Complexity Lower Bounds for Computing the Approximately-Commuting Operator Value of Non-Local Games to High Precision

Matthew Coudron
Institute for Quantum Computing, University of Waterloo, Waterloo, Canada
mcoudron@uwaterloo.ca

William Slofstra
Institute for Quantum Computing and Department of Pure Mathematics, University of Waterloo, Waterloo, Canada
weslofst@uwaterloo.ca

Abstract

We study the problem of approximating the commuting-operator value of a two-player non-local game. It is well-known that it is NP-complete to decide whether the classical value of a non-local game is 1 or $1-\epsilon$, promised that one of the two is the case. Furthermore, as long as ϵ is small enough, this result does not depend on the gap ϵ. In contrast, a recent result of Fitzsimons, Ji, Vidick, and Yuen shows that the complexity of computing the quantum value grows without bound as the gap ϵ decreases. In this paper, we show that this also holds for the commuting-operator value of a game. Specifically, in the language of multi-prover interactive proofs, we show that the power of $\mathrm{MIP}^{co}(2,1,1,s)$ (proofs with two provers, one round, completeness probability 1, soundness probability s, and commuting-operator strategies) can increase without bound as the gap $1-s$ gets arbitrarily small.

Our results also extend naturally in two ways, to perfect zero-knowledge protocols, and to lower bounds on the complexity of computing the approximately-commuting value of a game. Thus we get lower bounds on the complexity class $\text{PZK-MIP}^{co}_{\delta}(2,1,1,s)$ of perfect zero-knowledge multi-prover proofs with approximately-commuting operator strategies, as the gap $1-s$ gets arbitrarily small. While we do not know any computable time upper bound on the class MIP^{co}, a result of the first author and Vidick shows that for $s = 1 - 1/\operatorname{poly}(f(n))$ and $\delta = 1/\operatorname{poly}(f(n))$, the class $\mathrm{MIP}^{co}_{\delta}(2,1,1,s)$, with constant communication from the provers, is contained in $\mathrm{TIME}(\exp(\operatorname{poly}(f(n))))$. We give a lower bound of $\mathrm{coNTIME}(f(n))$ (ignoring constants inside the function) for this class, which is tight up to polynomial factors assuming the exponential time hypothesis.

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases Quantum complexity theory, Non-local game, Multi-prover interactive proof, Entanglement

Digital Object Identifier 10.4230/LIPIcs.CCC.2019.25

Funding *Matthew Coudron*: Supported at the IQC by Canada's NSERC and the Canadian Institute for Advanced Research (CIFAR), and through funding provided to IQC by the Government of Canada and the Province of Ontario.
William Slofstra: Supported by NSERC DG 2018-03968.

Acknowledgements We thank Zhengfeng Ji, Alex Bredariol Grilo, Anand Natarajan, Thomas Vidick, and Henry Yuen for helpful discussions.

1 Introduction

Non-local games are a subject of converging interest for quantum information theory and computational complexity theory. A central question in both fields is the complexity of approximating the optimal winning probability of a non-local game. Quantum mechanics

allows non-local strategies in which the players share entanglement, and in quantum complexity theory we are interested in understanding the optimal winning probability over these entangled strategies. Answering this question is necessary for understanding the power of multi-prover interactive proof systems with entangled provers and a classical verifier.

For classical strategies (i.e. strategies without entanglement), it is NP-hard to decide whether a non-local game has winning probability 1. The PCP theorem implies that it is NP-hard to decide whether a non-local game has winning probability 1 or winning probability $1 - \epsilon$, where ϵ is constant, promised that one of the two is the case [1, 2]. Therefore, for classical games, the complexity of computing the winning probability is the same for constant error as for zero error.

Two models for quantum strategies have historically been used when defining the entangled value of a nonlocal game: the tensor product model and the commuting-operator model. The optimal winning probability of a non-local game over tensor product strategies is called the quantum value, and optimal winning probability over all commuting-operator strategies is called the commuting-operator value.

A number of lower bounds on approximating the quantum value of a non-local game are known. In particular, Ji has shown that it is NEXP-hard to compute the quantum value of a non-local game with inverse polynomial precision, and NEEXP-hard to compute the entangled value with inverse exponential precision [11]. Fitzsimons, Ji, Vidick, and Yuen continue this line of results by showing, roughly, that for any computable function $f(n) : \mathbb{N} \to \mathbb{N}$, it is NTIME(exp($f(n)$)) hard to compute the quantum value of a nonlocal game with $1/f(n)$ precision (here n is the input size) [7]. In particular, this implies that the quantum value of a game behaves very differently from the classical winning probability, since the complexity of computing the quantum value increases without bound as the required precision increases.

It is also natural to ask whether one might be able to approximate the commuting-operator value of a game efficiently. The study of the commuting-operator value goes back to [9], where it is shown that it is NP-hard to distinguish whether the commuting operator value is 1 or 1-1/poly(n). The complexity of the commuting operator value does not seem to be explicitly studied in more recent work.

In this paper, we look at lower bounds on the complexity of approximating the commuting-operator value of linear system nonlocal games, a type of nonlocal game closely connected with the theory of finitely-presented groups [3]. We show that group-theoretic methods can be used to lower bound the complexity of approximating the commuting-operator value of a linear system nonlocal game. In particular we show that, just as with the quantum value of a game, the complexity of computing the commuting operator value of a non-local game to precision ϵ grows arbitrarily large as ϵ decreases. Because our results are based on group-theoretic methods, we observe that they naturally extend to lower bounds on approximately-commuting-operator strategies for games, a generalization of commuting-operator strategies in which Alice and Bob's strategies can interact slightly, but in such a way that the interaction is bounded by a parameter δ. Thus we show:

▶ **Theorem 1.** *There is a universal constant k such that for every language $L \subset A^*$ over a finite alphabet A and contained in* coNTIME($f(n)$)*, where $f(n)$ is at least polynomial, there is a constant $C > 0$ and a family of two-player non-local games $(\mathcal{G}_w)_{w \in A^*}$ of size* poly($|w|$) *and computable in* poly($|w|$)*-time, such that for any $\delta = o(1/f(Cn)^{2k})$, deciding whether $\omega^{co}_\delta(\mathcal{G}_w) = 1$, or*

$$\omega^{co}_\delta(\mathcal{G}_w) \leq 1 - 1/f(Cn)^k + O(\sqrt{\delta}),$$

promised that one of the two is the case, is as hard as deciding membership of w in L.

Here $\omega_\delta^{co}(\mathcal{G}_w)$ denotes the supremum of all winning probabilities for all δ-commuting-operator strategies for the game $\mathcal{G}_w$ (see Definition 9). The proof of Theorem 1 is given in Section 5. Setting $\delta = 0$ gives a hardness result for approximating the commuting-operator value $\omega^{co} := \omega_0^{co}$ of two-player non-local games.

The proof of Theorem 1 relies on a deep group theory result of Sapir, Birget, and Rips, which shows that the acceptance problem for any Turing machine can be encoded in the word problem of a finitely-presented group, in such a way that the Dehn function of the group is equivalent to the running time of the Turing machine [16]. We then use [17] to embed this group into linear system non-local games. In the case that a word $w \in A^*$ does not belong to L, the provers demonstrate this fact by showing that a certain word in the corresponding group is not equal to the identity. In this case, the representation of the group forms the proof that the word is not equal to the identity, and this representation is used to build the provers' quantum strategy. The reason that we use commuting-operator strategies in Theorem 1, and again in Theorem 2 below, is that this representation might not be finite-dimensional.

Little is known about upper bounds on the complexity of computing the value of non-local games. Most existing proposals for an algorithm are based on a hierarchy of semi-definite programs [13, 14, 6]. It remains open whether such an algorithm can approximate the commuting-operator value of a game to any precision ϵ in finite time. However, the first author and Vidick have shown that the SDP hierarchy of [13, 14, 6] can be used to estimate (with explicit convergence bounds) the optimal value of a non-local game over approximately-commuting strategies [5]. In particular [5] gives an algorithm which, given a description of a non-local game as a truth-table of size n, can decide whether the game has commuting-operator value equal to 1, or has no δ-commuting-operator strategy with winning probability higher than $1 - \epsilon$ (for constant ϵ, promised that one of the two is the case), in time $n^{O(\mathrm{poly}(\ell,1/\delta))}$, where ℓ is the size of the output alphabet for the game. Theorem 1 shows that the dependence of this algorithm on δ is necessary. For the games $\mathcal{G}_w$ in Theorem 1, $\ell = O(1)$, and $\omega_\delta^{co}(\mathcal{G}_w) = 1$ if and only if $\omega^{co}(\mathcal{G}_w) = 1$. According to the exponential time hypothesis, we might expect that the best deterministic upper bound for coNTIME($f(n)$) is TIME($2^{\mathrm{poly}(f(n))}$). Thus, if we assume the exponential time hypothesis, the non-deterministic lower bound in Theorem 1 matches the deterministic upper bound in [5] up to polynomial factors (for families of games with a constant number of outputs).

Results about the complexity of non-local games have direct and natural implications for the power of multi-prover interactive proofs. Multi-prover interactive proofs were originally defined and studied in a purely classical setting. A seminal result of Babai, Fortnow, and Lund, which studies the class MIP of languages which admit a multi-prover interactive proof with polynomial time verifier, states that MIP = NEXP. Once again, this equality is independent of the completeness-soundness gap, as long as this gap is a large enough constant. For entangled strategies, there are, a priori, two analogs of the class MIP to consider, the class MIP* of multi-prover interactive proofs in which provers may use finite-dimensional entangled strategies, and MIP^{co}, the equivalent class with commuting-operator strategies. A result of Ito and Vidick states that the class MIP$^*(4, 1, 1, 1 - 1/\mathrm{poly}(n))$ with four provers, one round, completeness probability 1, and soundness probability 1-1/poly(n) contains NEXP [10]. Ji's result mentioned earlier for computing the quantum value of game shows that with a sufficient number of provers k, MIP$^*(k, 1, 1, 1 - 1/\exp(n))$ contains NEEXP, in contrast again to the classical case [11]. Ji's result is based on a compression theorem for non-local games, which also shows that the problem of computing the quantum value of a game is complete for MIP*.

Theorem 1 can be translated into lower bounds on $\mathrm{MIP}^{co}_{\delta}$, the class of languages with a multiprover interactive proof sound against approximately commuting strategies. Furthermore, these lower bounds also apply to the class PZK-$\mathrm{MIP}^{co}_{\delta}$ of languages which admit a perfect zero knowledge multiprover interactive proof sound against approximately commuting strategies. In a perfect zero knowledge interactive proof the provers must reveal nothing to the verifier except the proven statement itself. The formal definition of these two classes is given in Definitions 18 and 20.

