

The complexity of constraint system games with entanglement

by

Kieran Mastel

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Pure Mathematics (Quantum Information)

Waterloo, Ontario, Canada, 2025

© Kieran Mastel 2025

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner: Anand Natarajan
Assistant Professor, Dept. of Electrical Engineering
and Computer Science,
Massachusetts Institute of Technology

Supervisor: William Slofstra
Associate Professor, Dept. of Pure Mathematics,
University of Waterloo

Internal Member: Michael Brannan
Associate Professor, Dept. of Pure Mathematics,
University of Waterloo

Internal-External Member: David Gosset
Associate Professor, Dept. of Combinatorics and Optimization,
University of Waterloo

Other Member(s): Ben Webster
Associate Professor, Dept. of Pure Mathematics,
University of Waterloo

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

This thesis is based on the following works:

- [48] Kieran Mastel and William Slofstra. Two prover perfect zero knowledge for MIP*. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing*, STOC 2024, page 991–1002, New York, NY, USA, 2024. Association for Computing Machinery. <https://doi.org/10.1145/3618260.3649702>
- [18] Eric Culf and Kieran Mastel. RE-completeness of entangled constraint satisfaction problems. *arXiv preprint arXiv:2410.21223*, 2024. (To appear in FOCS 2025).

The techniques in [18] generalize and improve the analysis in [48]. As such, Chapter 3 is an amalgamation of content from [48] and [18]. The contents of Chapter 4 are based entirely on [18]. Finally, Chapter 5 is based primarily on [48] but uses the improvements from [18].

Abstract

Constraint satisfaction problems (CSPs) are a natural class of decision problems where one must decide whether there is an assignment to variables that satisfies a given formula. Schaefer’s dichotomy theorem and its extension to all alphabets due to Bulatov and Zhuk, shows that CSP languages are either efficiently decidable or NP-complete. It is possible to extend CSP languages to quantum assignments using the formalism of nonlocal games. The recent $\text{MIP}^* = \text{RE}$ theorem of Ji, Natarajan, Vidick, Wright, and Yuen shows that the complexity class MIP^* of multiprover proof systems with entangled provers contains all recursively enumerable languages. As a consequence, general succinctly presented CSPs are RE-complete. We show that a wide range of NP-complete CSPs become RE-complete when the players are allowed entanglement, including all boolean CSPs, such as 3SAT and 3-colouring. This implies that these CSP languages remain undecidable even when not succinctly presented.

Prior work of Grilo, Slofstra, and Yuen shows (via a technique called simulatable codes) that every language in MIP^* has a perfect zero knowledge (PZK) MIP^* protocol. The $\text{MIP}^* = \text{RE}$ theorem uses two-prover one-round proof systems. Hence, such systems are complete for MIP^* . However, the construction in Grilo, Slofstra, and Yuen uses six provers, and there is no obvious way to get perfect zero knowledge with two provers via simulatable codes. This leads to a natural question: are there two-prover PZK- MIP^* protocols for all of MIP^* ?

In this work, we show that every language in MIP^* has a two-prover one-round PZK- MIP^* protocol, answering the question in the affirmative. For the proof, we use a new method based on a key consequence of the $\text{MIP}^* = \text{RE}$ theorem, which is that every MIP^* protocol can be turned into a family of boolean constraint system (BCS) nonlocal games. This makes it possible to work with MIP^* protocols as boolean constraint systems. In particular, it allows us to use a variant of a CSP due to Dwork, Feige, Kilian, Naor, and Safra that gives a classical MIP protocol for 3SAT with perfect zero knowledge.

To prove our results, we develop a toolkit for analyzing the quantum soundness of reductions between constraint system (CS) games, which we expect to be useful more broadly. In this formalism, synchronous strategies for a nonlocal game correspond to tracial states on an algebra. We equip the algebra with a finitely supported weight that allows us to gauge the players’ performance in the corresponding game using a weighted sum of squares. The soundness of our reductions hinges on guaranteeing that specific measurements for the players are close to commuting when their strategy performs well. To this end, we construct commutativity gadgets for all boolean CSPs and show that the

commutativity gadget for graph 3-colouring due to Ji is sound against entangled provers. We define a broad class of CSPs that have simple commutativity gadgets. We show a variety of relations between the different ways of presenting CSPs as games. This toolkit also applies to commuting operator strategies, and our argument shows that every language with a commuting operator BCS protocol has a two prover PZK commuting operator protocol.

Acknowledgements

I would first like to thank my excellent supervisor and collaborator, William Slofstra, for his guidance and support over the years. Thank you for encouraging my curiosity, and for pushing me to take on hard problems. I am also grateful to David Gosset for his generosity with his time and for sharing papers, stories, and climbs.

My time in graduate school was enriched by many dear friends and colleagues who helped make Kitchener-Waterloo feel like home. Thank you to Colter MacDonald for starting this journey with me from across the country during a pandemic. I am also thankful to Adina Goldberg, Joaco Prandi, John Sawatzky, Paul Lawrence, Nolan Shaw, Katie Pita, Yuming Zhao, Connor Paddock, Ben Lovitz, Adam Bene Watts, Archishna Bhattacharyya, Eric Culf, Jack Davis, Amit Anand, Sanchit Srivastava, Zachary Man, Emiliia Dyrenkova, Alex Frei, Amolak Ratan Kalra, Pulkit Sinha, and Calvin Liu. Thank you for the discussions (mathematical and otherwise) for rocks climbed, games played, meals shared, and journeys taken. A special thank you to Collin and Hannah Epstein, Maeve Wentland, and Irene Rodriguez for many days of food and fellowship. You've all become like family.

Thank you to my sisters, Petra and Annika Mastel, and to my parents, Ken and Gail Mastel, for their constant love and support; without it, I wouldn't be here.

Finally, I am grateful for the financial support provided by the Natural Sciences and Engineering Research Council of Canada (NSERC) through the Canada Graduate Scholarship–Doctoral (CGS D) and Canada Graduate Scholarship–Master's (CGS M) programs.

Dedication

To Emma, all troubles are less with you.

Table of Contents

Examining Committee Membership	ii
Author's Declaration	iii
Statement of Contributions	iv
Abstract	v
Acknowledgements	vii
Dedication	viii
List of Figures	xi
1 Introduction	1
1.1 Nonlocal games and multiprover interactive proofs	2
1.2 Constraint satisfaction problems	4
1.3 The weighted algebra formalism	7
1.4 Perfect zero knowledge	9
2 Background and Preliminaries	12
2.1 General notation	12
2.2 Quantum states and Hilbert space	12

2.3	*-algebras and representations	15
2.4	Decidability and complexity	17
2.5	Nonlocal games and MIP*	19
3	The weighted algebra formalism	23
3.1	Constraint system games	23
3.2	Weighted algebras for CS games	28
3.3	Relations between CS algebras	34
3.4	Subdivision and stability	43
4	RE-completeness of entangled CSPs	51
4.1	Constraint system languages	51
4.2	Quantum CSPs	54
4.3	Hardness of non-TVF CSPs	57
4.4	Hardness of boolean TVF CSPs	61
4.4.1	The basic commutativity gadget	61
4.4.2	Compression and simulation: building the needed constraints	63
4.4.3	The general commutativity gadget	75
4.4.4	Oracularizability of boolean TVF CSPs	82
4.5	Hardness of 2-CSPs	83
4.5.1	The case of 3-colouring	83
4.5.2	The case of 2-CSP(k)	90
4.6	Constraint-variable to constraint-constraint for CSPs	92
5	Two prover perfect zero knowledge for MIP*	94
5.1	Definitions	94
5.2	Parallel repetition	95
5.3	The tableau construction	97
5.4	Perfect zero knowledge	102
	References	116

List of Figures

3.1	C -homomorphisms (solid arrows) and trace-dependent mappings (dashed arrows) between the different weighted algebras for a k -ary CS $S = (X, \{(V_i, C_i)\}_{i=1}^m)$. Here $\pi'(i) = \sum_j \pi(i, j)$, $L = \max_i V_i $, $P = \max_{i,j. V_i \cap V_j \neq \emptyset} \frac{\pi'(i)}{\pi(i,j)}$, and $B(S)$ is the BCS defined in Definition 3.3.1.	34
4.1	The basic commutativity gadget for TVF boolean constraint systems. Exactly one variable in each triangle must be assigned value 1. These constraints bound the commutator $[x, y]$, and any assignment to x and y may be extended to an assignment to all three constraints.	61
4.2	Basic commutativity gadgets with one, two, or three variables per constraint negated. The white vertices indicate the negated variables: note that negated variables must be only connected amongst themselves to construct the gadget.	63
4.3	The compressible cycle TVF graph from Lemma 4.4.5.iii.	66
4.4	The triangular constraint system in Lemma 4.5.3. Each vertex corresponds to a variable and each edge corresponds to a 3-colouring constraint.	85
4.5	The triangular prism constraint system in Lemma 4.5.4. Each vertex corresponds to a variable and each edge correspond to a 3-colouring constraint.	88
4.6	Transitions in complexity based on soundness parameter for entangled 3-colouring	89

Chapter 1

Introduction

This thesis establishes results in the theory of nonlocal games and computational complexity. We introduce the weighted algebra formalism for analyzing the soundness of reductions between constraint system (CS) games. This formalism generalizes the well-studied correspondence between perfect strategies for synchronous CS games and tracial states on the synchronous algebra. Imperfect strategies correspond to traces on the weighted algebra that are nonzero on the support of the weight. The weighted algebra can be thought of as a model of the synchronous algebra. We define a trace-dependent quantity called the defect that quantifies how well traces on weighted algebras approximate traces on the synchronous algebra. Since the weight of an observable depends on the probability of the verifier asking the question it appears in, the defect gives us an algebraic tool to discuss imperfect strategies for the corresponding game.

Applying the weighted algebra formalism, we show that some broad classes of classical reductions between CS games remain sound even when the provers are allowed entanglement. We divide these reductions into two types of transformations: classical homomorphisms and subdivisions. Classical homomorphisms map constraints in one CS to constraints in another CS and thus are simple to analyze with weighted algebras. Subdivision involves separating one large constraint into many subconstraints and is more complex to analyze.

The players' measurements for variables that appear in the same constraint must commute as they are simultaneously measured. When a constraint is subdivided, there is no longer a guarantee that all the variables appearing in different subconstraints commute. For our soundness argument, we use commutativity gadgets like those introduced in [37] to force these variables to commute even though they appear in different subcon-

straints. Using weighted algebras, we show that these commutativity gadgets are sound against entangled provers and that they enforce approximate commutation in approximate strategies. By reducing from the $\text{MIP}^*=\text{RE}$ protocol for the halting problem, we show RE-completeness for the entangled version of any NP-complete CS game with unentangled players, provided that we have a commutativity gadget. Specifically, we show that NP-complete CS games become RE-complete with entanglement provided that they are boolean, a graph 3-colouring game, or not two-variable falsifiable. As an application of these results, we prove that MIP^* admits two prover perfect zero knowledge proofs by showing the quantum soundness of a variant of the two prover perfect zero knowledge protocol for NP due to Dwork, Feige, Kilian, Naor and Safra [21].

1.1 Nonlocal games and multiprover interactive proofs

In a nonlocal game, two spatially separated and non-communicating players, Alice and Bob, receive questions from a referee. The players' questions i and j are sampled by the referee from a finite set of questions I according to a probability distribution π known to both players. The players then respond with answers a and b from the finite answer sets O_i and O_j respectively. Often the players are defined to have separate question and answer sets, but it is convenient for us to assume without loss of generality that the players have the same question and answer sets. The referee then applies one of a family of functions $V(\cdot, \cdot | i, j) : O_i \times O_j \rightarrow \{0, 1\}$ for each question pair $(i, j) \in I \times I$. The players win if $V(a, b | i, j) = 1$ and lose otherwise. The function V is called the decision predicate.

Alice and Bob cannot communicate with each other during the interaction, but they can coordinate a strategy beforehand. This strategy can be deterministic or probabilistic, employing shared randomness. We call such strategies classical. A strategy is called perfect if it allows the players to satisfy the decision predicate for any question pair. It is clear that not all games admit perfect strategies. If they have access to quantum resources, the players may share a quantum state and condition their answers on the results of local measurements. Their choice of state and the collection of local measurements for each player make up a quantum strategy. Quantum strategies allow the players to correlate their answers in ways that are not possible with classical resources [22, 7]. In fact, there are nonlocal games that have perfect quantum strategies, but no perfect classical strategies, such as the Mermin-Peres magic square game [59, 49]. The Mermin-Peres magic square consists of a three-by-three grid. The referee asks Alice for an assignment of ± 1 to each square in one row, and asks Bob for an assignment of ± 1 to each square in one column. Alice's row must contain -1 an even number of times, and Bob's column must contain

-1 an odd number of times. They win the game if their assignments agree on the square where the row and column overlap; otherwise, they lose. There is no classical assignment to the whole grid that satisfies all of the constraints, but there is an assignment of ± 1 valued observables and a quantum state that allows the players to win on any question pair.

The Mermin-Peres magic square game is an example of a broad class of games called constraint system (CS) games. In a CS game, Alice and Bob receive constraints C_i and C_j over a set of variables X from a set of constraints $\{C_i\}_{i=1}^m$. The players must respond with satisfying assignments $\phi_i : V_i \rightarrow \Sigma$ and $\phi_j : V_j \rightarrow \Sigma$ over an alphabet Σ to the set of variables V_i and V_j that are in the scope of the constraint C_i and C_j respectively. Like in the magic square game, the winning condition is a consistency check between the players' assignments. The players lose when $\phi_i|_{V_i \cap V_j} \neq \phi_j|_{V_i \cap V_j}$, and win otherwise. What we have just described is the constraint-constraint (c-c) form of the CS game studied in, *e.g.*, [1, 58, 57, 45]. Alternatively, we can define the constraint-variable (c-v) form of the CS game. In the c-v game, Alice is sent a constraint C_i and responds with a satisfying assignment ϕ_i as before, but Bob is only sent a variable $x \in X$ and responds with an assignment to $\phi(x) \in \Sigma$ to that variable. The win condition is still a consistency check. The players lose if $\phi_i(x) \neq \phi(x)$ and win otherwise. This form of CS game was studied in, *e.g.*, [16, 37]. Finally, there is a third type of game for 2-CSs — where each constraint has only two variables. The verifier asks each player for an assignment to a single variable, and they win if the assignments satisfy the corresponding constraint. This type of game was studied in, *e.g.*, [25, 32, 19]. If the alphabet Σ has size k , then we call the CS k -ary, and in the particular case where $k = 2$, we say the CS is a boolean constraint system (BCS). Associated with each CS game is a finitely presented $*$ -algebra, called the CS algebra. Every perfect quantum strategy for a CS game is a representation of the CS algebra. In this way, the CS algebra generalizes the well-studied BCS algebra discussed in [16, 37]. It will be a central tool in proving our main results.

The family of synchronous games is another important class studied in the nonlocal game literature. The players' question and answer sets are the same in a synchronous game, and the synchronous condition states that if they receive the same question, they must respond with the same answer. An important example of a synchronous game is the graph k -colouring game first described in [15]. Given a graph $G = (V, E)$ with vertices V and edges E , the referee sends vertex x to Alice and y to Bob. The players respond with colours a and b respectively from the set $\{1, \dots, k\}$. If (x, y) is an edge in E , then the players lose if $a = b$ and win otherwise. If $x = y$, the players win if and only if $a = b$, which is the synchronous condition. Note that the graph colouring game with these questions and answers is also a 2-CS game. There is a perfect classical strategy if and only if there is a

k -colouring of the graph G . For some graphs however, entangled players can play the graph k -colouring game perfectly even though no actual k -colouring exists [10, 12]. Such graphs are called quantum k -colourable. Similarly to CS games, perfect quantum strategies for a synchronous game correspond to representations of a finitely presented $*$ -algebra called the synchronous algebra of the game [34, 45].

In a nonlocal game, the players can use unbounded computational power to construct their strategies, yet the referee is computationally bounded. Interactive proofs allow us to study and quantify the computational power the referee can access in such interactions. In an interactive proof protocol, a prover tries to convince a verifier that a string x belongs to a language \mathcal{L} . Interactive proof systems can be more powerful than non-interactive systems; famously, the class IP of interactive proofs with a polynomial time verifier and a single prover is equal to PSPACE [63], and the class MIP with a polynomial time verifier and multiple non-communicating provers is equal to NEXP [4]. The proof systems used in [4] are very efficient and require only two provers and one round of communication. Thus, they can be thought of as a family of nonlocal games indexed by the input string x , with the players as provers and the referee as the verifier. Since the provers in an MIP protocol are not allowed to communicate, it is natural to ask what happens if they are allowed to share entanglement. This leads to the complexity class MIP*, first introduced by Cleve, Hoyer, Toner, and Watrous (where the $*$ denotes that the players are entangled) [15]. Entanglement allows the provers to break some classical proof systems by coordinating their answers, but the improved ability of the provers also allows the verifier to set harder tasks. As a result, figuring out the power of MIP* has been difficult, and there have been successive lower bounds in [41, 35, 36, 66, 67, 38, 54, 39, 53, 23]. Most recently, the landmark work of Ji, Natarajan, Vidick, Wright, and Yuen showed that MIP* = RE, the class of languages equivalent to the halting problem [40]. A multiprover interactive proof protocol can be thought of as a decision problem, where the verifier must determine whether or not an instance x is in \mathcal{L} . In the accept instance, where $x \in \mathcal{L}$, there is a strategy that the provers can employ that wins the nonlocal game with probability c , called the completeness parameter, which in this work we will always take to be 1. In the reject instance, where $x \notin \mathcal{L}$, any strategy the provers employ can win the nonlocal game with probability at most s , called the soundness parameter.

1.2 Constraint satisfaction problems

As a result of the proof that MIP = NEXP in [5], MIP is equivalent to the class of CS – MIP proof systems, which are two-prover one-round proof systems in which the nonlocal games

are CS games. In a CS – MIP proof system, the CS games are succinctly presented. In other words, for a given instance, there may be exponentially many questions to sample or answers to verify (in the instance size $|x|$), but these operations are implemented efficiently. This can increase the complexity of the games significantly. For any family of constraints Γ over an alphabet Σ , the constraint satisfaction problem (CSP) language $\text{CSP}(\Gamma)_{1,1}$ consists of all constraint systems that can be expressed as the conjunction of constraints from Γ , where the yes instances are those for which all of the constraints can be satisfied, and the no instances are those where at least one constraint must be unsatisfied. Due to the CSP dichotomy theorem, it is well-understood that CSP languages are either solvable in polynomial time or NP-complete. Moreover, the theorem completely classifies which CSP languages are NP-complete [61, 11, 70]. The succinct version of this language is the promise problem $\text{SuccinctCSP}(\Gamma)_{1,s}$ consisting of efficiently sampleable CSs with constraints from Γ , where the yes instances are those for which all the constraints can be simultaneously satisfied, and the no instances are those where there is a probability less than s of sampling a satisfied constraint for any assignment. If $\text{CSP}(\Gamma)_{1,1}$ is NP-complete, then there is a constant $s \in [0, 1)$, such that $\text{SuccinctCSP}(\Gamma)_{1,s}$ is NEXP-complete. Thus, the class of CS games corresponding to $\text{SuccinctCSP}(\Gamma)_{1,s}$ is complete for MIP.

It is natural to ask if an analogous dichotomy holds when the provers are allowed entanglement. In this thesis, we provide a partial answer to this question. As a result of the $\text{MIP}^* = \text{RE}$ theorem, MIP^* is equivalent to the class of CS – MIP^* proof systems, the class of two-prover one-round proof systems in which the nonlocal games are CS games with entangled players. We prove that a family Γ of constraints with an NP-complete $\text{CSP}(\Gamma)_{1,1}$ is RE-complete with entangled provers by reducing from the CS – MIP^* proof system for RE to a CS – MIP^* proof system using only constraints from Γ . We use a classical reduction between NP-complete constraint satisfaction problems, modified to ensure that the reduction preserves the soundness of the entangled protocol.

