The Quantum Smooth Label Cover Problem Is Undecidable

Eric Culf
Department of Applied Mathematics, University of Waterloo, Canada
Institute for Quantum Computing, University of Waterloo, Canada
Perimeter Institute for Theoretical Physics, Waterloo, Canada

Kieran Mastel
Department of Mathematics and Statistics, University of Ottawa, Canada
Department of Pure Mathematics, University of Waterloo, Canada
Institute for Quantum Computing, University of Waterloo, Canada

Connor Paddock[1]
Department of Computer Science, University of Calgary, Canada
Department of Mathematics and Statistics, University of Ottawa, Canada

Taro Spirig
Department of Mathematical Sciences and QMATH, University of Copenhagen, Denmark

Abstract

We show that the quantum smooth label cover problem is undecidable and RE-hard. This sharply contrasts the quantum unique label cover problem, which can be decided efficiently by a result of Kempe, Regev, and Toner (FOCS'08). On the other hand, our result aligns with the RE-hardness of the quantum label cover problem, which follows from the celebrated MIP* = RE result of Ji, Natarajan, Vidick, Wright, and Yuen (ACM'21). Additionally, we show that the quantum oracularized smooth label cover problem is RE-hard. Our second result fits with the alternative quantum unique games conjecture recently proposed by Mousavi and Spirig (ITCS'25) on the RE-hardness of the quantum oracularized unique label cover problem. Our proof techniques include a quantum version of Feige's reduction from 3SAT to 3SAT5 (STOC'96) for BCS-MIP*-protocols, which may be of independent interest.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases Complexity Theory, Constraint Satisfaction Problems, Hardness of Approximation, Quantum Computing

Digital Object Identifier 10.4230/LIPIcs.ICALP.2026.71

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: https://arxiv.org/abs/2510.03477

Funding *Eric Culf*: CGS D scholarship from Canada's NSERC.
Kieran Mastel: CGS D scholarship from Canada's NSERC.
Connor Paddock: NSERC ALLRP-578455-2022
Taro Spirig: ERC grant QInteract (grant No 101078107), VILLUM FONDEN via the QMATH Centre of Excellence (grant No 10059), and Villum Young Investigator (grant No 37532).

Acknowledgements The authors thank the anonymous ICALP reviewers for valuable comments and feedback. They also thank Anand Natarajan and Honghao Fu for helpful discussions. TS thanks Hamoon Mousavi for initial discussions that motivated this work.

[1] Corresponding author

53rd International Colloquium on Automata, Languages, and Programming (ICALP 2026).
Editors: Sayan Bhattacharya, Danupon Nanongkai, Michael Benedikt, and Gabriele Puppis;
Article No. 71; pp. 71:1–71:23
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

The unique games conjecture is one of the most important open problems in theoretical computer science. Provided that P $\neq$ NP, it imposes strict limitations on the existence of efficient algorithms to approximate the value of many NP-hard problems. As an important example, under the unique games conjecture the Goemans and Williamson polynomial-time approximation algorithm for MAXCUT is optimal [17, 24]. Despite enormous attention and effort [36, 27, 26, 11], the conjecture remains unresolved. The conjecture, attributed to Khot [24], informally states that for any $\varepsilon, \delta > 0$, given a large enough alphabet, the unique label cover problem (*i.e.* a unique game), it is NP-hard to decide if the classical value is at least $1 - \varepsilon$ or no more than δ.

Surprisingly, Kempe, Regev, and Toner showed that the quantum analogue of the unique games conjecture is false. In particular, they proved that solutions of a semi-definite program (SDP) relaxation of any unique game can be rounded to entangled strategies that well-approximate the quantum value of the game [22]. At the time, this furthered the idea that the quantum value might be easier to approximate than the classical value of many two-player games. However, in light of the recent MIP* = RE result [20], which implies that even approximating the quantum value of two-player (nonlocal) games is RE-complete, the fact that an efficient approximation algorithm exists for the quantum value of unique games is even more peculiar.

In the classical setting, a source of evidence for the unique games conjecture is the NP-hardness of the label cover and smooth label cover problems [15, 23, 18]. Hardness of the label cover problem – a less structured variant of unique label cover – follows from the PCP theorem [1, 2]. While smooth label cover is more akin to unique label cover, the smooth variant is not an interpolation between the other two. Nonetheless, the NP-hardness of smooth label cover has been a powerful tool in its own right, used to establish hardness of approximation results for many problems in NP [18], even in cases where the unique games conjecture was thought to be necessary.

In the quantum setting of interactive proofs with entangled provers, the story is quite different. Mousavi and Spirig recently pointed out that, by the MIP* = RE result, the analogous quantum label cover problem is *undecidable* and RE-hard [32]. Together with the polynomial time approximation result for quantum unique label cover [22], this leaves us with a natural question: is the quantum smooth label cover problem easy, undecidable, or somewhere in between?

1.1 Results Overview

A two-player one referee game is akin to a two-prover interactive proof system. Here, non-communicating players (or provers) interact with a referee (or verifier) through a single round of question and answer. The players are cooperating and their goal is to win the game by satisfying the rule predicate, a condition checked by the referee. In the quantum setting the players have access to the additional resource of entangled quantum states, and the ability to measure them locally between receiving their questions and returning their answers.

An instance of smooth label cover consists of a bipartite graph G with partitioned vertices $U \cup V$, edges E, and an alphabet of size n. Given an instance, the goal is to assign labels to each of the vertices, maximizing the number of valid assignments; the validity of a label assignment is determined by constraints imposed by the edges of the graph. Smoothness is a certain global-to-local property on the constraints, which ensures the assignments satisfying any near-optimal number of constraints have specific local features. Each smooth label cover problem is an instance of a 2-ary constraint satisfaction problem (2-CSP), with degree n variables corresponding to the vertices, and constraints imposed by edges.

This smooth label cover game can be played as a constraint system nonlocal game $\mathcal{G}_{\mathrm{SLC}(n)}$. In the game, the referee samples a constraint e_{uv} according to some distribution, and sends one variable u to the first player, and the other variable v to the second player. The first player responds with an assignment to variable u, and the second player with an assignment to variable v. Finally, the referee checks that the assignments to u and v satisfy the constraint e_{uv}. The completeness of the game is clear: if they have an assignment which satisfies a large number of constraints, then they are likely to win the game. The notion of soundness is similar: players with a good strategy will be able to construct a good assignment.

The optimal classical value of the smooth label cover game $\mathcal{G}_{\mathrm{SLC}(n)}$, denoted by $\omega_c(\mathcal{G}_{\mathrm{SLC}(n)})$, will be directly related to the optimal fraction of satisfied constraints in the smooth label cover instance. The classical hardness of the smooth label cover problem is the statement that for all $0 < s < 1$, there is an n large enough, such that the following decision problem is NP-hard: given $\mathcal{G}_{\mathrm{SLC}(n)}$ decide if $\omega_c(\mathcal{G}_{\mathrm{SLC}(n)}) = 1$ or $\omega_c(\mathcal{G}_{\mathrm{SLC}(n)}) < s$, promised that one holds[2]. By allowing quantum strategies for the corresponding 2-CSP game, the quantum smooth label cover problem is the following: given $\mathcal{G}_{\mathrm{SLC}(n)}$ and $0 < s < 1$, decide if $\omega_q(\mathcal{G}_{\mathrm{SLC}(n)}) = 1$ or $\omega_q(\mathcal{G}_{\mathrm{SLC}(n)}) < s$, promised that one holds. Where the quantum value $\omega_q(\mathcal{G}_{\mathrm{SLC}(n)})$ is the supremum over all quantum strategies for $\mathcal{G}_{\mathrm{SLC}(n)}$. We establish the following result.

▶ **Theorem** (Informal). *For any soundness parameter, there exists a sufficiently large alphabet such that the quantum smooth label cover problem is* RE*-hard.*

With this result, it seems as though quantum unique label cover problem is the most strange, effectively switching from hard to easy. See Theorem 6 for the formal statement. One can take the perspective that the MIP* = RE result is a quantum version of the classical PCP theorem[3] for succinctly presented constraint satisfaction problems (CSPs). Unlike with the classical MIP = NEXP result [4], computable reductions between problems in RE are not sensitive to whether we have a compact presentation of the CSP. In other words, if a succinctly presented CSP is undecidable, then so is the non-succinct version of the problem. In this sense, the MIP* = RE result of [20] puts significant limitations on the existence of *any* approximation algorithms for the quantum value of nonlocal games. Still, understanding which games admit approximation algorithms for the quantum value is a very interesting problem. Here, little is known compared to the classical case, beyond the setting of XOR games [7, 39, 5], and some other partial results [9].

Mousavi and Spirig point out that the MIP* = RE result implies that the *quantum oracularized* label cover problem is also RE-hard. Oracularizable quantum strategies are a restricted class of quantum synchronous strategies which have convenient properties in the context of reductions between MIP* protocols for CSPs. In particular, they are necessary for the answer reduction portion of the MIP* = RE proof [20, 13]. This led Mousavi and Spirig to provide an alternative quantum unique games conjecture, based on the quantum oracularized value, for which the approximation algorithm of [22] does not readily apply. Assuming their conjecture, they established a hardness of approximation result for the quantum oracularized value for families of CSPs from 2-LIN and MAXCUT which resemble those established in the classical case under the unique games conjecture by Khot and others [24, 25].

Our results for quantum smooth label cover aligns with their quantum unique games conjecture. Moreover, we were able to extend our result to the oracularized setting. Accordingly, our second main result extends the hardness of quantum smooth label cover to

[2] Unlike for unique label cover, smooth label cover with perfect completeness is not necessarily easy.
[3] Although, there are some good reasons to not compare these statements directly, see [33].

the quantum oracularized value. In other words, we show that the problem of deciding if $\omega_{qo}(\mathcal{G}_{\mathrm{SLC}(n)}) = 1$ or $\omega_{qo}(\mathcal{G}_{\mathrm{SLC}(n)}) < s$, is RE-hard, where $\omega_{qo}(\mathcal{G}_{\mathrm{SLC}(n)})$ is the supremum over all quantum oracularizable strategies. For concreteness, one can consider the smooth label cover game, where one of the players must reply with an assignment to both vertices u and v, and the other player obtains only one of u or v chosen at random by the verifier, and replies with an assignment to that vertex.

▶ **Theorem** (Informal)**.** *For any soundness parameter, there is a sufficiently large alphabet such that the quantum oracularized smooth label cover problem is* RE*-hard.*

The formal result can be found in Corollary 13. Oracularizable quantum strategies for 2-CSPs, like smooth label cover, are more akin to quantum strategies for CSPs where each constraint contains more than two variables. In two-player protocols for CSPs where the constraints contain more than two variables, at least one of the players must be forced to respond with an assignments to multiple variables, otherwise the protocol lacks soundness. As such, that player's measurement operators for the variables in that constraint must commute. On the other hand, in a non-oracularizable strategy for a 2-CSP protocol, each player is sent only one variable, and yet the protocol is still sound. Unlike before, the player's measurements for these variables need not commute, since a single player only measures one variable. Such sound noncommutative quantum protocols only exist for 2-CSPs! Making this *constraint-noncommutativity* a behaviour unique to entangled 2-CSP protocols.

In an oracularizable strategy for a 2-CSP game, measurements for the variables in each constraint always commute. This is similar to the quantum strategies for CSP protocols with constraints involving more than 2 variables, such as the constraint-constraint and constraint-variable games, see for example [6, 19, 3, 34, 8]. The existence of non-oracularizable strategies for 2-CSPs is one of the reasons why the authors of [32] suggest a unique games conjecture based on the hardness of the quantum oracularized value, as it fits with the general theory of entangled CSP protocols of higher arity. In particular, the entangled strategies produced by efficient procedure of [22] are not oracularizable and therefore the RE-hardness of the quantum oracularized label cover problem is entirely plausible.

1.2 Technical Outline

Several of the proof ideas in this work are inspired by classical reductions between decision problems. Establishing completeness of a classical reduction in the quantum case is often more straightforward, and it is establishing the quantum soundness which often presents a new challenge. To overcome this obstacle in many cases, we make use of the powerful algebraic formalism for handling soundness arguments for MIP* protocols for CSPs that was recently introduced in [31] and expanded on in [8]. The weighted algebra formalism examines how reductions change the quantum *synchronous* value. One of the main observations used throughout this work (as well as, for instance [31, 8, 16]) is that we can reduce questions about the hardness of the quantum value to the hardness of the quantum synchronous value.