▶ **Theorem 2.** *There is a universal constant k such that for any language L in* $\mathrm{NTIME}(f(n))$*, where $f(n)$ is at least polynomial, there is a constant C such that for any $\delta = o(1/f(Cn)^{2k})$,*

$$\overline{L} \in \text{PZK-}\mathrm{MIP}^{co}_{\delta}(2, 1, 1, 1 - 1/f(Cn)^k),$$

where $\overline{L}$ is the complement of L.

Note that, since the containment PZK-$\mathrm{MIP}^{co}_{\delta} \subseteq \mathrm{MIP}^{co}_{\delta}$ is immediate (see Definition 20), Theorem 2 represents both a lower bound for PZK-$\mathrm{MIP}^{co}_{\delta}$ and for $\mathrm{MIP}^{co}_{\delta}$ itself. Similarly to Theorem 1, when $\delta = 0$, we get a lower bound on the class $\mathrm{MIP}^{co} := \mathrm{MIP}^{co}_{0}$ of multi-prover interactive proofs with commuting-operator strategies, which is the direct analog of the complexity class MIP^* in the commuting operator setting (indeed, the term MIP^* has been used to denote MIP^{co} in some previous works).

One reason we are interested in the class $\mathrm{MIP}^{co}_{\delta}$ is that the algorithm of [5] mentioned above gives a (deterministic) time upper bound for $\mathrm{MIP}^{co}_{\delta}$. For protocols with constant-sized outputs, this upper bound is stated in Theorem 23. In contrast, no computable upper bounds for MIP^* or MIP^{co} are known. Combining Theorem 2 with the upper bound in of [5] gives the following series of containments (written here with constants, polynomial factors, and some parameters of the $\mathrm{MIP}^{co}_{\delta}$ notation suppressed for conciseness, including a parameter requiring a constant number of outputs). For any $\delta = o(1/f(Cn)^{2k})$:

$$\begin{aligned} \mathrm{coNTIME}(f(n)) &\subseteq \text{PZK-}\mathrm{MIP}^{co}_{\delta}(1, 1 - 1/\operatorname{poly}(f(n))) \\ &\subseteq \mathrm{MIP}^{co}_{\delta}(1, 1 - 1/\operatorname{poly}(f(n))) \\ &\subseteq \mathrm{TIME}(\exp(1/\operatorname{poly}(\delta))) \end{aligned} \tag{1.1}$$

Just as for the decision problem in Theorem 1, if we assume the exponential time hypothesis then we can consider the left hand side and right hand side of Equation 1.1 above to be matching up to polynomial factors.

Our results are complementary to the results of Fitzsimons, Ji, Vidick, and Yuen, who show qualitatively similar lower bounds for computing the quantum value of k-player games and for $\mathrm{MIP}^*(k, 1, 1, s)$, where $k \geq 15$. Their results show that MIP^* with $1/f(n)$ completeness-soundness gap contains $\mathrm{NTIME}(2^{f(n)})$, matching the pattern seen in [11] for inverse polynomial and inverse exponential gaps. In contrast, in our result the scaling of the lower bound relative to the gap is weaker, requiring gap of order $1/f(n)$ to get a lower bound of $\mathrm{coNTIME}(f(n))$, and applying to commuting-operator strategies rather than quantum strategies. However, our results apply to two-player protocols, while the results of [7] apply to protocols with 15 or more players. That we get a lower bound of $\mathrm{coNTIME}(f(n))$ rather than $\mathrm{NTIME}(2^{f(n)})$ can be explained by the fact that our lower bound extends to $\mathrm{MIP}^{co}_{\delta}$, which, with the restriction to protocols with constant-sized outputs, is contained in $\mathrm{TIME}(2^{f(n)})$. Thus our results highlight the importance of considering soundness to approximately-commuting strategies when seeking lower bounds on MIP^* and MIP^{co}. It seems to be an interesting open problem to determine whether the improved bounds of [7] can be done with algebraic methods.

2 Group theory preliminaries

Recall that a *finitely-presented group* is a group G with a fixed presentation $G = \langle S : R \rangle$, meaning that G is the quotient of the free group $\mathcal{F}(S)$ generated by a finite set S, by the normal subgroup generated by a finite set of relations $R \subseteq \mathcal{F}(S)$. If $G = \langle S : R \rangle$, and $R' \subseteq \mathcal{F}(S \cup S')$, then the notation $\langle G, S' : R' \rangle$ refers to the presentation $\langle S \cup S' : R \cup R' \rangle$. A *(group) word of length k* over the generators S is a string $s_1^{a_1} \cdots s_k^{a_k}$ where $s_i \in S$ and $a_i \in \{\pm 1\}$ for all $1 \leq i \leq k$. Such a word is said to be *reduced* if $s_i = s_{i+1}$ implies that $a_i = a_{i+1}$ for all $1 \leq i \leq k-1$. Every element $w \in \mathcal{F}(S)$ is represented by a unique reduced word, and the *length* $|w|$ *of* w is defined to be the length of this reduced word. The *word problem* for G is the problem of deciding whether the image of a given element $w \in \mathcal{F}(S)$ is equal to the identity in G, or in other words, whether the word is in the normal subgroup of $\mathcal{F}(S)$ generated by R. Since the reduced form of any non-reduced word over S can be found in time linear in the length of that non-reduced word, we can ask that inputs to the word problem be represented either as reduced or non-reduced words without changing the problem.

A *(unitary) representation* of a group G is a homomorphism $\phi : G \to \mathcal{U}(\mathcal{H})$, where $\mathcal{U}(\mathcal{H})$ is the unitary group of a Hilbert space $\mathcal{H}$. If $G = \langle S : R \rangle$ is a finitely-presented group, then a representation $\phi : G \to \mathcal{U}(\mathcal{H})$ can be specified by giving a homomorphism $\widetilde{\phi} : \mathcal{F}(S) \to \mathcal{U}(\mathcal{H})$ such that $\widetilde{\phi}(r) = 1$ for every $r \in R$. If G is a group, then $\ell^2 G$ is the Hilbert space with Hilbert basis $\mathcal{B} = \{|g\rangle : g \in G\}$. This means that every element of $\mathcal{H}$ is of the form $\sum_{g \in G} c_g |g\rangle$, where $\sum_{g \in G} |c_g|^2 \leq +\infty$. Since every group G acts on itself by both left and right multiplication, G also acts by left and right multiplication on $\mathcal{B}$. Thus G acts unitarily on $\ell^2 G$ by left and right multiplication. The resulting representations $L, R : G \to \mathcal{U}(\ell^2 G)$ are called the *left* and *right regular representations* of G, respectively.

If $w \in \mathcal{F}(S)$ is a word which is equal to the identity in G, we let $\mathrm{Area}_G(w)$ be the minimum $t \geq 1$ such that

$$w = z_1 r_1^{a_1} z_1^{-1} \cdots z_t r_t^{a_t} z_t^{-1}$$

for some $r_1, \ldots, r_t \in R$, $z_1, \ldots, z_t \in \mathcal{F}(S)$, and $a_1, \ldots, a_t \in \{\pm 1\}$.[1] The *Dehn function* Dehn_G of G is the function $\mathbb{N} \to \mathbb{N}$ defined by

$$\mathrm{Dehn}_G(n) = \max\{\mathrm{Area}_G(w) : w \in \mathcal{F}(S) \text{ has } |w| \leq n \text{ and } w = 1 \text{ in G}\}.$$

If the word problem of G is decidable, then Dehn_G is computable. Conversely, the word problem of G belongs to $\mathrm{NTIME}(\mathrm{Dehn}_G(n))$ [16]. An easy way to see that the complexity of the word problem is bounded by the Dehn function (albeit with the slightly worse upper bound of $\mathrm{NTIME}(\mathrm{poly}(\mathrm{Dehn}_G(n)))$) is through the following lemma:

▶ **Lemma 3** ([8], Lemma 2.2). *Let $G = \langle S : R \rangle$ be a finitely-presented group, and let ℓ be the length of the longest relation in R. If $w \in \mathcal{F}(S)$ is equal to the identity in G and $k = \mathrm{Area}_G(w)$, then*

$$w = z_1 r_1^{a_1} z_1^{-1} \cdots z_k r_k^{a_k} z_k^{-1}$$

where $r_1, \ldots, r_k \in R$, $z_1, \ldots, z_k \in \mathcal{F}(S)$, $a_1, \ldots, a_k \in \{\pm 1\}$, and $|z_i| \leq k\ell + \ell + |w|$ for all $1 \leq i \leq k$.

[1] $\mathrm{Area}_G(w)$ can also be defined as the minimum number of regions in a van Kampen diagram with boundary word w, and this is where the name comes from.

In general, the Dehn function can be much larger than the time-complexity of the word problem of G. However, Sapir, Birget, and Rips have shown that every recursive language can be reduced to the word problem of a finitely-presented group for which the Dehn function is polynomially equivalent to the time-complexity of the word problem. For the statement of the theorem, recall that two functions $T, T' : \mathbb{N} \to \mathbb{N}$ are said to be *(asymptotically) equivalent* if there are constants C, C' such that $T(n) \leq CT'(Cn) + Cn + C$ and $T'(n) \leq C'T(C'n) + C'n + C'$ for all $n \geq 1$.

▶ **Theorem 4** ([16], Theorem 1.3)**.** *Let A be a finite alphabet, and $L \subset A^*$ a language over A contained in* NTIME$(T(n))$*, where $T(n)$ is computable and $T(n)^4$ is at least superadditive (i.e. $T(n+m)^4 \geq T(n)^4 + T(m)^4$. Then there exists a finitely-presented group $G = \langle S : R\rangle$ and an injective function $\kappa : A^* \to \mathcal{F}(S)$, such that*
(a) *$|\kappa(u)| = O(|u|)$ and $\kappa(u)$ is computable in time $O(|u|)$,*
(b) *$u \in L$ if and only if $\kappa(u) = 1$ in G, and*
(c) *$\mathrm{Dehn}_G(n)$ is bounded by a function equivalent to $T(n)^4$.*

A group over $\mathbb{Z}_2$ is a pair (G, J) where J is a central involution, i.e. an element of the center of G with $J^2 = 1$. Usually we just write G for the pair, and refer to $J = J_G$ in the same way we refer to the identity $1 = 1_G$ of a group. When $J_G \neq 1_G$, it can be used as a substitute for -1. Theorem 4 implies that any recursive decision problem can be encoded in the word problem of a group. We want an embedding of this type where the word w is a central involution. For this, we use the following trick:

▶ **Definition 5.** *Let $G = \langle S : R\rangle$ be a finitely-presented group, and let x, J, t be indeterminates not in S. Given $w \in \mathcal{F}(S)$, let*

$$\widetilde{G}_w := \langle G, x, J, t \ : \ J^2 = 1, [g, J] = 1 \text{ for all } g \in G, \\ [x, J] = 1, [t, J] = 1, [t, [x, w]] = J\rangle,$$

where $[a, b] := aba^{-1}b^{-1}$ is the group commutator.

