the characteristic vectors of the orbits of $\text{Qut}(G)$ as linear combinations of homomorphism matrices are the same as the expressions for the characteristic vectors of the orbits of $\text{Qut}(H)$, only with H -homomorphism matrices replaced by G -homomorphism matrices. In particular, this implies there is a bijection of the orbits of $\text{Qut}(G)$ and $\text{Qut}(H)$ that preserves cardinality.

We now consider the disjoint union $X = G \cup H$ with an aim to apply Theorem III.1. As with quantum isomorphism, planar isomorphism is preserved when taking complements [22, Lemma 7.12], and planar isomorphic graphs must have the same number of connected components [22, Lemma 7.11]. Thus, taking complements if necessary, we may assume that G and H were both connected, and so we can apply Theorem III.1. Now let $R = \sum_i \alpha_i T^{\mathbf{K}_i \rightarrow G}$ be the characteristic vector of an orbit of $\text{Qut}(G)$. Then, as outlined above, $R' = \sum_i \alpha_i T^{\mathbf{K}_i \rightarrow H}$ is the characteristic vector of a corresponding orbit of $\text{Qut}(H)$. Pick $u \in V(G)$, $u' \in V(H)$ in these orbits, i.e., such that $R_u = 1 = R'_{u'}$. Now let K be a *connected* planar graph and $a \in V(K)$. Then $\mathbf{K} = (K, (a), \emptyset) \in \mathcal{P}(1, 0)$. We will show that $\text{hom}((K, a), (X, u)) = \text{hom}((K, a), (X, u'))$ and thus u, u' are in the same orbit of $\text{Qut}(X)$ by Corollary III.4.

Let $T = T^{\mathbf{K} \rightarrow G}$ and $T' = T^{\mathbf{K} \rightarrow H}$. Since K is connected, any homomorphism from K to X mapping a to u has its image completely contained in $V(G)$, and similarly for homomorphisms mapping a to u' . In other words

$$\begin{aligned} \text{hom}((K, a), (X, u)) &= \text{hom}((K, a), (G, u)) = T_u \\ \text{hom}((K, a), (X, u')) &= \text{hom}((K, a), (H, u')) = T'_{u'}, \end{aligned}$$

and so we desire to show that $T_u = T'_{u'}$. For this, consider $R \bullet T$ and $R' \bullet T'$. As $T \in C_q^G(1, 0)$, it is constant on the orbits of $\text{Qut}(G)$, and similarly for T' and $\text{Qut}(H)$. Thus $R \bullet T = \alpha R$ and $R' \bullet T' = \alpha' R'$, where $\alpha = T_u$ and $\alpha' = T'_{u'}$. In fact, we must have $\alpha = \alpha'$, since $R \bullet T$ and $R' \bullet T'$ must have the same multiset of entries as outlined above. Thus $\text{hom}((K, a), (X, u)) = \text{hom}((K, a), (X, u'))$ for the arbitrary connected planar K and thus u, u' are in the same orbit of $\text{Qut}(X)$ by Corollary III.4. Therefore, by Theorem III.1, G and H are quantum isomorphic.

*C. Proof of $\mathcal{P} = \langle \mathbf{M}^{1,0}, \mathbf{M}^{1,2}, \mathbf{A} \rangle_{\circ, \otimes, *}$*

The combinatorial characterization of the intertwiners of $\text{Qut}(G)$ presented in Theorem III.2 requires us to show that the planar bi-labeled graphs introduced in Section II-A are precisely the bi-labeled graphs generated by $\mathbf{M}^{1,0}, \mathbf{M}^{1,2}$, and \mathbf{A} using the operations of composition, tensor product, and transposition. It is easy to see that $\mathbf{M}^{1,0}, \mathbf{M}^{1,2}, \mathbf{A} \in \mathcal{P}$,