To prove our first main result, we give a reduction from the halting problem to the quantum smooth label cover problem. Our starting point is the RE-hardness for 3SAT* established by Culf and Mastel [8], where 3SAT* denotes the MIP* protocol for 3SAT. Next, we exhibit a reduction from this RE-hard problem to 3SAT5*, the corresponding MIP* protocol where every variable appears in (exactly) 5 constraints. From here, we employ the classical construction of Guruswami, Raghavendra, Saket, and Wu [18], building on the work of Khot [23], to obtain a smooth label cover instance. One of our technical contributions is ensuring the completeness and soundness of the Guruswami, Raghavendra,

Saket, and Wu construction in the quantum case. This involves considering a variant of the (J, R)-dummy variable game (see Definition 7) where the questions and answers are ordered, since completeness of parallel repetition in the quantum setting depends on the players knowing which round of the repeated game they are playing. However, the primary technical contribution of this work is establishing the reduction from 3SAT* to 3SAT5*.

Our approach is inspired by the classical reductions from 3SAT to 3SATB (where every variable appears in B constraints for some constant B) by Papadimitriou and Yannakakis [35] and Arora et al. [1], and then from 3SATB to 3SAT5 by Feige [14]. This involves modifying the 3SAT system by first labelling each variable by its constraint and then adding equality constraints between these labelled variables. One can think of the labelled variables and equality constraints between them as the vertices and edges of a graph. Unfortunately, this reduction does not decrease the degree of the 3SAT instance. To overcome this issue, as in the classical reductions, we modify the construction by adding equality constraints only along the edges of an expander graph. Unlike in the classical setting, analysing the soundness of the expander construction is less straightforward, however we are able to prove that it remains sound in the quantum setting. The key observation is that we can employ a nice result from [21] (Lemma 22) to obtain a constant soundness drop-off in the quantum case.

The soundness proof of the expander graph construction hinges on a certain uniformity assumption for the question distribution of the initial 3SAT* protocol. Hence, to complete the proof of RE-hardness for 3SAT5*, we need to establish RE-hardness for such uniform 3SAT* protocols. Unfortunately, this is not clear a priori via polynomial-time reductions. Nevertheless, we modify the halting problem so that the required uniformity condition is baked in. This modified halting problem is RE-hard by an exponential-time reduction from the halting problem. Then, the RE-hardness of uniform 3SAT* follows by a polynomial-time reduction from the modified halting problem.

Our second main result, establishing undecidability in the oracularized case, follows by a modification of the 3SAT5* protocol described above. This modification results in a 3SAT10* protocol, in which the quantum assignments are what we call *2-oracularizable*. That is, pairs of observables corresponding to variables, say x and y, will commute not only if they are in the same constraint, but also whenever another variable z appears in separate constraints with each of x and y, respectively. The 2-oracularizable property of the quantum assignments to the 3SAT10* protocol ensures that in the construction of the smooth label cover game $\mathcal{G}_{\mathrm{SLC}(n)}$, the quantum strategies will be fully oracularizable. The construction is such that quantum assignments to the 3SAT5* protocol will give assignments to the 3SAT10* protocols. Lastly, the soundness of the construction follows by observing that the entire modification is an instance of a "constraint subdivision" transformation, and therefore quantum soundness follows by a result of [8], using the weighted algebra formalism.

2 Decision Problems for Constraint Satisfaction Games

2.1 Nonlocal Games

A two-player **nonlocal game** $\mathcal{G} = (I, \{O_i\}_{i\in I}, \pi, V)$ consists of a finite set of questions I, a collection of finite answer sets $\{O_i\}_{i\in I}$, a probability distribution π on $I \times I$, and a family of functions $V(\cdot,\cdot|i,j) : O_i \times O_j \to \{0,1\}$ for $(i,j) \in I \times I$. In the game, the players, often called Alice and Bob, receive questions i and j, respectively, from I with probability $\pi(i,j)$, and, without communicating with one another, reply with answers $a \in O_i$ and $b \in O_j$, respectively. They win if $V(a,b|i,j) = 1$ and lose otherwise. We have assumed without loss of generality that the players' question and answer sets are the same.

Although they cannot communicate during the game, the players can coordinate a strategy beforehand. The players' strategy in a nonlocal game is described by the probability distributions $p := p(a, b|i, j)$ of their answers conditioned on the question pair they receive. Such distributions are known as **correlations** for the game $\mathcal{G}$. Depending on the resources available to the players, they can only employ certain correlations. A correlation p is **quantum** if there are:

(i) finite-dimensional Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$,
(ii) a projective measurement $\{M_a^i\}_{a\in O_i}$ on $\mathcal{H}_A$ for every $i \in I$,
(iii) a projective measurement $\{N_a^j\}_{a\in O_j}$ on $\mathcal{H}_B$ for every $j \in I$, and
(iv) a state $|v\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$

such that $p(a, b|i, j) = \langle v|M_a^i \otimes N_b^j|v\rangle$ for all $i, j \in I$, $a \in O_i$, and $b \in O_j$. The collection $(\mathcal{H}_A, \mathcal{H}_B, \{M_a^i\}, \{N_a^j\}, |v\rangle)$ is called a **quantum strategy**. Quantum strategies capture the scenario where the players share some bipartite quantum state and sample their answers via measurements of that state. If the players instead only have access to classical resources such as shared randomness, their correlation and strategy are called **classical**.

The **winning probability** of a correlation p in a nonlocal game $\mathcal{G} = (I, \{O_i\}, \pi, V)$ is

$$\omega(\mathcal{G}; p) := \sum_{i,j\in I} \sum_{a\in O_i, b\in O_j} \pi(i, j)V(a, b|i, j)p(a, b|i, j).$$

Sometimes we denote a strategy by $\mathcal{S}$, and write $\omega(\mathcal{G}; S)$ for the winning probability. Let C_q be the set of quantum correlations, then the **quantum value** of $\mathcal{G}$ is $\omega_q(\mathcal{G}) := \sup_{p\in C_q}(\mathcal{G}; p)$.

A nonlocal game $\mathcal{G} = (I, \{O_i\}, \pi, V)$ is **synchronous** if $V(a, b|i, i) = 0$ for all $i \in I$ and $a \neq b \in O_i$. A correlation p is **synchronous** if $p(a, b|i, i) = 0$ for all $i \in I$ and $a \neq b \in O_i$. The set of synchronous classical and quantum correlations are denoted C_c^s and C_q^s, respectively. We define the synchronous quantum value ω_q^s analogously to ω_q by replacing C_q with C_q^s. If p is a synchronous quantum correlation, then there is a single projective measurement $\{M_a^i\}_{a\in O_i}$ for each $i \in I$ on a single finite dimensional Hilbert space $\mathcal{H}$, and the state $|v\rangle$ is tracial, in the sense that $\langle v|\alpha\beta|v\rangle = \langle v|\beta\alpha|v\rangle$ for all α and β in the $*$-algebra generated by the operators M_a^i, $i \in I$, $a \in O_i$. Moreover, the correlation can be written as $p(a, b|i, j) = \langle v|M_a^i M_b^j|v\rangle$ for all $i, j \in I$, $a \in O_i$, and $b \in O_j$. The collection $(\mathcal{H}, \{M_a^i\}, |v\rangle)$ is called a **synchronous quantum strategy**. A synchronous strategy is called **oracularizable** if $M_a^i M_b^j = M_b^j M_a^i$ for all $i, j \in I$, $a \in O_i$, and $b \in O_j$ with $\pi(i, j) > 0$. Similarly, we call a correlation **oracularizable** if it admits an oracularizable strategy. Let C_{qo} be the set of oracularizable quantum correlations, then the **quantum oracularizable value** of $\mathcal{G}$ is $\omega_{qo}(\mathcal{G}) := \sup_{p\in C_{qo}}(\mathcal{G}; p)$. In some previous work on 2-CSP games, namely [32], the quantum synchronous value of a synchronous game has been called the non-commutative value, and the oracularized synchronous quantum value has been called the quantum value. Every quantum correlation that is approximately synchronous, in the sense that $p(a, b|i, i) \approx 0$ for all $i \in I$ and $a \neq b \in O_i$, is close to an exactly synchronous quantum correlation [38, 28, 30]. Because of this, we work mainly with the synchronous value.

2.2 Constraint System Games and Weighted Algebras

A **constraint system (CS)** over an alphabet Σ consists of a set of variables $\mathcal{X}$ and a collection of constraints $\{\mathcal{C}_i\}_{i=1}^n$, where $\mathcal{C}_i = (\mathcal{U}_i, \mathcal{R}_i)$ for $1 \leq i \leq n$. Each $\mathcal{U}_i$, called a **context**, is an ordered subset of $\mathcal{X}$, and the **relations** $\mathcal{R}_i$ are nonempty subsets of $\Sigma^{\mathcal{U}_i}$. When $|\mathcal{U}_i| = k$ we say that $\mathcal{C}$ is a k-ary constraint, and we write k-CS to denote a CS where each constraint has arity k. An **assignment** to $\mathcal{U}_i$ is an element of $\Sigma^{\mathcal{U}_i}$. Given a constraint $\mathcal{C}_i$ the **satisfying assignments** are the elements of $\mathcal{R}_i$, and assignments not

in $\mathcal{R}_i$ are **unsatisfying assignments**. In the special case where the alphabet Σ has size two, we say S is a boolean constraint system (BCS). A **constraint language** is the family of constraint systems where each constraint comes from a restricted set of relations. An important example of a constraint language is 3SAT, where each context contains three variables and each constraint is a disjunction of three literals. Another example we mention here is the language of fixed degree 3SAT, which we denote by 3SATk for $k > 0$, where each variable appears in exactly k contexts.

A **constraint satisfaction problem** with completeness c and soundness s, for $0 < s < c \leq 1$, is a promise problem:

Given a CS $S = (\mathcal{X}, \{\mathcal{C}_i\}_{i=1}^n)$ where $|\mathcal{U}_i| = k$ for all $1 \leq i \leq n$ and $|\Sigma| = m$:
(*yes*) there exists a $\phi \in \Sigma^{\mathcal{X}}$ which satisfies at least cn constraints,
(*no*) each assignment $\phi \in \Sigma^{\mathcal{X}}$ satisfies less than sn constraints. (k-CSP$(m)_{c,s}$)

Similarly, a **succinct constraint satisfaction problem** with completeness c and soundness s, for $0 < s < c \leq 1$, is a promise problem:

Given a probabilistic Turing machine M that samples the constraints of a CS $S = (X, \{(\mathcal{U}_i, \mathcal{R}_i)\}_{i=1}^n)$ according to some probability distribution $\pi : [n] \to [0, 1]$:
(*yes*) there is an assignment $f : X \to \Sigma$ such that $\Pr_{i \leftarrow \pi}[f|_{\mathcal{U}_i} \in \mathcal{R}_i] \geq c$,
(*no*) for every assignment $\Pr_{i \leftarrow \pi}[f|_{\mathcal{U}_i} \in \mathcal{R}_i] < s$ (SuccinctCSP$(m)_{c,s}$)

For convenience later, we write SuccinctCSP(3SAT)$_{c,s}$ and CSP(3SAT)$_{c,s}$ for the succinct (and non-succinct) constraint satisfaction problems where each CS is a 3SAT instance. Additionally, it will be convenient to make use of the following graph associated with a CS, which exhibits which variables appear in which constraints in a constraint system.

▶ **Definition 1.** *Let $S = (\mathcal{X}, \{\mathcal{C}_i\}_{i=1}^n)$ be a CS. The **connectivity graph** of S is the bipartite graph $\Gamma(S)$ with left vertices $V_L(\Gamma(S)) = \mathcal{X}$, right vertices $V_R(\Gamma(S)) = [n]$, and edges $(x, i) \in E(\Gamma(S))$ if and only if $x \in \mathcal{U}_i$.*

Given a distribution $\mu : [n] \to \mathbb{R}_{\geq 0}$ and a CS S, the **constraint-variable CS game** is the nonlocal game $\mathcal{G}(S, \mu) = ([n] \cup \mathcal{X}, \{\mathcal{O}_i\}_{i \in [n] \cup \mathcal{X}}, \pi, V)$, where $O_i = \Sigma^{\mathcal{U}_i}$ when $i \in [n]$ and $O_i = \Sigma$ otherwise (i.e. $i \in \mathcal{X}$), $\pi(i, x) = \mu(i)/|\mathcal{U}_i|$ for $x \in \mathcal{U}_i$ and is 0 on other inputs, and $V(\phi, a|i, x) = 1$ if $\phi \in \mathcal{C}_i$ and $\phi(x) = a$ and is 0 otherwise. Note that although the players have different question sets, the game can be symmetrized without changing the synchronous winning probability. In this game, one player receives $i \in [n]$ sampled from μ and the other receives a uniformly random variable $x \in \mathcal{U}_i$. To win, the first player must answer with a satisfying assignment $\phi \in \mathcal{C}_i$ and the second must answer $a \in \Sigma$ such that $a = \phi(x)$. When S is boolean, we call that $\mathcal{G}(S, \pi)$ a BCS game.