Note that if G is finitely-presented, then we only need to include the relations $[g, J] = 1$ for g in a generating set of G, and this gives a finite presentation of $\widetilde{G}_w$.

▶ **Lemma 6.** *Given a group $G = \langle S : R\rangle$ and a word $w \in \mathcal{F}(S)$, let $\widetilde{G}_w$ be the group defined in Definition 5. Then*
(a) *J is a central involution in $\widetilde{G}_w$,*
(b) *$w = 1$ in G if and only if $J = 1$ in $\widetilde{G}_w$, and*
(c) *if $w = 1$ in G then $\mathrm{Area}_{\widetilde{G}_w}(J) \leq 4\,\mathrm{Area}_G(w) + 1$.*

Proof. Part (a) is clear. For part (b), let

$$G' := \langle G, x, J : J^2 = [x, J] = [g, J] = 1 \text{ for all } g \in G\rangle^2,$$

The element $y = [x, w]$ is equal to 1 in G' if and only if $w = 1$. If $w \neq 1$ then y has infinite order. Hence the subgroup $\langle y, J\rangle$ is equal to $\mathbb{Z} \times \mathbb{Z}_2$ if $w \neq 1$, and $\mathbb{Z}_2$ if $w = 1$. In both cases, the homomorphism induced by $y \mapsto Jy$ and $J \mapsto J$ is an automorphism of this subgroup, and

$$\widetilde{G}_w = \langle G', t : tyt^{-1} = Jy, tJt^{-1} = J\rangle$$

is the Higman-Neumann-Neumman (HNN) extension of G' by this automorphism (we refer to [12, Chapter IV] for the properties of HNN extensions). As a result, G' is a subgroup of $\widetilde{G}_w$, and part (b) follows.

For part (c), if $w = 1$ in G, then $\mathrm{Area}_G(w^{-1}) = \mathrm{Area}_G(w)$, so

$$\mathrm{Area}_{G'}([x, w]) \leq 2\,\mathrm{Area}_G(w)$$

and similarly

$$\mathrm{Area}_{\widetilde{G}_w}([t, [x, w]]) \leq 2\,\mathrm{Area}_{G'}([x, w]) \leq 4\,\mathrm{Area}_G(w).$$

Thus we can use the relation $J = [t, [x, w]]$ to conclude that $\mathrm{Area}_{\widetilde{G}_w}(J) \leq 4\,\mathrm{Area}_G(w)+1$. ◀

The last result we include in this section is a lemma which will be used to translate area calculations into bounds on distances between vectors in Hilbert spaces. If u and v are two vectors in a Hilbert space $\mathcal{H}$, we write $u \approx_\epsilon v$ to mean that $\|u - v\| \leq \epsilon$. We use the standard terminology and notation of quantum information, so for instance, a state in a Hilbert space $\mathcal{H}$ is a unit vector $|\psi\rangle$ in $\mathcal{H}$.

▶ **Definition 7.** *Let $G = \langle S : R\rangle$ be a finitely-presented group. A (δ, ϵ)-*bipartite representation of G with respect to a state $|\psi\rangle$ in a Hilbert space $\mathcal{H}$ *is a pair of homomorphisms $\Phi, \Phi' : \mathcal{F}(S) \to \mathcal{U}(\mathcal{H})$ such that*

(i) $\Phi(r)\,|\psi\rangle \approx_\epsilon |\psi\rangle$ *for all $r \in R$,*
(ii) $\Phi(s)^{-1}\,|\psi\rangle \approx_\epsilon \Phi'(s)\,|\psi\rangle$ *for all $s \in S$, and*
(iii) $\|[\Phi(s), \Phi'(t)] - \mathbb{1}\| \leq \delta$ *for all $s, t \in S$ (here $\mathbb{1}$ represents the identity operator in $\mathcal{U}(\mathcal{H})$).*

In Part (iii) and throughout this paper the notation $\|A\|$ for an operator A refers to the operator norm of A. Part (i) of Definition 7 essentially says that Φ is an approximate representation of G with respect to the state $|\psi\rangle$. Parts (ii) and (iii) are less straightforward, but these conditions arise naturally in the theory of non-local games.

▶ **Lemma 8.** *Let (Φ, Φ') be a (δ, ϵ)-bipartite representation of a finitely-presented group $G = \langle S : R\rangle$ with respect to a state $|\psi\rangle \in \mathcal{H}$, and let ℓ be the length of the longest relation in R. If $w \in \mathcal{F}(S)$ is equal to the identity in G, then*

$$\Phi(w)\,|\psi\rangle \approx_{A(w)\cdot(\epsilon+\delta)} |\psi\rangle\,,$$

where $A(w) \leq 5\ell^2\,\mathrm{Area}_G(w)^2 + 2\ell|w|\,\mathrm{Area}_G(w)$.

Proof. If $r \in R$, then $\Phi(r)\,|\psi\rangle \approx_\epsilon |\psi\rangle$, and consequently $\Phi(r)^{-1}\,|\psi\rangle \approx_\epsilon |\psi\rangle$. Thus for any $r \in R$, $z \in \mathcal{F}(S)$, and $a \in \{\pm 1\}$,

$$\begin{aligned}\Phi(zr^a z^{-1})\,|\psi\rangle &= \Phi(z)\Phi(r)^a\Phi(z)^{-1}\,|\psi\rangle \approx_{|z|\epsilon} \Phi(z)\Phi(r)^a\Phi'(z)^{-1}\,|\psi\rangle \\ &\approx_{|r||z|\delta} \Phi(z)\Phi'(z)^{-1}\Phi(r)^a\,|\psi\rangle \approx_\epsilon \Phi(z)\Phi'(z)^{-1}\,|\psi\rangle \\ &\approx_{|z|\epsilon} \Phi(z)\Phi(z)^{-1}\,|\psi\rangle = |\psi\rangle\,.\end{aligned}$$

We conclude that $\Phi(zr^a z^{-1})\,|\psi\rangle \approx_{(2|z|+1)\epsilon+\ell|z|\delta} |\psi\rangle$. The result follows from Lemma 3. ◀

3 Approximately-commuting operator strategies and linear system games

A *two-party Bell scenario* $(\mathcal{I}_A, \mathcal{I}_B, \mathcal{O}_A^*, \mathcal{O}_B^*)$ consists of finite input sets $\mathcal{I}_A, \mathcal{I}_B$, a finite set of outputs $\mathcal{O}_A^x$ for every $x \in \mathcal{I}_A$, and a finite set of outputs $\mathcal{O}_B^y$ for every $y \in \mathcal{I}_B$.[3] The

[3] The sets $\mathcal{O}_A^x$ and $\mathcal{O}_B^y$ are often assumed to be independent of the inputs x and y. However, this assumption is not essential, since we can make the output sets independent of the input sets by adding filler answers to make all output sets the same size, and stipulating that Alice and Bob lose if they output one of the filler answers. When working with linear system games, it is more convenient to have the output sets depend on the inputs.

number of outputs in a Bell scenario is the maximum of $|\mathcal{O}_A^x|$ and $|\mathcal{O}_B^y|$ over $x \in \mathcal{I}_A$ and $y \in \mathcal{I}_B$. A *two-player non-local game* consists of a Bell scenario $(\mathcal{I}_A, \mathcal{I}_B, \mathcal{O}_A^*, \mathcal{O}_B^*)$, a function $V(\cdot,\cdot|x,y) : \mathcal{O}_A^x \times \mathcal{O}_B^y \to \{0,1\}$ for every $x \in \mathcal{I}_A$ and $y \in \mathcal{I}_B$, and a probability distribution π on $\mathcal{I}_A \times \mathcal{I}_B$. In the operational interpretation of the game, the referee sends players Alice and Bob inputs $x \in \mathcal{I}_A$ and $y \in \mathcal{I}_B$ with probability $\pi(x,y)$, the players reply with outputs $a \in \mathcal{O}_A^x$ and $b \in \mathcal{O}_B^y$, and the players win if and only if $V(a,b|x,y) = 1$.

In a non-local game, the players are not usually allowed to communicate while the game is in progress. Thus, in a quantum strategy for a game, it's assumed that each player determines their output by measuring their own local system. Locality can be enforced in one of two ways: by requiring that the joint system is the tensor product of the subsystems, or by requiring that measurement operators for different players commute with each other. Strategies of the former type are called tensor-product strategies, while strategies of the latter type are called commuting-operator strategies. Tensor-product strategies are commuting-operator strategies by definition, and finite-dimensional commuting-operator strategies can be turned into equivalent tensor-product strategies. In infinite dimensional Hilbert spaces, there are commuting-operator strategies for which the corresponding correlations do not have a tensor-product model [17]. However, it's still an open question as to whether all correlations arising from commuting-operator strategies can be realized as a limit of tensor-product strategies. By a theorem of Ozawa, this question is equivalent to the Connes embedding problem. In [15, 5], the notion of a quantum strategy has been generalized to approximately-commuting strategies, where Alice and Bob's systems are allowed to interact slightly. In this paper, we focus on the case of approximately-commuting operator strategies. Unlike [5], we use projective measurements rather than the more general POVM measurements in this definition. We refer to Remark 19 for some of the consequences of this difference.