The provers' strategies in a CS – MIP^* protocol are an operator assignment to the variables in the constraint system and a choice of state. In general, operator assignments to CSPs are non-commuting, but to guarantee the validity of some reductions between CSPs, we need a way to guarantee commutativity (or near-commutativity) between variables. A canonical way to do that is using empty constraints, which contain variables but impose no relations between them, except that they are simultaneously measurable. However, we have no guarantee that a given set of constraints Γ includes an empty constraint. To remedy this, we construct commutativity gadgets: subsystems of constraints that behave like an empty constraint when restricted to two of the variables.

To characterize the commutativity gadgets we can construct, we use a property of some constraints we call two-variable falsifiability (TVF). A constraint is TVF if, for any

pair of variables, there is an assignment to that pair for which the constraint is false no matter what value the other variables are assigned. We call a CSP TVF if all of its constraints are. Non-TVF CSPs have simple one-constraint commutativity gadgets. If a CSP contains a non-TVF constraint, we can replace empty constraints with instances of the non-TVF constraint and replace all but two variables in the non-TVF constraint with dummy variables. If all constraints in Γ are TVF, this precludes the construction of these generic commutativity gadgets built from one constraint. However, many important CSPs, such as 3SAT, are not TVF. In particular, note that any CSP augmented by an empty constraint is non-TVF and that this does not change its classical complexity.

For TVF constraints, the situation is more complicated, but in the boolean case, we are able to construct a generic but more complicated commutativity gadget. The basic example of an NP-complete TVF constraint satisfaction problem is 1-in-3-SAT, generated by the three-variable boolean constraint where exactly one of the variables must be assigned value 1, $C = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$. This has a commutativity gadget constructed by connecting three copies of C in a triangular arrangement (see Figure 4.1), first studied by Ji [37]. For the remaining NP-complete boolean TVF CSPs, we show that there is a structure similar to C hidden within them, and hence that an analogous commutativity gadget may be constructed.

The final case where we are able to construct a commutativity gadget is for graph 3-colouring. Here, we make use of a triangular prism gadget (see Figure 4.5), also introduced by Ji [37]. The work of Ji shows that the gadget guarantees commutativity in the case of perfect completeness; for our purposes, we extend this to the case of imperfect completeness, showing soundness of the gadget.

In full, we show that if $\text{CSP}(\Gamma)_{1,1}$ is NP-complete, and Γ is boolean, non-TVF, or 3-colouring, then there exists a constant $s \in [0, 1)$ such that the entangled CSP language $\text{SuccinctCSP}(\Gamma)_{1,s}^*$ is RE-complete.

As a direct consequence, we find that the non-succinct version of the language, $\text{CSP}(\Gamma)_{1,s}^*$, is also RE-complete under Turing reductions, although it may not be polynomial-time reducible from the halting problem. Completeness in NEXP is with respect to polynomial-time Karp reductions; in fact, with exponential-time reductions, NP-complete problems become complete for NEXP. In the same way, if $\text{SuccinctCSP}(\Gamma)_{1,s}^*$ is RE-complete with a polynomial-time reduction, then $\text{CSP}(\Gamma)_{1,s}^*$ is RE-complete with an exponential-time reduction. In particular, there is a computable function that reduces the halting problem to $\text{CSP}(\Gamma)_{1,s}^*$, so it must be RE-complete.

For the proof of our main results, we begin with the output of the $\text{MIP}^* = \text{RE}$ theorem rather than encoding an arbitrary MIP^* protocol. The proof that $\text{MIP}^* = \text{RE}$ in [40] is very

involved, but has as an important consequence that only two-prover one-round proof systems are required to attain the full complexity of MIP^* . Dong, Fu, Natarajan, Qin, Xu, and Yao show in [20] that these proof systems can be reduced in size to have polynomial-length questions and constant-length answers. In both [40] and [20], the games are synchronous and admit oracularizable optimal strategies in the case of perfect completeness, meaning that the two players' measurement operators for any pair of questions asked at the same time commute. One-round MIP^* proof systems in which the games are synchronous and oracularizable are equivalent to the class of BCS- MIP^* proof systems, which are one-round two-prover proof systems in which the nonlocal games are BCS games, that is, CS games with boolean constraint systems. Since the games in [20] have constant answer size, so do the corresponding BCS games. We prove the RE-hardness of NP-complete CS protocols with entanglement by showing that the classical reduction from the BCS form of the protocol from [20] to a CS protocol is sound against quantum provers. Reductions between MIP^* protocols translate one proof system into another while preserving completeness and controlling the degradation of soundness. In some cases, reductions between MIP^* proof systems have polynomial soundness dropoff and use parallel repetition to recover constant soundness. Parallel repetition does not preserve many classes of CS games (for example, graph colouring games), so we require our transformations to have a constant soundness dropoff.

1.3 The weighted algebra formalism

In Section 1.1, we saw examples of the correspondence between perfect strategies for nonlocal games and representations of finitely presented $*$ -algebras. In the context of MIP^* protocols, this correspondence allows us to use transformations between the algebras to examine reductions between the games that preserve perfect completeness as in [37]. That is, where the completeness parameter remains equal to 1. These transformations are typically homomorphisms or other syntactic operations that map one finitely presented $*$ -algebra to another, often reflecting a structural reduction between the corresponding nonlocal games. For instance, one may rewrite a game by imposing new relations on the generators, such as identifying certain observables or introducing auxiliary variables, thereby inducing a quotient algebra. A perfect quantum strategy corresponds to a finite-dimensional C^* -algebra representation of the associated game algebra, so if such a transformation preserves the existence of the representation, then we can conclude that the second game is at least as hard as the first. However, this correspondence does not let us track how transformations between game algebras affect the soundness parameter. In the worst case, a transformation

of this type could take the soundness parameter to 1, meaning the players can perform arbitrarily well even in reject instances of the MIP^* protocol. From a complexity theory standpoint, gapless protocols can still be valuable because they provide insight into the expressive power of interactive proof systems and the kinds of languages that can be captured within MIP^* , even when the verifier cannot distinguish between accepting and rejecting instances. That is, such protocols may still characterize undecidable or recursively enumerable languages, helping us understand the upper bounds of provability or computability in quantum interactive models. In actual implementations of these protocols in the lab, we want the verifier to have a non-negligible chance of learning whether the instance is an accept or a reject. Furthermore, in the classical complexity theory literature, many hardness results are proven by reducing one CSP to another and showing that this transformation preserves the completeness and soundness of the corresponding MIP protocol. To show our main results, we focus on these classical CSP reductions and develop tools to show that, under certain circumstances, they preserve the soundness parameter even if the provers are allowed to share entanglement.

In general, it's difficult to determine if a classical transformation of constraint systems (of which there are many) remains sound (meaning that it preserves the soundness of protocols) in the quantum setting. For instance, one of the key parts of the $\text{MIP}^* = \text{RE}$ theorem is the construction of a PCP of proximity that is quantum sound. On the other hand, some transformations lift fairly easily to the quantum setting. We identify two such classes of transformations, “classical transformations,” which are applied constraint by constraint, and “context subdivision transformations,” in which each constraint is split into several subclauses. Both types of transformations are used implicitly throughout the literature on nonlocal games, including in [37], which was the first paper to consider reductions between quantum strategies in BCS games. In this thesis, we systematically investigate the quantum soundness of these transformations. In Section 3.2, we show that classical transformations preserve soundness by a relatively simple argument. In subdivision, each subclause becomes a different question in the associated BCS game. As a result, a strategy for the subdivided game has many more observables than the original game. Since these new observables don't need to commute with each other, subdivision is more complex to analyze. We show that if we have a way of enforcing commutativity among these new variables, and if the subclauses have a bounded number of variables, then subdivision preserves soundness with a constant dropoff. We prove this in Section 3.4.

While reductions between nonlocal games have been important in previous work, they are difficult to reason about, since it's necessary to keep track of how strategies for the input game map to those for the output game. One advantage of working with constraint systems in the classical setting is that it's more convenient to work with assignments (and think

about the fraction of constraints in the system that can be satisfied) than to work with strategies and winning probabilities. In the quantum setting, it isn't possible to work with assignments to the variables because strategies involve observables that don't necessarily commute with each other. However, we can achieve a similar conceptual simplification by replacing assignments with representations of the CS algebra of the constraint system. The CS algebra is the same as the synchronous algebra of the CS game; we refer to [58] for more background on the boolean case. With this approach, reductions between CS games correspond to homomorphisms between CS algebras, which are much easier to describe and work with than mappings between strategies. Near-perfect strategies correspond to approximate representations of the CS algebra (See [57] for the boolean case). We can thus prove that our reductions preserve the soundness of MIP* protocols by examining how mappings between CS algebras treat these approximate representations. Previous work using this idea (see e.g. [57, 30]) has focused on reductions between single games, and the definitions are not suitable for working with protocols, as they do not incorporate question distributions. To solve this problem, we introduce a notion of weighted algebras and weighted homomorphisms, which allows us to keep track of the soundness of reductions between games using completely algebraic arguments involving sums of squares.

Another advantage of the weighted algebras framework is that arguments can be made simultaneously for both quantum and commuting operator strategies. As a result, our proof methods extend to commuting operator strategies. However, our results here are not as conclusive, as the exact characterization of the corresponding complexity class MIP^{co} is not known. There is a conjecture that $\text{MIP}^{\text{co}} = \text{coRE}$, and with that conjecture, we expect that it would be possible to extend our main results and show that NP-complete CSPs for which we can construct commutativity gadgets are coRE-complete when the players can share commuting operator entanglement.

1.4 Perfect zero knowledge

Interactive proof systems allow for zero knowledge protocols, in which the prover demonstrates that $x \in \mathcal{L}$ without revealing any other information to the verifier. As a result, interactive proof systems are important to cryptography in addition to complexity theory. The first zero knowledge proof systems go back to the invention of interactive proof systems by Goldwasser, Micali, and Rackoff [27], and every language in MIP admits a two-prover one-round perfect zero knowledge proof system by a result of Ben-Or, Goldwasser, Kilian, and Wigderson [8]. Perfect means that absolutely no information is revealed to the verifier, in contrast to statistical zero knowledge (in which the amount of knowledge gained by the

verifier is small but bounded), or computational zero knowledge (in which zero knowledge relies on some computational intractability assumption).

Since any language in MIP admits a perfect zero knowledge proof system, it is natural to wonder if the same can be done when the provers have access to entanglement. Chiesa, Forbes, Gur, and Spooner showed that every language in NEXP (and hence in classical MIP) has a perfect zero knowledge MIP* proof system, or in other words belongs to PZK-MIP* [14]. Subsequently, Grilo, Slofstra, and Yuen show that all of MIP* belongs to PZK-MIP* [29]. Combining PZK-MIP* = MIP* with MIP* = RE shows that there are one-round perfect zero-knowledge MIP* proof systems for all languages that can be reduced to the halting problem, a very large class. However, the construction in [29] is involved. The idea behind the proof is to encode a circuit for an arbitrary MIP verifier in a “simulatable” quantum error correcting code, and then hide information from the verifier by splitting the physical qubits of this code between different provers. The resulting proof systems in [29] require 6 provers, and because the core concept of the proof is to split information between provers, bringing this down to 2 provers (as can be done with perfect zero-knowledge for MIP) seems to require new ideas.

In this work, we apply our main results on the hardness of entangled CSPs to show that MIP* does indeed admit two-prover one-round perfect zero knowledge protocols. In [21], the authors construct a two-prover one-round perfect zero knowledge protocol for NP. We show a reduction from the classes of CSP that are RE-hard with entanglement to a slight variant of this protocol that is sound against entangled provers. We also demonstrate that this variant protocol retains the zero knowledge property when the players are allowed entanglement. Consequently, we show that every language in MIP* admits a two-prover one-round perfect zero knowledge proof system with polynomial length questions and constant length answers. Since we care only about the zero knowledge property and not the exact CSP we get at the end, we can apply a constant amount of parallel repetition to recover soundness $s = 1/2$.

Additionally, we prove that MIP* admits a two-prover one-round perfect zero knowledge protocol with polynomial length questions and answers, completeness probability $c = 1$ and soundness probability $s = 1/2$, in which the verifier chooses questions uniformly at random. For the proof, we use a different version of the MIP* = RE protocol as the input of our CSP reduction. Natarajan and Zhang have shown that MIP* proof systems require only a constant number of questions and polylog length answers from the provers [55]. This shows that MIP* = AM*(2), the complexity class of languages with two-prover MIP* protocols in which the verifier chooses their messages to the prover uniformly at random. When we reduce from this protocol to our perfect zero knowledge protocol, we get a polynomial soundness dropoff, but a polynomial amount of parallel repetition allows us to recover

soundness $1/2$.

Since the weighted algebra formalism applies to commuting operator strategies as well, if we had a characterization of MIP^{co} such as a confirmation that $\text{MIP}^{\text{co}} = \text{coRE}$ and a parallel repetition theorem for commuting operator strategies, we should be able to show that all languages in MIP^{co} have a perfect zero knowledge commuting operator protocol. Without these ingredients, we are limited to showing that $\text{BCS-MIP}^{\text{co}} = \text{PZK-BCS-MIP}^{\text{co}}$. Previous work on perfect zero knowledge for commuting operator protocols does not preserve soundness gaps [17].

Our results also have applications for the membership problem for quantum correlations. For exact membership, the cohalting problem is many-one reducible to membership in the set of quantum-approximable correlations C_{qa} , and to membership in the set of commuting operator correlations C_{qc} [64, 17, 24]. It follows from $\text{MIP}^* = \text{RE}$ that the halting problem is Turing reducible to approximate membership in C_q , the set of quantum correlations, but this is not a many-one reduction. The proof of Theorem 5.4.8 immediately implies that there is a many-one reduction from the halting problem to approximate membership in C_q .

Chapter 2

Background and Preliminaries

2.1 General notation

For $n \in \mathbb{N}$, we write the set $[n] = \{1, 2, \dots, n\}$. We assume the logarithm \log is base 2 unless otherwise specified.

We identify \mathbb{Z}_k with the subset $\{0, \dots, k-1\}$ of \mathbb{N} , and denote the primitive k -th root of unity $\omega_k = e^{2\pi i/k}$, dropping the subscript if clear from context.

For a graph $G = (V, E)$ and $U \subseteq V$, write $G|_U$ for the subgraph on vertices in U and $G \setminus U$ for the subgraph on vertices in $V \setminus U$.

We deal only with probability distributions on finite sets. Hence, we present any probability distribution π on a set A by a function $\pi : A \rightarrow [0, 1]$ such that $\sum_{a \in A} \pi(a) = 1$. We say a probability distribution π on $A \times A$ is **symmetric** if $\pi(a, b) = \pi(b, a)$ for all $a, b \in A$. Following [47], we say that a probability distribution π on $A \times A$ is **C -diagonally dominant** if $\pi(a, a) \geq C \sum_{b \in A} \pi(a, b)$ and $\pi(a, a) \geq C \sum_{b \in A} \pi(b, a)$ for all $a \in A$. We write \mathfrak{u}_n for the uniform distribution on $[n]$, *i.e.* $\mathfrak{u}_n(i) = \frac{1}{n}$ for all $i \in [n]$.

2.2 Quantum states and Hilbert space

Let \mathbb{C} be the field of complex numbers. Given a set X , we define the **free complex vector space** generated by X , denoted by $\mathbb{C}X$, as the set of all finite linear combinations of formal symbols $|x\rangle$ for $x \in X$. That is, every element $|v\rangle \in \mathbb{C}X$ can be written as a sum $|v\rangle = \sum_{x \in S} a_x |x\rangle$, for some finite subset $S \subseteq X$ and coefficients $a_x \in \mathbb{C}$. In this way, the

collection $|x\rangle_{x \in X}$ forms a basis for $\mathbb{C}X$. When X is finite, then $\mathbb{C}X$ is isomorphic to \mathbb{C}^X , the space of functions $X \rightarrow \mathbb{C}$.

For each $x \in X$, define a linear functional $\langle x| : \mathbb{C}X \rightarrow \mathbb{C}$ by $\langle x|(|y\rangle) = 1$ if $y = x$ and is 0 otherwise for all $y \in X$. These functionals form a basis for the **dual space** of complex linear functionals from $\mathbb{C}X \rightarrow \mathbb{C}$ when X is finite. Elements $|v\rangle$ of $\mathbb{C}X$ are known as **kets**, and elements of the dual space $\langle u| \in \mathbb{C}^X$ are known as **bras**. This is called **bra-ket** notation and is commonly used in quantum physics to represent vectors and their duals in a form aligned with the structure of inner products.

A complex **inner product space** V is a vector space with a map $\langle \cdot | \cdot \rangle : V \times V \rightarrow \mathbb{C}$ called an **inner product** that satisfies the following properties for all vectors $|x\rangle, |y\rangle, |z\rangle \in V$ and scalars $\lambda, \nu \in \mathbb{C}$:

- $\langle x|y\rangle = \overline{\langle y|x\rangle}$ (conjugate symmetry),
- $\langle z|\lambda x + \nu y\rangle = \lambda \langle z|x\rangle + \nu \langle z|y\rangle$ (linearity in the second argument),
- $\langle x|x\rangle > 0$ for all $|x\rangle \neq 0$ (positive definiteness).

If there is ambiguity about which inner product we mean, we will denote the inner product on V by $\langle \cdot | \cdot \rangle_V$. The inner product induces a norm on V defined by $\| |x\rangle \|^2 = \langle x|x\rangle$ for all $|x\rangle \in V$. With this norm, V becomes a **normed vector space**. If there is ambiguity about which norm we are discussing, we will denote the norm induced by the inner product on V by $\| \cdot \|_V$.

A sequence $|x_1\rangle, |x_2\rangle, |x_3\rangle, \dots$ in a normed vector space is called a **Cauchy sequence** if for every positive real number ϵ , there is a positive integer N such that for all natural numbers $m, n > N$, $\| |x_m\rangle - |x_n\rangle \| < \epsilon$. A normed vector space V is called **complete** if every Cauchy sequence in V converges to an element of V . An inner product space that is complete is called a **Hilbert space**.

A linear operator on $T : \mathcal{H} \rightarrow \mathcal{H}'$ that maps from a Hilbert space \mathcal{H} to a Hilbert space \mathcal{H}' is called **bounded** if there exists a constant C such that $\| T|x\rangle \|_{\mathcal{H}'} \leq C \| |x\rangle \|_{\mathcal{H}}$ for all $|x\rangle \in \mathcal{H}$. The set of all bounded linear operators from \mathcal{H} to \mathcal{H}' is denoted $\mathcal{B}(\mathcal{H}, \mathcal{H}')$ and for convenience we write $\mathcal{B}(\mathcal{H}) := \mathcal{B}(\mathcal{H}, \mathcal{H})$. If $T \in \mathcal{B}(\mathcal{H}, \mathcal{H}')$ then there exists a map $T^* \in \mathcal{B}(\mathcal{H}', \mathcal{H})$ called the **adjoint** map, satisfying $\langle y|Tx\rangle_{\mathcal{H}'} = \langle T^*y|x\rangle_{\mathcal{H}}$ for all $|x\rangle \in \mathcal{H}, |y\rangle \in \mathcal{H}'$. An operator $T \in \mathcal{B}(\mathcal{H})$ is called **unitary** if $TT^* = T^*T = Id_{\mathcal{H}}$, where $Id_{\mathcal{H}}$ is the identity map on \mathcal{H} . If $T^* = T$, then T is called self-adjoint. Taking the adjoint is an **antilinear involution** $*$: $\mathcal{B}(\mathcal{H}) \rightarrow \mathcal{B}(\mathcal{H})$ sending $T \rightarrow T^*$, with the property that

$(T^*)^* = T$, $(TS)^* = S^*T^*$, and $(\lambda T)^* = \bar{\lambda}T^*$ for all $\lambda \in \mathbb{C}$, and $T, S \in \mathcal{B}(\mathcal{H})$. An operator $P \in \mathcal{B}(\mathcal{H})$ is called projective if $P^2 = P$.