The constraint-variable game is not naturally synchronous. To make it synchronous, the verifier must randomly choose which player to ask the constraint question and which player to ask the variable question, and also ask consistency check questions with some constant probability. This transformation preserves constant completeness-soundness gaps and guarantees that the players can win near-optimally with synchronous strategies due to Theorem 0.1 of [30], so we do not need to worry about it in practice.

Let S be a 2-CS over an alphabet Σ and π' be a probability distribution on $[m]$, define the **2-CS game** $G_a(S, \pi')$ as the nonlocal game $(\mathcal{X}, \Sigma, \nu_a, V_a)$, where $\nu_a(x, y) = \frac{\pi'(i)}{2}$ if $\mathcal{U}_i = \{x, y\}$ and 0 otherwise, and $V_a(a, b|x, y) = 1$ iff there exists $\phi \in \mathcal{R}_i$ such that $\phi(x) = a$

and $\phi(y) = b$. In this game, the referee samples i from π', and then asks each of the players one variable from the constraint. Each player responds with an assignment to the variable she received. They win if they have answered a satisfying assignment.

If the players are classical, then their strategy for either type of CS game can be taken to be deterministic. The responses of the player who received the variable question constitute an assignment to the constraint system. Similarly, the variable player's measurements in a synchronous strategy constitute a **quantum (synchronous) assignment** to the CS. If the corresponding strategy is oracularizable, then the assignment is called an **oracularizable assignment**. One can reason about reductions between classical strategies for CS games by working with global assignments. When the players have access to quantum resources, we can get a similar simplification when thinking about synchronous strategies using the weighted algebra formalism developed in [31] and [8].

For a set of variables $\mathcal{U}$, let $\mathbb{C}\mathbb{Z}_k^{*\mathcal{U}}$ be the free algebra generated by order-k unitaries labelled by the elements $x \in \mathcal{U}$, and let $\mathbb{C}\mathbb{Z}_k^{\mathcal{U}}$ be its abelianization. For any order-k unitary x and $a \in \mathbb{Z}_k$ with eigenvalues $\{\omega_k^a\}_{a\in\mathbb{Z}_k}$, write $\Pi_a^{(k)}(x)$ for the projector onto the ω_k^a-eigenspace of x; where k is clear, we suppress the superscript (k). Given $\phi \in \mathbb{Z}_k^{\mathcal{U}}$, we define the element $\Phi_{\mathcal{U},\phi}^{(k)} \in \mathbb{C}\mathbb{Z}_k^{*\mathcal{U}}$ as $\Phi_{\mathcal{U},\phi}^{(k)} = \prod_{x\in V} \Pi_{\phi(x)}^{(k)}(x)$, where as before the superscript k is suppressed where clear. We use the same notation for the image of $\Phi_{\mathcal{U},\phi}^{(k)}$ under any homomorphism when the homomorphism is clear. Given a constraint $(\mathcal{U}, \mathcal{R})$ over the alphabet $\mathbb{Z}_k$, we write $\mathcal{A}(\mathcal{U}, \mathcal{R}) = \mathbb{C}\mathbb{Z}_k^{\mathcal{U}} / \langle \Phi_{\mathcal{U},\phi} : \phi \notin \mathcal{R} \rangle$. The algebra $\mathcal{A}(\mathcal{U}, \mathcal{R})$ is isomorphic to the C^*-algebra of functions on the finite set $\mathcal{R}$. Consequently, if $\rho : \mathcal{A}(\mathcal{U}, \mathcal{R}) \to \mathcal{B}(\mathcal{H})$ is a $*$-representation, then $\{\rho(\Phi_{\mathcal{U},\phi})\}_{\phi\in\mathcal{R}}$ is a projective measurement on $\mathcal{H}$, and conversely if $\{M_\phi\}_{\phi\in\mathcal{R}}$ is a projective measurement on $\mathcal{H}$ then there is a $*$-representation $\rho : \mathcal{A}(\mathcal{U}, \mathcal{R}) \to \mathcal{B}(\mathcal{H})$ with $\rho(\Phi_{\mathcal{U},\phi}) = M_\phi$.

Let $S = (\mathcal{X}, (\mathcal{U}_i, \mathcal{R}_i)_{i=1}^n)$ be a CS. The **constraint-variable algebra** is the free product $\mathcal{A}_{c-v}(S) = \ast_{i=1}^m \mathcal{A}(\mathcal{U}_i, \mathcal{R}_i) * \mathbb{C}\mathbb{Z}_k^{*\mathcal{X}}$. Let $\sigma_i : \mathcal{A}(\mathcal{U}_i, \mathcal{R}_i) \to \mathcal{A}_{c-v}(S)$ be the inclusion on the i^{th} factor, and let $\sigma' : \mathbb{C}\mathbb{Z}_k^{*\mathcal{X}} \to \mathcal{A}_{c-v}(S)$ be the inclusion on the final factor. To keep our formulas tidy, we'll often omit the σ_i and σ' when it's clear what subalgebra $\mathcal{A}(\mathcal{U}_i, \mathcal{R}_i)$ the element belongs to. By the GNS representation theorem, there is a correspondence between finite dimensional tracial states on the constraint variable algebra and synchronous quantum strategies for the constraint-variable game. A correlation $p \in C_q^s$ if and only if there is a finite dimensional tracial state τ on $\mathcal{A}_{c-v}(S)$ such that $p(\phi, a|i, x) = \tau(\Phi_{\mathcal{U}_i,\phi}\Pi_a(\sigma'(x)))$. A tracial state τ on $\mathcal{A}_{c-v}(S)$ is **perfect** if it corresponds to a perfect correlation. This happens if and only if $\tau(\Phi_{\mathcal{U}_i,\phi}\Pi_a(\sigma'(x))) = 0$ whenever $\phi(x) \neq a$. We are now ready to define the weight and defect, which gives us a way of tracking how close players' strategies are to being perfect at the level of tracial states on algebras.

▶ **Definition 2** ([31]). *A **(finitely-supported) weight function** on a set X is a function $\mu : X \to [0, +\infty)$ such that $\mathrm{supp}(\mu) := \mu^{-1}((0, +\infty))$ is finite. A **weighted $*$-algebra** is a pair $(\mathcal{A}, \mu)$ where $\mathcal{A}$ is a $*$-algebra and μ is a weight function on $\mathcal{A}$.*

*If τ is a tracial state on $\mathcal{A}$, then the **defect of τ** is*

$$\mathrm{def}(\tau; \mu) := \sum_{a\in\mathcal{A}} \mu(a)\|a\|_\tau^2,$$

where $\|a\|_\tau := \sqrt{\tau(a^*a)}$ is the τ-norm. When the weight function is clear, we just write $\mathrm{def}(\tau)$.

Since μ is finitely supported, the sum in the definition of the defect is finite, and hence is well-defined. The support of the weight is chosen so that if $\mathrm{def}(\tau; \mu) = 0$, then τ is perfect.

For a probability distribution π' on $[m]$, define the **weighted constraint-variable algebra** to be $\mathcal{A}_{c-v}(S,\pi') = (\mathcal{A}_{c-v}(S), \mu_{c-v,\pi'})$ where the weight function $\mu_{c-v}(\Phi_{V_i,\phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))) = \frac{\pi'(i)}{|\mathcal{U}_i|}$ for all $x \in \mathcal{U}_i$, $\phi \in \mathcal{R}_i$, and 0 on all other elements. Similarly, Given a probability distribution π' on $[m]$, we define the **weighted assignment algebra** to be $\mathcal{A}_a(S) = \mathbb{C}\mathbb{Z}_k^{*\mathcal{X}}$, along with the weight function $\mu_{a,\pi'}(\Phi_{\mathcal{U}_i,\phi}) = \pi'(i)$ for all $\phi \notin \mathcal{R}_i$, and 0 on all other elements.

The point of these algebras is the following lemma.

▶ **Lemma 3** ([8]). *Let $S = (\mathcal{X}, \{\mathcal{C}_i\}_{i=1}^m)$ be a CS, and let μ be a probability distribution on $[m]$. A tracial state τ on $\mathcal{A}_{c-v}(S)$ is an ε-perfect strategy for $\mathcal{G}(S,\pi)$ if and only if $\mathrm{def}(\tau) \leq \varepsilon$. If S is a 2-CS, then a tracial state τ_a on $\mathcal{A}_a(S)$ is an ε-perfect strategy for $G_a(S,\pi)$ if and only if $\mathrm{def}(\tau_a) \leq \varepsilon$*

See [8] for a more complete treatment of the weighted algebra formalism.

3 The quantum smooth label cover game

A **two-prover one-round MIP protocol** consists of a probabilistic Turing machine Q and another Turing machine V, along with a family of nonlocal games $G_x = (I_x, \{O_{xi}\}_{i\in I_x}, \pi_x, V_x)$ for $x \in \{0,1\}^*$, such that

- on input x, the Turing machine Q outputs $(i,j) \in I \times I$ with probability $\pi_x(i,j)$, and
- on input (x,a,b,i,j), the Turing machine V outputs $V_x(a,b|i,j)$.

If the bit lengths of the questions and answers are bounded by functions $q(n)$ and $a(n)$ of $n = |x|$, respectively, and $\{G_x\}_x$ is a set of games indexed by strings x, we denote the problem of deciding whether the quantum value of the game $\omega_q(G_x) \geq c$ or $\omega_q(G_x) < s$ by $\mathrm{MIP}^*(q(n),a(n))_{c,s}$. In general for MIP protocols, the completeness parameter c, and the soundness parameter s, may also depend on n in an efficiently computable way. When we restrict the nonlocal games to be BCS games, we denote the decision problem by $\text{BCS-MIP}^*(q(n),a(n))_{c,s}$. When we further restrict the BCS games to be constraint-variable 3SAT games we denote the decision problem by $\mathrm{CSP}_{c-v}(3\mathrm{SAT})^*_{c,s}$ and we write $\mathrm{SuccinctCSP}_{c-v}(3\mathrm{SAT})^*_{c,s}$ for the succinctly presented version.

As mentioned in the previous section, the synchronous quantum value and the quantum value of synchronous nonlocal games are closely related. In our proofs of hardness, we deal with the synchronous quantum value as it is easier to reason about mathematically. This simplification is okay, since by Theorem 0.1 of [30], for constant s, the hardness of deciding if $\omega_q^s(G) = 1$, or $\omega_q^s(G) < s$ for a synchronous game will be the same as that of deciding whether $\omega_q(G) = 1$, or $\omega_q(G) < s'$ where $1 > s' > 0$ is also a constant.

Let $\mathcal{P}$ be an MIP protocol, and let α and β be functions from instances of $\mathcal{P}$ to $[0,1]$. Denote by $(\alpha,\beta,\mathcal{P})_{c,s}$ the decision problem where G is a yes instance if $\alpha(G) \geq c$ and a no instances if $\beta(G) < s$ and $\alpha(G) < c$, with the promise that one of these is the case. For a synchronous nonlocal game G, we have

$$\omega_c(G) \leq \omega_{qo}(G) \leq \omega_q^s(G), \tag{1}$$

which implies the trivial reductions in Figure 1.

▶ Remark 4. Note that there are always trivial reductions from relaxing β in the definition of no instances to $\beta' \leq \beta$. However, we don't have trivial reductions from relaxing α to $\alpha' \geq \alpha$ without the assumption that $\alpha' \leq \beta$, as it could be possible that $\alpha \leq \beta$ and $\alpha' \geq c$. In particular, the intersection of the no instances of $(\alpha,\beta,\mathcal{P})_{c,s}$ and the yes instances of $(\alpha',\beta,\mathcal{P})_{c,s}$ could be non-empty. The definitions of gapped decision problems appearing in [32] are incorrect, as we need to ensure that the yes and no instances are disjoint.