and thus to show that $\langle \mathbf{M}^{1,0}, \mathbf{M}^{1,2}, \mathbf{A} \rangle_{\circ, \otimes, *} \subseteq \mathcal{P}$ it suffices to show that \mathcal{P} is closed under these three operations. For transposition, this is straightforward, since $K^\circ(\mathbf{a}, \mathbf{b})$ is isomorphic to $K^\circ(\mathbf{b}, \mathbf{a})$. In order to show that \mathcal{P} is closed under composition, we consider $\mathbf{H}_1 = (H_1, \mathbf{a}, \mathbf{b}) \in \mathcal{P}(\ell, k)$, $\mathbf{H}_2 = (H_2, \mathbf{c}, \mathbf{d}) \in \mathcal{P}(k, m)$, and let $\mathbf{H} = (H, \mathbf{a}', \mathbf{b}')$ be the bi-labeled graph such that $\mathbf{H} = \mathbf{H}_1 \circ \mathbf{H}_2$. To show $\mathbf{H} \in \mathcal{P}$, we carefully construct H° from the disjoint union of H_1° and H_2° , ensuring that at every step we remain planar. The full proof is given in [22] but the essential idea is to perform a sort of ‘‘reverse mitosis’’ procedure illustrated in Figure 3.

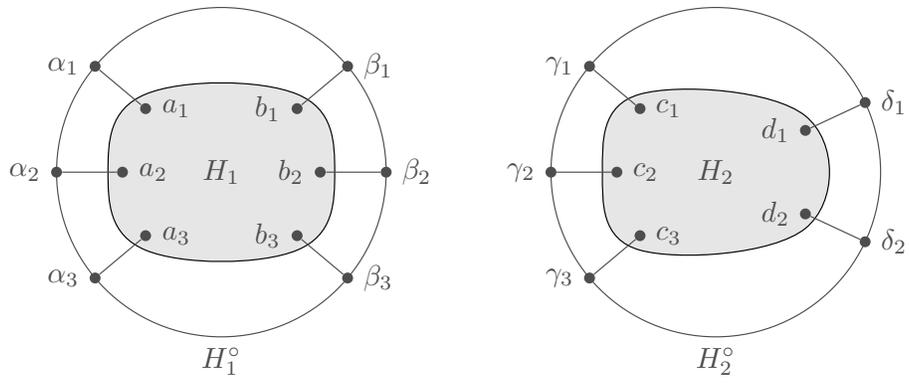
The proof that \mathcal{P} is closed under tensor products is similar. So we see that \mathcal{P} is indeed closed under composition, tensor product, and transpose, and thus $\langle \mathbf{M}^{1,0}, \mathbf{M}^{1,2}, \mathbf{A} \rangle_{\circ, \otimes, *} \subseteq \mathcal{P}$.

The proof of the other containment is more challenging. In this case, instead of constructing a bi-labeled graph from two separate bi-labeled graphs, we must pull apart a single bi-labeled graph into two parts. There is some choice in how to do this, and we must choose carefully. More precisely, given a bi-labeled graph $\mathbf{K} \in \mathcal{P}$, we select a vertex $v \in V(K)$ that we can ‘‘pluck’’ out of \mathbf{K} to obtain $\mathbf{K}' \in \mathcal{P}$ with one fewer vertex. By induction, we have that $\mathbf{K}' \in \langle \mathbf{M}^{1,0}, \mathbf{M}^{1,2}, \mathbf{A} \rangle_{\circ, \otimes, *}$. We can then express \mathbf{K} in terms of \mathbf{K}' and some simple bi-labeled graphs that can be shown to be in $\langle \mathbf{M}^{1,0}, \mathbf{M}^{1,2}, \mathbf{A} \rangle_{\circ, \otimes, *}$ directly. It follows that $\mathbf{K} \in \langle \mathbf{M}^{1,0}, \mathbf{M}^{1,2}, \mathbf{A} \rangle_{\circ, \otimes, *}$. The vertex v that we pluck out of $\mathbf{K} = (K, (a_1, \dots, a_\ell), (b_1, \dots, b_k))$ must be chosen so that its occurrences in $a_1, \dots, a_\ell, b_k, \dots, b_1$ are cyclically consecutive. In other words, the neighbors of v on the enveloping cycle of K° appear consecutively in this cycle (such a vertex always exists [22, Corollary 6.5]). The main difficulty of the proof is in showing that the bi-labeled graph \mathbf{K}' obtained by removing v from \mathbf{K} is indeed an element of \mathcal{P} . This is done by carefully constructing K'° from K° , showing that the former has a planar embedding in which its enveloping cycle is the boundary of a face. The rough idea of the proof is illustrated in Figure 3 (for full details, see [22, Lemma 6.6]). Having shown that an arbitrary $\mathbf{K} \in \mathcal{P}$ must be an element of $\langle \mathbf{M}^{1,0}, \mathbf{M}^{1,2}, \mathbf{A} \rangle_{\circ, \otimes, *}$, we obtain our desired equality, completing the proof of Theorem III.3.