▶ **Definition 9.** *A* δ-approximately-commuting operator strategy $\mathcal{S}$ *(or* δ-AC operator strategy *for short) for a Bell scenario* $(\mathcal{I}_A, \mathcal{I}_B, \mathcal{O}_A^*, \mathcal{O}_B^*)$ *consists of a Hilbert space* $\mathcal{H}$*, a projective measurement* $\{P_a^x\}_{a \in \mathcal{O}_A^x}$ *on* $\mathcal{H}$ *for every* $x \in \mathcal{I}_A$*, a projective measurement* $\{Q_b^y\}_{b \in \mathcal{O}_B^y}$ *on* $\mathcal{H}$ *for every* $y \in \mathcal{I}_B$*, and a state* $|\psi\rangle \in \mathcal{H}$ *such that*

$$\|P_a^x Q_b^y - Q_b^y P_a^x\| \leq \delta$$

for all $(x,y) \in \mathcal{I}_A \times \mathcal{I}_B$ *and* $(a,b) \in \mathcal{O}_A^x \times \mathcal{O}_B^y$*. A* δ-approximately-commuting quantum *(or* δ-AC quantum*)* strategy *is a* δ*-AC operator strategy in which* $\mathcal{H}$ *is finite-dimensional.*

Let $\mathcal{G} = (\mathcal{I}_A, \mathcal{I}_B, \mathcal{O}_A^*, \mathcal{O}_B^*, V, \pi)$ *be a non-local game. The* winning probability of $\mathcal{G}$ with strategy $\mathcal{S}$ *is*

$$\omega(\mathcal{G};\mathcal{S}) = \left| \sum_{x \in \mathcal{I}_A, y \in \mathcal{I}_B} \pi(x,y) \sum_{a \in \mathcal{O}_A, b \in \mathcal{O}_B} V(a,b|x,y) \langle\psi| P_a^x Q_b^y |\psi\rangle \right|.$$

The δ*-AC operator value* $\omega_\delta^{co}(\mathcal{G})$ *(resp.* δ*-AC quantum value* $\omega_\delta^*(\mathcal{G})$*) of* $\mathcal{G}$ *is defined to be the supremum of* $\omega(\mathcal{G};\mathcal{S})$ *across* δ*-AC operator strategies (resp.* δ*-AC quantum strategies).*

With this definition, a *commuting-operator strategy* is simply a 0-AC operator strategy, and the usual *commuting-operator value* of a game is $\omega^{co}(\mathcal{G}) := \omega_0^{co}(\mathcal{G})$. Since commuting-operator strategies are the same as tensor product strategies in finite dimensions, a *(tensor-product) quantum strategy* is simply a 0-AC quantum strategy, and the usual *quantum value* of a game is $\omega^*(\mathcal{G}) := \omega_0^*(\mathcal{G})$. Note that when $\delta = 0$, the absolute value can be dropped in the definition of $\omega(\mathcal{G};\mathcal{S})$. When $\delta > 0$, the values $\langle\psi| P_a^x Q_b^y |\psi\rangle$ can be complex, and the absolute value is necessary. This also means that $\omega(\mathcal{G},\mathcal{S})$ cannot necessarily be interpreted as a probability when $\mathcal{S}$ is approximately but not exactly commuting.

We look at a specific class of non-local games called linear system games. Let $Mx = c$ be an $m \times n$ linear system over $\mathbb{Z}_2$, so $M \in \mathbb{Z}_2^{m\times n}$ and $c \in \mathbb{Z}_2^m$. For each $1 \leq i \leq m$, let $V_i = \{1 \leq j \leq n : M_{ij} \neq 0\}$. The linear system game $\mathcal{G}_{Mx=c}$ is the non-local game with

$$\mathcal{I}_A = \{1, \ldots, m\}, \quad \mathcal{I}_B = \{1, \ldots, n\},$$

$$\mathcal{O}_A^i = \left\{ a \in \mathbb{Z}_2^{V_i} : \sum_{j\in V_i} a_j = c_i \right\}, \quad \mathcal{O}_B^j = \mathbb{Z}_2,$$

$$V(a, b|i, j) = \begin{cases} 1 & j \notin V_i \text{ or } a_j = b \\ 0 & \text{otherwise} \end{cases},$$

and π the uniform distribution over pairs (i, j) such that $j \in V_i$. In other words, Alice receives the index i of an equation and Bob receives the index j of a variable, chosen uniformly at random from pairs (i, j) with $j \in V_i$. Alice replies with a satisfying assignment to the variables which appear in the ith equation, and Bob replies with an assignment for the jth variable. The players win if Alice and Bob both give the same assignment to variable j.

For linear system games, it is often convenient to express strategies in terms of observables, rather than measurement operators (see, for instance, [4, 3]). If $\mathcal{S} = (\mathcal{H}, \{P_a^i\}_{a\in\mathcal{O}_A^i}, \{Q_b^j\}_{b\in\mathbb{Z}_2}, |\psi\rangle)$ is a δ-AC strategy for $\mathcal{G}_{Mx=c}$, the corresponding observables are

$$A_{ij} := \sum_{a\in\mathcal{O}_A^i} (-1)^{a_j} P_a^i \text{ for } 1 \leq i \leq m, j \in V_i, \tag{3.1}$$

and

$$B_j := Q_0^j - Q_1^j \text{ for } 1 \leq j \leq n. \tag{3.2}$$

These operators are ± 1-valued observables (meaning, self-adjoint unitary operators) satisfying the equations

$$\prod_j A_{ij}^{M_{ij}} = (-\mathbb{1})^{c_i} \text{ for all } 1 \leq i \leq m, \tag{3.3}$$

$$[A_{ij}, A_{ij'}] = \mathbb{1} \text{ whenever } j, j' \in V_i \text{ for some } 1 \leq i \leq m, \text{ and} \tag{3.4}$$

$$\|[A_{ij}, B_k] - \mathbb{1}\| \leq 2^{|V_i|+1}\delta \text{ for all } 1 \leq i \leq m, j \in V_i, \text{ and } 1 \leq k \leq n. \tag{3.5}$$

We can recover the projections P_a^i, $a \in \mathcal{O}_A^i$, and Q_b^j, $b \in \mathcal{O}_B^j$, from the observables A_{ik} and B_j via the formulas

$$P_a^i = \prod_{k\in V_i} \left(\frac{\mathbb{1} + (-1)^{a_k} A_{ik}}{2} \right) \text{ and } Q_b^j = \frac{1 + (-1)^b B_j}{2}. \tag{3.6}$$

We define *bias of strategy* $\mathcal{S}$ to be

$$\beta(\mathcal{G}_{Mx=c}; \mathcal{S}) := \sum_{1\leq i\leq m} \sum_{j\in V_i} \pi(i, j) \langle\psi| A_{ij} B_j |\psi\rangle .$$

It is not hard to see that

$$\omega(\mathcal{G}_{Mx=c};\mathcal{S}) = \frac{1}{2}|\beta(\mathcal{G}_{Mx=c},\mathcal{S})+1|,$$

so we can work with the winning probability using observables as well.

It follows from [3] that when $\delta = 0$, perfect commuting-operator strategies of $\mathcal{G}_{Mx=c}$ can be understood using the following group.

▶ **Definition 10.** *Let $Mx = c$ be an $m \times n$ linear system over $\mathbb{Z}_2$. Then the* solution group *of the system is the finitely presented group $\Gamma_{Mx=c}$ generated by $x_1, \ldots, x_n, J$, and satisfying relations*

1. $[x_i, J] = x_i^2 = J^2 = 1$ *for all* $1 \le i \le n$,
2. $\prod_j x_j^{M_{ij}} = J^{c_i}$ *for all* $1 \le j \le m$, *and*
3. $[x_j, x_k] = 1$ *if there is some* $1 \le i \le m$ *with* $M_{ij}, M_{ik} \neq 0$.

We consider $\Gamma = \Gamma_{Mx=c}$ to be a group over $\mathbb{Z}_2$ with J_Γ equal to the generator J.

In particular, we can characterize when the optimal winning probability of the game is equal to 1 using this group.

▶ **Theorem 11** ([3, 18])**.** *Let $Mx = c$ be a linear system over $\mathbb{Z}_2$. Then*

(a) $\omega^{co}(\mathcal{G}_{Mx=c}) = 1$ *if and only if* $J \neq 1$ *in* $\Gamma_{Mx=c}$, *and*

(b) $\omega^*(\mathcal{G}_{Mx=c}) = 1$ *if and only if J is non-trivial in approximate representations of* $\Gamma_{Mx=c}$.

For the definition of *non-trivial in approximate representations*, we refer to [18].

Near-perfect finite-dimensional strategies of $\mathcal{G}_{Mx=c}$ correspond to approximate representations of $\Gamma_{Mx=c}$ [18]. We want to develop this theory when $\delta > 0$.

▶ **Proposition 12.** *Let $Mx = c$ be an $m \times n$ linear system, let $V_i := \{1 \le j \le n : M_{ij} \neq 0\}$, let $r := \max_i |V_i|$ be the maximum number of non-zero entries in any row, and let $K := \sum_{i=1}^m |V_i|$ be the number of non-zero entries in M. Suppose $\mathcal{S} = (\mathcal{H}, \{P_a^x\}, \{Q_b^y\}, |\psi\rangle)$ is a δ-AC operator strategy with $\omega(\mathcal{G}_{Mx=c};\mathcal{S}) \ge 1-\epsilon$ for some $\epsilon, \delta \ge 0$. Let A_{ij}, B_k be the corresponding observables defined in Equations* (3.1) *and* (3.2)*. Then*

(a) $A_{ij}|\psi\rangle \approx_{2\sqrt{K(\epsilon+2^{r-1}\delta)}} B_j|\psi\rangle$ *for all* $1 \le i \le m$ *and* $j \in V_i$,

(b) $\prod_{j=1}^m B_{ij}^{M_{ij}}|\psi\rangle \approx_{2r\sqrt{K(\epsilon+2^{r-1}\delta)}+\binom{r}{2}2^{r+1}\delta} (-\mathbb{1})^{c_i}|\psi\rangle$ *for all* $1 \le i \le m$, *and*

(c) $[B_j, B_k]|\psi\rangle \approx_{8\sqrt{K(\epsilon+2^{r-1}\delta)}+6\cdot 2^{r+1}\delta} |\psi\rangle$ *whenever there is* $1 \le i \le m$ *with* $j, k \in V_i$.

Proof. For part (a), any two unit vectors $|\psi\rangle$ and $|\phi\rangle$ satisfy $|\psi\rangle \approx_2 |\phi\rangle$, so we can assume that $\epsilon + 2^{r-1}\delta \le 1$. Write β for $\beta(\mathcal{G}_{Mx=c},\mathcal{S})$, and observe that

$$\begin{aligned}|2\operatorname{Im}\beta| = |\beta - \overline{\beta}| &= \left|\sum_{i,j}\pi(i,j)\langle\psi|A_{ij}B_j - B_jA_{ij}|\psi\rangle\right| \\ &\le \sum_{i,j}\pi(i,j)\|A_{ij}B_j - B_jA_{ij}\| \le 2^{r+1}\delta\end{aligned}$$

by Equation (3.5). Since $\omega(\mathcal{G}_{Mx=c};\mathcal{S}) \ge 1-\epsilon$, we have that

$$(1-\epsilon)^2 \le \left|\frac{\beta+1}{2}\right|^2 = \frac{(\operatorname{Re}\beta+1)^2 + (\operatorname{Im}\beta)^2}{4} \le \frac{(\operatorname{Re}\beta+1)^2 + (2^r\delta)^2}{4}.$$

Since $A_{ij}|\psi\rangle$ and $B_j|\psi\rangle$ are unit vectors, $-1 \le \operatorname{Re}\beta \le 1$, and in particular $\operatorname{Re}\beta + 1 \ge 0$. Thus

$$\operatorname{Re}\beta + 1 \ge \sqrt{4(1-\epsilon)^2 - (2^r\delta)^2} = \sqrt{(2-2\epsilon-2^r\delta)(2-2\epsilon+2^r\delta)} \ge 2 - 2\epsilon - 2^r\delta,$$

where the last inequality holds because of the assumption $2\epsilon + 2^r\delta \leq 2$. We conclude that $\operatorname{Re}\beta \geq 1 - 2\epsilon - 2^r\delta$, or $1 - \operatorname{Re}\beta \leq 2\epsilon + 2^r\delta$.