A **projective measurement** is a set of linear operators $\{M_a\}_{a \in O}$ in $\mathcal{B}(\mathcal{H})$ such that M_a is self adjoint and projective for all $a \in O$, $\sum_{a \in O} M_a = Id_{\mathcal{H}}$ and $M_a M_b = 0$ when $a \neq b$. In quantum mechanics, the **state** of a quantum system is a unit vector in Hilbert space. O is the set of possible outcomes when measuring a given quantity. If $|x\rangle$ is the state before measurement, then the probability of getting outcome a is $p(a) = \langle x | M_a | x \rangle$. The final state after measuring and getting outcome a will be $M_a | x \rangle / \langle x | M_a | x \rangle$. This probabilistic collapse of the state to one measurement outcome is called **Born's rule** and is the only way that quantum mechanics makes predictions about the outcome of experiments. From a measurement, we can construct an operator $P = \sum_{a \in O} \lambda_a M_a$ in $\mathcal{B}(\mathcal{H})$ called an observable, with real eigenvalues λ_a . The average value of a measurement is then $\langle x | P | x \rangle$. The state $M_a | x \rangle / \langle x | M_a | x \rangle$ after measurement of P is then a unit eigenvector of P called an **eigenstate**. If two observables commute, then exchanging the order in which they are measured has no effect on the outcome. Otherwise, the measurements cannot be meaningfully performed simultaneously. Since a state cannot simultaneously be an eigenstate for two non-commuting observables, one measurement disturbs the outcome of the other. The disturbance of the measurement of one observable P_1 by another P_2 is given by the magnitude of their commutator $[P_1, P_2] = P_1 P_2 - P_2 P_1$. This is Heisenberg's famous uncertainty principle.

Suppose \mathcal{H}_1 and \mathcal{H}_2 are Hilbert spaces with $|x_1\rangle, |x_2\rangle \in \mathcal{H}_1$ and $|y_1\rangle, |y_2\rangle \in \mathcal{H}_2$. Using the inner products on \mathcal{H}_1 and \mathcal{H}_2 , we define the inner product on the tensor product vector space $\mathcal{H}_1 \otimes \mathcal{H}_2$ by $\langle x_1 \otimes y_1 | x_2 \otimes y_2 \rangle = \langle x_1 | x_2 \rangle \langle y_1 | y_2 \rangle$. The completion of $\mathcal{H}_1 \otimes \mathcal{H}_2$ with respect to this inner product is called the **tensor product** of the Hilbert spaces \mathcal{H}_1 and \mathcal{H}_2 . Since we will not further discuss the vector space tensor product, we will denote the tensor product Hilbert space by $\mathcal{H}_1 \otimes \mathcal{H}_2$. If \mathcal{H}_1 and \mathcal{H}_2 have bases $\{|x\rangle\}_{x \in X}$ and $\{|y\rangle\}_{y \in Y}$ respectively, then $\mathcal{H}_1 \otimes \mathcal{H}_2$ is spanned by the elements $\{|x\rangle \otimes |y\rangle\}_{x \in X, y \in Y}$. A quantum state $|\psi\rangle \in \mathcal{H}_1 \otimes \mathcal{H}_2$ is called **bipartite**. If there are vectors $|\phi_1\rangle \in \mathcal{H}_1$ and $|\phi_2\rangle \in \mathcal{H}_2$ such that $|\psi\rangle = |\phi_1\rangle \otimes |\phi_2\rangle$, then $|\psi\rangle$ is called a **product state**. States that are not product states are called **entangled**. Measurements involving entangled states can create correlations between spatially separated experiments that are impossible in classical physics.

For some of the definitions in the next section, we need to equip $\mathcal{B}(\mathcal{H})$ with a topology. For a Hilbert space \mathcal{H} , the **weak operator topology** is the weakest topology on $\mathcal{B}(\mathcal{H})$ that makes the functional $\mathcal{B}(\mathcal{H}) \rightarrow \mathbb{C} : T \mapsto \langle x | T y \rangle$ continuous for all $|x\rangle, |y\rangle \in \mathcal{H}$, and $T \in \mathcal{B}(\mathcal{H})$. In this topology, the adjoint map is continuous. To see this, note that if $\{T_i\}_{i \in \mathbb{N}}$

is a sequence of operators converging to T , then

$$\lim_{i \rightarrow \infty} \langle (T_i^* - T^*)x | y \rangle = \lim_{i \rightarrow \infty} \langle x | (T_i - T)y \rangle = 0$$

for all $|x\rangle$ and $|y\rangle \in \mathcal{H}$.

2.3 *-algebras and representations

In this section, we recall some key concepts from the theory of *-algebras. These tools will allow us to model the algebras of observables used by players in nonlocal games. See [56, 62] for a more complete background.

A complex *-algebra \mathcal{A} is a unital algebra over \mathbb{C} with an antilinear involution

$$\mathcal{A} \rightarrow \mathcal{A} : a \mapsto a^*,$$

such that $(ab)^* = b^*a^*$. In a *-algebra, we denote the **hermitian square** as $|a|^2 := a^*a$. Let $\mathbb{C}^*\langle X \rangle$ denote the free complex *-algebra generated by the set X . If $R \subseteq \mathbb{C}^*\langle X \rangle$, let $\mathbb{C}^*\langle X : R \rangle$ denote the quotient of $\mathbb{C}^*\langle X \rangle$ by the two-sided ideal generated by R . If X and R are finite then $\mathbb{C}^*\langle X : R \rangle$ is called a **finitely presented** *-algebra.

A ***-homomorphism** $\phi : \mathcal{A} \rightarrow \mathcal{B}$ between *-algebras is an algebra homomorphism such that $\phi(x^*) = \phi(x)^*$ for all $x \in \mathcal{A}$. A ***-representation** of \mathcal{A} is a *-homomorphism $\rho : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{H})$ from \mathcal{A} to the *-algebra of bounded operators on a Hilbert space \mathcal{H} . If \mathcal{A} and \mathcal{B} are *-algebras, and $\mathbb{C}^*\langle X : R \rangle$ is a presentation of \mathcal{A} , then *-homomorphisms $\mathcal{A} \rightarrow \mathcal{B}$ correspond to homomorphisms $\phi : \mathbb{C}^*\langle X \rangle \rightarrow \mathcal{B}$ such that $\phi(r) = 0$ for all $r \in R$. Thus, a *-representation is an assignment of operators to the elements of X that satisfies the defining relations R .

If \mathcal{A} is a *-algebra, we write $a \geq b$ if $a - b$ is a sum of hermitian squares, *i.e.* there is $k \geq 0$ and $c_1, \dots, c_k \in \mathcal{A}$ such that $a - b = \sum_{i=1}^k c_i^*c_i$. A finitely presented *-algebra \mathcal{A} is called **archimedean** if for all $a \in \mathcal{A}$ there exists a $\lambda > 0$ such that $a^*a \leq \lambda 1$. The algebras we consider in this thesis are all archimedean. We will need the following lemma when handling sums of hermitian squares.

Lemma 2.3.1. *Let $a_i \in \mathcal{A}$, where \mathcal{A} is a *-algebra. Then, we have that $|\sum_{i=1}^k a_i|^2 \leq 2^{\lceil \log k \rceil} \sum_{i=1}^k |a_i|^2$.*

Proof. Since $|a + b|^2 + |a - b|^2 = 2|a|^2 + 2|b|^2$, we see that $|a + b|^2 \leq 2|a|^2 + 2|b|^2$. Thus $h(\sum_{i=1}^n a_i) \leq 2|\sum_{i=1}^{\lfloor n/2 \rfloor} a_i|^2 + 2|\sum_{i=\lfloor n/2 \rfloor + 1}^n a_i|^2$, and repeated applications gives the desired inequality. \square

If $f : \mathcal{A} \rightarrow \mathbb{C}$ is a linear functional then f is **positive** if $f(a) \geq 0$ whenever $a \geq 0$. A **state** on \mathcal{A} is a positive unital hermitian linear functional $\tau : \mathcal{A} \rightarrow \mathbb{C}$, that is $\tau(a^*a) \geq 0$, $\tau(1) = 1$, and $\tau(a^*) = \overline{\tau(a)}$ for all $a \in \mathcal{A}$. A state is **tracial** if $\tau(ab) = \tau(ba)$ for all $a, b \in \mathcal{A}$, and **faithful** if $\tau(a^*a) > 0$ for all $a \neq 0$. A tracial state τ induces the **trace norm** $\|a\|_\tau := \sqrt{\tau(a^*a)}$, also called the τ -norm. A directed set is a non-empty set D together with a **preorder** \leq , which is a binary relation on D that is reflexive and transitive. A **net** in \mathcal{A} is a function from a directed set D to \mathcal{A} taking d to x_d . whose domain D is a directed set. We say that a net is monotone increasing if $x_d \leq x_e$ whenever $d \leq e$. A state τ is called **normal** if for every monotone increasing net x_d of operators with a least upper bound x , $\tau(x_d)$ converges to $\tau(x)$. Similarly to the case of bounded linear operators on Hilbert spaces, an element $u \in \mathcal{A}$ is called **unitary** if $u^*u = 1 = uu^*$. Trace norms are unitarily invariant, meaning that $\|uav\|_\tau = \|a\|_\tau$ for all $a \in \mathcal{A}$, and all unitaries u and v .

If $\rho : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{H})$ is a $*$ -algebra representation, then a vector $|v\rangle \in \mathcal{H}$ is **cyclic** for ρ if the closure of $\rho(\mathcal{A})|v\rangle$ with respect to the Hilbert space norm is equal to \mathcal{H} . A **cyclic representation** of \mathcal{A} is a tuple $(\rho, \mathcal{H}, |v\rangle)$, where ρ is a representation of \mathcal{A} on \mathcal{H} and $|v\rangle$ is a cyclic vector for ρ . If $\tau : \mathcal{A} \rightarrow \mathbb{C}$ is a positive linear functional on \mathcal{A} , then there is a cyclic representation ρ_τ of \mathcal{A} , called the **GNS representation** of τ , such that $\tau(a) = \langle \xi_\tau | \rho_\tau(a) | \xi_\tau \rangle$ for all $a \in \mathcal{A}$. Two representations $\rho : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{H})$ and $\pi : \mathcal{A} \rightarrow \mathcal{B}(\mathcal{K})$ of \mathcal{A} are **unitarily equivalent** if there is a unitary operator $U : \mathcal{H} \rightarrow \mathcal{K}$ such that $U\rho(a)U^* = \pi(a)$ for all $a \in \mathcal{A}$. If τ is the state defined by $\tau(a) = \langle \xi | \rho(a) | \xi \rangle$ for all $a \in \mathcal{A}$ and some cyclic representation $(\rho, \mathcal{H}, |\xi\rangle)$, then $(\rho, \mathcal{H}, |\xi\rangle)$ is unitarily equivalent to the GNS representation. So states on $*$ -algebras correspond to states on Hilbert spaces up to unitary equivalence. A state τ is **finite-dimensional** if the Hilbert space \mathcal{H}_τ in the GNS representation $(\rho_\tau, \mathcal{H}_\tau, |\xi_\tau\rangle)$ is finite-dimensional. A state τ on \mathcal{A} is called **Connes-embeddable** if there is a trace-preserving embedding of \mathcal{A} into the ultrapower of the hyperfinite II_1 factor.

If \mathcal{A} is a $*$ -algebra then two elements $a, b \in \mathcal{A}$ are said to be **cyclically equivalent** if there is $k \geq 0$ and $f_1, \dots, f_k, g_1, \dots, g_k \in \mathcal{A}$ such that $a - b = \sum_{i=1}^k [f_i, g_i]$, where $[f, g] = fg - gf$. We say that $a \gtrsim b$ if $a - b$ is cyclically equivalent to a sum of squares. If τ is a tracial state on \mathcal{A} then $\tau(c_i^*c_i) \geq 0$ and $\tau([f_j, g_j]) = 0$. Thus if $a \gtrsim b$ then $\tau(a) \geq \tau(b)$, and if a and b are cyclically equivalent then $\tau(a - b) = 0$.

The $*$ -algebras we use in this work are built out of the group algebras of the finitely presented groups

$$\mathbb{Z}_q^{*V} = \langle V : x^q = 1 \text{ for all } x \in V \rangle \text{ and } \mathbb{Z}_q^V = \langle V : x^q = 1, xy = yx \text{ for all } x, y \in V \rangle.$$

Note that the group \mathbb{Z}_q^V is in bijection with the functions $V \rightarrow \mathbb{Z}_q$, via $\phi \mapsto \prod_{x \in V} x^{\phi(x)}$; we often make use of this identification implicitly. The group algebra $\mathbb{C}\mathbb{Z}_q^{*V}$ is the $*$ -algebra

generated by variables $x \in V$ with the defining relations from \mathbb{Z}_q^{*V} , along with the relations $x^*x = xx^* = 1$ for all $x \in V$. Similarly \mathbb{CZ}_q^V is the $*$ -algebra generated by variables $x \in V$ with the defining relations of \mathbb{Z}_q^V , along with the relations $x^*x = xx^* = 1$ for all $x \in V$. Notice that \mathbb{CZ}_q^V is the quotient of \mathbb{CZ}_q^{*V} by the relations $xy = yx$ for all $x, y \in V$. If \mathcal{A} and \mathcal{B} are complex $*$ -algebras, then we let $\mathcal{A} * \mathcal{B}$ denote their free product, and $\mathcal{A} \otimes \mathcal{B}$ denote their tensor product. Both are again complex $*$ -algebras.

When working with \mathbb{CZ}_q^V , a **monomial** in V is an element of the form $\prod_{x \in V} x^{a_x}$, where $0 \leq a_x < q$. We say that the monomial contains a variable $y \in V$ if $a_y > 0$. The degree of a monomial is $\sum_x a_x$. If \mathcal{A}_1 and \mathcal{A}_2 are $*$ -algebras for which we have a defined notion of monomial, then a monomial in $\mathcal{A}_1 \otimes \mathcal{A}_2$ is an element of the form $v_1 v_2$, where v_i is a monomial in \mathcal{A}_i . The degree of $v_1 v_2$ is the sum of the degrees of v_1 and v_2 , and a variable y is contained in $v_1 v_2$ if y is contained in v_1 or v_2 . For instance, a monomial in $\mathbb{CZ}_{q_1}^{V_1} \otimes \mathbb{CZ}_{q_2}^{V_2}$ is an element of the form $\prod_{x \in V_1} x^{a_x} \cdot \prod_{y \in V_2} y^{b_y}$, where $0 \leq a_x < q_1$ and $0 \leq b_y < q_2$. Similarly, a monomial in $\mathcal{A}_1 * \mathcal{A}_2$ is an element of the form $v_1 \cdots v_k$, where v_j is a monomial in \mathcal{A}_{i_j} for all $1 \leq j \leq k$, and $i_j \neq i_{j+1}$ for all $1 \leq j < k$. In this case, the degree of $v_1 \cdots v_k$ is the sum of the degrees of v_1, \dots, v_k , and a variable y is contained in $v_1 \cdots v_k$ if y is contained in one of the monomials v_1, \dots, v_k . In any $*$ -algebra where we have a defined notion of monomial, a polynomial is a linear combination of monomials.

A **C^* -algebra** \mathcal{A} is a complex $*$ -algebra with a submultiplicative Banach norm that satisfies the C^* identity $\|aa^*\| = \|a\|^2$ for all $a \in \mathcal{A}$. Every C^* -algebra can be realized as a norm-closed $*$ -subalgebra of the algebra of bounded operators $\mathcal{B}(\mathcal{H})$ on some Hilbert space \mathcal{H} . A C^* -algebra is a **von Neumann algebra** if it can be realized as a $*$ -subalgebra of $\mathcal{B}(\mathcal{H})$ which is closed in the weak operator topology. A **tracial von Neumann algebra** is a von Neumann algebra \mathcal{M} equipped with a faithful normal tracial state τ . We denote the unitary group on \mathcal{M} by $\mathcal{U}(\mathcal{M})$. If τ is a tracial state on a $*$ -algebra \mathcal{A} , and $(\rho, \mathcal{H}, |v\rangle)$ is the GNS representation, then the closure $\mathcal{M} = \overline{\rho(\mathcal{A})}$ of $\rho(\mathcal{A})$ in the weak operator topology is a von Neumann algebra, and $\tau_0(a) = \langle v | a | v \rangle$ is a faithful normal tracial state on \mathcal{M} . See [9] for more background on C^* -algebras and von Neumann algebras.

2.4 Decidability and complexity

We recall some of the basics of computational complexity and computability theory. See [2] for a more complete introduction. In computability and complexity theory, we often study questions that have a yes-or-no answer for each possible input string over an alphabet Σ . These are called **decision problems**. A set of strings called a language \mathcal{L} over an alphabet Σ is said to be **decidable** if there exists a Turing machine with input alphabet Σ that halts

on all inputs and correctly determines membership in \mathcal{L} . If there is a Turing machine M that accepts every string in \mathcal{L} but may run indefinitely on inputs not in \mathcal{L} , then we say that \mathcal{L} is **recursively enumerable (RE)**. We say that M **recognizes** \mathcal{L} . For a language \mathcal{L} , being recognizable is a weaker property than being decidable, as the Turing machine that recognizes \mathcal{L} is not required to halt on strings that are not in \mathcal{L} . A language belongs to **coRE** if its complement is in RE. If \mathcal{L} belongs to both RE and coRE, then it is decidable. A fundamental example of an RE language that is not decidable is the **halting problem**: the set of pairs (M, x) of Turing machines M and inputs x such that M halts on input x . The complement of this set, consisting of pairs (M, x) where the Turing machine does not halt on input x , belongs to coRE but not to RE.

The notion of reducibility helps us compare the difficulty of different decision problems. A language \mathcal{L} over an alphabet Σ is said to be **RE-complete** if every RE language can be reduced to it via a computable function. In particular, \mathcal{L} is RE-complete if there is a computable function f from pairs (M, x) of Turing machines and inputs to strings over Σ such that $f(M, x) \in \mathcal{L}$ if and only if M halts on input x . This expresses that \mathcal{L} is as hard as any language in RE under many-one reductions.

When a language is decidable, we classify its complexity by how many steps it takes to decide. For a Turing machine M , the **time function** $T : \mathbb{N} \rightarrow \mathbb{N}$ is defined by taking $T(n)$ to be the smallest integer such that if M halts on x , and x is an input string of length $\leq n$, then some computation path of M halts in $\leq T(n)$ steps. If there exists a Turing machine M that recognizes a language \mathcal{L} and has a time function bounded above by $T : \mathbb{N} \rightarrow \mathbb{N}$, then we say that \mathcal{L} can be recognized in time $T : \mathbb{N} \rightarrow \mathbb{N}$. A language \mathcal{L} is in P (resp. EXP) if it is decidable by a deterministic Turing machine that recognizes both $x \in \mathcal{L}$ and $x \in \bar{\mathcal{L}}$ in time $T = O(\text{poly}(n))$ (resp. $T = O(\text{exp}(n))$). A language \mathcal{L} is in NP (resp. NEXP) if it is decidable by a nondeterministic Turing machine in time $T = O(\text{poly}(n))$ (resp. $T = O(\text{exp}(n))$). A language \mathcal{L} is called **NP-hard** (resp. **NEXP-hard**) if, for any language \mathcal{L}' in NP (resp. NEXP) there is a polynomial-time computable function $f : \Sigma'^* \rightarrow \Sigma^*$ from strings over the alphabet for \mathcal{L}' to strings over the alphabet for \mathcal{L} such that $f(x) \in \mathcal{L}$ if and only if $x \in \mathcal{L}'$. The function f is called a polynomial-time reduction. Intuitively, a language \mathcal{L} is hard for a complexity class if any language in that class can be solved by first reducing to an instance of \mathcal{L} and then running a Turing machine that decides \mathcal{L} . A language \mathcal{L} is called **NP-complete** (resp. **NEXP-complete**) if it is both NP-hard (resp. NEXP-hard) and in NP (resp. NEXP).