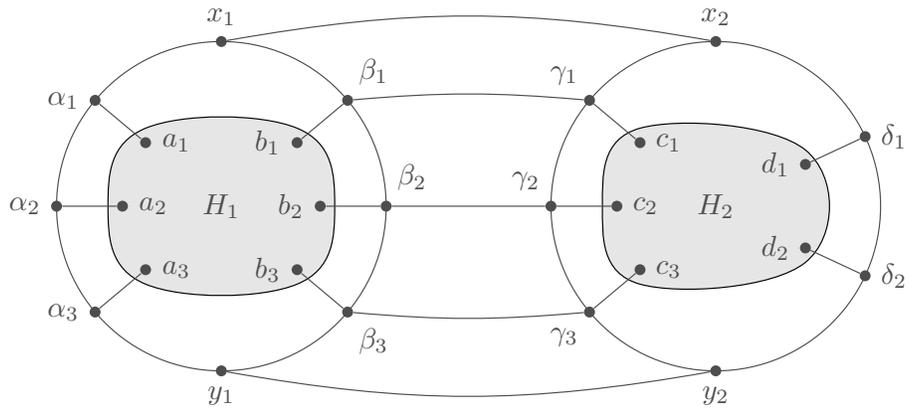
ACKNOWLEDGMENT

LM: Supported by the Villum Fonden via the QMATH Centre of Excellence (Grant No. 10059).

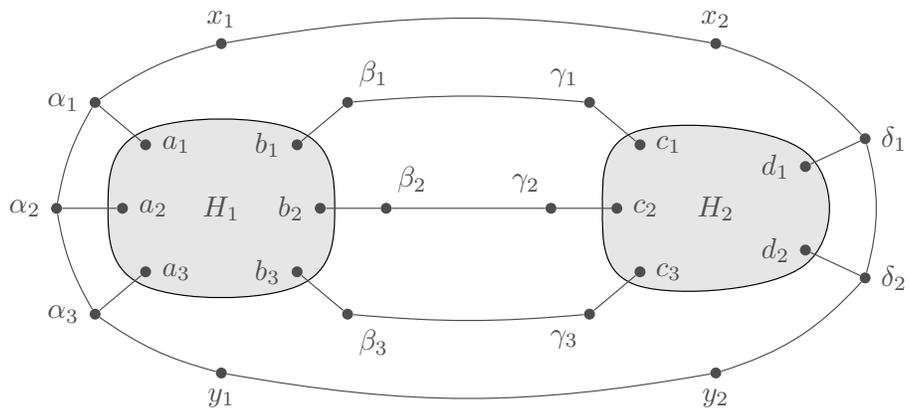
DR: I would like to acknowledge the helpful discussions with Arthur Chassaniol concerning intertwiners of quantum automorphism groups which took place at the ‘Cohomology of quantum groups and quantum automorphism groups of finite graphs’ Workshop at Saarland University in October of 2018. The research leading to these results has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement no. 713683 (COFUNDfellowsDTU).