Now $\pi(i,j) = 1/K$ for all $1 \leq i \leq m$, $j \in V_i$, so

$$1 - \operatorname{Re}\beta = \frac{1}{K}\sum_{i,j}(1 - \operatorname{Re}\langle\psi| A_{ij}B_j |\psi\rangle) \leq 2\epsilon + 2^r\delta.$$

Since $\operatorname{Re}\langle\psi| A_{ij}B_j |\psi\rangle \leq 1$, we have that $1 - \operatorname{Re}\langle\psi| A_{ij}B_j |\psi\rangle \leq 2K(\epsilon + 2^{r-1}\delta)$ for all $1 \leq i \leq m$ and $j \in V_i$. So

$$\|A_{ij}|\psi\rangle - B_j|\psi\rangle\|^2 = 2 - 2\operatorname{Re}\langle\psi| A_{ij}B_j |\psi\rangle \leq 4K(\epsilon + 2^{r-1}\delta),$$

finishing the proof of part (a).

For parts (b) and (c), let $\tau = 2\sqrt{K(\epsilon + 2^{r-1}\delta)}$. Given $1 \leq i \leq m$, let $V_i = \{j_1, \ldots, j_k\}$, where $1 \leq j_1 < \ldots < j_k \leq n$. Then

$$B_{j_1}\cdots B_{j_k}|\psi\rangle \approx_\tau B_{j_1}\cdots B_{j_{k-1}}A_{ij_k}|\psi\rangle \approx_{(k-1)2^{r+1}\delta} A_{ij_k}B_{j_1}\cdots B_{j_{k-1}}|\psi\rangle.$$

Continuing this pattern, we see that

$$B_{j_1}\cdots B_{j_k} \approx_{k\tau + \binom{k}{2}2^{r+1}\delta} A_{ij_k}A_{ij_{k-1}}\cdots A_{ij_1}|\psi\rangle = (-\mathbb{1})^{c_i}|\psi\rangle,$$

where the last equality is Equation (3.3). Part (c) follows similarly from Equation (3.4). ◀

▶ **Corollary 13.** *Using the notation and hypotheses of Proposition 12, if we define* $\Phi, \Phi' : \mathcal{F}(x_1, \ldots, x_n, J) \to \mathcal{U}(\mathcal{H})$ *by*

$$\Phi(x_j) = B_j \text{ for all } 1 \leq j \leq n, \quad \Phi(J) = -\mathbb{1}$$

and

$$\Phi'(x_j) = \begin{cases} A_{ij} & \text{any } i \text{ such that } j \in V_i \\ \mathbb{1} & \text{if no such } i \text{ exists} \end{cases}, \quad \Phi'(J) = -\mathbb{1}$$

then (Φ, Φ') *is a* (τ, κ)*-bipartite representation of* $\Gamma_{Mx=c}$ *with respect to* $|\psi\rangle$*, where*

$$\tau = 2\max(r,4)\sqrt{K(\epsilon + 2^{r-1}\delta)} + \binom{\max(r,4)}{2}2^{r+1}\delta$$

and $\kappa = 2^{r+1}\delta$.

Proof. Follows immediately from Proposition 12, Equation (3.5), and the fact that $B_j^2 = \mathbb{1}$. ◀

4 Embedding finitely-presented groups in solution groups

By Theorem 4, every recursive language can be efficiently encoded as the word problem of a finitely-presented group. By Lemma 6, the word problem for finitely-presented groups reduces to the problem of determining whether $J_G = 1$ in finitely-presented groups G over $\mathbb{Z}_2$. By Theorem 11, if $G = \Gamma_{Mx=c}$ is a solution group, then $J_G = 1$ if and only if $\omega^{co}(\mathcal{G}_{Mx=c}) = 1$.

The main result of [17] is that the problem of determining whether $J_G = 1$ for general finitely-presented groups G over $\mathbb{Z}_2$ reduces to the problem of determining whether $J_\Gamma = 1$ for solution groups $\Gamma = \Gamma_{Mx=c}$. In this paper, we use the following version of this result:

▶ **Theorem 14** ([17]). *Let $G = \langle S : R\rangle$ be a finitely presented group over $\mathbb{Z}_2$, such that $J_G \in S$, and let $N = |S| + \sum_{r\in R} |r|$ be the size of the presentation. Then there is an $m \times n$ linear system $Mx = c$ and a map $\phi : \mathcal{F}(S) \to \mathcal{F}(x_1, \ldots, x_n, J)$ such that*

(a) *$\phi(J_G) = J_\Gamma$, and ϕ descends to an injection $G \to \Gamma_{Mx=c}$ (in other words, for all $w \in \mathcal{F}(S)$, $\phi(w)$ is trivial in $\Gamma_{Mx=c}$ if and only if w is trivial in G);*

(b) *for all $w \in \mathcal{F}(S)$, $|\phi(w)| \leq 4|w|$, and if w is trivial in G, then $\mathrm{Area}_\Gamma(\phi(w)) = O(N \cdot \mathrm{Area}_G(w))$; and*

(c) *M has exactly three non-zero entries in every row, the dimensions m and n of M are $O(N)$, and M and b can be constructed from $\langle S : R\rangle$ in time polynomial in N.*

Note that if $G = \langle S : R\rangle$ and $G' = \langle S' : R'\rangle$ are finitely-presented groups, and $\phi : \mathcal{F}(S) \to \mathcal{F}(S')$ is a homomorphism which descends to a homomorphism $G \to G'$, then $\mathrm{Area}_{G'}(\phi(w)) = O(\mathrm{Area}_G(w))$, with a constant which depends on G, G', and ϕ. The statement in part (b) of Theorem 14 is stronger, in that the constant is independent of G (so the only dependence on G comes from N).

Proof of Theorem 14. Part (a) is Theorem 3.1 of [17]. For the complexity statements in parts (b) and (c), we need to analyze the construction of M and b, which occurs in Proposition 4.3, Corollary 4.8, and Theorem 5.1 of [17]. For this purpose, suppose that $G = \langle S : R\rangle$ is a finitely presented group over $\mathbb{Z}_2$. For simplicity, we assume that $J_G = J \in S$, and that all relations containing J are of the form $J \cdot r = 1$ for some word $r \in \mathcal{F}(S \setminus \{J\})$. This assumption can always be satisfied by adding an extra generator.

For the first step of the construction, we also need some notation. If $x \in \mathcal{F}(S')$ is equal to the reduced word $s_1^{a_1} \cdots s_k^{a_k}$, where $s_i \in S'$ and $a_i \in \{\pm 1\}$ for all $1 \leq i \leq k$, let $x^+ = s_1 \cdots s_k$. Note that this word is still reduced, and that x and x^+ represent the same element in the group

$$\langle S' : s^2 = 1 \text{ for all } s \in S'\rangle.$$

Now, starting from $G = \langle S : R\rangle$, we take a new set of indeterminates $S' = \{u_s, v_s : s \in S \setminus \{J\}\}$, and define $\phi_1 : \mathcal{F}(S) \to \mathcal{F}(S' \cup \{J\})$ by $\phi_1(s) = u_s v_s u_s v_s$ for all $s \in S \setminus \{J\}$ and $\phi_1(J) = J$. We then let

$$G' = \langle S' \cup \{J\} : R' \cup \{u_s^2 = v_s^2 = 1 : s \in S \setminus \{J\}\} \cup \{J^2 = 1\}\rangle,$$

where $R' = \{\phi_1(r)^+ : r \in R\}$. Since $u_s^2 = v_s^2 = J^2 = 1$ in G', we conclude that ϕ_1 descends to a homomorphism $\phi_1 : G \to G'$. It is not hard to see that this morphism is injective (see, for instance, [17, Proposition 4.3]), and clearly $|\phi_1(w)| \leq 4|w|$. If $r \in R$, then $\phi_1(r)$ can be turned into $\phi_1(r)^+$ in at most $4|r|$ applications of the relations $u_s^2 = v_s^2 = 1$, $s \in S \setminus \{J\}$. (In particular, $\mathrm{Area}_{G'}(\phi_1(w)) \leq 4N \,\mathrm{Area}_G(w)$, although we use a more refined calculation for bound on Area_Γ in part (b).) The size of the presentation of G' is

$$N' = |S'| + 1 + \sum_{r\in R'} |r| + 4|S| - 2 \leq 6|S| + 4\sum_{r\in R} |r| \leq 6N,$$

and the presentation can be constructed from $\langle S : R\rangle$ in $O(N)$ time.

To finish the construction of $Mx = c$, we apply the wagon wheel construction from Section 5 of [17] to the group G'. This construction is best understood pictorially. An $m \times n$ matrix M with entries in $\mathbb{Z}_2$ can be represented graphically by drawing a hypergraph with a vertex for each row of M, and an edge for each column, such that the jth hyperedge is incident to the ith vertex if and only if $M_{ij} = 1$. With this representation, a vector $b \in \mathbb{Z}_2^m$ is

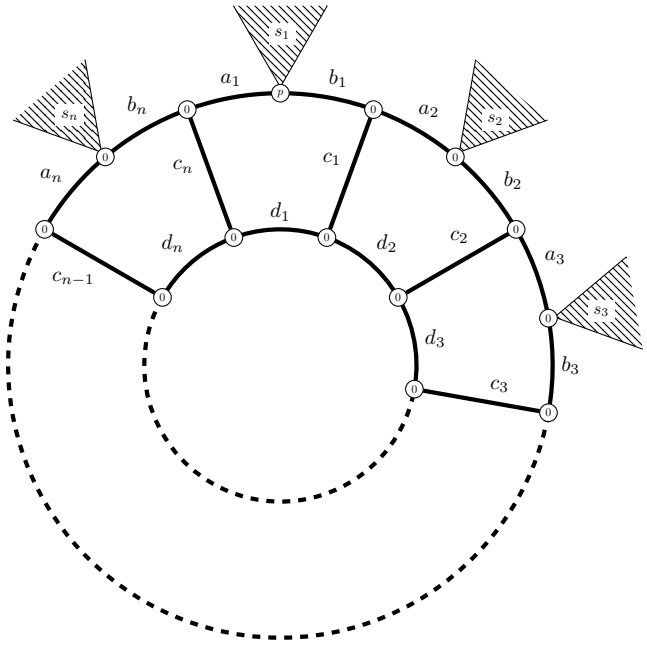

Figure 1 Pictorial depiction of the linear system associated to each relation in the wagon wheel embedding as described in the proof of Theorem 14. Figure reproduced from [17, Figure 2].

the same as function from the vertices of the hypergraph to $\mathbb{Z}_2$. So a linear system $Mx = c$ can thus be represented by a hypergraph with a (not necessarily proper) $\mathbb{Z}_2$-vertex colouring, where the edges correspond to the variables of the system, and the vertices to the equations.