Promise problems are a generalization of decision problems where the input is promised to be in some specific subset of inputs. That is, the accept and reject instances of the problem need not exhaust all possible inputs. When faced with an input that is not in the promised subset, an algorithm for a promise problem is allowed to give any output

and is not required to halt.

2.5 Nonlocal games and MIP*

A two-player **nonlocal game** $G = (I, \{O_i\}_{i \in I}, \pi, V)$ 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 $V(\cdot, \cdot | i, j) : O_i \times O_j \rightarrow \{0, 1\}$ for $(i, j) \in I \times I$. In the game, the players (commonly called Alice and Bob) receive questions from i and j from I with probability $\pi(i, j)$, and 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. For the sake of convenience, we have assumed that the players have the same question and answer sets. This assumption can be made without loss of generality. We often think of the question and answer sets as subsets of $\{0, 1\}^n$ and $\{0, 1\}^{m_i}$ for $i \in I$ respectively. In this case, we say that the questions have length n and the answers have length $\max_{i \in I} m_i$.

A **correlation** for a set of inputs and outputs $(I, \{O_i\}_{i \in I})$ is a family p of probability distributions $p(\cdot, \cdot | i, j)$ for all $(i, j) \in I \times I$. Correlations describe the players' behaviour in a nonlocal game. The probability $p(a, b | i, j)$ is interpreted as the probability that the players answer (a, b) on questions (i, j) . A correlation p is **classical** if there is a set Λ with a probability measure μ , and if for each $\lambda \in \Lambda$ there are functions $f_1^\lambda, f_2^\lambda : I \rightarrow \cup_i O_i$ such that $f_1^\lambda(i), f_2^\lambda(i) \in O_i$ for all $i \in I$, and $p(a, b | i, j) = \Pr_{\lambda \sim \mu}((f_1^\lambda(i) = a) \wedge (f_2^\lambda(j) = b))$ for all $i, j \in I, a \in O_i, b \in O_j$. The collection $(\Lambda, \mu, \{f_1^\lambda\}, \{f_2^\lambda\})$ is called a **classical strategy**. This captures the notion that a strategy for classical unentangled provers consists of some shared randomness that is independent of the verifier's questions and player strategies that are deterministic for a given state λ of the shared randomness. A correlation p is **quantum** if there are

- (a) finite-dimensional Hilbert spaces \mathcal{H}_A and \mathcal{H}_B ,
- (b) a projective measurement $\{M_a^i\}_{a \in O_i}$ on \mathcal{H}_A for every $i \in I$,
- (c) a projective measurement $\{N_a^j\}_{a \in O_j}$ on \mathcal{H}_B for every $j \in I$, and
- (d) 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 $i, j \in I, a \in O_i, b \in O_j$. 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 condition their answers

on measurements of that state. In quantum field theory, there is a more general notion of entanglement called the commuting operator model. Players in a nonlocal game may also employ commuting operator entanglement. A correlation is **commuting operator** if there exists

- (i) a Hilbert space \mathcal{H} ,
- (ii) projective measurements $\{M_a^i\}_{a \in O_i}$ and $\{N_a^j\}_{a \in O_i}$ on \mathcal{H} for every $i \in I$, and
- (iii) a state $|v\rangle \in \mathcal{H}$

such that $M_a^i N_b^j = N_b^j M_a^i$ and $p(a, b|i, j) = \langle v | M_a^i N_b^j | v \rangle$ for all $i, j \in I$, $a \in O_i$, and $b \in O_j$. The collection $(\mathcal{H}, \{M_a^i\}, \{N_a^j\}, |v\rangle)$ is called a **commuting operator strategy**. The set of classical correlations for a set of inputs and outputs $(I, \{O_i\})$ is denoted $C_c(I, \{O_i\})$. Similarly, the set of quantum and commuting operator correlations are denoted $C_q(I, \{O_i\})$ and $C_{qc}(I, \{O_i\})$, respectively. If the inputs and outputs are clear from context we denote the sets by C_c , C_q , and C_{qc} , respectively. It follows from the definitions that $C_c \subseteq C_q \subseteq C_{qc}$.

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

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

Given a strategy S for \mathbf{G} with corresponding correlation p , the winning probability is also denoted $\mathfrak{w}(\mathbf{G}; S) = \mathfrak{w}(\mathbf{G}; p)$. The **classical value** of \mathbf{G} is

$$\mathfrak{w}_c(\mathbf{G}) := \sup_{p \in C_c} \mathfrak{w}(\mathbf{G}; p).$$

The **quantum value** is

$$\mathfrak{w}_q(\mathbf{G}) := \sup_{p \in C_q} \mathfrak{w}(\mathbf{G}; p).$$

The **commuting operator value** is

$$\mathfrak{w}_{qc}(\mathbf{G}) := \sup_{p \in C_{qc}} \mathfrak{w}(\mathbf{G}; p).$$

A correlation p is **perfect** for \mathbf{G} if $\mathfrak{w}(\mathbf{G}, p) = 1$, and ε -**perfect** if $\mathfrak{w}(\mathbf{G}, p) \geq 1 - \varepsilon$. A strategy is ε -perfect if its corresponding correlation is ε -perfect. Since C_c (resp. C_{qc}) is closed and compact, \mathbf{G} has a perfect classical (resp. commuting operator) strategy if and only if $\mathfrak{w}_c(\mathbf{G}) = 1$ (resp. $\mathfrak{w}_{qc}(\mathbf{G}) = 1$). The set of quantum correlations C_q is not necessarily closed.

There are games for which $\mathfrak{w}_q(\mathbf{G}) = 1$, but which do not have a perfect quantum correlation. A correlation p is **quantum approximable** if it is an element of the closure $C_{qa} = \overline{C_q}$, and a game has a perfect quantum approximable strategy if and only if $\mathfrak{w}_q(\mathbf{G}) = 1$.

A nonlocal game $\mathbf{G} = (I, \{O_i\}, \pi, V)$ 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, quantum and commuting operator correlations are denoted C_c^s , C_q^s , and C_{qc}^s , respectively. A correlation is in C_{qc}^s (resp. C_q^s) if and only if there is

- (A) a Hilbert space \mathcal{H} (resp. finite dimensional Hilbert space \mathcal{H}),
- (B) a projective measurement $\{M_a^i\}_{a \in O_i}$ on \mathcal{H} for all $i \in I$, and
- (C) a state $|v\rangle \in \mathcal{H}$

such that $|v\rangle$ 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 , $i \in I$, $a \in O_i$, and $p(a, b|i, j) = \langle v|M_a^i M_b^j|v\rangle$ 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 commuting operator strategy**. If the Hilbert space \mathcal{H} is finite dimensional, then the collection $(\mathcal{H}, \{M_a^i\}, |v\rangle)$ is called a **synchronous quantum strategy**. The synchronous quantum and commuting operator values $\mathfrak{w}_q^s(\mathbf{G})$ and $\mathfrak{w}_{qc}^s(\mathbf{G})$ of a game \mathbf{G} are defined similarly to $\mathfrak{w}_q(\mathbf{G})$ and $\mathfrak{w}_{qc}(\mathbf{G})$ by replacing C_q and C_{qc} with C_q^s and C_{qc}^s , respectively. 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 $\pi(i, j) > 0$.

Every quantum correlation that is close to being 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 a synchronous quantum correlation [68]. This is also true for commuting operator correlations [46]. As a result, the synchronous quantum and commuting operator values of a synchronous game are polynomially related to the non-synchronous quantum and commuting operator values. We use a version of this statement from [47]:

Theorem 2.5.1 ([47]). *Suppose \mathbf{G} is a synchronous game with a C -diagonally dominant question distribution. If $\mathfrak{w}_q(\mathbf{G})$ (resp. $\mathfrak{w}_{qc}(\mathbf{G})$) is $\geq 1 - \varepsilon$, then $\mathfrak{w}_q^s(\mathbf{G})$ (resp. $\mathfrak{w}_{qc}^s(\mathbf{G})$) is $\geq 1 - O((\varepsilon/C)^{1/4})$.*

Note that the probability distribution in any synchronous nonlocal game can be made C -diagonally dominant by mixing the question distribution π with a distribution on symmetric question pairs (i, i) . This will only perturb the synchronous quantum and commuting operator values slightly, but it may change the general quantum and commuting operator values significantly.

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

- for all $x \in \{0, 1\}^*$ and $i \in I_x$, there are integers n_x and m_{xi} such that $I_x = \{0, 1\}^{n_x}$ and $O_{xi} = \{0, 1\}^{m_{xi}}$,
- 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 $V_x(a, b|i, j)$.

Let $c, s : \{0, 1\}^* \rightarrow [0, 1]$ be computable functions with $c(x) > s(x)$ for all $x \in \{0, 1\}^*$. A language $\mathcal{L} \subset \{0, 1\}^*$ belongs to $\text{MIP}(2, 1, c, s)$ if there is a MIP protocol $(\{\mathbf{G}_x\}, Q, V)$ such that n_x and m_{xi} are polynomial in $|x|$, Q and V run in polynomial time in $|x|$, if $x \in \mathcal{L}$ then $\mathfrak{w}_c(\mathbf{G}_x) \geq c$, and if $x \notin \mathcal{L}$ then $\mathfrak{w}_c(\mathbf{G}_x) \leq s$. The function c is called the **completeness probability**, and s is called the **soundness probability**. The functions n_x and m_{xi} are called the **question length** and **answer length** respectively. The classes $\text{MIP}^*(2, 1, c, s)$ and $\text{MIP}^{co}(2, 1, c, s)$ are defined equivalently to $\text{MIP}(2, 1, c, s)$, but with \mathfrak{w}_c replaced by \mathfrak{w}_q and \mathfrak{w}_{qc} , respectively. The protocols in these cases are called MIP^* and MIP^{co} protocols. When we make the additional restriction that the nonlocal games must be synchronous, we denote this by replacing MIP with SynMIP . The class of classical MIP protocols were characterized in [5], which found that $\text{MIP}(2, 1, 1, 1/2) = \text{NEXP}$ with polynomial-sized questions and constant-sized answers. The $\text{MIP}^* = \text{RE}$ theorem of Ji, Natarajan, Vidick, Wright, and Yuen states that $\text{MIP}^*(2, 1, 1, 1/2) = \text{RE}$ [40]. We use the following stronger version of the $\text{MIP}^* = \text{RE}$ theorem due to [20].

Theorem 2.5.2. ([20]) *There is a two-prover one round MIP^* protocol $(\{\mathbf{G}_x\}, Q, V)$ for the halting problem with completeness $c = 1$ and soundness $s = 1/2$, such that \mathbf{G}_x is a synchronous game with $\text{poly}(|x|)$ length questions, and constant length answers. Furthermore, if \mathbf{G}_x has a perfect strategy, then it has a perfect oracularizable synchronous quantum strategy.*

The constant-length answers will be key to preserving constant soundness in our reductions without requiring parallel repetition.

Chapter 3

The weighted algebra formalism

In this chapter, we introduce the weighted algebra formalism. We begin by introducing constraint system games in Section 3.1 and defining their corresponding weighted algebras in Section 3.2. We define the defect of a weighted algebra as a way of keeping track of the performance of the players' strategies in constraint system games. In section Section 3.3, we examine mappings between different versions of the constraint system algebra and how they affect the defect. Finally, we introduce the context subdivision transformation in Section 3.4 and show that it preserves the defect when the contexts have constant size.

3.1 Constraint system games

We now formally introduce constraint system games. If V is a set of variables, a **constraint on V** is a subset C of \mathbb{Z}_k^V . We call \mathbb{Z}_k the **alphabet** of the constraint, and think of it as the multiplicative group of the k^{th} roots of unity, since this is more convenient when working with observables and measurements. An **assignment to V** is an element $\phi \in \mathbb{Z}_k^V$, and we refer to the elements of C as **satisfying assignments for C** . For convenience, we assume every constraint is non-empty, i.e. has a satisfying assignment. A k -ary **constraint system (CS)** S is a pair $(X, \{(V_i, C_i)\}_{i=1}^m)$, where X is an ordered set of variables, V_i is a nonempty subset of X for all $1 \leq i \leq m$, and C_i is a constraint on the variables V_i with alphabet \mathbb{Z}_k . If $k = 2$, then we call S a **boolean constraint system (BCS)**, the name “binary constraint system” is also used in other works. When working with nonlocal games, the sets V_i are sometimes called the **contexts** of the system. If all the contexts have $|V_i| = 2$ then we call S a **2-CS**. A CS is **satisfiable** if there exists an assignment $f : X \rightarrow \mathbb{Z}_k$ such that $f|_{V_i} \in C_i$ for all i . The order on X induces an order on the contexts

V_i , which we use for some specific models of the weighted BCS algebra in Section 3.4. This is the only thing we use the order on X for, so it can be ignored otherwise. A **satisfying assignment** for S is an assignment ϕ to X such that $\phi|_{V_i} \in C_i$ for all $1 \leq i \leq m$. Also, if $V = \bigcup_{i=1}^k V_i$ and C_i is a constraint on V_i , then the **conjunction** $\bigwedge_{i=1}^k C_i$ is the constraint C on variables V such that $\phi \in C$ if and only if $\phi|_{V_i} \in C_i$ for all $1 \leq i \leq k$.

Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a CS, and let π be a probability distribution on $[m] \times [m]$. The **CS game** $\mathbf{G}(S, \pi)$ is the nonlocal game $([m], C_{i \in m}, \pi, V)$, where $V(\phi_i, \phi_j | i, j) = 1$ if $\phi_i|_{V_i \cap V_j} = \phi_j|_{V_i \cap V_j}$, and is 0 otherwise. In other words, in $\mathbf{G}(S, \pi)$, the players are given integers $i, j \in [m]$ according to the distribution π , and must reply with satisfying assignments $\phi_i \in C_i$ and $\phi_j \in C_j$ respectively. They win if their assignments agree on the variables in $V_i \cap V_j$. With this definition, $\mathbf{G}(S, \pi)$ has questions of length $\lceil \log m \rceil$, and answer sets of length $|V_i|$. This is also called the **constraint-constraint** version of the CS game to contrast it with the other versions we will now discuss.

Given a constraint system S as above, we can define other variants of the CS game. Let π' be a probability distribution on $[m]$, and define the probability distribution $\nu(i, x) = \pi'(i)/|V_i|$ for all $i \in [m]$ and $x \in V_i$, and 0 otherwise. The **constraint-variable CS game** $\mathbf{G}_{c-v}(S, \pi')$ is the nonlocal game $([m] \times X, \{C_i \times \mathbb{Z}_k\}_{i \in [m], x \in X}, \nu, V')$, where $V'(\phi, a | i, x) = 1$ if $\phi(x) = a$, and is 0 otherwise. Note that in this case, the players have different question sets, but the game can be symmetrized without changing the synchronous winning probability. In this game, one player receives $i \in [m]$ sampled from π' and the other receives a uniformly random variable $x \in V_i$. To win, the first player must answer a satisfying assignment $\phi \in C_i$ and the second must answer $a \in \mathbb{Z}_k$ such that $a = \phi(x)$.

For a 2-CS $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ and a probability distribution π' on $[m]$, define the **2-CS game** $\mathbf{G}_a(S, \pi')$ as the nonlocal game $(X, \mathbb{Z}_k, \nu_a, V_a)$, where $\nu_a(x, y) = \frac{\pi'(i)}{2}$ if $V_i = \{x, y\}$ and 0 otherwise, and $V_a(a, b | x, y) = 1$ iff there exists $\phi \in C_i$ 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.

A CS-MIP **protocol** is a family of CS games $\mathbf{G}(S_x, \pi_x)$, where $S_x = (X_x, \{(V_i^x, C_i^x)\}_{i=1}^{m_x})$, along with a probabilistic Turing machine Q and another Turing machine C , such that

1. on input x , Q outputs $(i, j) \in [m_x] \times [m_x]$ with probability $\pi_x(i, j)$, and
2. on input (x, ϕ, i) , C outputs true if $\phi \in C_i^x$ and false otherwise.

Technically, this definition should also include some way of computing the sets X_x and V_i^x . For instance, we might say that the integers $|N_x|$ and $|V_i^x|$ are all computable, and there are

computable order-preserving injections $[[V_i^x]] \rightarrow [[X_x]]$. However, for simplicity, we ignore this aspect of the definition going forward, and assume that in any CS-MIP* protocol, we have some efficient way of working with the sets X_x and V_i^x , the intersections $V_i^x \cap V_j^x$, and assignments $\phi \in \mathbb{Z}_k^{V_i^x}$. A language \mathcal{L} belongs to the complexity class CS-MIP*(s) if there is a CS-MIP protocol as above such that $\lceil \log m_x \rceil$ and $|V_i^x|$ are polynomial in $|x|$, Q and C run in polynomial time, if $x \in \mathcal{L}$ then $\mathfrak{w}_q^s(\mathbf{G}_x) = 1$, and if $x \notin \mathcal{L}$ then $\mathfrak{w}_q^s(\mathbf{G}_x) \leq s$. The parameter s is called the soundness. Any CS-MIP* protocol for \mathcal{L} can be transformed into a SynMIP* protocol by playing the game \mathbf{G}_x with the answer sets C_i replaced by $\mathbb{Z}_k^{V_i^x}$, and on input (x, ϕ, ψ, i, j) , asking the verifier V to first check that $\phi \in C_i$ and $\psi \in C_j$ using C , and then checking that $\phi|_{V_i \cap V_j} = \psi|_{V_i \cap V_j}$. Hence CS-MIP*(s) is contained in SynMIP*(2, 1, 1, s). Notice that in this modified version of the CS game, the players can answer with non-satisfying assignments, but they always lose if they do so. Thus, any strategy for the modified game can be converted into a strategy for the original game with the same winning probability, and perfect strategies for both types of games (ignoring questions that aren't in the support of π) are identical, so the SynMIP* protocol has the same completeness and soundness as the CS-MIP* protocol. Replacing \mathfrak{w}_q with \mathfrak{w}_{qc} in the definition gives us the class CS-MIP^{co}(s), which is contained in SynMIP^{co}(2, 1, 1, s). We can also define subclasses of CS-MIP* and CS-MIP^{co}. For instance, we let 3SAT-MIP* be the class of languages with a CS-MIP* protocol $(\{\mathbf{G}(S_x, \pi_x)\}, Q, C)$, in which every constraint of S_x is a 3SAT clause, i.e. a disjunction $x \vee y \vee z$, where x, y, z are either variables from S_x , negations of said variables, or constants.

An analogous protocol can be constructed using the constraint-variable version of the CS games, which we call a **constraint-variable CS-MIP protocol**. Here, since the constraint-variable game is not naturally synchronous, the only difference is that 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 gap and guarantees that the players can win near-optimally with synchronous strategies due to Theorem 2.5.1, so we do not need to worry about it in practice.

In the constraint-constraint game, if the players receive the same question $i \in [m]$, then they must reply with the same assignment ϕ to win. Consequently, if $\pi(i, i) > 0$ for all i then $\mathbf{G}(S, \pi)$ is a synchronous game. As per the previous section, a synchronous strategy for $\mathbf{G}(S, \pi)$ consists of projective measurements $\{M_\phi^i\}_{\phi \in \mathbb{Z}_2^{V_i}, i \in [m]}$, on a Hilbert space \mathcal{H} , along with a state $|v\rangle \in \mathcal{H}$ which is tracial on the algebra generated by M_ϕ^i .