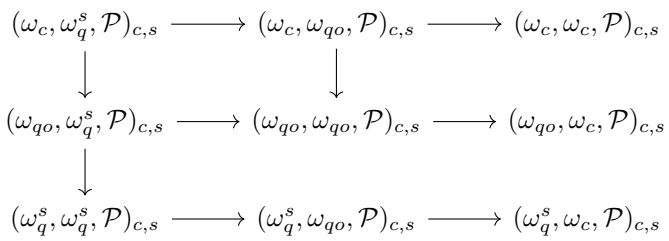

Figure 1 Trivial reductions between decision problems for the MIP protocol $\mathcal{P}$.

The main decision problems that we study centre around the 2-CS smooth label cover. Although we introduced the smooth label cover problem in the context of bipartite graphs, the problem can be generalized to any regular connected graph. That is, an instance of **Smooth Label Cover** is given by a tuple $(G, [n], [k], \Sigma)$ consisting of a (undirected) regular connected graph $G = (V, E)$, a label set $[n]$ (for positive integer n), and a set

$$\Sigma = \big((\pi_{ev}, \pi_{ew}) : e = (v, w) \in E\big)$$

consisting of pairs of maps both from $[n]$ to $[k]$ associated with the endpoints of the edges in E. Given an **assignment** $A : V \to [n]$, we say that an edge $e = (v, w) \in E$ is **satisfied** if $\pi_{ev}\big(A(v)\big) = \pi_{ew}\big(A(w)\big)$. For a given $J \geq 1$, the maps π_{ev} are ***J*-smooth** (or just **smooth** when $J \gg 1$) in the sense that, for every vertex v and every pair $a, a' \in [n]$ with $a \neq a'$, we have

$$\Pr_{w \sim v}[\pi_{ev}(a) \neq \pi_{ew}(a')] \geq 1 - \frac{1}{J}.$$

Let $\mathrm{SLC}_{\alpha,\beta}(c, s)$ be the MIP protocol in which every instance is a 2-CS game $\mathcal{G}$ for a smooth label cover instance. The game $\mathcal{G}$ is a yes instance if $\alpha(\mathcal{G}) \geq c$, and a no instance if $\beta(\mathcal{G}) < s$. When $\alpha = \beta$, we write $\mathrm{SLC}_{\alpha}(c, s)$.

3.1 The Hardness of Quantum Smooth Label Cover

In this section, we establish the RE-hardness of smooth label cover. Our reduction is based on the following hardness result involving the quantum value of a 3SAT5* instance.

▶ **Theorem 5.** *There exists a constant $0 < s < 1$ such that it is* RE*-hard to distinguish whether the clause-variable game associated to a* 3SAT5 *instance has quantum value 1 or at most s.*

We reserve the proof until Section 4. Using Theorem 5, we will prove the following theorem.

▶ **Theorem 6.** *For any $0 < s < 1$, there exists large enough alphabets over which $\mathrm{SLC}_q(1, s)$ is* RE*-hard.*

The proof can be seen as a quantum analog of Theorem 3.5 in [18]. In particular, the proof of Theorem 6 is based on the construction in Section 2.2 of [23], and in Appendix A of [18]. Before we give the proof, we require several intermediary results and definitions.

▶ **Definition 7.** *Let $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a BCS. The* (J, R)-dummy clause-variable game *is the nonlocal game $\mathcal{G}(B, J, R)$ with question set*

$$I = [m]^{(J+1)R} \cup \left\{ \vec{i} \in (X \cup [m])^{(J+1)R} \mid |\{q \mid i_q \in X\}| = R \right\},$$

answer set

$$O_{\vec{i}} = \bigtimes_{q=1}^{(J+1)R} \begin{cases} \{0,1\} & i_q \in X \\ C_{i_q} & i_q \in [m] \end{cases},$$

probability distribution $\pi(\vec{i},\vec{j})$*, obtained by sampling* $\vec{i} \in [m]^{(J+1)R}$ *uniformly and then sampling* $\vec{j}$ *by sampling a subset* $L \subseteq [(J+1)R]$ *of size* R *and* $x_q \in V_{i_q}$ *for all* $q \in L$ *uniformly and then letting* $j_q = x_q$ *if* $q \in L$ *and* $j_q = i_q$ *if* $q \notin R$*, and predicate* $V(\vec{\sigma},\vec{\tau}|\vec{i},\vec{j}) = 1$ *iff* $\sigma_q = \tau_q$ *if* $i_q = j_q$ *and* $\sigma_q(j_q) = \tau_q$ *else.*

Verifying the predicate function can be simplified by introducing the projection $\pi^{\vec{j},\vec{i}} : O_{\vec{i}} \to O_{\vec{j}}$ *defined as*

$$\pi^{\vec{j},\vec{i}}(\vec{\sigma})_q = \begin{cases} \sigma_q & j_q = i_q \\ \sigma_q(j_q) & j_q \in V_{i_q} \end{cases}.$$

Then $V(\vec{\sigma},\vec{\tau}|\vec{i},\vec{j}) = \delta_{\vec{\tau},\pi^{\vec{j},\vec{i}}(\vec{\sigma})}$.

The purpose of the (J,R)-dummy variable game is twofold. It both amplifies the soundness of the original BCS game and ensures the smoothness property with respect to any assignments of the BCS. We start by establishing how the (J,R)-dummy variable game exponentially amplifies the soundness with respect to the quantum synchronous value.

▶ **Lemma 8.** *If the quantum value of the constraint-variable game of a BCS* B *is* < 1*, there exists* $\eta(R) = O(\exp(-R))$ *such that the quantum value of* $\mathcal{G}(B,J,R)$ *is upper-bounded by* $\eta(R)$*. If the synchronous quantum value of the constraint-variable game of a BCS* B *is* 1*, then the synchronous quantum value of* $\mathcal{G}(B,J,R)$ *is* 1*.*

Proof. Consider the modified version of $\mathcal{G}(B,J,R)$ where the predicate only checks that $\sigma_q(j_q) = \tau_q$ for those q such that $j_q \in X$, and does not check for consistency on the remaining q. Then, it is easy to see that the quantum value of this game is an upper bound on the quantum value of $\mathcal{G}(B,J,R)$; and that this game is equivalent to the R-fold parallel repetition of the constraint-variable game of B. Hence, if the quantum value of a constraint-variable game of a BCS B is less than 1, then by the parallel repetition theorem for entangled projection games [12], the quantum value of the parallel repetition is upper-bounded by $O(\exp(-R))$. Hence, the same bound applies to $\mathcal{G}(B,J,R)$.

Now, consider the case that the synchronous quantum value of the constraint-variable game of B is 1. Then there exists a perfect (quantum-approximate) strategy for this game consisting of Alice's PVMs $\{A_i^\sigma\}_{\sigma \in C_i}$ for $i \in [m]$ and Bob's PVMs $\{B_x^b\}_{b \in \{0,1\}}$, and a tracial state tr. Then, construct the following strategy for $\mathcal{G}(B,J,R)$: let the PVMs $P_{\vec{i}}^{\vec{\sigma}} = P_{i_1}^{\sigma_1} \otimes \cdots \otimes P_{i_{(J+1)R}}^{\sigma_{(J+1)R}}$ where $P_{i_q}^{\sigma_q} = A_{i_q}^{\sigma_q}$ if $i_q \in [m]$ and $P_{i_q}^{\sigma_q} = B_{i_q}^{\sigma_q}$ if $j_q \in X$. With respect to the tracial state $\mathrm{tr}^{\otimes(J+1)R}$, it is clear that this is a perfect synchronous quantum approximate strategy, since it passes each of the consistency checks perfectly term-by-term. ◀

▶ **Corollary 9.** *There exists an inverse exponential* $\eta(R)$ *such that it is* RE*-hard to decide whether* $\mathcal{G}(B,J,R)$ *for a* 3SAT5 *instance* B *has synchronous quantum value* 1 *or quantum value that is* $\leq \eta(R)$*.*

Proof. By Theorem 5, there exists $s < 1$ such that it is RE-hard to decide if the synchronous quantum value of the constraint-variable game of a 3SAT5 instance B is 1 or $< s$. By Theorem 0.1 of [30], there is a constant s' such that it is RE-hard to decide if the quantum value of a 3SAT5 is 1 or $< s'$. Using Lemma 8 gives the wanted reduction. ◀

▶ **Definition 10.** *Let $\mathcal{G} = (I, \{O_i\}, \mu, V)$ be a nonlocal game. The* connectivity graph of $\mathcal{G}$ *is the bipartite graph $\Gamma(\mathcal{G})$ with left vertices $V_L(\Gamma(\mathcal{G})) = \left\{i \in I \mid \sum_j \mu(i,j) > 0\right\}$, right vertices $V_R(\Gamma(\mathcal{G})) = \{j \in I \mid \sum_i \mu(i,j) > 0\}$, and edges $(i,j) \in E(\Gamma(\mathcal{G}))$ if and only if $\pi(i,j) > 0$.*

Let u_n be the uniform distribution on the constraints of a CS S. The connectivity graph of the constraint-variable game $\mathcal{G}(S, \mathrm{u}_n)$ is equal to the connectivity graph of the CS S (see Definition 1). Hence, the connectivity graph $\Gamma(\mathcal{G}(B,J,R))$ for B a 3SAT5 instance with variables X has $n = m^{(J+1)R}$ left vertices and $k = \binom{(J+1)R}{R}|X|^R m^{JR}$ right vertices. In the following, we often equate the left vertices of $\Gamma(\mathcal{G}(B,J,R))$ with $[n]$ and the right vertices with $[k]$.

As mentioned earlier the other role of J and R is to ensure J-smoothness of the assignments to the newly constructed game $\mathcal{G}(B,J,R)$. The proof follows along the lines of the lemma from [23].

▶ **Lemma 11** ([23], Lemma 2.4)**.** *For fixed $\vec{i} \in [m]^{(J+1)R}$ and any two distinct assignments $\vec{\sigma}_1$ and $\vec{\sigma}_2$ to $\vec{i}$,*

$$\Pr_{\vec{j}}\left[\pi^{\vec{j},\vec{i}}(\vec{\sigma}_1) \neq \pi^{\vec{j},\vec{i}}(\vec{\sigma}_2)\right] \geq 1 - \frac{1}{J},$$

where $\vec{j}$ is distributed according to the conditional distribution of μ for fixed $\vec{i}$, i.e. $\Pr[\vec{j} = \vec{j}_0] = \mu(\vec{i},\vec{j}_0)/\sum_{\vec{k}} \pi(\vec{i},\vec{k})$.

Proof. By assumption, there exists some $q_0 \in [(J+1)R]$ such that $(\sigma_1)_{i_{q_0}} \neq (\sigma_2)_{i_{q_0}}$. As such, if $j_{q_0} = i_{q_0}$, then $\pi^{\vec{j},\vec{i}}(\vec{\sigma}_1) \neq \pi^{\vec{j},\vec{i}}(\vec{\sigma}_2)$, since the projection preserves the q_0-th term. Let $L = \{q \mid j_q \neq i_q\} \subseteq [(J+1)R]$ be a set-valued random variable. Then, we have the lower bound $\Pr_{\vec{j}}\left[\pi^{\vec{j},\vec{i}}(\vec{\sigma}_1) \neq \pi^{\vec{j},\vec{i}}(\vec{\sigma}_2)\right] \geq \Pr[q_0 \notin L]$. The number of subsets of size R in $[(J+1)R]$ is $\binom{(J+1)R}{R}$ and the number that do no include q_0 is $\binom{(J+1)R-1}{R}$, so the probability

$$\Pr[q_0 \notin L] = \frac{\binom{(J+1)R-1}{R}}{\binom{(J+1)R}{R}} = \frac{((J+1)R-1)!(JR)!R!}{(JR-1)!R!((J+1)R)!} = \frac{JR}{(J+1)R} = 1 - \frac{1}{J+1} \geq 1 - \frac{1}{J}. \quad ◀$$

We state the following lemma for later use.