In the wagon wheel construction, $Mx = c$ is defined as a union of subsystems $M^r x^r = c^r$, each corresponding to a relation $r \in R'$. The variables of $Mx = c$ consist of the indeterminates S', as well as an additional set of ancillary variables S''. Each ancillary variable appears in exactly one of the subsystems $M^r x^r = c^r$, while the variables S' are shared. If $r = J^p s_1 \cdots s_n$, where $p \in \mathbb{Z}_2$ and $s_1, \ldots, s_n \in S'$, then the portion of the hypergraph of $Mx = c$ corresponding to $M^r x^r = c^r$ is shown in Figure 1, with the ancillary variables denoted by a_i, b_i, c_i, d_i, $1 \leq i \leq n$. The vertex colouring is also shown in Figure 1: one vertex is given colour p, and the remaining vertices are coloured 0.

As can be seen from Figure 1, the number of ancillary variables added for subsystem $M^r x^r = c^r$ is $4|r|$, and the number of equations added is $3|r|$. Since every vertex in the hypergraph has degree three, every row of M^r has exactly three non-zero entries. Theorem 5.1 of [17] then states that the natural inclusion $\phi_2 : \mathcal{F}(S' \cup \{J\}) \to \mathcal{F}(S' \cup \{J\} \cup S'') : s \mapsto s$ descends to an injection $G' \to \Gamma_{Mx=c}$.

Recall from Definition 10 that every linear equation in $Mx = c$ becomes a defining relation of $\Gamma := \Gamma_{Mx=c}$. The wagon wheel construction is designed so that if $r \in R'$, then $\phi_2(r)$ can be turned into the identity by applying each defining relation from $M_r x = c_r$ exactly once, so $\mathrm{Area}_\Gamma(\phi_2(r)) \leq 3|r|$ for all $r \in R'$. This is easiest to see using pictures of the group, for which we refer to Section 7 of [17]; with this formalism, Figure 1 is itself a proof that $\phi_2(r) = 1$, with each vertex corresponding to a use of the corresponding relation. For relations $r = J^p s_1 \cdots s_n$ with $p \neq 1$, we start with the relation coloured by p, after which J no longer appears in the word. If $r \in R$, then $\phi_2(\phi_1(r))$ can be turned into the identity with at most $7|r|$ applications of the relations of Γ, by first changing $\phi_2(\phi_1(r))$ to $\phi_2(\phi_1(r)^+)$ using the relations $s^2 = 1$, $s \in S'$, and then applying the linear relations of Γ. It follows

that $\mathrm{Area}_\Gamma(\phi_2(w)) = O(N \mathrm{Area}_G(w))$ for all $w \in \mathcal{F}(S)$ which are trivial in G. It should also be clear from Figure 1 that $M^r x^r = c^r$ can be constructed in time polynomial in $|r|$. We conclude that $Mx = c$ is an $m \times n$ linear system with m and n equal to $O(N)$, and that M and b can be constructed in time polynomial in N, so the theorem holds with $\phi = \phi_2 \circ \phi_1$. ◀

Theorem 14 is sufficient to prove Theorem 1. However, to get perfect zero-knowledge protocols for MIP^{co}_δ, we need to prove an additional fact about the embedding in Theorem 14.

▶ **Lemma 15.** *Let $Mx = c$ be an $m \times n$ linear system from the wagon wheel construction in the proof of Theorem 14. In the solution group $\Gamma_{Mx=c}$, the generator x_i is not equal to 1 or J for all $1 \leq i \leq n$, and similarly the product $x_i x_j$ is not equal to 1 or J for all $1 \leq i \neq j \leq n$.*

Proof. We revisit the wagon wheel construction in the proof of Theorem 14. We need to show that $x_i \neq 1$ and $x_i \neq x_j$ in $\Gamma_0 := \Gamma_{Mx=c}/\langle J \rangle$ for all $1 \leq i \neq j \leq n$. This is the same as showing that $x_i \neq 1$ and $x_i \neq x_j$ in $\Gamma_{Mx=0} = \Gamma_0 \times \mathbb{Z}_2$. Recall that the generators of $\Gamma_{Mx=0}$ are split into two sets, the generators S' of G', and the ancillary variables S''. The group $G'_0 := G'/\langle J \rangle$ has a presentation where every generator $s \in S'$ occurs an even number of times in every relation. Thus for any $s \in S'$, we can define a representation $G' \to \mathbb{C}^x$ by sending $s \in S'$ to -1, and $t \in S' \setminus \{s\}$ to 1. It follows that $s \neq 1$ and $s \neq t$ in G'_0 for every $s \neq t \in S'$. Since $G'_0 \to \Gamma_0$ is an injection, we conclude that the same holds in Γ_0.

For the ancillary variables, consider the hypergraph description of the system $Mx = 0$. Given a subset of edges C, let $y \in \mathbb{Z}_2^n$ be the vector with $y_i = 1$ if and only if the ith edge is in C. Then y is a classical solution to $Mx = 0$ if and only if every vertex of the hypergraph is incident with an even number of edges from C. The classical solutions of $Mx = 0$ correspond to 1-dimensional representations of Γ_0; if y is a solution of $Mx = 0$, then the corresponding 1-dimensional representation of Γ_0 sends $x_i \mapsto (-1)^{y_i}$.

Inspecting the wagon wheel hypergraph in Figure 1, we see that every ancilla variable $s \in S''$ belongs to a cycle C which does not contain any edges from S'. Using the corresponding representation of Γ_0, we see that $s \neq 1$ and $s \neq t$ in Γ_0 for all $s \in S''$ and $t \in S'$. Similarly, if $s \neq t \in S''$, and $\{s, t\}$ is not one of the pairs $\{a_i, b_i\}$, then there is a cycle C containing s and not containing t, so $s \neq t$ in Γ_0.

For the pairs $\{a_i, b_i\}$, fix $s \in S''$, and recall that if $r = s_1 \cdots s_n$ is a relation of G', where $s_1, \ldots, s_n \in S'$, then s occurs an even number of times in r. Let $1 \leq i_1 < \cdots < i_{2k} \leq n$ be the indices such that $s_{i_j} = s$, and let

$$\begin{aligned} C_r := \{&s_{i_1}, b_{i_1}, a_{i_1+1}, b_{i_1+1}, \ldots, a_{i_2}, s_{i_2}, \\ &s_{i_3}, b_{i_3}, a_{i_3+1}, b_{i_3+1}, \ldots, a_{i_4}, s_{i_4}, \\ &\ldots, \\ &s_{i_{2k-1}}, b_{i_{2k-1}}, a_{i_{2k-1}+1}, b_{i_{2k-1}+1}, \ldots, a_{i_{2k}}, s_{i_{2k}}\}. \end{aligned}$$

be the collection of paths along the outer cycle of the wagon wheel connection s_{i_1} with s_{i_2}, s_{i_3} with s_{i_4}, and so on. Let $C := \bigcup_{r \in R'} C_r$. Then every vertex of the hypergraph of $Mx = 0$ is incident to an even number of edges in C. If we look at a particular relation r, then for every $1 \leq j \leq 2k$, exactly one of the edges a_{i_j}, b_{i_j} belongs to C_r, so $a_{i_j} \neq b_{i_j}$ in Γ_0. It follows that all of the pairs of ancillary generators a_i, b_i are distinct in Γ_0. ◀

▶ **Proposition 16.** *Let $Mx = c$ be an $m \times n$ linear system from the wagon wheel construction in the proof of Theorem 14, and suppose $J \neq 1$ in $\Gamma_{Mx=c}$. Then $\mathcal{G}_{Mx=c}$ has a commuting-operator strategy $\mathcal{S} = (\mathcal{H}, \{P_a^i\}_{a \in \mathcal{O}_a^i}, \{Q_b^j\}_{b \in \mathbb{Z}_2}, |\psi\rangle)$ such that $\omega(\mathcal{G}_{Mx=c}; \mathcal{S}) = 1$, and*

$$\langle\psi| P_a^i Q_b^j |\psi\rangle = \begin{cases} \frac{1+(-1)^{a_j+b}}{8} & j \in V_i \\ \frac{1}{8} & j \notin V_i \end{cases}$$

for all $1 \leq i \leq m$, $a \in \mathcal{O}_A^i$, $1 \leq j \leq m$, $b \in \mathbb{Z}_2$.

Proof. Suppose $J \neq 1$ in $\Gamma_{Mx=c}$. We recall the construction of a perfect commuting-operator strategy for $\mathcal{G}_{Mx=c}$ from [3]. Let $\mathcal{H} = \ell^2\Gamma_{Mx=c}$ be the regular representation of $\Gamma_{Mx=c}$, and given $g \in \Gamma_{Mx=c}$, let $L(g)$ (resp. $R(g)$) denote left (resp. right) multiplication by g. Then $L(g)$ and $R(g)$ are unitaries for all $g \in \Gamma_{Mx=c}$, and we can get a perfect strategy for $\mathcal{G}_{Mx=c}$ by taking $A_{ij} = L(X_j)$ for all $1 \leq i \leq m$, $j \in V_i$, $B_j = R(X_j)$ for all $1 \leq j \leq n$, and $|\psi\rangle = \frac{1-J}{\sqrt{2}}$ considered as an element of $\mathcal{H}$. Since J is central of order 2, we have that

$$\langle\psi| A_{ik}B_j |\psi\rangle = \langle\psi| L(X_k)R(X_j) |\psi\rangle = \langle\psi| R(X_kX_j) |\psi\rangle = \begin{cases} 1 & X_kX_j = 1 \\ -1 & X_kX_j = J \\ 0 & \text{otherwise.} \end{cases}.$$

Recall from Equation (3.6) that

$$P_a^i = \prod_{k \in V_i} \left(\frac{\mathbb{1} + (-1)^{a_k} A_{ik}}{2} \right)$$

for all $a \in \mathcal{O}_A^i$ and $Q_b^j = \frac{1+(-1)^b B_j}{2}$ for all $b \in \mathcal{O}_B^j$. Using the fact that $\prod_{k \in V_i} A_{ik} = (-\mathbb{1})^{c_i}$ in perfect strategies, and that $|V_i| = 3$ in the linear system constructed in Theorem 14, we get that

$$P_a^i = \prod_{k \in V_i} \left(\frac{\mathbb{1} + (-1)^{a_k} A_{ik}}{2} \right) = \frac{\mathbb{1}}{4} + \frac{1}{4} \sum_{k \in V_i} (-1)^{a_k} A_{ik}.$$

By Lemma 15

$$\langle\psi| P_a^i Q_b^j |\psi\rangle = \frac{1}{8} + \frac{1}{8} \langle\psi| \sum_{k \in V_i} (-1)^{a_k+b} A_{ik}B_j |\psi\rangle = \begin{cases} \frac{1}{8} & j \notin V_i \\ \frac{1+(-1)^{a_j+b}}{8} & j \in V_i. \end{cases}. \quad \blacktriangleleft$$

5 Proof of Theorem 1

In this section we prove Theorem 1, by proving the main technical result of the paper.