Conversely, it is well-known that every synchronous game $\mathbf{G} = (I, \{\mathcal{O}_i\}, \pi, V)$ can be turned into a BCS game. One way to do this (see, e.g. [58, 57]) is to make a constraint

system with variables x_{ia} for $i \in I$ and $a \in \mathcal{O}_i$, and constraints $\bigvee_{a \in \mathcal{O}_i} x_{ia} = \text{true}$ for all $i \in [m]$ and $x_{ia} \wedge x_{jb} = \text{false}$ whenever $V(a, b|i, j) = 0$. The variable x_{ia} represents whether the player answers a on input i , and the constraints express the idea that the players must choose an answer for every question, and that they should reply with winning answers (the synchronous condition on V implies that $x_{ia} \wedge x_{ib} = \text{false}$ is a constraint for all i and $a \neq b$, which means that the players should choose a single answer for question i). The BCS game \mathbf{G}' associated with this constraint system has a perfect quantum (resp. quantum approximable, commuting operator) strategy if and only if \mathbf{G} has a perfect quantum (resp. quantum approximable, commuting operator) strategy. Unfortunately, if the size of the answer set is allowed to vary, this construction results in a game with answer sets $\{\pm 1\}^{\mathcal{O}_i}$. This means that the bit-length of the answers increases exponentially from \mathbf{G} . If $\mathfrak{w}_q(\mathbf{G}) = 1 - \varepsilon$, then $\mathfrak{w}_q(\mathbf{G}') = 1 - O(\varepsilon/|\mathcal{O}_i|)$, meaning that if this construction is used in a MIP* protocol, soundness can drop off exponentially.

To fix this, we look at the oracularization \mathbf{G}^{orac} of \mathbf{G} . There are several versions of \mathbf{G}^{orac} in the literature, all closely related. We use the version from [51], in which the verifier picks a question pair $(i_1, i_2) \in I$ according to π . The verifier then picks $a, b, c \in \{1, 2\}$ uniformly at random. When $a = 1$, they send player b both questions (i_1, i_2) , and the other player question (i_c) . Player b must respond with $a_j \in \mathcal{O}_j$ such that $V(a_1, a_2|i_1, i_2) = 1$, and the other player responds with $b \in \mathcal{O}_{i_c}$. The players win if $a_c = b$. If $a = 2$, both players are sent (i_1, i_2) and must respond with (a_1, a_2) and (b_1, b_2) in $\mathcal{O}_{i_1} \times \mathcal{O}_{i_2}$. They win if $(a_1, a_2) = (b_1, b_2)$. If \mathbf{G} has questions of length q and answers of length a , then \mathbf{G}^{orac} has questions of length $2q$ and answers of length $2a$, so this construction only increases the question and answer length polynomially. The following lemma shows that this construction is sound, in the sense that $\mathfrak{w}_q(\mathbf{G}^{orac})$ cannot be much larger than $\mathfrak{w}_q(\mathbf{G})$.

Lemma 3.1.1 ([51, 40]). *Let \mathbf{G} be a synchronous game. If \mathbf{G} has a perfect oracularizable synchronous strategy, then \mathbf{G}^{orac} has a perfect synchronous strategy. Conversely, if $\mathfrak{w}_q(\mathbf{G}^{orac}) = 1 - \varepsilon$, then $\mathfrak{w}_q(\mathbf{G}) \geq 1 - \text{poly}(\varepsilon)$.*

Proof. This is asserted in Definition 17.1 of [51]. Although a proof isn't supplied, the proof follows the same lines as Theorem 9.3 of [40]. \square

Given a synchronous game $\mathbf{G} = (I, \{O_i\}, \pi, V)$ where $I \subseteq \{0, 1\}^n$ and $O_i \subseteq \{0, 1\}^{m_i}$, construct a constraint system B as follows. Take X to be the set of variables x_{ij} , where $i \in I$ and $1 \leq j \leq m_i$. Let $V_i = \{x_{ij}, 1 \leq j \leq m_i\}$, and identify $\mathbb{Z}_2^{V_i}$ with bit strings $\{0, 1\}^{m_i}$, where the assignment to x_{ij} corresponds to the j th bit, and let $C_i \subseteq \mathbb{Z}_2^{V_i}$ be the subset corresponding to O_i . Let $P = \{(i, j) \in I \times I : \pi(i, j) > 0\}$. For $(i, j) \in P$, let $V_{ij} = V_i \cup V_j$, and let $C_{ij} \subseteq \mathbb{Z}_2^{V_{ij}} = \mathbb{Z}_2^{V_i} \times \mathbb{Z}_2^{V_j}$ be the set of pairs of strings (a, b) such that

$a \in O_i, b \in O_j$, and $V(a, b|i, j) = 1$. Then B is the constraint system with variables X and constraints $\{(V_i, C_i)\}_{i \in I}$ and $\{(V_{ij}, C_{ij})\}_{(i,j) \in P}$. Let $I' = I \cup P$ and π^{orac} be the probability distribution on $I' \times I'$ such that

$$\pi^{orac}(i', j') = \begin{cases} \frac{1}{8}\pi(i, j) & i' = (i, j), j' = i \\ \frac{1}{8}\pi(i, j) & i' = (i, j), j' = j \\ \frac{1}{8}\pi(i, j) & i' = i, j' = (i, j) \\ \frac{1}{8}\pi(i, j) & i' = j, j' = (i, j) \\ \frac{1}{2}\pi(i, j) & i' = j' = (i, j) \\ 0 & \text{otherwise} \end{cases}$$

Then $\mathbf{G}(B, \pi^{orac}) = \mathbf{G}^{orac}$, so the oracularization of a synchronous game is a BCS game. As a result, Theorem 2.5.2 has the following corollary:

Corollary 3.1.2. *There is a BCS – MIP* protocol $(\mathbf{G}(B_x, \pi_x), S, C)$ for the halting problem with constant soundness $s < 1$, where B_x has exponentially many contexts of constant size. Furthermore, if $\mathbf{G}(B_x, \pi_x)$ has a perfect quantum strategy, then it has a perfect oracularizable synchronous quantum strategy.*

Proof. Let $(\{\mathbf{G}_x\}, Q, V)$ be the protocol from Theorem 2.5.2. Then \mathbf{G}_x^{orac} is a BCS game in which the underlying BCS has exponentially many contexts of constant size. If \mathbf{G}_x has a perfect strategy, then it has a perfect oracularizable synchronous strategy, so the transformation to \mathbf{G}_x^{orac} preserves perfect completeness. The probability distribution π^{orac} and the constraints of \mathbf{G}_x^{orac} can be computed in polynomial time from Q and V , so by Lemma 3.1.1 there is a BCS-MIP* protocol for the halting problem with constant soundness $s' < 1$. If \mathbf{G}_x^{orac} has a perfect quantum strategy, then \mathbf{G}_x has a perfect oracularizable synchronous strategy \mathcal{S} constructed using the players' measurements for the question in \mathbf{G}_x^{orac} where they are asked both questions (i, j) from the original game \mathbf{G}_x , and the same state as their strategy for \mathbf{G}_x^{orac} . From \mathcal{S} we can construct a strategy \mathcal{S}' for \mathbf{G}_x^{orac} that is synchronous and oracularizable. The measurements in \mathcal{S} for any pair of questions (i, j) asked in the \mathbf{G}_x commute and satisfy the decision predicate for the \mathbf{G}_x , so when a player is asked for both questions (i, j) in \mathbf{G}_x^{orac} they can simultaneously measure the observables for i and j from \mathcal{S} . Otherwise, if the player is asked only i or j , they measure the observable from \mathcal{S} for this question. The strategy \mathcal{S}' constructed this way is perfect, synchronous, and oracularizable. \square

3.2 Weighted algebras for CS games

It is often worth thinking about synchronous strategies more abstractly. For any finite set of variables X , we can suppose that there is an ordering on the set and that the ordering induces an ordering on any subset of variables $V \subseteq X$. For a set of variables V , recall that \mathbb{CZ}_k^{*V} is the free algebra generated by order- k unitaries labelled by the elements $x \in V$, and \mathbb{CZ}_k^V is its abelianization. For any order- k unitary x and $a \in \mathbb{Z}_k$, write $\Pi_a^{(k)}(x) = \frac{1}{k} \sum_{n=0}^{k-1} \omega_k^{-na} x^n$ for the projector onto the ω_k^a -eigenspace of x ; where k is clear, we suppress the superscript (k) . Given $\phi \in \mathbb{Z}_k^V$, we define the element $\Phi_{V,\phi}^{(k)} \in \mathbb{CZ}_k^{*V}$ as

$$\Phi_{V,\phi}^{(k)} = \prod_{x \in V} \Pi_{\phi(x)}^{(k)}(x),$$

where the product is ordered according to the ordering of V , and as before the superscript (k) is suppressed where clear. We use the same notation for the image of $\Phi_{V,\phi}^{(k)}$ under any homomorphism when the homomorphism is clear. In particular, in \mathbb{CZ}_k^V , $\{\Phi_{V,\phi}\}_{\phi \in \mathbb{Z}_k^V}$ forms a generating PVM for the commutative algebra. Given a constraint (V, C) over the alphabet \mathbb{Z}_k , we write

$$\mathcal{A}(V, C) = \mathbb{CZ}_k^V / \langle \Phi_{V,\phi} : \phi \notin C \rangle.$$

Since \mathbb{CZ}_k^V is commutative, the image of $\Phi_{V,\phi}$ is independent of the order of V ; however we will work with \mathbb{CZ}_q^{*V} in Section 3.4. The algebra $\mathcal{A}(V, C)$ is isomorphic to the C^* -algebra of functions on the finite set C . Consequently, if $\rho : \mathcal{A}(V, C) \rightarrow \mathcal{B}(\mathcal{H})$ is a $*$ -representation, then $\{\rho(\Phi_{V,\phi})\}_{\phi \in C}$ is a projective measurement on \mathcal{H} , and conversely if $\{M_\phi\}_{\phi \in C}$ is a projective measurement on \mathcal{H} then there is a $*$ -representation $\rho : \mathcal{A}(V, C) \rightarrow \mathcal{B}(\mathcal{H})$ with $\rho(\Phi_{V,\phi}) = M_\phi$.

If $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ is a CS, then we define the **constraint-constraint CS algebra** $\mathcal{A}_{c-c}(S)$ by the free product $\mathcal{A}_{c-c}(S) := *_{i \in [m]} \mathcal{A}(V_i, C_i)$. We let $\sigma_i : \mathcal{A}(V_i, C_i) \rightarrow \mathcal{A}(B)$ denote the natural inclusion of the i th factor, so $\mathcal{A}_{c-c}(S)$ is generated by the involutions $\sigma_i(x)$ for $i \in [m]$ and $x \in V_i$. Equivalently, $\mathcal{A}_{c-c}(S)$ is generated by the projections $\sigma_i(\Phi_{V_i,\phi})$ for $i \in [m]$ and $\phi \in C_i$. To avoid clogging up formulas with symbols, we'll often write $\Phi_{V_i,\phi}$ instead of $\sigma_i(\Phi_{V_i,\phi})$ when it's clear what subalgebra $\mathcal{A}(V_i, C_i)$ the element belongs to. As with $\mathcal{A}(V, C)$, representations α of $\mathcal{A}_{c-c}(S)$ are in bijective correspondence with families of projective measurements $\{M_\phi^i\}_{\phi \in C_i, i \in [m]}$ via the relation $M_\phi^i = \alpha(\Phi_{V_i,\phi})$. If $(\{M_\phi^i\}, |v\rangle, \mathcal{H})$ is a synchronous commuting operator strategy for $\mathbf{G}(S, \pi)$, and $\alpha : \mathcal{A}_{c-c}(S) \rightarrow \mathcal{B}(\mathcal{H})$ is the representation with $\alpha(\Phi_{V_i,\phi}) = M_\phi^i$, then $a \mapsto \langle v | \alpha(a) | v \rangle$ is a tracial state on $\mathcal{A}_{c-c}(S)$. Conversely, if τ is a tracial state on $\mathcal{A}_{c-c}(S)$, then the GNS representation theorem implies that there is a synchronous commuting operator strategy

$\mathcal{S} = (\{M_\phi^i\}, |v\rangle, \mathcal{H})$ such that $\tau(a) = \langle v | \alpha(a) | v \rangle$ where α is the representation corresponding to $\{M_\phi^i\}$. Note that the trace is faithful on the image of the GNS representation. As a result, synchronous commuting operator strategies for $\mathbb{G}(S, \pi)$ and tracial states on $\mathcal{A}_{c-c}(S)$ can be used interchangeably, and in particular $p \in C_{qc}$ if and only if there is a tracial state τ with $p(\phi, \psi | i, j) = \tau(\Phi_{V_i, \phi} \Phi_{V_j, \psi})$ for all i, j, ϕ , and ψ . Finite-dimensional tracial states on $\mathcal{A}_{c-c}(S)$ can be used interchangeably with synchronous quantum strategies for $\mathbb{G}(S, \pi)$, and $p \in C_q$ if and only if there is a finite-dimensional tracial state τ with $p(\phi, \psi | i, j) = \tau(\Phi_{V_i, \phi} \Phi_{V_j, \psi})$ for all i, j, ϕ , and ψ . Similarly, $p \in C_{qa}$ if and only if there is a Connes-embeddable tracial state τ such that $p(\phi, \psi | i, j) = \tau(\Phi_{V_i, \phi} \Phi_{V_j, \psi})$ for all i, j, ϕ , and ψ [45].

A correlation p is perfect for a CS game $\mathcal{G}(S, \pi)$ if $p(\phi, \psi | i, j) = 0$ whenever $\pi(i, j) > 0$ and (ϕ, ψ) is a losing answer to questions (i, j) . As a result, a tracial state τ on $\mathcal{A}_{c-c}(S)$ is **perfect** (aka. corresponds to a perfect correlation) if and only if $\tau(\Phi_{V_i, \phi} \Phi_{V_j, \psi}) = 0$ whenever $\phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}$. Consequently a tracial state on $\mathcal{A}_{c-c}(S)$ is perfect for $\mathcal{G}(S, \pi)$ if and only if it is the pullback of a tracial state on the **synchronous algebra** of $\mathcal{G}(S, \pi)$, which is the quotient

$$\text{SynAlg}(S, \pi) = \mathcal{A}_{c-c}(S) / \langle \Phi_{V_i, \phi} \Phi_{V_j, \psi} = 0 \text{ for all } i, j \in [m] \text{ with } \pi(i, j) > 0 \\ \text{and } \phi \in C_i, \psi \in C_j \text{ with } \phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j} \rangle.$$

For the special case of BCS games, this result about perfect strategies is due to Kim, Paulsen, and Schafhauser [45]. The general notion of a synchronous algebra is due to [34]. In [26, 58], it is shown that the synchronous algebra of a BCS game is isomorphic to the so-called BCS algebra of the game, and the general CS case follows by similar arguments. In working with MIP* protocols, we also need a way to keep track of ε -perfect strategies. In [57], it is shown that ε -perfect strategies for a BCS game correspond to ε -representations of the BCS algebra, where an ε -representation is a representation of $\mathcal{A}_{c-c}(S)$ such that all the defining relations of $\text{SynAlg}(B, \pi)$ are bounded by ε in the normalized Frobenius norm. In this prior work, the focus was on the behaviour of ε -perfect strategies for a fixed game, so the number of questions and answers was constant. For MIP* protocols, the game size is not constant, and we need to work with approximate representations where the average, rather than the maximum, of the norms of the defining relations is bounded. For this, we introduce the following algebraic structure:

Definition 3.2.1. A *(finitely-supported) weight function* on a set X is a function $\mu : X \rightarrow [0, +\infty)$ such that $\text{supp}(\mu) := \mu^{-1}((0, +\infty))$ is finite. A *weighted *-algebra* is a pair (\mathcal{A}, μ) 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

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

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

Since μ is finitely supported, the sum in the definition of the defect is finite, and hence is well-defined. Note that traces τ on a weighted algebra (\mathcal{A}, μ) with $\text{def}(\tau) = 0$ correspond to traces on the algebra $\mathcal{A}/\langle \text{supp}(\mu) \rangle$. In general, $\text{def}(\tau)$ is a measure of how far τ is from being a trace on \mathcal{A} . Thus we can think of a weighted algebra (\mathcal{A}, μ) as a presentation or model for the algebra $\mathcal{A}/\langle \text{supp}(\mu) \rangle$ that allows us to talk about approximate traces on this algebra.

Definition 3.2.2. Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a CS, and let π be a probability distribution on $[m] \times [m]$. The **(weighted) constraint-constraint CS algebra $\mathcal{A}(S, \pi)$** is the $*$ -algebra $\mathcal{A}_{c-c}(S)$, with weight function $\mu_{c-c, \pi}$ defined by

$$\mu_{c-c, \pi}(\Phi_{V_i, \phi} \Phi_{V_j, \psi}) = \pi(i, j)$$

for all $i, j \in [m]$ and $\phi \in C_i, \psi \in C_j$ with $\phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}$, and $\mu_{c-c, \pi}(r) = 0$ for all other $r \in \mathcal{A}_{c-c}(S)$.

Note that $\mathcal{A}_{c-c}(S)/\langle \text{supp}(\mu_{c-c, \pi}) \rangle$ is the synchronous algebra $\text{SynAlg}(S, \pi)$ defined above, so $\mathcal{A}_{c-c}(S, \pi)$ is a model of this synchronous algebra, and perfect strategies for $\mathcal{G}(S, \pi)$ correspond to tracial states τ on $\mathcal{A}_{c-c}(S, \pi)$ with $\text{def}(\tau) = 0$. The following lemma is an immediate consequence of the definitions:

Lemma 3.2.3. Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a CS, and let π be a probability distribution on $[m] \times [m]$. A tracial state τ on $\mathcal{A}_{c-c}(S)$ is an ε -perfect strategy for $\mathcal{G}(S, \pi)$ if and only if $\text{def}(\tau) \leq \varepsilon$.

Proof. Let p be the correlation corresponding to τ , so $p(\phi, \psi | i, j) = \tau(\Phi_{V_i, \phi} \Phi_{V_j, \psi})$. Then

$$\text{def}(\tau) = \sum \pi(i, j) \tau(\Phi_{V_i, \phi} \Phi_{V_j, \psi}),$$

where the sum is across $i, j \in [m]$ and $\phi \in C_i, \psi \in C_j$ with $\phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}$. So $\text{def}(\tau) = 1 - \mathfrak{w}(\mathcal{G}(B, \pi); p)$. \square

In addition to the constraint-constraint algebra, we define the following weighted algebras for constraint systems:

Definition 3.2.4. *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a k -ary CS. We define the following algebras:*

- **The *inter-contextual algebra*:** *Given a probability distribution π on $[m] \times [m]$, define the weighted algebra $\mathcal{A}_{inter}(S, \pi) = (\mathcal{A}_{c-c}(S), \mu_{inter, \pi})$, with weight function $\mu_{inter, \pi}(\sigma_i(x)^l - \sigma_j(x)^l) = \pi(i, j)$ for all $i \neq j \in [m]$, $l \in [k-1]$ and $x \in V_i \cap V_j$, and 0 on all other elements.*
- **The *constraint-variable algebra*** $\mathcal{A}_{c-v}(S) = \bigstar_{i=1}^m \mathcal{A}(V_i, C_i) * \mathbb{CZ}_k^{*X}$: *As above, write the inclusion $\sigma_i : \mathcal{A}(V_i, C_i) \rightarrow \mathcal{A}_{c-v}(S)$, and write also the inclusion $\sigma' : \mathbb{CZ}_k^{*X} \rightarrow \mathcal{A}_{c-v}(S)$. For a probability distribution π' on $[m]$, define the weighted algebra $\mathcal{A}_{c-v}(S, \pi') = (\mathcal{A}_{c-v}(S), \mu_{c-v, \pi'})$ where the weight function $\mu_{c-v}(\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))) = \frac{\pi'(i)}{|V_i|}$ for all $x \in V_i$, $\phi \in C_i$, and 0 on all other elements.*
- **The *assignment algebra*** $\mathcal{A}_a(S) = \mathbb{CZ}_k^{*X}$: *Given a probability distribution π' on $[m]$, define the weighted algebra $\mathcal{A}_a(S, \pi') = (\mathcal{A}_a(S), \mu_{a, \pi'})$ where the weight function $\mu_{a, \pi'}(\Phi_{V_i, \phi}) = \pi'(i)$ for all $\phi \notin C_i$, and 0 on all other elements.*
- **The *assignment algebra with commutation*:** *Define $\mathcal{A}_{a+comm}(S, \pi') = (\mathcal{A}_a(S), \mu_{a, \pi'} + \mu_{comm, \pi'})$, where the weight function $\mu_{comm, \pi'}([\Pi_a(x), \Pi_b(y)]) = \sum_{i. x, y \in V_i} \pi'(i)$ for all $a, b \in \mathbb{Z}_k$ and $x, y \in X$, and 0 on all other elements.*

The goal of defining these algebras is that their traces correspond to synchronous strategies for other variants of the CS game, except for $\mathcal{A}_{inter}(S, \pi)$. Instead, the inter-contextual algebra is used as an intermediary step in the proof of the soundness of the subdivision transformation in Section 3.4.