▶ **Proposition 12.** *For any $\vec{i} \in [m]^{(J+1)R}$ and $\vec{j}$ such that $\mu(\vec{i},\vec{j}) > 0$, we have for any assignment $\vec{\tau}$ to $\vec{j}$ that $|(\pi^{\vec{j},\vec{i}})^{-1}(\tau)| \leq 4^R$.*

Proof. For any $x \in \{0,1\}$, there are at most 4 ways to extend it to a satisfying 3SAT clause. Hence, since there are R indices of τ to be extended to a satisfying assignment to $\vec{i}$, there are at most 4^R ways to do it. ◀

Proof of Theorem 6. Our goal is to reduce 3SAT5* to $\mathrm{SLC}_q(1,s)$. We start by constructing a Smooth Label Cover instance $(G,[n],[k],\Sigma)$ along the lines of [18, Appendix A]. As above, consider a bipartite graph $\Gamma(\mathcal{G}(B,J,R))$. Associated to each edge $(\vec{i},\vec{j})$ there is a map $\pi^{\vec{j},\vec{i}}$ which satisfies the smoothness property (Lemma 11). Moreover, the graph $\Gamma(\mathcal{G}(B,J,R))$ is bi-regular as the 3SAT5 instances we started with have the property that every variable appears exactly in 5 clauses, giving that right degree is 5^R. For each fixed left vertex, one must choose which of R coordinates are chosen from $(J+1)R$ and then which of the 3 variables are chosen from each clause. Hence, the left degree is $\binom{(J+1)R}{R}3^R$.

Now, let V be the set of left vertices of $\Gamma(\mathcal{G}(B,J,R))$, and connect $\vec{i},\vec{i'} \in V$ by an edge labelled by $\vec{j}$ if $\vec{j}$ is a right vertex that is adjacent to both of them in $\Gamma(\mathcal{G}(B,J,R))$. For each edge $e=\{\vec{i},\vec{i'}\}$, define the projection maps $\pi_{e,\vec{i}} = \pi_{\vec{j},\vec{i}}$ and $\pi_{e,\vec{i'}} = \pi_{\vec{j},\vec{i'}}$, where $\vec{j}$ is the corresponding right vertex. By Lemma 11, we see that smoothness property holds. That is, for every vertex $v \in V$ and distinct $i,j \in [n]$, we have

$$\Pr_{w\sim v}\left[\pi_{ev}(i) = \pi_{ew}(j)\right] \leq \frac{1}{J}\,. \tag{2}$$

We move on to show that it is RE-hard to distinguish between the following two cases: $\omega_q(\mathfrak{S}) = 1$, or $\omega_q^s(\mathfrak{S}) \leq 1/\exp(R)$. For completeness, let $\mathcal{S}$ denote a quantum strategy winning perfectly at the synchronous version of the clause variable (J,R)-dummy game associated to B as in Section 2.1. Since the game is synchronous, there must exist a perfect quantum synchronous assignment $\mathcal{S}_s$. Let us denote it $\mathcal{S}_s = (P_u, Q_v)$, $P_u = \{P_u^l\}_{l\in[k]}$ is a PVM for every $u \in U$ and $Q_v = \{Q_v^i\}_{i\in[n]}$ is a PVM for every $v \in V$. Then we have

$$1 = \omega_q(\mathcal{G}(B,J,R),\mathcal{S}_s) = \mathbb{E}_{(u,v)\in E}\sum_{i\in[n]} \operatorname{tr}\left(P_u^{\pi_{v,u}(i)} Q_v^i\right) \tag{3}$$

Now consider the quantum assignment for $\mathfrak{S}$ given by the operators $\mathcal{S}' = \{Q_v : v \in V\}$. We compute its value:

$$\omega_q(\mathfrak{S},\mathcal{S}') = \mathbb{E}_{u\in U, v_1,v_2\in N_u} \sum_{\substack{i,j\in[n]:\\ \pi_{v_1,u}(i)=\pi_{v_2,u}(j)}} \operatorname{tr}\left(Q_{v_1}^i Q_{v_2}^j\right) = \mathbb{E}_{u\in U, v_1,v_2\in N_u}\sum_{l\in[k]} \operatorname{tr}\left(\widetilde{Q}_{v_1}^l \widetilde{Q}_{v_2}^l\right),$$

where we defined $\widetilde{Q}_{v_1}^l = \sum_{i\in\pi_{v_1,u}^{-1}(l)} Q_{v_1}^i$ and similarly for $\widetilde{Q}_{v_2}^l$.

Now we note that $\{\widetilde{Q}_{v_1}^l\}_{l\in[k]}$ defines a PVM. Indeed, $\sum_{l\in[k]} \widetilde{Q}_{v_1}^l = \sum_l \sum_{i\in\pi_{v_1,u}^{-1}(l)} Q_{v_1}^i = 1$. Moreover, since $\{Q_v^i\}_{i\in[n]}$ forms an orthonormal set of projections, it is clear that $\{\widetilde{Q}_{v_1}^l\}_{l\in[k]}$ will as well form an orthonormal set of projections. Therefore, we can use Lemma 34 of [32]:

$$\omega_q(\mathfrak{S},S) \geq 2\mathbb{E}_{u\in U, v_1,v_2\in N_u}\sum_{l\in[k]} \operatorname{tr}\left(\widetilde{Q}_{v_1}^l P_u^l + \widetilde{Q}_{v_2}^l P_u^l\right) - 3 = 1,$$

where N_u is the set of vertices adjacent to u, and we used that

$$\mathbb{E}_{u\in U, v_1,v_2\in N_u}\sum_{l\in[k]} \operatorname{tr}\left(\widetilde{Q}_{v_1}^l P_u^l\right) = \mathbb{E}_{(u,v_1)\in E}\sum_{i\in[n]} \operatorname{tr}\left(Q_{v_1}^i P_u^{\pi_{v_1,u}(i)}\right),$$

along with the assumption from equation (3). This completes the proof of completeness.

For soundness, let $\mathcal{L} = (\{Q_v^i\}_{i\in[n]})$ be a quantum synchronous assignment for $\mathfrak{S}$ which satisfies:

$$\omega_q(\mathfrak{S},\mathcal{L}) = \mathbb{E}_{u\in U, v_1,v_2\in N_u} \sum_{\substack{i,j\in[n]:\\ \pi_{v_1,u}(i)=\pi_{v_2,u}(j)}} \operatorname{tr}\left(Q_{v_1}^i Q_{v_2}^j\right) > s, \tag{4}$$

for some $s < 1$. Define the operators $P_u^l = \mathbb{E}_{v\in N_u}\sum_{i\in\pi_{e,v}^{-1}(l)} Q_v^i$. These operators form a POVM as they are clearly positive and satisfy

$$\sum_{l\in[k]} P_u^l = \mathbb{E}_{v\in N_u}\sum_l \sum_{i\in\pi_{e,v}^{-1}(l)} Q_v^i = 1,$$

as $\{Q_v^i\}_i$ is a PVM. Using Naimark's dilation theorem, we can replace P_u with a PVM that performs at least as well. Now, consider the strategy $\mathcal{S} = (P_u, Q_v)$ for $\mathcal{G}(B, J, R)$. This assignment has value:

$$\omega_q(\mathcal{G}(B,J,R),\mathcal{S}) = \mathbb{E}_{(u,v_1)\in E} \sum_{i\in[n]} \operatorname{tr}\left(P_u^{\pi_{v_1,u}(i)} Q_{v_1}^i\right) = \mathbb{E}_{v_1,v_2\in N_u \atop u\in U} \sum_{\substack{i,j\in[n]: \\ \pi_{v_1,u}(i)=\pi_{v_2,u}(j)}} \operatorname{tr}\left(Q_{v_1}^i Q_{v_2}^j\right) > s,$$

using the assumption from equation (4). By Lemma 8 we conclude that the $s \leq 1/\exp(R)$ and the claim follows. ◀

3.2 The Hardness of Quantum Oracularized Smooth Label Cover

In this section, we show an analog of Theorem 6 where we consider the quantum oracularized value instead of the quantum synchronous value.

▶ **Corollary 13.** *For any $0 < s < 1$, there exists large enough alphabets over which* $\mathrm{SLC}_{qo}(1, s)$ *is* RE*-hard.*

This result follows directly from the following theorem and the trivial reduction from $\mathrm{SLC}_{qo,q}(1, s)$.

▶ **Theorem 14.** *For any $0 < s < 1$, there exists large enough alphabets over which* $\mathrm{SLC}_{qo,q}(1, s)$ *is* RE*-hard.*

To prove this theorem, we start from a slightly altered version of Theorem 5. In particular, the following lemma shows that we can alter the 3SAT5 game so that the resulting optimal strategies for the SLC instance obtained by the reduction outlined in the proof of Theorem 6 will be oracularizable.

▶ **Lemma 15.** *There exists a constant $0 < s < 1$ such that it is* RE*-hard to distinguish whether the clause-variable game associated with a* 3SAT10 *instance has synchronous quantum value 1 or at most s. Moreover, when the value is 1, there exists binary observables $\{A_x\}_{x\in X}$ corresponding to each variable $x \in X$ in the instance which satisfy:*

1. *the observables satisfy the constraints*
2. *$[A_x, A_y] = 0$ if there exists $i \in [m]$ such that $x, y \in V_i$,*
3. *$[A_x, A_y] = 0$ if there exist $i, j \in [m]$ and $z \in X$ such that $x, z \in V_i$ and $y, z \in V_j$.*

Proof. The only non-trivial part of the statement is the third property. Let S be a 3SAT5 instance and $\Gamma = \Gamma(S)$ be its constraint graph with constraint (left) vertices $u \in U$ and (right) variables vertices $x \in X$. Starting from Γ consider the following graph Γ' which is constructed as follows: Let $N(u) \subset X$ be the neighbourhood of u. These represent the variables in the corresponding constraint. For each $u \in U$, add a "local" copy x_u to X for each "global" vertex $x \in N(u)$ (i.e. 5 copies for each variable vertex, as S is a 3SAT5 instance). For each copy of $N(u)$ added to X, we remove u and add 7 copies of the vertex u to U and, for each new copy u', an edge from u' to each vertex x_u in the copy of $N(u)$. Now, for each added variable vertex x_u, add 2 vertices w_{x_u} and v_{x_u} to the constraint vertices U, and edges connecting both x_u and x to w_{x_u} and v_{x_u}. We now observe that each vertex in X has degree 10. For convenience later, we let $U' \subset U$ denote the vertices in Γ' incident to only "local" variable vertices. Likewise, we let $V \subset U$ and $W \subset U$ denote the subsets of vertices w_{x_u} and v_{x_u}, respectively. Notably, $U = U' \sqcup V \sqcup W$ are the constraint vertices of Γ'.

Now, consider the CS S and the connectivity graph Γ' built from Γ as above. From S and Γ' we construct a new CS S' as follows. For each $u' \in U'$, there is a 3SAT clause u from S that was copied to create u. Add to S' a copy of that constraint with x_u taking the role of x for each $x \in N(u)$. Next, for the vertices $v_{x_u} \in V$, add the constraint $x \vee x \vee \neg x_u$ to S'. Finally, for the vertices $w_{x_u} \in W$, add the constraint $x_u \vee x_u \vee \neg x$ to S'.

We now observe that for the constructed CS S' we have that $\Gamma(S') = \Gamma'$. Next, suppose there exists a quantum perfect strategy for the constraint-variable game corresponding to S. Denote A_x for the observable corresponding to $x \in X$ in this perfect strategy. Now, extend this quantum assignment to the variables of S' by taking the observables for x_u to be $A_{x_u} = A_x$. By construction, this satisfies both the 3CNF constraints on the labelled variables and the new equality constraints. Hence, it gives a quantum satisfying assignment to S' – hence the reduction is complete. Moreover, we note that this assignment also satisfies conditions (2) and (3) of the statement. First, if x and y belong to the same context of S', then either $x = a_u$ and $y = b_u$ for some a, b belonging to the same context of S and hence $[A_x, A_y] = [A_a, A_b] = 0$; or $y = x_u$ (or vice-versa) so $A_x = A_y$ and hence they commute. Second, suppose x and y are variables of S' that each share a constraint with a variable z. By construction, either the constraint containing x and z or the constraint containing y and z must contain a "global" variable w. Without loss of generality, let us assume we are in the first case. That implies that $x = w$ or $x = w_u$, and $z = w$ or $z = w_u$, so $A_x = A_w = A_z$. As z is in the same constraint as y, this means $[A_x, A_y] = [A_z, A_y] = 0$.