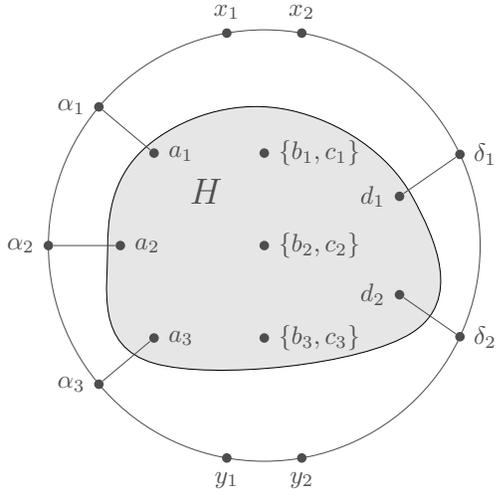
(a) $H_1^\circ \cup H_2^\circ$



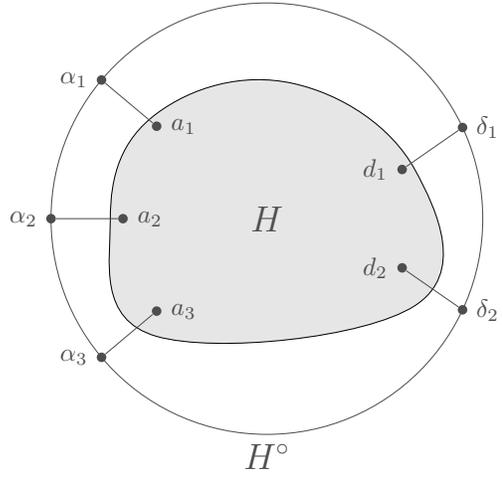
(b) Subdivide with x_1, x_2, y_1, y_2 and add edges $\beta_i\gamma_i$ for $i = 1, 2, 3$ and x_1x_2, y_1, y_2 .



(c) Remove edges $x_1\beta_1, \beta_3y_1, x_2\gamma_1, \gamma_3y_2$, and $\beta_i\beta_{i+1}$ and $\gamma_i\gamma_{i+1}$ for $i = 1, 2$.

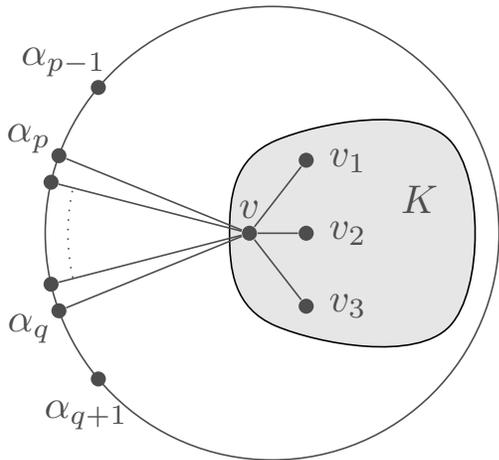


(d) Unsubdivide the paths $b_i, \beta_i, \gamma_i, c_i$ and then contract the edges $b_i c_i$ for $i = 1, 2, 3$.

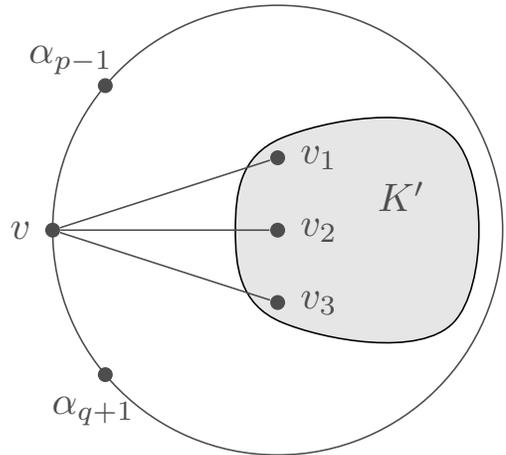


(e) Unsubdivide $\alpha_1, x_1, x_2, \delta_1$ and $\alpha_3, y_1, y_2, \delta_2$.

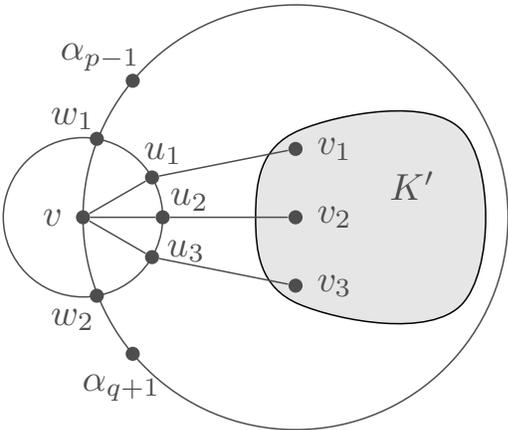
Figure 3. Illustration of the proof that \mathcal{P} is closed under composition.