As we saw, a tracial state on $\mathcal{A}_{c-c}(S)$ is perfect for $\mathbf{G}(S, \pi)$ if and only if it is the pullback of a tracial state on the synchronous algebra of $\mathbf{G}(S, \pi)$. Thus, $\mathcal{A}_{c-c}(S, \pi)$ can be thought of as a model of $\text{SynAlg}(S, \pi)$. In fact, assuming that for all $x \in V_i$ there exists some j such that $x \in V_j$ and $\pi(i, j) > 0$, each of the algebras defined above — except in general the assignment algebra — are models for $\text{SynAlg}(S, \pi)$, since

$$\begin{aligned} \text{SynAlg}(S, \pi) &= \mathcal{A}_{c-c}(S) / \langle \text{supp}(\mu_{c-c, \pi}) \rangle = \mathcal{A}_{c-c}(S) / \langle \text{supp}(\mu_{inter, \pi}) \rangle \\ &= \mathcal{A}_{c-v}(S) / \langle \text{supp}(\mu_{c-v, \pi'}) \rangle = \mathcal{A}_a(S) / \langle \text{supp}(\mu_{a, \pi'} + \mu_{comm, \pi'}) \rangle, \end{aligned}$$

where $\pi'(i) = \sum_j \pi(i, j)$. Thus, if one of these algebras has a perfect trace, then all the others do.

In addition to looking at CS games, we also want to consider transformations between constraint systems and examine how they affect the corresponding games. To keep track of how near-perfect strategies change, we introduce a notion of homomorphism for weighted algebras.

Definition 3.2.5. *Let (\mathcal{A}, μ) and (\mathcal{B}, ν) be weighted $*$ -algebras, and let $C > 0$. A **C -homomorphism** $\alpha : (\mathcal{A}, \mu) \rightarrow (\mathcal{B}, \nu)$ is a $*$ -homomorphism $\alpha : \mathcal{A} \rightarrow \mathcal{B}$ such that*

$$\alpha\left(\sum_{a \in \mathcal{A}} \mu(a) a^* a\right) \lesssim C \sum_{b \in \mathcal{B}} \nu(b) b^* b.$$

The point of this definition is the following:

Lemma 3.2.6. *Suppose $\alpha : (\mathcal{A}, \mu) \rightarrow (\mathcal{B}, \nu)$ is a C -homomorphism. If τ is a trace on (\mathcal{B}, ν) , then $\text{def}(\tau \circ \alpha) \leq C \text{def}(\tau)$.*

Proof. Let $A = \alpha(\sum_{a \in \mathcal{A}} \mu(a) a^* a)$ and $B = \sum_{b \in \mathcal{B}} \nu(b) b^* b$. Note that

$$\text{def}(\tau \circ \alpha) = \sum_{a \in \mathcal{A}} \mu(a) \|a\|_{\tau \circ \alpha} = \sum_{a \in \mathcal{A}} \mu(a) \tau(\alpha(a^* a)) = \tau(A),$$

By the definition of \lesssim , there are c_1, \dots, c_k and $f_1, \dots, f_\ell, g_1, \dots, g_\ell \in \mathcal{B}$ such that

$$CB - A = \sum_{i=1}^k c_i^* c_i + \sum_{j=1}^{\ell} [f_j, g_j].$$

Since τ is a tracial state, $\tau(c_i^* c_i) \geq 0$ and $\tau([f_j, g_j]) = 0$ for all i and j . Hence $C\tau(B) \geq \tau(A)$ as required. \square

We can now easily handle transformations of constraint systems which apply a homomorphism to each context. Note that a homomorphism $\sigma : \mathcal{A}(V, C) \rightarrow \mathcal{A}(W, D)$ between finite abelian C^* -algebras is equivalent to a function $f : D \rightarrow C$. Indeed, given a function $f : D \rightarrow C$, we can define a homomorphism σ by $\sigma(\Phi_{V, \phi}) = \sum_{W, \psi \in f^{-1}(\phi)} \Phi_{W, \psi}$, and it is not hard to see that all homomorphisms have this form. We extend this notion to CS algebras in the following way.

Definition 3.2.7. *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ and $S' = (X', \{(W_i, D_i)\}_{i=1}^m)$ be constraint systems. A homomorphism $\sigma : \mathcal{A}(S) \rightarrow \mathcal{A}(S')$ is a **classical homomorphism** if*

1. $\sigma(\mathcal{A}(V_i, C_i)) \subseteq \mathcal{A}(W_i, D_i)$ for all $1 \leq i \leq m$, and
2. if $\sigma(\Phi_{V_i, \phi_i}) = \sum_k \Phi_{W_i, \psi_{ik}}$, $\sigma(\Phi_{V_j, \phi_j}) = \sum_k \Phi_{W_j, \psi_{jk}}$, and $\phi_i|_{V_i \cap V_j} \neq \phi_j|_{V_i \cap V_j}$ then $\psi_{ik}|_{W_i \cap W_j} \neq \psi_{jl}|_{W_i \cap W_j}$ for all k, l .

To explain this definition, note that condition (1) implies that σ restricts to a homomorphism $\mathcal{A}(V_i, C_i) \rightarrow \mathcal{A}(W_i, D_i)$, and hence gives a collection of functions $f_i : D_i \rightarrow C_i$ for all $1 \leq i \leq m$. Condition (2) states that if $f_i(\phi)|_{V_i \cap V_j} \neq f_j(\psi)|_{V_i \cap V_j}$ for some $\phi \in D_i$, $\psi \in D_j$, then $\phi|_{W_i \cap W_j} \neq \psi|_{W_i \cap W_j}$. Conversely, any collection of functions $f_i : D_i \rightarrow C_i$ satisfying this condition can be turned into a classical homomorphism $\sigma : \mathcal{A}(S) \rightarrow \mathcal{A}(S')$.

Lemma 3.2.8. *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ and $S' = (Y, \{(W_i, D_i)\}_{i=1}^m)$ be constraint systems, and let π be a probability distribution on $[m] \times [m]$. If $\sigma : \mathcal{A}(S) \rightarrow \mathcal{A}(S')$ is a classical homomorphism, then σ is a 1-homomorphism $\mathcal{A}(S, \pi) \rightarrow \mathcal{A}(S', \pi)$.*

Proof. Suppose σ arises from a family of functions $f_i : D_i \rightarrow C_i$ as above. For any $1 \leq i, j \leq m$, let $R_{ij} = \{(\phi, \psi) \in C_i \times C_j : \phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}\}$, and let $T_{ij} = \{(\phi, \psi) \in D_i \times D_j : \phi|_{W_i \cap W_j} \neq \psi|_{W_i \cap W_j}\}$. Then

$$\begin{aligned} \sigma \left(\sum_{i,j} \sum_{(\phi, \psi) \in R_{ij}} \pi(i, j) \Phi_{V_i, \phi} \Phi_{V_j, \psi} \right) &= \sum_{i,j} \sum_{\phi' \in f_i^{-1}(\phi), \psi' \in f_j^{-1}(\psi)} \pi(i, j) \Phi_{W_i, \phi'} \Phi_{W_j, \psi'} \\ &\leq \sum_{i,j} \sum_{(\phi, \psi) \in T_{ij}} \pi(i, j) \Phi_{W_i, \phi} \Phi_{W_j, \psi}. \end{aligned}$$

□

One situation where we get a classical homomorphism is the following:

Corollary 3.2.9. *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a CS, and let $S' = (X', \{(W_i, D_i)\}_{i=1}^m)$ be a CS with $X \subset X'$, $V_i \subseteq W_i$ for all $1 \leq i \leq m$, and $W_i \cap W_j = V_i \cap V_j$ for all $1 \leq i, j \leq m$. Suppose that for all $i \in [m]$, $\phi \in C_i$ if and only if there exists $\psi \in D_i$ with $\psi|_{V_i} = \phi$. Then for any probability distribution π on $[m] \times [m]$, the homomorphism*

$$\sigma : \mathcal{A}(S) \rightarrow \mathcal{A}(S') : \sigma_i(x) \mapsto \sigma_i(x) \text{ for } i \in [m], x \in V_i$$

defined by the inclusions $V_i \subseteq W_i$ is a 1-homomorphism $\mathcal{A}(S, \pi) \rightarrow \mathcal{A}(S', \pi)$, and there is another 1-homomorphism $\sigma' : \mathcal{A}(S', \pi) \rightarrow \mathcal{A}(S, \pi)$.

In the next section we will see the relationship between imperfect traces on the different CS algebras.




Figure 3.1: C -homomorphisms (solid arrows) and trace-dependent mappings (dashed arrows) between the different weighted algebras for a k -ary CS $S = (X, \{(V_i, C_i)\}_{i=1}^m)$. Here $\pi'(i) = \sum_j \pi(i, j)$, $L = \max_i |V_i|$, $P = \max_{i,j. V_i \cap V_j \neq \emptyset} \frac{\pi'(i)}{\pi(i,j)}$, and $B(S)$ is the BCS defined in Definition 3.3.1.

3.3 Relations between CS algebras

In this section, we study the relationships between the weighted algebras coming from a CS. Fig. 3.1 summarises the C -homomorphisms we find between these weighted algebras. First, we note that any constraint-constraint algebra is equivalent to a related BCS algebra.

Definition 3.3.1. *Given a k -ary CS $S = (X, \{(V_i, C_i)\}_{i=1}^m)$, the **boolean form** of S is defined as $B(S) = (X', \{(V'_i, C'_i)\}_{i=1}^m)$, where $X' = \{(x, a) | x \in X, a \in \mathbb{Z}_k\}$, $V'_i = \{(x, a) | x \in V_i, a \in \mathbb{Z}_k\}$, $C'_i = \{\phi' | \phi \in C_i\}$ where*

$$\phi'(x, a) = \begin{cases} 1 & \phi(x) = a \\ 0 & \text{else} \end{cases}.$$

This corresponds to replacing each variable in X with k indicator variables indicating which value in \mathbb{Z}_k is assigned to x .

Lemma 3.3.2. *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a k -ary CS and π be a probability distribution on $[m] \times [m]$. There is a $*$ -isomorphism $\alpha : \mathcal{A}_{c-c}(B(S)) \rightarrow \mathcal{A}_{c-c}(S)$ such that α and α^{-1} are 1-homomorphisms between $\mathcal{A}_{c-c}(B(S), \pi)$ and $\mathcal{A}_{c-c}(S, \pi)$.*

Proof. Write $B(S) = (X', \{(V'_i, C'_i)\}_{i=1}^m)$. Consider first the map $\alpha_i : \mathbb{C}\mathbb{Z}_2^{V'_i} \rightarrow \mathbb{C}\mathbb{Z}_k^{V_i}$ defined as the homomorphism such that

$$\alpha((x, a)) = 1 - 2\Pi_a^{(k)}(x).$$

Since for any $\phi \in \mathbb{Z}_2^{V'_i}$, $\Phi_{V'_i, \phi} = \prod_{(x,a) \in V'_i} \frac{1 + (-1)^{\phi(x,a)}(x,a)}{2} = \prod_{x \in V_i} \prod_{a \in \mathbb{Z}_k} \frac{1 + (-1)^{\phi(x,a)}(x,a)}{2}$,

$$\begin{aligned} \alpha_i(\Phi_{V'_i, \phi}) &= \prod_{x \in V_i} \prod_{a \in \mathbb{Z}_k} \frac{1 + (-1)^{\phi(x,a)}(1 - 2\Pi_a^{(k)}(x))}{2} = \prod_{x \in V_i} \prod_{\phi(x,a)=0} (1 - \Pi_a^{(k)}(x)) \prod_{\phi(x,a)=1} \Pi_a^{(k)}(x) \\ &= \begin{cases} \prod_{x \in V_i} \Pi_{\psi(x)}^{(k)}(x) & \exists \psi \in \mathbb{Z}_k^{V_i}. \phi(x, b) = 1 \iff b = \psi(x) \forall x \in V_i \\ 0 & \text{otherwise} \end{cases}. \end{aligned}$$

This means that, if $\alpha_i(\Phi_{V'_i, \phi}) \neq 0$, there exists $\psi \in \mathbb{Z}_k^{V_i}$ such that $\phi = \psi'$. Further, in this case $\alpha_i(\Phi_{V'_i, \phi}) = \Phi_{V_i, \psi}^{(k)}$. As such, the map α_i factors through to a map $\bar{\alpha}_i : \mathcal{A}(V'_i, C'_i) \rightarrow \mathcal{A}(V_i, C_i)$. Since we have $\bar{\alpha}_i(\Phi_{V'_i, \phi'}) = \Phi_{V_i, \phi}^{(k)}$ for all $\phi \in C_i$, $\bar{\alpha}_i$ is an isomorphism. Finally, we take $\alpha : \mathcal{A}_{c-c}(B(S)) \rightarrow \mathcal{A}_{c-c}(S)$ to be given by $\bar{\alpha}_i$ on the corresponding term of the free product. By construction, α is an isomorphism, and α and α^{-1} both preserve the weight function as they exchange $\Phi_{V'_i, \phi'}$ and $\Phi_{V_i, \phi}^{(k)}$. \square

Now we examine the homomorphisms between the $\mathcal{A}_{c-c}(S, \pi)$ and $\mathcal{A}_{inter}(S, \pi)$ algebras.

Proposition 3.3.3. *Suppose $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ is a k -ary CS and π is a probability distribution on $[m] \times [m]$. Then the identity map $\mathcal{A}_{c-c}(S) \rightarrow \mathcal{A}_{c-c}(S)$ gives an $O(1)$ -homomorphism $\mathcal{A}_{c-c}(S, \pi) \rightarrow \mathcal{A}_{inter}(S, \pi)$, and an $O(L)$ -homomorphism $\mathcal{A}_{inter}(S, \pi) \rightarrow \mathcal{A}_{c-c}(S, \pi)$, where $L = \max_{i,j} |V_i \cap V_j|$.*

Recall that $\sigma_i : \mathcal{A}(V_i, C_i) \rightarrow \mathcal{A}_{c-c}(S)$ is the natural inclusion of the i th factor.

Proof. Fix $1 \leq i, j \leq m$. Since $\Phi_{V_i, \phi}$ is a projection in $\mathcal{A}(V_i, C_i)$, $(\Phi_{V_i, \phi} \Phi_{V_j, \psi})^* (\Phi_{V_i, \phi} \Phi_{V_j, \psi})$ is cyclically equivalent to $\Phi_{V_i, \phi} \Phi_{V_j, \psi}$ for all $\phi \in C_i, \psi \in C_j$. For $x \in V_i \cap V_j$, let R_x be the pairs $(\phi, \psi) \in C_i \times C_j$ such that $\phi(x) \neq \psi(x)$. Then

$$\sum_{\phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}} \Phi_{V_i, \phi} \Phi_{V_j, \psi} \lesssim \sum_{x \in V_i \cap V_j} \sum_{(\phi, \psi) \in R_x} \Phi_{V_i, \phi} \Phi_{V_j, \psi},$$

and since $\phi|_{V_i \cap V_j}$ and $\psi|_{V_i \cap V_j}$ can disagree in at most $|V_i \cap V_j|$ places,

$$\sum_{x \in V_i \cap V_j} \sum_{(\phi, \psi) \in R_x} \Phi_{V_i, \phi} \Phi_{V_j, \psi} \lesssim |V_i \cap V_j| \sum_{\phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}} \Phi_{V_i, \phi} \Phi_{V_j, \psi}.$$

Fix $x \in V_i \cap V_j$, and let $V'_i = V_i \setminus \{x\}$, $V'_j = V_j \setminus \{x\}$.

$$\begin{aligned} \sum_{(\phi, \psi) \in R_x} \Phi_{V_i, \phi} \Phi_{V_j, \psi} &= \sum_{\substack{\phi \in \mathbb{Z}_k^{V'_i}, \psi \in \mathbb{Z}_k^{V'_j} \\ a \neq b}} \Phi_{V'_i, \phi} \left[\Pi_a^{(k)}(\sigma_i(x)) \Pi_b^{(k)}(\sigma_j(x)) \right] \Phi_{V'_j, \psi} \\ &= \sum_{a \neq b} \Pi_a^{(k)}(\sigma_i(x)) \Pi_b^{(k)}(\sigma_j(x)), \end{aligned}$$

where the last equality holds because $\sum_{\phi \in \mathbb{Z}_k^{V'_i}} \Phi_{V'_i, \phi}$ and $\sum_{\psi \in \mathbb{Z}_k^{V'_j}} \Phi_{V'_j, \psi}$ are both equal to 1.

Notice that

$$\begin{aligned} \sum_{a \neq b} \Pi_a^{(k)}(\sigma_i(x)) \Pi_b^{(k)}(\sigma_j(x)) &= \sum_a \Pi_a^{(k)}(\sigma_i(x)) (1 - \Pi_a^{(k)}(\sigma_j(x))) \\ &\lesssim \frac{1}{2} \sum_a \left| \Pi_a^{(k)}(\sigma_i(x)) - \Pi_a^{(k)}(\sigma_j(x)) \right|^2. \end{aligned}$$

Thus, we get

$$\begin{aligned} \sum_{(\phi, \psi) \in R_x} \Phi_{V_i, \phi} \Phi_{V_j, \psi} &\lesssim \frac{1}{2} \sum_a \left| \Pi_a^{(k)}(\sigma_i(x)) - \Pi_a^{(k)}(\sigma_j(x)) \right|^2 \\ &= \frac{1}{2} \sum_a \left| \frac{1}{k} \sum_{\ell=0}^{k-1} a^{-\ell} (\sigma_i(x)^\ell - \sigma_j(x)^\ell) \right|^2 \\ &\leq 2^{\lceil \log k \rceil - 1} \sum_a \frac{1}{k^2} \sum_{\ell=0}^{k-1} \left| \sigma_i(x)^\ell - \sigma_j(x)^\ell \right|^2 \\ &\leq 2^{\lceil \log k \rceil - 1} \frac{1}{k} \sum_{\ell=0}^{k-1} \left| \sigma_i(x)^\ell - \sigma_j(x)^\ell \right|^2, \end{aligned}$$

where the second last inequality is due to Lemma 2.3.1. Finally, notice that $\sum_{\ell=0}^{k-1} |\sigma_i^\ell(x) - \sigma_j^\ell(x)|^2$ is cyclically equivalent to $\sum_{a \in \mathbb{Z}_k} 2 \Pi_a^{(k)}(\sigma_i(x)) (\Pi_a^{(k)}(\sigma_i(x)) - \Pi_a^{(k)}(\sigma_j(x)))$. Then

$$\Pi_a^{(k)}(\sigma_i(x)) (\Pi_a^{(k)}(\sigma_i(x)) - \Pi_a^{(k)}(\sigma_j(x))) = \Pi_a^{(k)}(\sigma_i(x)) (1 - \Pi_a^{(k)}(\sigma_j(x)))$$

gives the result. \square

Next, we look at homomorphisms between the constraint-constraint and constraint-variable algebras.