For soundness, one can observe that the transformation between S to S' is an instance of subdivision, and therefore we can appeal to [8, Theorem 5.2]. In particular, this implies that if the constraint-variables CS games corresponding to S has quantum synchronous value less than s, then the constraint variable game corresponding to S' has quantum synchronous value less than $C_0 s$, for some universal constant $C_0 > 1$. ◀

Proof of Theorem 14. This proof is similar to the proof of Theorem 6. The only difference is we start the reduction from the instances in Lemma 15 rather than Theorem 5. The proof of soundness is exactly the same. We only need to show that in the completeness argument that the assignment we construct for the smooth label cover instance is a quantum oracularizable assignment.

Consider an instance B of 3SAT $-$ 10 from Lemma 15 which has quantum value 1. Let A_x denote the operators as in the statement of Lemma 15, and let $\{A_x^b\}_{b\in\{0,1\}}$ denote the corresponding PVMs. Then, by the construction in the proof of Lemma 8, Alice's operators in the corresponding perfect strategy for $\mathcal{G}(B, J, R)$ are $P_{\vec{i}}^{\vec{\sigma}} = \bigotimes_{q=1}^{(J+1)R} \prod_{x\in V_{i_q}} A_x^{\sigma_q(x)}$. We have that, for questions $\vec{i}, \vec{i}'$ of $\mathcal{G}(B, J, R)$ such that there exists $\vec{j}$ with $\pi(\vec{i}, \vec{j}), \pi(\vec{i}', \vec{j}) > 0$, $[P_{\vec{i}}^{\vec{\sigma}}, P_{\vec{i}'}^{\vec{\sigma}'}] = 0$. This is because, for all $q \in [(J+1)R]$, either $j_q = i_q = i'_q$ or $j_q \in V_{i_q} \cap V_{i'_q}$, and therefore for each $x \in V_{i_q}$ and $y \in V_{i'_q}$ either x and y are in the same context, or there exists z that shares a context with both x and y. Hence, by properties (2) and (3) of the strategy induced by the A_x from Lemma 15, $[A_x^{\sigma_q(x)}, A_y^{\sigma'_q(y)}] = 0$. Then, using the construction of Theorem 6, the strategy $\mathcal{S}'$ for $\mathfrak{S}$ is given by $Q_v^i = P_{\vec{i}}^{\vec{\sigma}}$, where $v \in V$ corresponding to $\vec{i}$ and $i \in [n]$ corresponding to $\vec{\sigma}$. It follows from the above commutation relation on the $P_{\vec{i}}^{\vec{\sigma}}$ that this strategy is oracularizable. This is a perfect strategy for $\mathfrak{S}$ by the same argument as in the proof of Theorem 6. ◀

4 RE-hardness of 3SAT* with fixed degree

In this section, we establish the RE-hardness of 3SAT5* and prove Theorem 5. Our starting point is the following undecidability result for succinctly presented 3SAT* protocols.

▶ **Theorem 16** ([8]). *There exists a constant $0 < s < 1$ such that there is a polynomial-time reduction from the halting problem to* $\mathrm{SuccinctCSP}_{c-v}(3\mathrm{SAT})^*_{1,s}$.

There are two technical issues we need to resolve before we can prove Theorem 5. Firstly, we need to resolve the issue of succinctness. In particular, our result requires a direct presentation of the 3SAT* protocol, rather than the succinct presentation from the BCS-MIP* protocol in Theorem 16. Secondly, due to technical limitations of our proof techniques (see e.g. in the proof of Theorem 28), we require that the question distribution to our 3SAT* protocol be uniform, another property not guaranteed a priori from the protocols in Theorem 16. Fortunately, we can resolve both of these issues.

▶ **Definition 17.** *Fix a universal Turing machine M. The* halting problem *is the language* HALTING *of all strings $x \in \{0,1\}^*$ such that M halts on input x. For a function $f : \mathbb{N} \to \mathbb{N}$ such that $f(n) \geq n$, the* f-padded halting problem *is the language* $\mathrm{PADDEDHALTING}_f = \{x0^{f(|x|)-|x|} \mid x \in \mathrm{HALTING}\}$.

In particular, there is a $O(f)$-time reduction from the halting problem to the padded halting problem by padding with 0, and a polynomial-time reduction from the padded-halting problem to the halting problem by truncating. Therefore $\mathrm{PADDEDHALTING}_f$ is RE-complete. Hence, to deal with the succinctness issue, we push the exponential-time reduction to the realm of RE, where the computation time is immaterial. In fact, we can exploit this idea further, and use it to enforce our uniformity conditions on the resulting 3SAT3 instances. More explicitly, we show that constraint-variable BCS nonlocal games with uniform distribution are RE-complete in the following sense.

▶ **Lemma 18.** *There exists $s \in (0,1)$ such that there is a polynomial-time reduction from the* exp*-padded halting problem to* $\mathrm{CSP}_{c-v}(3\mathrm{SAT})^*_{1,s}$ *with uniform question distribution.*

Proof. From Theorem 16 there exists a polynomially-bounded function p such that the reduction to succinct entangled 3SAT maps from a string x – seen as input to a universal Turing machine M – to Turing machines that sample a probability distribution on constraints from a 3SAT instance $S_x = (X, \{(V_i, C_i)\}_{i=1}^{2^{p(|x|)}})$. Now consider the reduction from the 2^p-padded halting problem defined as follows. For any string y, truncate it to the first $\lceil p^{-1}(\log |y|)\rceil$ bits and call that substring x. If $y \neq x0^{2^{p(|x|)}-|x|}$, then map y to a fixed unsatisfiable constraint system $S'_y = S_0$. Else, since x is an instance of the halting problem, the reduction of Theorem 16 maps it to a probability distribution π on constraints of a 3SAT instance S_x with $2^{p(|x|)} = |y|$ constraints. Now, by deterministically enumerating over random seeds for the Turing machine that samples π, we obtain the distribution in $\mathrm{poly}(|y|)$ time. Next, let $S'_y = (X, \{(V_{i,j}, C_i)\}_{i,j})$ be the 3SAT instance defined from S_x by repeating each constraint i $\lceil |y|^2\pi(i)\rceil$ times. S'_y has $m'_y = \sum_i \lceil |y|^2\pi(i)\rceil$ constraints, which is bounded as $m'_y \leq \sum_i (|y|^2\pi(i) + 1) = |y|^2 + |y|$ and $m'_y \geq \sum_i |y|^2\pi(i) = |y|^2$.

We claim that the map $y \mapsto S'_y$ is a polynomial-time reduction from the 2^p-padded halting problem to $\mathrm{CSP}_{c-v}(3\mathrm{SAT})^*_{1,(1+s)/2}$. It is clear that the reduction is polynomial-time in $|y|$. If y is a yes instance of the 2^p-padded halting problem, then x is a yes instance of the halting problem, so S_x has a perfect quantum satisfying assignment. Since S'_y has the same constraints as S_x, it is also perfectly quantum satisfiable. Now, if y is a no instance of the halting problem, either $y \neq x0^{2^{p(|x|)}-|x|}$ or x is a no instance of the halting problem. In the former case, $S'_y = S_0$ can be chosen to have large enough defect. In the latter case, $\mathrm{def}(\tau) \geq 1 - s$ for any finite-dimensional tracial state on $\mathcal{A}_{c-v}(S_x, \pi)$. Now, let τ' be a finite-dimensional tracial state on $\mathcal{A}_{c-v}(S'_y, \mathfrak{u}_{m'_y})$, and let $\varphi' : \mathcal{A}_{c-v}(S'_y) \to \mathcal{M}$ be the GNS representation, with tracial state ρ. Now, let φ be the representation of $\mathcal{A}_{c-v}(S_x)$

defined as $\varphi(\Pi_b(\sigma'(a))) = \varphi'(\Pi_b(\sigma'(a)))$ and $\varphi(\Phi_{V_i,\phi}) = \varphi'(\Phi_{V_{i,j},\phi})$ for the j that minimizes $\sum_{\phi\in C_i}\sum_{a\in V_i}\rho(\varphi'(\Phi_{V_{i,j},\phi})(1-\varphi(\Pi_{\phi(a)}(\sigma'(a)))))$; let $\tau = \rho\circ\varphi$. Since τ is a tracial state on $\mathcal{A}_{c-v}(S_x,\pi)$, and the PVMs are extremal within the POVMs, the defect is bounded by

$$\begin{aligned}\operatorname{def}(\tau) &\leq \sum_{i=1}^{|y|}\frac{\pi(i)}{|V_i|}\sum_{\varphi\in C_i}\sum_{a\in V_i}\rho\Big(\frac{1}{\lceil|y|^2\pi(i)\rceil}\sum_{j=1}^{\lceil|y|^2\pi(i)\rceil}\varphi'(\Phi_{V_{i,j},\phi})(1-\varphi'(\Pi_{\phi(a)}(\sigma'(a))))\Big)\\ &\leq \sum_{i,j}\frac{\pi(i)}{\lceil|y|^2\pi(i)\rceil}\frac{1}{|V_{i,j}|}\sum_{\varphi\in C_i}\sum_{a\in V_i}\tau'(\Phi_{V_{i,j},\phi}(1-\Pi_{\phi(a)}(\sigma'(a)))).\end{aligned}$$

Now, $\lceil|y|^2\pi(i)\rceil \geq |y|^2\pi(i)$, so

$$\frac{\pi(i)}{\lceil|y|^2\pi(i)\rceil} \leq \frac{1}{|y|^2} = \frac{m'_y}{|y|^2}\mathfrak{u}_{m'_y}(i,j) \leq \left(1+\frac{1}{|y|^2}\right)\mathfrak{u}_{m'_y}(i,j).$$

Hence, $\operatorname{def}(\tau) \leq \left(1+\frac{1}{|y|^2}\right)\operatorname{def}(\tau') \leq 2\operatorname{def}(\tau')$. Hence, for every τ', $\operatorname{def}(\tau') \geq \frac{1-s}{2}$, giving the desired gap. ◀

Now, we show that the 3SAT* protocol can be reduced to an instance of entangled 3SAT5, which denotes a 3SAT CSP where every variable appears exactly in 5 clauses, equivalently the degree of every left vertex in its connectivity graph (see Definition 1) is 5. In other words, we show that we can fix the degree of 3SAT* while preserving a constant soundness gap. Although our approach is inspired by the classical proof in [14], establishing the proof in the quantum case requires some additional technical results, in particular a novel use of expander graphs, which is not needed to establish the classical analogue.

▶ **Definition 19.** *A graph $G=(V,E)$ is an* (n,d,λ)-spectral expander *if $|V|=n$, the degree of G is d, G is connected, and the second-largest eigenvalue of the adjacency matrix $A(G)$ is at most $d(1-\lambda)$. A* (d,λ)-expander family *is a sequence of graphs (G_n) such that G_n is a (n,d,λ)-spectral expander for each n, and a description of G_n can be generated in time polynomial in n.*

▶ **Theorem 20** (*e.g.* [29, 37])**.** *There are constants $d\in\mathbb{N}$ and $\lambda>0$ such that there exists a (d,λ)-expander family.*

It will also be useful to know the expansion of the cycle graph, which is a very weak expander, but has minimal degree and nonzero expansion.