▶ **Theorem 17.** *Let $L \subset A^*$ be a language over a finite alphabet A, and contained in* NTIME$(T(n))$*, where $T(n)^4$ is superadditive. Then for any string $w \in A^*$, there is a non-local game $\mathcal{G}_w$ such that*

(a) *the game $\mathcal{G}_w$ has question sets of size $O(|w|)$ and output sets of size at most 8,*
(b) *the function $w \mapsto \mathcal{G}_w$ is computable in $O(|w|^k)$-time, where k is some universal constant,*
(c) *if $w \notin L$ then $\omega^{co}(\mathcal{G}_w) = 1$, and*
(d) *if $w \in L$ then*

$$\omega_\delta^{co}(\mathcal{G}_w) \leq 1 - \frac{1}{T(O(|w|))^{k'}} + O(\delta)$$

for some universal constant k'.

While the constants k, k' in Theorem 17 are independent of L, the other constants appearing in the big-O can depend on L. The game $\mathcal{G}_w$ will be a linear system game $\mathcal{G}_{M(w)x=c(w)}$, where $M(w)x = c(w)$ is an $O(|w|) \times O(|w|)$-linear system. Since the linear system game of an $m \times n$ linear system $Mx = c$ can be constructed in $O(mn)$-time from M and b, the goal in proving Theorem 17 will be to show that the linear system $M(w)x = c(w)$ can be constructed in time polynomial in $|w|$. Theorem 1 is an immediate corollary of Theorem 17.

Proof of Theorem 17. Given the language L, let $G = \langle S : R \rangle$ be the group from Theorem 4, and let κ be the function $A^* \to \mathcal{F}(S)$. Given $w \in A^*$, we let $\widetilde{G}_{\kappa(w)}$ be the group over $\mathbb{Z}_2$ constructed in Definition 5, and $M(w)x = c(w)$ be the linear system constructed from $\widetilde{G}_{\kappa(w)}$ in Theorem 14. Finally, we let $\mathcal{G}_w := \mathcal{G}_{M(w)x=c(w)}$ and $\Gamma_w := \Gamma_{M(w)x=c(w)}$. The only part of the presentation of $\widetilde{G}_{\kappa(w)}$ that changes with w is the relation $[t, [x, \kappa(w)]] = 1$, so the presentation of $\widetilde{G}_{\kappa(w)}$ has size $O(|\kappa(w)|) = O(|w|)$, and $M(w)x = c(w)$ is an $O(|w|) \times O(|w|)$ linear system. Because $M(w)$ has only three non-zero entries per equation, Alice's output sets in $\mathcal{G}_w$ will have size $2^3 = 8$, while Bob's output sets will have size 2. Thus parts (a) and (b) of Theorem 17 follow from part (c) of Theorem 14.

By Theorem 4 and Lemma 6, if $w \notin L$ then $\kappa(w) \neq 1$ in G, and hence $J \neq 1$ in $\widetilde{G}_{\kappa(w)}$. Since the inclusion $\widetilde{G}_{\kappa(w)} \hookrightarrow \Gamma_w$ sends $J_{\widetilde{G}_{\kappa(w)}} \mapsto J_{\Gamma_w}$, we conclude that $J \neq 1$ in Γ_w. By Theorem 11, $\omega^{co}(\mathcal{G}_w) = 1$, proving part (c).

This leaves part (d). Suppose $w \in L$. Then $\kappa(w) = 1$ in G, $J = 1$ in $\widetilde{G}_{\kappa(w)}$, and hence $J = 1$ in Γ_w. Suppose $\mathcal{S}$ is a δ-AC operator strategy for $\mathcal{G}_w$ with $\omega(\mathcal{G}_w; \mathcal{S}) \geq 1 - \epsilon$. Since M has only three non-zero entries per row, the parameters r and K appearing in Corollary 13 are $O(1)$ and $O(|w|)$ respectively. Also, because we are interested in $\delta \leq 2$, we can say that $\delta = O(\sqrt{\delta})$. Thus Corollary 13 states that there is a $(O(\sqrt{|w|(\epsilon + \delta)}), O(\delta))$-bipartite representation (Φ, Φ') of $\mathcal{G}_w$ with respect to the state $|\psi\rangle$ used in $\mathcal{S}$. By construction, this bipartite representation has $\Phi(J) = -\mathbb{1}$. The length of the longest relation in Γ_w is 4, and the length of J in Γ_w is 1, so Lemma 8 implies that

$$-|\psi\rangle = \Phi(J)|\psi\rangle \approx_{O\left(\mathrm{Area}_{\Gamma_w}(J)^2 \sqrt{|w|(\epsilon+\delta)}\right)} |\psi\rangle . \tag{5.1}$$

By Theorem 14, part (b) and Lemma 6, part (c),

$$\mathrm{Area}_{\Gamma_w}(J) = O\left(|w| \cdot \mathrm{Area}_{\widetilde{G}_{\kappa(w)}}(J)\right) = O\left(|w| \cdot \mathrm{Area}_G(\kappa(w))\right).$$

Finally, by Theorem 4, $|\kappa(u)| = O(|u|)$ and Dehn_G is bounded by a function equivalent to $T(n)^4$, so there is a constant C such that $\mathrm{Area}_G(\kappa(w)) = O(T(C|w|)^4 + |w|)$. Since $T(n)^4$ is superadditive by assumption, $|w| = O(T(|w|)^4)$, and we can conclude that $\mathrm{Area}_{\Gamma_w}(J) = O(T(C|w|)^8)$. Returning to Equation (5.1), since $\|-|\psi\rangle - |\psi\rangle\| = 2$, we see that there is a constant $C_0 > 0$ such that

$$C_0 T(C|w|)^{18} \sqrt{\epsilon + \delta} \geq 2.$$

Hence

$$C_0^2 T(C|w|)^{36} (\epsilon + \delta) \geq 4.$$

so,

$$\epsilon \geq \frac{4}{C_0^2 T(C|w|)^{36}} - \delta,$$

So we conclude that

$$\omega(\mathcal{G}_w;\mathcal{S}) \leq 1 - \Omega\left(\frac{1}{T(C|w|)^{36}}\right) + O(\delta).$$

Because $T(n)^4$ is superadditive, $T(C_0 \cdot C|w|)^{36} \geq C_0 T(C|w|)^{36}$ for any integer C_0, so we can move the constant from the big-Ω inside T, proving part (d). ◀

6 Multi-prover interactive proofs

In this section we define the complexity class PZK-MIP$^{co}_{\delta}(2,1,1,1-1/f(n))$, and prove Theorem 2. We first recall the definition of MIP$^{co}_{\delta}$. The definition given here is a simple variant on Definition 8 of [5].

▶ **Definition 18.** *A language L over an alphabet A is in the class* MIP$^{co}_{\delta}(2,1,1,1-1/f(n))$ *of multi-prover interactive proofs with two provers, one round, completness probability* 1*, and soundness probability* $1-1/f(n)$*, if and only if there is family of two-player non-local games* $\mathcal{G}_w = (\mathcal{I}_A^w, \mathcal{I}_B^w, \mathcal{O}_A^{*,w}, \mathcal{O}_B^{*,w}, V_w, \pi_w)$ *indexed by strings* $w \in A^*$*, such that*

- *the input sets* $\mathcal{I}_A^w$*,* $\mathcal{I}_B^w$ *and output sets* $\mathcal{O}_A^{*,w}$*,* $\mathcal{O}_B^{*,w}$ *for* $\mathcal{G}_w$ *are subsets of strings of length* poly $|w|$ *(and hence can have size at most* $2^{\text{poly}\,|w|}$*).*
- *the function* V_w *can be computed in polynomial time in* $|w|$ *and the lengths of its inputs,*
- *the distribution* π_w *can be sampled in polynomial time in* $|w|$ *and the lengths of its inputs,*
- *(completeness) if* $w \in L$ *then* $\omega^{co}(\mathcal{G}_w) = 1$*, and*
- *(soundness) if* $w \notin L$ *then* $\omega^{co}_{\delta}(\mathcal{G}_w) \leq 1 - 1/f(|w|)$*.*

The family $\{\mathcal{G}_w\}$ *is referred to as a* protocol *for* L*.*

Here δ can also be a function of $|w|$. When $\delta = 0$, MIP$^{co}_0$ is the class of commuting-operator multi-prover interactive proofs, which dates back to [9]. Note that, in Definition 18, the protocol must be sound against δ-AC operator strategies, whereas the completeness condition requires a perfect commuting-operator strategy. As a result, MIP$^{co}_{\delta} \subset$ MIP^{co} for all δ. Similarly, MIP$^*_{\delta} \subset$ MIP*.

▶ **Remark 19.** *Our definition is slightly different from [5] in that we use δ-AC strategies with projective measurements, rather than POVMs. It's not clear how this changes the complexity class in general, since we are restricting the class of strategies that a protocol must be sound against (which potentially strengthens the class) and restricting the class of strategies that can be used for completeness (which potentially weakens the class). However, Claim 9 of [5] shows that projective measurements and POVMs are equivalent up to an increase in δ proportional to the size of the output sets. Since our lower bounds use protocols with a constant number of outputs, the lower bounds will also apply if we define* MIP$^{co}_{\delta}$ *using POVMs.*

Next we will define the perfect zero knowledge version of MIP$^{co}_{\delta}$, called PZK-MIP$^{co}_{\delta}$. Informally, a multi-prover interactive proof is perfect zero-knowledge if the verifier gains no new information from interacting with the provers. This is formalized by requiring that, for every yes instance, the provers have a strategy for which the verifier can efficiently simulate the provers' behaviour.

Let $\mathcal{G} = (\mathcal{I}_A, \mathcal{I}_B, \mathcal{O}_A^*, \mathcal{O}_B^*, V, \pi)$ be a non-local game. If the players use a commuting-operator strategy given by measurements $\{P_a^x\}$ and $\{Q_b^y\}$ and a state $|\psi\rangle$ in a Hilbert space $\mathcal{H}$, then to an outside party (such as the verifier), the players actions are described by the probabilities

$$p(a,b|x,y) = \langle\psi| P_a^x Q_b^y |\psi\rangle .$$

When x, y are fixed, $p(a, b|x, y)$ gives a probability distribution over outcomes $(a, b) \in \mathcal{O}_A^x \times \mathcal{O}_B^y$. The family of probability distributions $\underline{p} = \{p(a, b|x, y) : (x, y) \in \mathcal{I}_A \times \mathcal{I}_B, (a, b) \in \mathcal{O}_A^x \times \mathcal{O}_B^y\}$ is called the *correlation matrix* of the strategy.