Lemma 3.3.4. *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a k -ary CS and π be a symmetric probability distribution on $[m] \times [m]$. There is a $4L$ -homomorphism $\alpha : \mathcal{A}_{c-c}(S, \pi) \rightarrow \mathcal{A}_{c-v}(S, \pi')$ where $\pi'(i) = \sum_j \pi(i, j)$ and $L = \max_i |V_i|$.*

Proof. Let α be the inclusion of $\mathcal{A}_{c-c}(S)$ in $\mathcal{A}_{c-v}(S)$, then for all i, j

$$\begin{aligned}
\sum_{\substack{\phi \in C_i, \psi \in C_j \\ \phi|_{V_i \cap V_j} = \psi|_{V_i \cap V_j}}} |\Phi_{V_i, \phi} \Phi_{V_j, \psi}|^2 &\lesssim \sum_{\substack{\phi \in C_i, \psi \in C_j \\ \phi|_{V_i \cap V_j} = \psi|_{V_i \cap V_j}}} \Phi_{V_i, \phi} \Phi_{V_j, \psi} \lesssim \sum_{x \in V_i \cap V_j} \sum_{\substack{\phi \in C_i, \psi \in C_j \\ \phi(x) \neq \psi(x)}} \Phi_{V_i, \phi} \Phi_{V_j, \psi} \\
&= \sum_{x \in V_i \cap V_j} \sum_{\phi \in C_i, \psi \in C_j} \Phi_{V_i, \phi} \sum_{a \neq b} \Pi_a(\sigma_i(x)) \Pi_b(\sigma_j(x)) \Phi_{V_j, \psi} \\
&= \sum_{x \in V_i \cap V_j} \sum_{a \neq b} \Pi_a(\sigma_i(x)) \Pi_b(\sigma_j(x)).
\end{aligned}$$

Now, noting that

$$\sum_{a \neq b} \Pi_a(\sigma_i(x)) \Pi_b(\sigma_j(s)) = \sum_a \Pi_a(\sigma_i(x)) (1 - \Pi_a(\sigma_j(x))) \lesssim \frac{1}{2} \sum_a |\Pi_a(\sigma_i(x)) - \Pi_a(\sigma_j(x))|^2,$$

we get that

$$\begin{aligned}
\sum_{\substack{i, j \\ \phi \in C_i, \psi \in C_j \\ \phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}}} \pi(i, j) |\Phi_{V_i, \phi} \Phi_{V_j, \psi}|^2 &\lesssim \frac{1}{2} \sum_{\substack{i, j \\ x \in V_i \cap V_j}} \pi(i, j) \sum_a |\Pi_a(\sigma_i(x)) - \Pi_a(\sigma_j(x))|^2 \\
&\leq \sum_{\substack{i, j \\ x \in V_i \cap V_j}} \pi(i, j) \sum_a \left(|\Pi_a(\sigma_i(x)) - \Pi_a(\sigma'(x))|^2 + |\Pi_a(\sigma_j(x)) - \Pi_a(\sigma'(x))|^2 \right).
\end{aligned}$$

Next, using the symmetry of π and the fact that $L \geq |V_i|$ for all i ,

$$\begin{aligned}
\sum_{\substack{i, j \\ \phi \in C_i, \psi \in C_j \\ \phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}}} \pi(i, j) |\Phi_{V_i, \phi} \Phi_{V_j, \psi}|^2 &\lesssim 2 \sum_{i, x \in V_i} \pi'(i) \sum_a |\Pi_a(\sigma_i(x)) - \Pi_a(\sigma'(x))|^2 \\
&\leq 2L \sum_{i, x \in V_i} \frac{\pi'(i)}{|V_i|} \sum_a |\Pi_a(\sigma_i(x)) - \Pi_a(\sigma'(x))|^2 \\
&\lesssim 4L \sum_{i, x \in V_i} \frac{\pi'(i)}{|V_i|} \sum_a \Pi_a(\sigma_i(x)) (1 - \Pi_a(\sigma'(x))).
\end{aligned}$$

Finally, reintroducing the projectors $\Phi_{V_i, \phi}$,

$$\begin{aligned}
\sum_{\substack{i,j \\ \phi \in C_i, \psi \in C_j \\ \phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}}} \pi(i, j) |\Phi_{V_i, \phi} \Phi_{V_j, \psi}|^2 &\lesssim 4L \sum_i \frac{\pi'(i)}{|V_i|} \sum_{\substack{x \in V_i \\ \phi \in C_i}} \sum_a \Phi_{V_i, \phi} \Pi_a(\sigma_i(x)) (1 - \Pi_a(\sigma'(x))) \\
&= 4L \sum_i \frac{\pi'(i)}{|V_i|} \sum_{\substack{x \in V_i \\ \phi \in C_i}} \Phi_{V_i, \phi} (1 - \Pi_{\phi(x)}(\sigma'(x))) \\
&\lesssim 4L \sum_i \frac{\pi'(i)}{|V_i|} \sum_{\substack{x \in V_i \\ \phi \in C_i}} |\Phi_{V_i, \phi} (1 - \Pi_{\phi(x)}(\sigma'(x)))|^2. \quad \square
\end{aligned}$$

Lemma 3.3.5. *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a k -ary CS and π be a probability distribution on $[m] \times [m]$. There is a P -homomorphism $\beta : \mathcal{A}_{c-v}(S, \pi') \rightarrow \mathcal{A}_{c-c}(S, \pi)$ where $\pi'(i) = \sum_j \pi(i, j)$ and $P = \max_{i,j. V_i \cap V_j \neq \emptyset} \frac{\pi'(i)}{\pi(i, j)}$.*

Note that, for an arbitrary π , P can be arbitrarily large. This C -homomorphism will nevertheless be useful in the case of perfect completeness (defect 0).

Proof. For each $x \in X$, choose $i_x \in [m]$ such that $x \in V_{i_x}$. Let β be the $*$ -homomorphism defined by

$$\begin{aligned}
\beta(\sigma_i(x)) &= \sigma_i(x), \text{ and} \\
\beta(\sigma'(x)) &= \sigma_{i_x}(x)
\end{aligned}$$

for all $x \in X$. Then, for all $x \in X$ and $\phi \in C_i$

$$\begin{aligned}
\beta \left[|\Phi_{V_i, \phi} (1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \right] &= |\Phi_{V_i, \phi} (1 - \Pi_{\phi(x)}(\sigma_{i_x}(x)))|^2 \lesssim \Phi_{V_i, \phi} (1 - \Pi_{\phi(x)}(\sigma_{i_x}(x))) \\
&= \sum_{\substack{\psi \in \mathbb{Z}_k^{V_{i_x}} \\ \phi(x) \neq \psi(x)}} \Phi_{V_i, \phi} \Phi_{V_{i_x}, \psi} \lesssim \sum_{\substack{\psi \in \mathbb{Z}_k^{V_{i_x}} \\ \phi(x) \neq \psi(x)}} |\Phi_{V_i, \phi} \Phi_{V_{i_x}, \psi}|^2 \\
&\leq \sum_{\substack{\psi \in C_{i_x} \\ \phi|_{V_i \cap V_{i_x}} \neq \psi|_{V_i \cap V_{i_x}}}} |\Phi_{V_i, \phi} \Phi_{V_{i_x}, \psi}|^2.
\end{aligned}$$

Thus, we get

$$\begin{aligned}
\beta \left[\sum_i \frac{\pi'(i)}{|V_i|} \sum_{\substack{x \in V_i \\ \phi \in C_i}} |\Phi_{V_i, \phi}(1 - \pi_{\phi(x)}(\sigma_x(x)))|^2 \right] &\lesssim \sum_i \frac{\pi'(i)}{|V_i|} \sum_{x \in V_i} \sum_{\substack{\phi \in C_i, \psi \in C_{i_x} \\ \phi|_{V_i \cap V_{i_x}} \neq \psi|_{V_i \cap V_{i_x}}} |\Phi_{V_i, \phi} \Phi_{V_{i_x}, \psi}|^2 \\
&\leq \sum_{i, j} \pi'(i) \sum_{\substack{\phi \in C_i, \psi \in C_j \\ \phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}}} |\Phi_{V_i, \phi} \Phi_{V_j, \psi}|^2 \\
&\leq P \sum_{i, j} \pi(i, j) \sum_{\substack{\phi \in C_i, \psi \in C_j \\ \phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}}} |\Phi_{V_i, \phi} \Phi_{V_j, \psi}|^2. \quad \square
\end{aligned}$$

Now, we look at homomorphisms between the constraint-variable and assignment algebras. In order to preserve the oracularizability, we first consider the variant of the assignment algebra with commutation constraints added to the defect. In one direction, we find a C -homomorphism; in the other, we have a trace-dependent mapping.

Lemma 3.3.6. *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a k -ary CS and π be a probability distribution on $[m]$. There is a $20L^2$ -homomorphism $\alpha : \mathcal{A}_{a+comm}(S, \pi) \rightarrow \mathcal{A}_{c-v}(S, \pi)$, where $L = \max_i |V_i|$.*

Proof. Let α be the inclusion of $\mathcal{A}_a(S)$ in $\mathcal{A}_{c-v}(S)$. Fix i and write $V_i = \{x_1, \dots, x_n\}$ ordered according to the ordering of X . Also, for $j = 1, \dots, n$, write $x_{j,a} = \Pi_a(\sigma'(x_j))$ and $\bar{x}_{j,a} = \Pi_a(\sigma_i(x_j))$. Then, as $\sigma_i(\Phi_{V_i, \phi}) = 0$ for all $\phi \notin C_i$,

$$\begin{aligned}
\alpha \left[\sum_{\phi \notin C_i} |\Phi_{V_i, \phi}|^2 \right] &= \sum_{\phi \notin C_i} |x_{1, \phi(x_1)} \cdots x_{n, \phi(x_n)}|^2 = \sum_{\phi \notin C_i} |x_{1, \phi(x_1)} \cdots x_{n, \phi(x_n)} - \bar{x}_{1, \phi(x_1)} \cdots \bar{x}_{n, \phi(x_n)}|^2 \\
&\leq \sum_{a_1, \dots, a_n} |x_{1, a_1} \cdots x_{n, a_n} - \bar{x}_{1, a_1} \cdots \bar{x}_{n, a_n}|^2 \\
&= \sum_{a_1, \dots, a_n} \left| \sum_{l=1}^n x_{1, a_1} \cdots x_{l-1, a_{l-1}} (x_{l, a_l} - \bar{x}_{l, a_l}) \bar{x}_{l+1, a_{l+1}} \cdots \bar{x}_{n, a_n} \right|^2 \\
&\leq 2^{\lceil \log(n) \rceil} \sum_{a_1, \dots, a_n} \sum_{l=1}^n |x_{1, a_1} \cdots x_{l-1, a_{l-1}} (x_{l, a_l} - \bar{x}_{l, a_l}) \bar{x}_{l+1, a_{l+1}} \cdots \bar{x}_{n, a_n}|^2 \\
&\leq 2n \sum_a \sum_{j=1}^n |x_{j,a} - \bar{x}_{j,a}|^2 = 2|V_i| \sum_{x \in V_i} \sum_a |\Pi_a(\sigma'(x)) - \Pi_a(\sigma_i(x))|^2.
\end{aligned}$$

Also, we find that

$$\begin{aligned}
\alpha \left[\sum_{x,y \in V_i; a,b \in \mathbb{Z}_k} |[\Pi_a(x), \Pi_b(y)]|^2 \right] &= \sum_{j,k=1,\dots,n; a,b \in \mathbb{Z}_k} |[x_{j,a}, x_{k,b}]|^2 \\
&= \sum_{j,k=1,\dots,n; a,b \in \mathbb{Z}_k} |[x_{j,a}, x_{k,b}] - [\bar{x}_{j,a}, \bar{x}_{k,b}]|^2 \\
&= \sum_{j,k=1,\dots,n; a,b \in \mathbb{Z}_k} |[x_{j,a} - \bar{x}_{j,a}, x_{k,b}] - [\bar{x}_{j,a}, \bar{x}_{k,b} - x_{k,b}]|^2 \\
&\leq 2 \sum_{j,k=1,\dots,n; a,b \in \mathbb{Z}_k} |[x_{j,a} - \bar{x}_{j,a}, x_{k,b}]|^2 + |[\bar{x}_{j,a}, \bar{x}_{k,b} - x_{k,b}]|^2 \\
&\lesssim 4 \sum_{j,k=1,\dots,n; a,b \in \mathbb{Z}_k} |(x_{j,a} - \bar{x}_{j,a})x_{k,b}|^2 + |\bar{x}_{j,a}(\bar{x}_{k,b} - x_{k,b})|^2 \\
&\leq 8n \sum_{j=1,\dots,n; a \in \mathbb{Z}_k} |x_{j,a} - \bar{x}_{j,a}|^2 \\
&= 8|V_i| \sum_{x \in V_i} \sum_a |\Pi_a(\sigma'(x)) - \Pi_a(\sigma_i(x))|^2.
\end{aligned}$$

By the proof of Lemma 3.3.4, we know

$$\sum_{x \in V_i} \sum_a |\Pi_a(\sigma(x)) - \Pi_a(\sigma_i(x))|^2 \lesssim 2 \sum_{\substack{x \in V_i \\ \phi \in C_i}} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2.$$

Thus, we get that

$$\begin{aligned}
\alpha \left(\sum_{r \in \mathcal{A}_a(S)} (\mu_{a,\pi}(r) + \mu_{comm,\pi}(r)) r^* r \right) &\lesssim \sum_i \pi(i) 20|V_i| \sum_{\substack{x \in V_i \\ \phi \in C_i}} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \\
&\leq 20L^2 \sum_i \frac{\pi(i)}{|V_i|} \sum_{\substack{x \in V_i \\ \phi \in C_i}} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2. \quad \square
\end{aligned}$$

Lemma 3.3.7. *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ and let π be a probability distribution on $[m]$. Then, for every trace τ on $\mathcal{A}_{a+comm}(S, \pi)$, there exists a trace τ' on $\mathcal{A}_{c-v}(S, \pi)$ such that $\text{def}(\tau') \leq \text{poly}(k^L) \text{def}(\tau)$, where $L = \max_i |V_i|$.*

Proof. Let $\tau = \rho \circ \varphi$ be the GNS representation, where $\varphi : \mathcal{A}_a(S) \rightarrow \mathcal{M} \subseteq \mathcal{B}(\mathcal{H})$ is a $*$ -homomorphism, and $\rho : \mathcal{M} \rightarrow \mathbb{C}$ is a tracial state. Let

$$\varepsilon_i = \sum_{x,y \in V_i} \sum_{a,b \in \mathbb{Z}_k} \|[\Pi_a(\sigma'(x)), \Pi_b(\sigma'(y))]\|_{\tau}^2.$$

We have $\text{def}(\tau) = \text{def}(\tau; \mu_{a,\pi}) + \sum_i \pi(i) \varepsilon_i$. Let $u_{x,a} = \varphi(1 - 2\Pi_a(x))$. Fix some $i \in [m]$. We have in particular $\| [u_{x,a}, u_{y,b}] \|_\rho^2 \leq \varepsilon_i$ for all $x, y \in V_i$ and $a, b \in \mathbb{Z}_k$. So, using Lemma 3.4.1, there exist commuting order-2 unitaries $v_{x,a} \in \mathcal{M}$ such that $\|v_{x,a} - u_{x,a}\|_\rho^2 \leq \text{poly}(k|V_i|)\varepsilon_i$. Next, note that

$$\begin{aligned} \left\| \sum_{a \in \mathbb{Z}_k} \Pi_1(v_{x,a}) - 1 \right\|_\rho^2 &= \frac{1}{4} \left\| \sum_{a \in \mathbb{Z}_k} (u_{x,a} - v_{x,a}) \right\|_\rho^2 \leq \frac{1}{4} 2^{\lceil \log_2 k \rceil} \sum_{a \in \mathbb{Z}_k} \|u_{x,a} - v_{x,a}\|_\rho^2 \\ &\leq k^2 \text{poly}(k|V_i|)\varepsilon_i = \text{poly}(|V_i|, k)\varepsilon_i. \end{aligned}$$

Let p be the projector onto the kernel of $\sum_{a \in \mathbb{Z}_k} \Pi_1(v_{x,a}) - 1$. Then, as the eigenvalues of $\sum_{a \in \mathbb{Z}_k} \Pi_1(v_{x,a}) - 1$ are integers, $\|1 - p\|_\rho^2 \leq \text{poly}(|V_i|, k)\varepsilon_i$. Take $w_{i,x} = \sum_{a \in \mathbb{Z}_k} \omega_k^a \Pi_1(v_{x,a})p + (1 - p)$. As the supports of the $\Pi_1(v_{x,a})p = \Pi_1(v_{x,a}) \wedge p$ are disjoint, $w_{i,x}$ is an order- k unitary. Further,

$$\begin{aligned} \|\Pi_a(w_{i,x}) - \Pi_a(\varphi(x))\|_\rho^2 &\leq 4(\|\Pi_1(v_{x,a})(p - 1)\|_\rho^2 + \|1 - p\|_\rho^2 + \|\Pi_1(v_{x,a}) - \Pi_1(u_{x,a})\|_\rho^2) \\ &\leq 4(2\|1 - p\|_\rho^2 + \frac{1}{4}\|v_{x,a} - u_{x,a}\|_\rho^2) \leq \text{poly}(|V_i|, k)\varepsilon_i. \end{aligned}$$

Now, take the PVM $\{P_{V_i, \phi}\}_{\phi \in \mathbb{Z}_k^{V_i}}$ as $P_{V_i, \phi} = \prod_{x \in V_i} \Pi_{\phi(x)}(w_{i,x})$. Fix some $\phi_i \in C_i$ and take the PVM $\{\tilde{P}_{V_i, \phi}\}_{\phi \in C_i}$ as $\tilde{P}_{V_i, \phi} = P_{V_i, \phi} + \delta_{\phi, \phi_i} \sum_{\psi \notin C_i} P_{V_i, \psi}$. Now, we can take the *-representation $\chi : \mathcal{A}_{c-v}(S, \pi) \rightarrow \mathcal{M}$ defined by $\chi(\sigma'(x)) = \varphi(x)$ and $\chi(\Phi_{V_i, \phi}) = \tilde{P}_{V_i, \phi}$; and take the tracial state $\tau' = \rho \circ \chi$. It remains to calculate the defect of τ' . First, note that

$$\begin{aligned} \text{def}(\tau') &= \sum_{i=1}^m \frac{\pi(i)}{|V_i|} \sum_{x \in V_i, \phi \in C_i} \tau'(\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))) \\ &= \sum_{i=1}^m \frac{\pi(i)}{|V_i|} \sum_{x \in V_i, \phi \in C_i} \rho(\tilde{P}_{V_i, \phi}(1 - \Pi_{\phi(x)}(\varphi(x)))) \\ &= \sum_{i=1}^m \frac{\pi(i)}{|V_i|} \left(\sum_{x \in V_i, \phi \in C_i} \rho \left(P_{V_i, \phi}(1 - \Pi_{\phi(x)}(\varphi(x))) + \sum_{x \in V_i, \phi \notin C_i} P_{V_i, \phi}(1 - \Pi_{\phi_i(x)}(\varphi(x))) \right) \right) \\ &= \sum_{i=1}^m \frac{\pi(i)}{|V_i|} \left(\sum_{x \in V_i, \phi \in \mathbb{Z}_k^{V_i}} \rho(P_{V_i, \phi}(1 - \Pi_{\phi(x)}(\varphi(x)))) \right. \\ &\quad \left. + \sum_{x \in V_i, \phi \notin C_i} \rho(P_{V_i, \phi}(\Pi_{\phi(x)}(\varphi(x)) - \Pi_{\phi_i(x)}(\varphi(x)))) \right) \\ &\leq \sum_{i=1}^m \frac{\pi(i)}{|V_i|} \left(\frac{1}{4} \sum_{x \in V_i, a \in \mathbb{Z}_k} \|\Pi_a(w_{i,x}) - \Pi_a(\varphi(x))\|_\rho^2 + 2|V_i| \sum_{\phi \notin C_i} \|P_{V_i, \phi}\|_\rho^2 \right). \end{aligned}$$