▶ **Lemma 21.** *Let C_d be the cycle on d vertices. C_d is a $(d,2,\frac{8}{d^2})$-spectral expander.*

Proof. Let ω be a primitive d-th root of unity. It is easy to see that for any $i=0,\dots,d-1$, the vector $(1,\omega^i,\omega^{2i},\dots,\omega^{(d-1)i})$ is an eigenvector of the adjacency matrix $A(C_d)$ with eigenvalue $2\cos\left(\frac{2\pi i}{d}\right)$, and that these eigenvectors provide an eigenbasis. As such, the largest eigenvalue is the degree 2 and the second-largest eigenvalue is $2\cos\left(\frac{2\pi}{d}\right)$. Thus, the expansion coefficient is $\lambda = \frac{2-2\cos\left(\frac{2\pi}{d}\right)}{2} = 2\sin^2\left(\frac{\pi}{d}\right) \geq \frac{8}{d^2}$. ◀

▶ **Lemma 22** (Lemma A.2 [21])**.** *Let $G=(V,E)$ be an (n,d,λ)-spectral expander, let $\mathcal{A}$ be a $*$-algebra, let τ be a state on $\mathcal{A}$, and let $\{a_u\}_{u\in V}$ be a collection of operators in $\mathcal{A}$. Then,*

$$\frac{1}{n^2}\sum_{u,v}\|a_u-a_v\|_\tau^2 \leq \frac{1}{\lambda}\frac{2}{dn}\sum_{\{u,v\}\in E}\|a_u-a_v\|_\tau^2$$

▶ **Corollary 23.** *For each $u \in V$, suppose $\{A_i^u\}_{i=1}^k$ is a PVM in some finite-dimensional von Neumann algebra $\mathcal{M}$ such that $\frac{2}{nd}\sum_{\{u,v\}\in E}\sum_i \|A_i^u - A_i^v\|_\tau^2 \leq \varepsilon$ for some tracial state τ on $\mathcal{M}$. Then, there exists a PVM $\{P_i\}_{i=1}^k$ such that $\frac{1}{n}\sum_{u,i}\|A_i^u - P_i\|_\tau^2 \leq 5\varepsilon/\lambda$.*

Proof. By Lemma 22 we know that $\frac{1}{n^2}\sum_{u,v}\sum_i \|A_i^u - A_i^v\|_\tau^2 \leq \frac{\varepsilon}{\lambda}$. Then, note that

$$\sum_i \|A_i^u - A_i^v\|_\tau^2 = \sum_i \tau(A_i^u + A_i^v - 2A_i^u A_i^v) = 2 - 2\sum_i \tau(A_i^u A_i^v),$$

and hence $\frac{1}{n^2}\sum_{u,v}\sum_i \tau(A_i^u A_i^v) \geq 1 - \frac{\varepsilon}{2\lambda}$. Writing the POVM $A_i' = \frac{1}{n}\sum_u A_i^u$, we have that

$$\sum_i \tau((A_i')^2) = \frac{1}{n^2}\sum_{u,v,i}\tau(A_i^u A_i^v) \geq 1 - \frac{\varepsilon}{2\lambda},$$

so by [10][Theorem 1], there exists a PVM $\{P_i\}_i$ such that $\sum_i \|P_i - A_i'\|_\tau^2 \leq \frac{9\varepsilon}{2\lambda}$. This gives that

$$\begin{aligned}
&\frac{1}{n}\sum_{u,i}\|A_i^u - P_i\|_\tau^2 = 2 - \frac{2}{n}\sum_{u,i}\tau(A_i^u P_i) = 2 - 2\sum_i \tau(A_i' P_i) \\
&\leq \frac{\varepsilon}{2\lambda} + \sum_i \tau(P_i) + \sum_i \tau((A_i')^2) - 2\sum_i \tau(A_i' P_i) = \frac{\varepsilon}{2\lambda} + \sum_i \|A_i' - P_i\|_\tau^2 \leq \frac{5\varepsilon}{\lambda}
\end{aligned}$$

◀

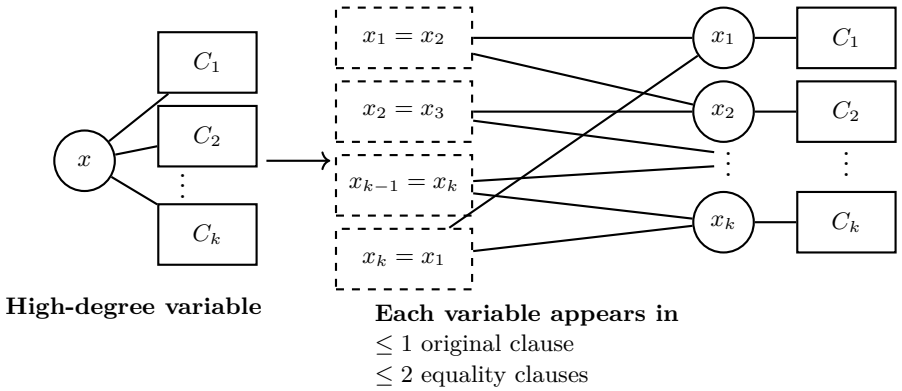

Figure 2 The reduction from 3SAT to 3SAT5. We first replace each variable with a labelled copy, then add equality constraints between the new variables. By replacing each equality constraint with two 3SAT constraints, we get a 3SAT5 instance.

With these results and definitions established we now move onto the reduction. To reduce the number of constraints that each variable appears in, we modify the constraint system by labelling each variable by the constraint it appears in and adding equality constraints between these labelled variables, one for each edge of an expander graph, see Figure 2. We use the properties of the expanders (defined above) to ensure equality while maintaining the low degree. We formalize this in the following definition.

▶ **Definition 24.** *Let $G = (G_n)$ be a sequence of graphs such that $G_n = ([n], E_n)$, and let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a k-ary CS. For each variable $x \in X$, let $n_x = |\{i \mid x \in V_i\}|$ and fix a bijection $r_x : \{i \mid x \in V_i\} \mapsto [n_x]$. We call the* G-replacement of S *the k-ary CS $S|_G$ with variables $X|_G = \{x_i \mid i \in [m], x \in V_i\}$ and two types of constraints: for each $i \in [m]$, the constraint (V_i', C_i) where $V_i' = \{x_i \mid x \in V_i\}$; and for each $x \in X$ and edge $\{u, v\} \in E_{n_x}$,*

the constraint $(V_{x,r_x^{-1}(u),r_x^{-1}(v)}, C_=)$*, where* $V_{x,i,j} = \{x_i, x_j\}$ *and* $C_= = \{(a,a) \mid a \in \mathbb{Z}_k\}$ *is the* k*-ary equality constraint. Given a probability distribution* π *on* $[m]$*, define the probability distribution* $\pi|_G$ *via* $\pi|_G(i) = \pi(i)/2$ *and* $\pi|_G(x,i,j) = \frac{\pi(x)}{2dn_x}$ *where* $\pi(x) = \sum_{i.\ x\in V_i} \frac{\pi(i)}{|V_i|}$.

▶ Remark 25. If G is a (d,λ)-expander family, then the left degree of the connectivity graph of $\Gamma(S|_G)$ is $d+1$; the right degree is simply the largest number of variables in a constraint, that is $\max\{2, \max_i |V_i|\}$. In particular, if S is a 3SAT instance, we have that the left vertices all have degree $d+1$ and the right vertices have degree 2 or 3. Also, the equality constraint between two boolean variables x, y can be represented by the gadget $(\neg x \vee y \vee y) \wedge (x \vee \neg y \vee \neg y)$. If B is a BCS which is a 3SAT instance, we call $B|_{G,3SAT}$ the 3SAT instance constructed by replacing all the equality constraints in $B|_G$ by the two constraints given in the above gadget. Then, the left vertices of $\Gamma(B|_{G,3SAT})$ all have degree $2d+1$ and the right vertices all have degree 3.

▶ **Lemma 26.** *Let* $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ *be a CS, let* π *be a probability distribution on* $[m]$*, and let* τ *be a tracial state on* $\mathcal{A}_{c-v}(S,\pi)$*. Then,*

$$\mathrm{def}(\tau) = \sum_{i,x\in V_i, a\in\mathbb{Z}_k} \frac{\pi(i)}{2|V_i|} \|\Pi_a(\sigma_i(x)) - \Pi_a(\sigma'(x))\|_\tau^2.$$

Proof. In fact,

$$\mathrm{def}(\tau) = \sum_{i,x\in V_i} \frac{\pi(i)}{|V_i|} \sum_{a\in\mathbb{Z}_k} \tau\Big(\Pi_a(\sigma_i(x))(1 - \Pi_a(\sigma'(x)))\Big).$$

To finish the proof, note that for any projectors P, Q, we have that

$$\|P - Q\|_\tau^2 = \tau(P + Q - 2PQ) = \tau(P - Q) + 2\tau(P(1-Q)),$$

and therefore

$$\begin{aligned}
&\sum_{a\in\mathbb{Z}_k} \frac{1}{2}\|\Pi_a(\sigma_i(x)) - \Pi_a(\sigma'(x))\|_\tau^2 \\
&= \sum_{a\in\mathbb{Z}_k} \frac{1}{2}\tau(\Pi_a(\sigma_i(x)) - \Pi_a(\sigma'(x))) + \tau(\Pi_a(\sigma_i(x))(1 - \Pi_a(\sigma'(x)))) \\
&= \frac{1}{2}\tau(1-1) + \sum_{a\in\mathbb{Z}_k} \tau(\Pi_a(\sigma_i(x))(1 - \Pi_a(\sigma'(x)))) = \sum_{a\in\mathbb{Z}_k} \tau(\Pi_a(\sigma_i(x))(1 - \Pi_a(\sigma'(x)))).
\end{aligned}$$ ◀

▶ **Lemma 27.** *Consider the BCS* $B_{x=y} = (\{x,y\}, \{(V_1, C_1), (V_2, C_2)\})$*, where* $V_1 = V_2 = \{x, y, z\}$*, and* C_1 *and* C_2 *the 3CNF constraints given by* $\neg x \vee y \vee y$ *and* $x \vee \neg y \vee \neg y$*, respectively. Then, for a state* τ *on* $\mathcal{A}_{c-v}(B_{x=y}, \mathfrak{u}_3)$*,* $\|\sigma'(x) - \sigma'(y)\|_\tau^2 \leq 64\,\mathrm{def}(\tau)$.

Proof. For a variable v, write $v^i = \sigma_i(v)$, $v_a^i = \sigma_i(\Pi_a(v))$, $v = \sigma'(v)$, and $v_a = \sigma'(\Pi_a(v))$. Due to the constraints, we have $x_{-1}^1 y_1^1 = 0$ and $x_1^2 y_{-1}^2 = 0$. Note that $\|v^i - v\|_\tau^2 = 4\|v_a^i - v_a\|_\tau^2$ and hence using Lemma 26,

$$\mathrm{def}(\tau) = \sum_{i,v\in V_i} \frac{1}{8|V_i|}\|v^i - v\|_\tau^2 = \frac{1}{16}\Big(\|x^1 - x\|_\tau^2 + \|x^2 - x\|_\tau^2 + \|y^1 - y\|_\tau^2 + \|y^2 - y\|_\tau^2\Big).$$

Next,

$$\begin{aligned}
\|x - y\|_\tau^2 = \|1 - xy\|_\tau^2 &= 4\|x_{-1}y_1 + x_1 y_{-1}\|_\tau^2 \\
&= 4\|x_{-1}^1(y_1 - y_1^1) + (x_{-1} - x_{-1}^1)y_1 + x_1^2(y_{-1} - y_{-1}^2) + (x_1 - x_1^2)y_{-1}\|_\tau^2 \\
&\leq 4\Big(\|y - y^1\|_\tau^2 + \|y - y^2\|_\tau^2 + \|x - x^1\|_\tau^2 + \|x - x^2\|_\tau^2\Big) \leq 64\,\mathrm{def}(\tau),
\end{aligned}$$

where the first inequality is $\|\sum_{i=1}^k a_i\|_\tau^2 \leq 2^{\lceil \log k\rceil} \sum_{i=1}^k \|a_i\|_\tau^2$ [31, Lemma 7.2]. ◀

▶ **Theorem 28.** *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a k-ary CS, and $G = (G_n)$ be a (d, λ)-expander family. For every tracial state τ on $\mathcal{A}_{c-v}(S|_G, \mathrm{u}_m|_G)$, there exists a tracial state τ' on $\mathcal{A}_{c-v}(S, \mathrm{u}_m)$ such that $\mathrm{def}(\tau') \leq \frac{80L}{\lambda}\mathrm{def}(\tau)$, where $L = \max_i |V_i|$.*

Proof. Write $\pi = \mathrm{u}_m$. Write $\varepsilon_i = \frac{1}{|V_i'|}\sum_{x_i \in V_i', a} \|\Pi_a(\sigma_i(x_i)) - \Pi_a(\sigma'(x_i))\|_\tau^2$ and

$$\varepsilon_x = \frac{1}{2dn_x} \sum_{\substack{i,j.\ x \in V_i \cap V_j,\\ \{r_x(i), r_x(j)\} \in E_{n_x}}} \sum_a \|\Pi_a(\sigma_{x,i,j}(x_i)) - \Pi_a(\sigma'(x_i))\|_\tau^2 + \|\Pi_a(\sigma_{x,i,j}(x_j)) - \Pi_a(\sigma'(x_j))\|_\tau^2.$$