In a interactive proof system, a record of interactions between verifier and provers is called a transcript. Let $\{\mathcal{G}_w\}$ be a $\mathrm{MIP}_\delta^{co}(2, 1, 1, s)$ protocol for a language L as in Definition 18. During the game $\mathcal{G}_w$, the transcript consists simply of the inputs $(x, y) \in \mathcal{I}_A^w \times \mathcal{I}_B^w$ sent to the provers, and the outputs $(a, b) \in \mathcal{O}_A^x \times \mathcal{O}_B^y$ received back. If the verifier asks questions x, y with probability $\pi(x, y)$, then the distribution over transcripts (x, y, a, b) is given by $\pi(x, y)p(a, b|x, y)$, where $\{p(a, b|x, y)\}$ is the correlation matrix of the provers' strategy. A strategy is said to be *perfect zero-knowledge against an honest verifier* if it is possible to sample from the distribution $\{\pi_w(x, y)p(a, b|x, y)\}_{(x,y,a,b)}$ in polynomial time. However, this assumes that the verifier chooses questions x, y according to the probability distribution π_w given in the protocol, something that the provers cannot validate themselves while the game is in progress. To be perfect zero-knowledge against a possibly dishonest verifier, it is necessary that the verifier be able to simulate $\pi(x, y)p(a, b|x, y)$ for any (simulable) distribution $\pi(x, y)$ on inputs. This is equivalent to being able to simulate the distributions $\{p(a, b|x, y)\}$, so we make the following definition:

▶ **Definition 20.** *Let $\{\mathcal{G}_w\}$ be a $\mathrm{MIP}_\delta^{co}(2, 1, 1, 1 - s)$-protocol for a language L. Then $\{\mathcal{G}_w\}$ is said to be* perfect zero-knowledge *if for each string w and pair $(x, y) \in \mathcal{I}_A \times \mathcal{I}_B$, there is a probability distribution $\{p_w(a, b|x, y) : (a, b) \in \mathcal{O}_A^x \times \mathcal{O}_B^y\}$ over $\mathcal{O}_A^x \times \mathcal{O}_B^y$ such that*

1. *the distribution $\{p_w(a, b|x, y)\}$ can be sampled in polynomial time in $|w|$, $|x|$, and $|y|$, and*
2. *if $w \in L$, then $\{p_w(a, b|x, y) : (x, y) \in \mathcal{I}_A \times \mathcal{I}_B, (a, b) \in \mathcal{O}_A^x \times \mathcal{O}_B^y\}$ is the correlation matrix of a commuting-operator strategy $\mathcal{S}$ with winning probability $\omega(\mathcal{G}_w; \mathcal{S}) = 1$.*

The class PZK-$\mathrm{MIP}_\delta^{co}(2, 1, 1, 1-1/f(n))$ *is the class of languages in* $\mathrm{MIP}_\delta^{co}(2, 1, 1, 1-1/f(n))$ *with a perfect zero-knowledge protocol.*

Proof of Theorem 2. Theorem 17 immediately implies that any language $L \in \mathrm{coNTIME}(f(n))$ has a protocol in $\mathrm{MIP}_\delta^{co}(2, 1, 1, 1 - 1/f(Cn)^k)$ for some constants C and k, where $\delta = o(1/f(Cn)^{2k})$. Since the games constructed in the proof of Theorem 17 come from the wagon wheel construction, Proposition 16 implies that when $w \in L$, the game $\mathcal{G}_w$ has a perfect commuting operator strategy with a correlation that can easily be simulated by the verifier. ◀

6.1 Upper bounds

As mentioned in the introduction, no upper bound on MIP^{co} is known, but an upper bound on MIP_δ^{co} follows from [5] as we will now describe. Consider the following theorem, which is a restatement of Theorem 2 in [5]:

▶ **Theorem** (Theorem 2 [5])**.** *Let G be a 2-prover non-local game with classical messages in which each prover has ℓ possible answers, and $\omega_{QCSDP}^N(G)$ is the optimum of the N-th level of the QC SDP hierarchy for G. Then there exists a $\delta = \Theta(\ell^2/\sqrt{N})$ such that $\omega_\delta^*(G) = \omega_{QCSDP}^N(G)$.*[4]

Here the QC SDP hierarchy for a non-local game G is as in Definition 10 of [5]. For our purposes the only properties of the QC SDP hierarchy that we will require are the following:

[4] The same statement could be made for non-local games with more players by raising the exponent on ℓ.

▶ **Fact 21.** *The QC SDP hierarchy gives an upper bound on the entangled winning probability of a game G at every level. That is, $\omega^N_{QCSDP}(G) \geq \omega^{co}(G)$ for all N. This is an elementary property of this hierarchy and is discussed in [5].*

▶ **Fact 22.** *The quantity $\omega^N_{QCSDP}(G)$ can be computed in time polynomial in $(Q\ell)^N$ where Q is the maximum number of questions to either prover in G, and ℓ is the maximum number of answers. This is because $\omega^N_{QCSDP}(G)$ is defined (in Definition 10 of [5]) to be the optimal value of an semi-definite program on $\mathrm{poly}((Q\ell)^N)$ dimensional space, with $\mathrm{poly}((Q\ell)^N)$ constraints.*

Now, suppose that one wishes to decide whether a non-local game G has $\omega^{co}(G) = 1$, or has $\omega^{co}_\delta(G) \leq 1 - 1/f$ promised that one of the two is the case. By Theorem 2 of [5], there exists $M = O(\ell^4/\delta^2)$ such that $\omega^*_\delta(G) = \omega^M_{QCSDP}(G)$. To resolve the decision problem it then suffices to compute the quantity $\omega^M_{QCSDP}(G)$. In the case that $\omega^{co}(G) = 1$ we know by Fact 21 that we will have $\omega^M_{QCSDP}(G) = 1$. On the other hand, in the case that $\omega^{co}_\delta(G) \leq 1 - 1/f$ we know that $\omega^M_{QCSDP}(G) = \omega^*_\delta(G) \leq \omega^{co}_\delta(G) \leq 1 - 1/f$. It follows by Fact 22 that this decision problem can be solved in time that is polynomial in $(Q\ell)^M = (Q\ell)^{O(\ell^4/\delta^2)}$ where Q and ℓ are the sizes of the question and answer sets in G respectively.

This upper bound uses strategies with POVM measurements, but if we restrict to protocols with constant size output sets, we can state this result for the class defined in Definition 18.

▶ **Theorem 23** ([5])**.** *If $L \in \mathrm{MIP}^{co}_\delta(2, 1, 1, 1 - 1/f(n))$ has a protocol with constant size output sets, and $\delta = o(1/f(n))$, then L is contained in $\mathrm{TIME}(\exp(\mathrm{poly}(n)/\delta^2))$.*

Proof. While the result in [5] is stated for MIP^*_δ which has completeness and soundness conditions stated for finite-dimensional strategies, the proof is still valid for the analogously defined MIP^{co}_δ, which has completeness and soundness conditions stated for commuting-operator strategies. ◀

References

1. Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof Verification and the Hardness of Approximation Problems. *JACM*, 45(3):501–555, 1998.
2. Sanjeev Arora and Shmuel Safra. Probabilistic Checking of Proofs: A New Characterization of NP. *JACM*, 45(1):70–122, 1998.
3. Richard Cleve, Li Liu, and William Slofstra. Perfect commuting-operator strategies for linear system games. *Journal of Mathematical Physics*, 2016. to appear (`arXiv:1606.02278`).
4. Richard Cleve and Rajat Mittal. Characterization of Binary Constraint System Games. In *Automata, Languages, and Programming*, number 8572 in Lecture Notes in Computer Science, pages 320–331. Springer Berlin Heidelberg, 2014. `arXiv:1209.2729`.
5. Matthew Coudron and Thomas Vidick. Interactive proofs with approximately commuting provers. In *International Colloquium on Automata, Languages, and Programming (ICALP 2015)*, pages 355–366, 2015.
6. Andrew C. Doherty, Yeong-Cherng Liang, Benjamin Toner, and Stephanie Wehner. The Quantum Moment Problem and Bounds on Entangled Multi-prover Games. In *Proc. 23rd IEEE Conf. on Computational Complexity (CCC'08)*, pages 199–210. IEEE Computer Society, 2008.
7. J. Fitzsimons, Z. Ji, T. Vidick, and H. Yuen. Quantum proof systems for iterated exponential time, and beyond. *preprint*, 2018. `arXiv:1805.12166`.
8. SM Gersten. Isoperimetric and isodiametric functions of finite presentations. In Graham A. Niblo and Martin A. Roller, editors, *Geometric group theory: Proceedings of the Symposium held in Sussex 1991, Volume 1*, volume 181 of *London Mathematical Society Lecture Note Series 181*, pages 79–97. Cambridge University Press, 1993.

9 T. Ito, H. Kobayashi, D. Preda, X. Sun, and A C.-C. Yao. Generalized Tsirelson Inequalities, Commuting-Operator Provers, and Multi-prover Interactive Proof Systems. In *23rd Annual IEEE Conference on Computational Complexity (CCC 2008)*, 2008.

10 Tsuyoshi Ito and Thomas Vidick. A multi-prover interactive proof for NEXP sound against entangled provers. In *Foundations of Computer Science (FOCS), 2012 IEEE 53rd Annual Symposium on*, pages 243–252. IEEE, 2012.

11 Zhengfeng Ji. Compression of Quantum Multi-prover Interactive Proofs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, pages 289–302, New York, NY, USA, 2017. ACM.

12 R.C. Lyndon and P.E. Schupp. *Combinatorial Group Theory*. Classics in Mathematics. Springer, 1977.

13 Miguel Navascués, Stefano Pironio, and Antonio Acín. Bounding the Set of Quantum Correlations. *Phys. Rev. Lett.*, 98:010401, January 2007. doi:10.1103/PhysRevLett.98.010401.

14 Miguel Navascués, Stefano Pironio, and Antonio Acín. A Convergent Hierarchy of Semidefinite Programs Characterizing the Set of Quantum Correlations. *New Journal of Physics*, 10(073013), 2008. arXiv:0803.4290v1.

15 Narutaka Ozawa. Tsirelson's problem and asymptotically commuting unitary matrices. *Journal of Mathematical Physics*, 54(3):032202, 2013.

16 Mark V. Sapir, Jean-Camille Birget, and Eliyahu Rips. Isoperimetric and Isodiametric Functions of Groups. *Annals of Mathematics*, 156(2):345–466, 2002.

17 William Slofstra. Tsirelson's problem and an embedding theorem for groups arising from non-local games. *CoRR*, 2016. preprint. arXiv:1606.03140.

18 William Slofstra and Thomas Vidick. Entanglement in non-local games and the hyperlinear profile of groups. *Annales Henri Poincare*, 19(10):2979–3005, 2018.