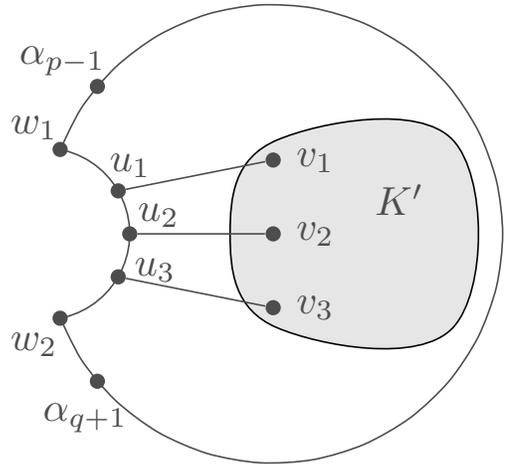
(a) The vertex v and its neighbors in K° .



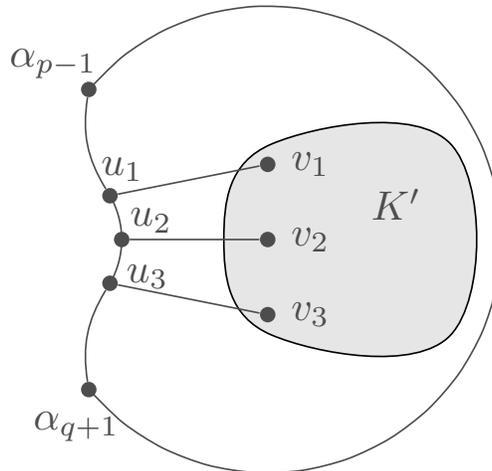
(b) Contract edges $v\alpha_i$ for $i = p, \dots, q$, and edges $\alpha_j\alpha_{j+1}$ for $j = p, \dots, q - 1$, to form new vertex v .



(c) Subdivide edges incident to v and add cycle through new vertices.



(d) Remove vertex v and edge $w_1 w_2$.



(e) Unsubdivide α_{p-1}, w_1, u_1 and u_3, w_2, α_{q+1} to obtain K'^o .

Figure 3. Illustration of the fact that $\mathbf{K}' \in \mathcal{P}$, from the proof of Theorem III.3.

REFERENCES

- [1] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)*, 45(3):501–555, 1998.
- [2] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM (JACM)*, 45(1):70–122, 1998.
- [3] Albert Atserias, Laura Mančinska, David E. Roberson, Robert Šámal, Simone Severini, and Antonios Varvitsiotis. Quantum and non-signalling graph isomorphisms. *Journal of Combinatorial Theory, Series B*, 136:289 – 328, 2019. doi:10.1016/j.jctb.2018.11.002.
- [4] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational complexity*, 1(1):3–40, 1991.
- [5] Teodor Banica, Julien Bichon, and Benoît Collins. Quantum permutation groups: a survey. *Banach Center Publications*, 78:13–34, 2007. doi:10.4064/bc78-0-1.
- [6] John S. Bell. On the Einstein Podolsky Rosen paradox. *Physics*, 1(3):195, 1964.
- [7] Arthur Chassaniol. Study of quantum symmetries for vertex-transitive graphs using intertwiner spaces. 2019. arXiv:1904.00455.
- [8] Richard Cleve, Peter Hoyer, Benjamin Toner, and John Watrous. Consequences and limits of nonlocal strategies. In *Proceedings of the 19th IEEE Annual Conference on Computational Complexity*, pages 236–249, 2004. arXiv:quant-ph/0404076, doi:10.1109/CCC.2004.1313847.
- [9] Richard Cleve, Li Liu, and William Slofstra. Perfect commuting-operator strategies for linear system games. *Journal of Mathematical Physics*, 58(1):012202, 2017. doi:10.1063/1.4973422.
- [10] Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming, ICALP '14*, pages 320–331. 2014. arXiv:1209.2729.
- [11] Matthew Coudron and William Slofstra. Complexity lower bounds for computing the approximately-commuting operator value of non-local games to high precision. 2019. arXiv:1905.11635.
- [12] Holger Dell, Martin Grohe, and Gaurav Rattan. Lovász Meets Weisfeiler and Leman. In *45th International Colloquium on Automata, Languages, and Programming (ICALP 2018)*, volume 107 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 40:1–40:14, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ICALP.2018.40.
- [13] Andrew C Doherty, Yeong-Cherng Liang, Ben Toner, and Stephanie Wehner. The quantum moment problem and bounds on entangled multi-prover games. In *2008 23rd Annual IEEE Conference on Computational Complexity*, pages 199–210, 2008.
- [14] Zdeněk Dvořák. On recognizing graphs by numbers of homomorphisms. *Journal of Graph Theory*, 64(4):330–342, 2010. doi:10.1002/jgt.20461.
- [15] Uriel Feige and László Lovász. Two-prover one-round proof systems: Their power and their problems (extended abstract). In *Proceedings of the Twenty-fourth Annual ACM Symposium on Theory of Computing, STOC '92*, pages 733–744, New York, NY, USA, 1992. ACM. doi:10.1145/129712.129783.
- [16] Tsuyoshi Ito and Thomas Vidick. A multi-prover interactive proof for nexp sound against entangled provers. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 243–252, 2012.
- [17] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP*=RE, 2020. arXiv:2001.04383.