We have $\mathrm{def}(\tau) = \sum_i \frac{\pi(i)}{4}\varepsilon_i + \sum_x \frac{\pi(x)}{4}\varepsilon_x$. Since $C_{x,i,j} = C_=$ is the equality constraint between x_i and x_j, $\sigma_{x,i,j}(x_i) = \sigma_{x,i,j}(x_j)$ so that

$$\frac{1}{dn_x} \sum_{\substack{i,j.\ x \in V_i \cap V_j,\\ \{r_x(i), r_x(j)\} \in E_{n_x}}} \sum_a \|\Pi_a(\sigma'(x_i)) - \Pi_a(\sigma'(x_j))\|_\tau^2 \leq 4\varepsilon_x.$$

Let $\varphi : \mathcal{A}_{c-v}(S|_G, \mathrm{u}_m|_G) \to \mathcal{M}$ be a $*$-representation and $\rho : \mathcal{M} \to \mathbb{C}$ be a state such that $\tau = \rho \circ \varphi$ is the GNS representation of τ. Using Corollary 23, there exists an order-k unitary $T_x \in \mathcal{M}$ such that

$$\frac{1}{n_x} \sum_{i.x \in V_i} \sum_a \|\varphi(\Pi_a(\sigma'(x_i))) - \Pi_a(T_x)\|_\rho^2 \leq \frac{20}{\lambda}\varepsilon_x.$$

Take $\chi : \mathcal{A}_{c-v}(S, \mathrm{u}_m) \to \mathcal{M}$ to be the $*$-homomorphism defined via $\chi(\sigma_i(x)) = \varphi(\sigma_i(x_i))$ and $\chi(\sigma'(x)) = T_x$. Take $\tau' = \rho \circ \chi$. Using the fact that π is uniform, note that $\pi(i) = \frac{1}{m}$ and $\frac{n_x}{Lm} \leq \pi(x) \leq \frac{n_x}{m}$. Then, the defect

$$\begin{aligned}
\mathrm{def}(\tau') &= \sum_i \frac{\pi(i)}{2|V_i|} \sum_{x \in V_i, a} \|\varphi(\Pi_a(\sigma_i(x_i))) - \Pi_a(T_x)\|_\rho^2 \\
&\leq \sum_i \frac{\pi(i)}{|V_i|} \sum_{x \in V_i, a} \|\varphi(\Pi_a(\sigma_i(x_i))) - \varphi(\Pi_a(\sigma'(x_i)))\|_\rho^2 + \|\varphi(\Pi_a(\sigma'(x_i))) - \Pi_a(T_x)\|_\rho^2 \\
&\leq \sum_i \pi(i)\varepsilon_i + \sum_x \frac{1}{m} \sum_{i.\ x \in V_i} \sum_a \|\varphi(\Pi_a(\sigma'(x_i))) - \Pi_a(T_x)\|_\rho^2 \\
&\leq \sum_i \pi(i)\varepsilon_i + \frac{20}{\lambda}\sum_x \frac{n_x}{m}\varepsilon_x \leq \sum_i \pi(i)\varepsilon_i + \frac{4L}{\lambda}\sum_x \pi(x)\varepsilon_x \leq \frac{80L}{\lambda}\mathrm{def}(\tau).
\end{aligned}$$

◀

Proof of Theorem 5. Due to Lemma 18, there is a polynomial-time reduction from the exp-padded halting problem, an RE-complete problem, to $\mathrm{CSP}_{c-v}(3\mathrm{SAT})^*_{1,s}$ for some $s \in (0, 1)$. Let $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a 3SAT instance and let $G = (G_n = ([n], E_n))_n$ be a (d, λ)-expander family. Next, using Theorem 28, for every trace τ on $\mathcal{A}_{c-v}(B|_G, \mathrm{u}_m|_G)$, there exists a trace τ' on $\mathcal{A}_{c-v}(B, \mathrm{u}_m)$ with defect $\mathrm{def}(\tau') \leq \frac{240}{\lambda}\mathrm{def}(\tau)$. We have $\mathrm{u}_m|_G(i) = \frac{1}{2m}$ and $\frac{1}{6dm} \leq \mathrm{u}_m|_G(x, i, j) \leq \frac{1}{2dm}$, hence we can replace $\mathrm{u}_m|_G$ by the uniform distribution on the m' constraints of $B|_G$ while incurring only a constant penalty. Explicitly, there exists a constant C such that for every trace τ on $\mathcal{A}_{c-v}(B|_G, \mathrm{u}_{m'})$, there exists a trace τ' on $\mathcal{A}_{c-v}(B, \mathrm{u}_m)$ with defect $\mathrm{def}(\tau') \leq C\,\mathrm{def}(\tau)$. Consider the family of graphs $H = (C_n)_n$ where C_n is the cycle on n vertices. By Lemma 21, C_n is a $(n, 2, \frac{8}{n^2})$ expander, so H is not an expander family. However, since each variable in $B|_G$ is contained in at most $d + 1$ constraints, H behaves like an expander family with $\lambda = \frac{8}{(d+1)^2}$ with respect to $B|_G$. As such, using Theorem 28 again, for every trace τ on $\mathcal{A}_{c-v}(B|_G|_H, \mathrm{u}_{m'}|_H)$, there exists a trace

τ' on $\mathcal{A}_{c-v}(B, \mathfrak{u}_m)$ with defect $\mathrm{def}(\tau') \leq 6(d+1)^2 C\,\mathrm{def}(\tau)$. To finish, note that $B|_G|_H$ has 3CNF or equality constraints. Using Lemma 27, we can replace all the equality constraints by gadgets built out of 3CNF constraints. As such, there exists a 3SAT-5 instance B', a probability distribution and a constant C' such that for every trace τ on $\mathcal{A}_{c-v}(B', \pi')$, there exists a trace τ' on $\mathcal{A}_{c-v}(B, \pi)$ such that $\mathrm{def}(\tau') \leq C'\,\mathrm{def}(\tau)$. ◀

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. doi:10.1145/278298.278306.

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. doi:10.1145/273865.273901.

3 Albert Atserias, Phokion G. Kolaitis, and Simone Severini. Generalized satisfiability problems via operator assignments. *Journal of Computer and System Sciences*, 105:171–198, 2019. doi:10.1016/J.JCSS.2019.05.003.

4 László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *computational complexity*, 1(1):3–40, March 1991. doi:10.1007/BF01200056.

5 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*, CCC '04, pages 236–249, Washington, DC, USA, 2004. IEEE Computer Society. doi:10.1109/CCC.2004.9.

6 Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *International Colloquium on Automata, Languages, and Programming (ICALP) 2012*, pages 320–331, 2012. doi:10.1007/978-3-662-43948-7_27.

7 Richard Cleve, William Slofstra, Falk Unger, and Sarvagya Upadhyay. Perfect parallel repetition theorem for quantum XOR proof systems. *Computational Complexity*, 17(2):282–299, 2008. doi:10.1007/S00037-008-0250-4.

8 Eric Culf and Kieran Mastel. RE-completeness of entangled constraint satisfaction problems. In *2025 IEEE 66th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2194–2230, 2025. doi:10.1109/FOCS63196.2025.00116.

9 Eric Culf, Hamoon Mousavi, and Taro Spirig. Approximation algorithms for noncommutative CSPs. In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 920–929. IEEE, 2024. doi:10.1109/FOCS61266.2024.00061.

10 Mikael de la Salle. Orthogonalization of positive operator valued measures. *Comptes Rendus. Mathématique*, 360(G5):549–560, 2022. doi:10.5802/crmath.326.

11 Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the 2-to-1 games conjecture? In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 376–389, New York, NY, USA, 2018. Association for Computing Machinery. doi:10.1145/3188745.3188804.

12 Irit Dinur, David Steurer, and Thomas Vidick. A parallel repetition theorem for entangled projection games. *computational complexity*, 24(2):201–254, 2015. doi:10.1007/S00037-015-0098-3.

13 Yangjing Dong, Honghao Fu, Anand Natarajan, Minglong Qin, Haochen Xu, and Penghui Yao. The computational advantage of MIP* vanishes in the presence of noise. *Journal of the ACM*, 2023.

14 Uriel Feige. A threshold of ln n for approximating set cover. *Journal of the ACM (JACM)*, 45(4):634–652, 1998. doi:10.1145/285055.285059.

15 Uriel Feige, Shafi Goldwasser, Laszlo Lovász, Shmuel Safra, and Mario Szegedy. Interactive proofs and the hardness of approximating cliques. *J. ACM*, 43(2):268–292, March 1996. doi:10.1145/226643.226652.

16 Honghao Fu, Kieran Mastel, and Xingjian Zhang. Succinct perfect zero-knowledge for MIP*. *arXiv preprint*, 2025. arXiv:2503.04517.

17 Michel X. Goemans and David P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, November 1995. doi:10.1145/227683.227684.

18 Venkatesan Guruswami, Prasad Raghavendra, Rishi Saket, and Yi Wu. Bypassing UGC from some optimal geometric inapproximability results. *ACM Trans. Algorithms*, 12(1), feb 2016. doi:10.1145/2737729.

19 Zhengfeng Ji. Binary constraint system games and locally commutative reductions. *arXiv preprint*, 2013. arXiv:1310.3794.

20 Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. $\mathsf{MIP}^* = \mathsf{RE}$. *Commun. ACM*, 64(11):131–138, October 2021. doi:10.1145/3485628.

21 Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. Quantum soundness of testing tensor codes. *arXiv preprint*, 2021. arXiv:2111.08131.

22 Julia Kempe, Oded Regev, and Ben Toner. Unique games with entangled provers are easy. *SIAM Journal on Computing*, 39(7):3207–3229, 2010. doi:10.1137/090772885.

23 S. Khot. Hardness results for coloring 3-colorable 3-uniform hypergraphs. In *The 43rd Annual IEEE Symposium on Foundations of Computer Science, 2002. Proceedings.*, pages 23–32, 2002. doi:10.1109/SFCS.2002.1181879.

24 S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings 17th IEEE Annual Conference on Computational Complexity*, pages 25–, 2002. doi:10.1109/CCC.2002.1004334.

25 Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007. doi:10.1137/S0097539705447372.

26 Subhash Khot, Dor Minzer, and Muli Safra. On independent sets, 2-to-2 games, and grassmann graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pages 576–589, New York, NY, USA, 2017. Association for Computing Machinery. doi:10.1145/3055399.3055432.

27 Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within $2-\varepsilon$. *Journal of Computer and System Sciences*, 74(3):335–349, 2008. Computational Complexity 2003. doi:10.1016/j.jcss.2007.06.019.

28 Junqiao Lin. Tracial embeddable strategies: Lifting MIP* tricks to MIPco. *arXiv preprint*, 2024. arXiv:2304.01940.

29 Grigorii Aleksandrovich Margulis. Explicit construction of concentrators. *Problem Inf. Trans.*, 9:325–332, 1973.

30 Amine Marrakchi and Mikael de la Salle. Almost synchronous correlations and Tomita-Takesaki theory. *arXiv preprint*, 2023. arXiv:2307.08129.

31 Kieran Mastel and William Slofstra. Two prover perfect zero knowledge for MIP*. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, pages 991–1002, 2024. doi:10.1145/3618260.3649702.

32 Hamoon Mousavi and Taro Spirig. A Quantum Unique Games Conjecture. In *16th Innovations in Theoretical Computer Science Conference (ITCS 2025)*, volume 325 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 76:1–76:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.ITCS.2025.76.

33 Anand Natarajan and Chinmay Nirkhe. The status of the quantum PCP conjecture (games version). *arXiv preprint*, 2024. doi:10.48550/arXiv.2403.13084.

34 Connor Paddock and William Slofstra. Satisfiability problems and algebras of boolean constraint system games. *Illinois Journal of Mathematics*, 69(1):81–107, 2025.

35 Christos Papadimitriou and Mihalis Yannakakis. Optimization, approximation, and complexity classes. In *Proceedings of the twentieth annual ACM symposium on Theory of computing*, pages 229–234, 1988.

36 Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, pages 245–254, New York, NY, USA, 2008. Association for Computing Machinery. doi:10.1145/1374376.1374414.

37 O. Reingold, S. Vadhan, and A. Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders and extractors. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 3–13, 2000. doi:10.1109/SFCS.2000.892006.

38 Thomas Vidick. Almost synchronous quantum correlations. *Journal of Mathematical Physics*, 63(2), February 2022. doi:10.1063/5.0056512.

39 Stephanie Wehner. Tsirelson bounds for generalized Clauser-Horne-Shimony-Holt inequalities. *Physical Review A – Atomic, Molecular, and Optical Physics*, 73(2):022110, 2006.