- [18] László Lovász. Operations with structures. *Acta Mathematica Hungarica*, 18(3-4):321–328, 1967. doi:[10.1007/BF02280291](https://doi.org/10.1007/BF02280291).
- [19] László Lovász. *Large networks and graph limits*, volume 60. American Mathematical Society, 2012. doi:[10.1090/coll/060](https://doi.org/10.1090/coll/060).
- [20] Martino Lupini, Laura Mančinska, and David E. Roberson. Nonlocal games and quantum permutation groups. 2017. arXiv:[1712.01820](https://arxiv.org/abs/1712.01820).
- [21] Laura Mancinska, David E. Roberson, Robert Samal, Simone Severini, and Antonios Varvitsiotis. Relaxations of Graph Isomorphism. In *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 76:1–76:14, 2017. doi:[10.4230/LIPIcs.ICALP.2017.76](https://doi.org/10.4230/LIPIcs.ICALP.2017.76).
- [22] Laura Mančinska and David E. Roberson. Quantum isomorphism is equivalent to equality of homomorphism counts from planar graphs. 2019. arXiv:[1910.06958](https://arxiv.org/abs/1910.06958).
- [23] William J. Martin. Scaffolds: A graph-based system for computations in Bose-Mesner algebras. 2020. arXiv:[2001.02346](https://arxiv.org/abs/2001.02346).
- [24] Hamoon Mousavi, Seyed Sajjad Nezhadi, and Henry Yuen. On the complexity of zero gap mip, 2020. arXiv:[2002.10490](https://arxiv.org/abs/2002.10490).
- [25] Anand Natarajan and John Wright. NEXP is contained in MIP. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 510–518, 2019.
- [26] Miguel Navascués, Stefano Pironio, and Antonio Acín. A convergent hierarchy of semidefinite programs characterizing the set of quantum correlations. *New Journal of Physics*, 10(7):073013, 2008.
- [27] Simon Schmidt and Moritz Weber. Quantum symmetries of graph C^* -algebras. *Canadian Mathematical Bulletin*, 61(4):848–864, 2018. doi:[10.4153/CMB-2017-075-4](https://doi.org/10.4153/CMB-2017-075-4).
- [28] William Slofstra. The set of quantum correlations is not closed. *Forum of Mathematics, Pi*, 7:e1, 2019. doi:[10.1017/fmp.2018.3](https://doi.org/10.1017/fmp.2018.3).
- [29] William Slofstra. Tsirelson’s problem and an embedding theorem for groups arising from non-local games. *Journal of the American Mathematical Society*, 2019. doi:[10.1090/jams/929](https://doi.org/10.1090/jams/929).