For fixed i and $\phi \in \mathbb{Z}_k^{V_i}$, write $V_i = \{x_1, \dots, x_{|V_i|}\}$, $p_j = \Pi_{\phi(x_j)}(x_j)$, and $\bar{p}_j = \Pi_{\phi(x_j)}(w_{i,x_j})$. So we have that

$$\begin{aligned}
\|P_{V_i, \phi}\|_\rho^2 &= \|\bar{p}_1 \cdots \bar{p}_{|V_i|}\|_\rho^2 \\
&= \left\| p_1 \cdots p_{|V_i|} + \sum_{j=1}^{|V_i|} p_1 \cdots p_{j-1} (\bar{p}_j - p_j) \bar{p}_{j+1} \cdots \bar{p}_{|V_i|} \right\|_\rho^2 \\
&\leq 2^{\lceil \log_2(|V_i|+1) \rceil} \left(\|p_1 \cdots p_{|V_i|}\|_\rho^2 + \sum_{j=1}^{|V_i|} \|\bar{p}_j - p_j\|_\rho^2 \right) \\
&\leq 2(|V_i| + 1) \left(\|\Phi_{V_i, \phi}\|_\tau^2 + \sum_{x \in V_i} \|\Pi_{\phi(x)}(w_{i,x}) - \Pi_{\phi(x)}(\varphi(x))\|_\rho^2 \right).
\end{aligned}$$

As such, we get that the defect

$$\begin{aligned}
\text{def}(\tau') &\leq \sum_{i=1}^m \frac{\pi(i)}{|V_i|} \left(\frac{1}{4} \sum_{x \in V_i, a \in \mathbb{Z}_k} \|\Pi_a(w_{i,x}) - \Pi_a(\varphi(x))\|_\rho^2 + 4|V_i|(|V_i| + 1) \sum_{\phi \notin C_i} \|\Phi_{V_i, \phi}\|_\tau^2 \right. \\
&\quad \left. + 4|V_i|(|V_i| + 1) \sum_{\phi \notin C_i} \sum_{x \in V_i} \|\Pi_{\phi(x)}(w_{i,x}) - \Pi_{\phi(x)}(\varphi(x))\|_\rho^2 \right) \\
&\leq 4(L + 1) \text{def}(\tau; \mu_{a, \pi}) \\
&\quad + \sum_{i=1}^m \pi(i) \left(\frac{1}{4|V_i|} + 4(|V_i| + 1)k^{|V_i|-1} \right) \sum_{x \in V_i, a \in \mathbb{Z}_k} \|\Pi_a(w_{i,x}) - \Pi_a(\varphi(x))\|_\rho^2 \\
&\leq 4(L + 1) \text{def}(\tau; \mu_{a, \pi}) + \sum_{i=1}^m \pi(i) \left(\frac{1}{4|V_i|} + 4(|V_i| + 1)k^{|V_i|-1} \right) k|V_i| \text{poly}(|V_i|, k) \varepsilon_i \\
&\leq \text{poly}(k^L) \text{def}(\tau). \quad \square
\end{aligned}$$

Finally, we consider the relationship between the assignment algebras with and without commutation. Here, we are only able to find a C -homomorphism in one direction, as a C -homomorphism in the other direction would imply oracularizability of the underlying assignment algebra, which does not hold for general CSs.

Lemma 3.3.8. *Let $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a k -ary CS and π be a probability distribution on $[m]$. There is a 1-homomorphism $\alpha : \mathcal{A}_a(S, \pi) \rightarrow \mathcal{A}_{a+comm}(S, \pi)$.*

Proof. Let α be the identity map. Noting that $\mu_{a, \pi} + \mu_{comm, \pi} \geq \mu_{a, \pi}$ gives the result. \square

3.4 Subdivision and stability

Suppose we have a BCS where each constraint is made up of subconstraints on subsets of the variables (for instance, a 3SAT instance made of 3SAT clauses). In this section, we look at what happens when we split up the contexts and constraints so that each subconstraint is in its own context. In the weighted BCS algebra, splitting up a context changes the commutative subalgebra corresponding to the context to a non-commutative subalgebra. To deal with this, we use a tool from the approximate representation theory of groups, namely the stability of \mathbb{Z}_2^k .

Lemma 3.4.1 ([13]). *Let (\mathcal{M}, τ) be a tracial von Neumann algebra, and suppose $f : [k] \rightarrow \mathcal{M}$ is a function such that $f(i)^2 = 1$ for all $i \in [k]$ and $\|[f(i), f(j)]\|_\tau^2 \leq \varepsilon$ for all $i, j \in [k]$, where $k \geq 1$ and $\varepsilon \geq 0$. Then there is a homomorphism $\psi : \mathbb{Z}_2^k \rightarrow \mathcal{U}(\mathcal{M})$ such that $\|\psi(x_i) - f(i)\|_\tau^2 \leq \text{poly}(k)\varepsilon$ for all $i \in [k]$, where the x_i generate \mathbb{Z}_2^k .*

A function f satisfying the conditions of Lemma 3.4.1 is called an ε -**homomorphism from \mathbb{Z}_2^k to $\mathcal{U}(\mathcal{M})$** .

We now formally define the subdivision of a BCS.

Definition 3.4.2. *Let $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a BCS. Suppose that for all $1 \leq i \leq m$ there exists a constant $m_i \geq 1$ and a set of constraints $\{D_{ij}\}_{j=1}^{m_i}$ on variables $\{V_{ij}\}_{j=1}^{m_i}$ respectively, such that*

1. $V_{ij} \subseteq V_i$ for all $i \in [m]$ and $j \in [m_i]$,
2. for every $i \in [m]$ and $x, y \in V_i$, there is a $j \in [m_i]$ such that $x, y \in V_{ij}$, and
3. $C_i = \bigwedge_{j=1}^{m_i} D_{ij}$ for all $i \in [m]$, where \wedge is conjunction.

The BCS $B' = (X, \{(V_{ij}, D_{ij})\}_{i,j})$ is called a **subdivision** of B . When working with subdivisions, we refer to D_{ij} as the **clauses** of constraint C_i , and we refer to m_i as the **number of clauses** in constraint i . A subdivision is **uniform** if $m_i = m_j$ for all i, j .

Given a subdivision B' of B with $M = \sum_{i=1}^m m_i$, we pick a bijection between $[M]$ and the set of pairs (i, j) with $1 \leq i \leq m$ and $1 \leq j \leq m_i$. If π is a probability distribution on $[m] \times [m]$, then let π_{sub} be the probability distribution on $[M] \times [M]$ with $\pi_{sub}(ij, kl) = \pi(i, k)/m_i m_k$. Note that if π is uniform and the subdivision is uniform, then π_{sub} is uniform as well.

One of the first things we notice about subdivision is that strategies for $\mathbf{G}(B, \pi)$ can be lifted to strategies for the subdivided game.

Proposition 3.4.3. *Let $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a BCS, and let $B' = (X, \{V_{ij}, D_{ij}\}_{i,j})$ be a subdivision. Let π be a probability distribution on $[m] \times [m]$, and let π_{sub} be the probability distribution defined from π as above. The homomorphism $\alpha : \mathcal{A}(B') \rightarrow \mathcal{A}(B)$ defined by $\sigma_{ij}(x) \mapsto \sigma_i(x)$ is a 1-homomorphism $\mathcal{A}_{inter}(B', \pi_{sub}) \rightarrow \mathcal{A}_{inter}(B, \pi)$, and also induces an isomorphism $\text{SynAlg}(B', \pi_{sub}) \cong \text{SynAlg}(B, \pi)$.*

Here $\sigma_{ij}(x)$ denotes the copy of x in $\mathcal{A}(W_{ij}, D_{ij}) \subseteq \mathcal{A}(B')$.

Proof. Recall that $|a|^2 = a^*a$ denotes the hermitian square of a . By definition, $\alpha(\sigma_{ij}(x) - \sigma_{kl}(x)) = \sigma_i(x) - \sigma_k(x)$. Hence

$$\begin{aligned} \alpha \left(\sum_{\substack{ij \neq kl \\ x \in V_{ij} \cap V_{kl}}} \pi_{sub}(ij, kl) |\sigma_{ij}(x) - \sigma_{kl}(x)|^2 \right) &= \sum_{\substack{ij \neq kl \\ x \in V_{ij} \cap V_{kl}}} \frac{\pi(i, k)}{m_i m_k} |\sigma_i(x) - \sigma_k(x)|^2 \\ &\leq \sum_{\substack{i \neq k \\ x \in V_i \cap V_k}} \pi(i, k) |\sigma_i(x) - \sigma_k(x)|^2, \end{aligned}$$

since each variable $x \in V_i$ appears in at most m_i subclauses V_{ij} . Hence $\alpha : \mathcal{A}_{inter}(B', \pi_{sub}) \rightarrow \mathcal{A}_{inter}(B, \pi)$ is a 1-homomorphism.

To show that the synchronous algebras are isomorphic, observe that since every pair of elements $x, y \in V_i$ belongs to some V_{ij} , there is an isomorphism

$$\text{SynAlg}(B', \pi_{sub}) \cong *_{i=1}^m \mathbb{Z}_2^{V_i} / \langle R \rangle,$$

where R is the set of relations $\sigma_i(\Phi_{V_{ij}, \phi}) \sigma_i(\Phi_{V_{kl}, \psi}) = 0$ for all ϕ and ψ which do not agree on $V_{ij} \cap V_{kl}$, and $\sigma_i(\Phi_{V_{ij}, \phi}) = 0$ for all $\phi \notin D_{ij}$. From these latter relations, it is possible to recover the relations $\Phi_{V_i, \phi} = 0$ for $\phi \notin C_i$, and then to recover all the relations of $\text{SynAlg}(B, \pi)$. \square

Proposition 3.4.3 implies that $\mathbf{G}(B, \pi)$ has a perfect quantum (resp. commuting operator) strategy if and only if $\mathbf{G}(B', \pi_{sub})$ has a perfect quantum (resp. commuting operator) strategy. The main result of this section is that near perfect strategies for $\mathbf{G}(B', \pi_{sub})$ can be pulled back to near perfect strategies for $\mathbf{G}(B, \pi)$.

Theorem 3.4.4. *Let $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a BCS and let $B' = (X, \{(V_{ij}, D_{ij})\}_{i,j})$ be a subdivision of B with m_i clauses in constraint C_i . Let π be a probability distribution on $[m] \times [m]$ that is C -diagonally dominant for some $C > 0$, and let π_{sub} be the probability distribution defined from π as above. If there is a trace τ on $\mathcal{A}_{c-c}(B', \pi_{sub})$, then there is a trace $\tilde{\tau}$ on $\mathcal{A}_{c-c}(B, \pi)$ with $\text{def}(\tilde{\tau}) \leq \text{poly}(2^L, M, K) \text{def}(\tau)$, where $L = \max_{i,j} |V_{ij}|$, $K = \max_i |V_i|$, and $M = \max m_i$.*

To prove the theorem, we consider other versions of the weighted BCS algebra, where $\mathcal{A}(V_i, C_i)$ is replaced by $\mathbb{C}\mathbb{Z}_2^{*V_i}$, and the defining relations of $\mathcal{A}(V_i, C_i)$ are moved into the weight function.

Definition 3.4.5. Let $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a BCS with a probability distribution π on $[m] \times [m]$, and let $B' = (X, \{(V_{ij}, D_{ij})\}_{i,j})$ be a subdivision, with m_i clauses in constraint C_i and probability distribution π_{sub} induced by π . Let $\sigma_i : \mathbb{C}\mathbb{Z}_2^{*V_i} \rightarrow *_{i=1}^m \mathbb{C}\mathbb{Z}_2^{*V_i}$ denote the inclusion of the i th factor. Let $\mathcal{A}_{free}(B) := *_{i=1}^m \mathbb{C}\mathbb{Z}_2^{*V_i}$, and define weight functions μ_{inter} , μ_{sat} , μ_{clause} , and μ_{comm} on $\mathcal{A}_{free}(B)$ by

$$\begin{aligned} \mu_{inter}(\sigma_i(x) - \sigma_j(x)) &= \pi(i, j) \text{ for all } i \neq j \in [m] \text{ and } x \in V_i \cap V_j, \\ \mu_{sat}(\Phi_{V_i, \phi}) &= \pi(i, i) \text{ for all } i \in [m] \text{ and } \phi \in \mathbb{Z}_2^{V_i} \setminus C_i, \\ \mu_{clause}(\Phi_{V_{ij}, \phi}) &= \pi(i, i) / m_i^2 \text{ for all } (i, j) \in [m] \times [m_i] \text{ and } \phi \in \mathbb{Z}_2^{V_{ij}} \setminus D_{ij}, \text{ and} \\ \mu_{comm}([\sigma_i(x), \sigma_i(y)]) &= \pi(i, i) \text{ for all } i \in [m] \text{ and } x, y \in V_i, \end{aligned}$$

and $\mu_{inter}(r) = 0$, $\mu_{sat}(r) = 0$, $\mu_{clause}(r) = 0$, and $\mu_{comm}(r) = 0$ for any elements r other than those listed. Let $\mathcal{A}_{free}(B, B', \pi)$ be the weighted algebra $(\mathcal{A}_{free}(B), \mu_{all})$, where $\mu_{all} := \mu_{inter} + \mu_{clause} + \mu_{comm}$.

Note that μ_{inter} is the same as the weight function of the algebra $\mathcal{A}_{inter}(B, \pi)$ defined in Definition 3.2.4, except that it's defined on $\mathcal{A}_{free}(B)$ rather than $\mathcal{A}_{c-c}(B)$. The weight function μ_{sat} comes from the defining relations for $\mathcal{A}_{c-c}(B)$, while μ_{clause} comes from the defining relations for $\mathcal{A}_{c-c}(B')$, so $\mathcal{A}_{free}(B, B', \pi)$ is a mix of relations from $\mathcal{A}_{inter}(B, \pi)$ and $\mathcal{A}_{inter}(B', \pi)$. As mentioned previously, the context V_i has an order inherited from X , and this is used for the order of the product when talking about $\Phi_{V_i, \phi}$ and $\Phi_{V_{ij}, \phi}$ in $\mathcal{A}_{free}(B)$. In particular, the order on V_{ij} is compatible with the order on V_i .

The weight functions μ_{inter} , μ_{sat} and μ_{clause} can also be defined on $*_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}$ using the same formulae as in Definition 3.4.5, and we use the same notation for both versions. The following lemma shows that we can relax $\mathcal{A}_{inter}(B, \pi)$ to $(*_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}, \mu_{inter} + \mu_{clause})$, as long as π is C -diagonally dominant.

Lemma 3.4.6. Let $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a BCS, and let π be a probability distribution on $[m] \times [m]$ which is C -diagonally dominant for some $C > 0$. Let $B' = (X, \{(V_{ij}, D_{ij})\}_{i,j})$ be a subdivision of B . Let μ_{inter} , μ_{sat} and μ_{clause} be the weight functions defined above with respect to π . Then there is an $O(K)$ -homomorphism $\mathcal{A}_{inter}(B, \pi) \rightarrow (*_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}, \mu_{inter} + \mu_{sat})$, where $K = \max_i |V_i|$. Furthermore, there is an M^2 -homomorphism $(*_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}, \mu_{inter} + \mu_{sat}) \rightarrow (*_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}, \mu_{inter} + \mu_{clause})$, where $M = \max_i m_i$ is the maximum number of clauses m_i in constraint i .

Proof. Since C_i is non-empty by convention, we can choose $\psi_i \in C_i$ for every $1 \leq i \leq m$. Define the homomorphism $\alpha : \mathcal{A}_{inter}(B, \pi) \rightarrow (*_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}, \mu_{inter} + \mu_{sat})$ by

$$\alpha(\sigma_i(x)) = \sum_{\varphi \in C_i} \Phi_{V_i, \varphi} \sigma_i(x) + \sum_{\varphi \in \mathbb{Z}_2^{V_i} \setminus C_i} \Phi_{V_i, \varphi} (-1)^{\psi_i(x)}.$$

Let $\Phi_i = \sum_{\varphi \in C_i} \Phi_{V_i, \varphi}$, and recall that $|a|^2 = a^*a$ denotes the hermitian square of a . Then

$$\begin{aligned} \alpha(|\sigma_i(x) - \sigma_j(x)|^2) &= |\Phi_i \sigma_i(x) + (1 - \Phi_i)(-1)^{\psi_i(x)} - \Phi_j \sigma_j(x) - (1 - \Phi_j)(-1)^{\psi_j(x)}|^2 \\ &\leq 4 |\Phi_i \sigma_i(x) + (1 - \Phi_i)(-1)^{\psi_i(x)} - \sigma_i(x)|^2 \\ &\quad + 4 |\Phi_j \sigma_j(x) + (1 - \Phi_j)(-1)^{\psi_j(x)} - \sigma_j(x)|^2 + 4 |\sigma_i(x) - \sigma_j(x)|^2. \end{aligned}$$

Observe that $\sigma_i(x) = \sum_{\varphi \in \mathbb{Z}_2^{V_i}} \Phi_{V_i, \varphi} (-1)^{\varphi(x)}$, so

$$|\Phi_i \sigma_i(x) + (1 - \Phi_i)(-1)^{\psi_i(x)} - \sigma_i(x)|^2 = \sum_{\varphi \in \mathbb{Z}_2^{V_i} \setminus C_i} \Phi_{V_i, \varphi} ((-1)^{\psi_i(x)} - (-1)^{\varphi(x)})^2 \leq 4 \sum_{\varphi \in \mathbb{Z}_2^{V_i} \setminus C_i} \Phi_{V_i, \varphi}.$$

Thus

$$\begin{aligned} \alpha\left(\sum_{\substack{1 \leq i \neq j \leq m \\ x \in V_i \cap V_j}} \pi(i, j) |\sigma_i(x) - \sigma_j(x)|^2\right) &\leq \sum_{\substack{1 \leq i \neq j \leq m \\ x \in V_i \cap V_j}} \pi(i, j) \left(16 \sum_{\varphi \in \mathbb{Z}_2^{V_i} \setminus C_i} \Phi_{V_i, \varphi} + 16 \sum_{\varphi \in \mathbb{Z}_2^{V_j} \setminus C_j} \Phi_{V_j, \varphi} + 4 |\sigma_i(x) - \sigma_j(x)|^2\right) \\ &\leq \sum_{a \in *_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}} 4\mu_{inter}(a)a^*a + \sum_{a \in *_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}} 32\frac{K}{C}\mu_{sat}(a)a^*a \\ &\leq O(K) \sum_{a \in *_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}} (\mu_{inter}(a) + \mu_{sat}(a))a^*a, \end{aligned}$$

since π is C -diagonally dominant.

Next, suppose B' is a subdivision of B . If $\phi \in \mathbb{Z}_2^{V_i} \setminus C_i$, then we can choose $j_\phi \in [m_i]$ such that $\phi|_{V_{ij_\phi}} \notin D_{ij_\phi}$. Since $\sum_{\phi: \phi|_{V_{ij}} = \phi'} \Phi_{V_i, \phi} = \Phi_{V_{ij}, \phi'}$,

$$\sum_{\phi \notin C_i} \Phi_{V_i, \phi} = \sum_{1 \leq j \leq m_i} \sum_{\phi: \phi|_{V_j} = j} \Phi_{V_i, \phi} \leq \sum_{1 \leq j \leq m_i} \sum_{\phi: \phi|_{V_{ij}} \notin D_{ij}} \Phi_{V_i, \phi} = \sum_{1 \leq j \leq m_i} \sum_{\phi' \notin D_{ij}} \Phi_{V_{ij}, \phi'}.$$

Hence

$$\sum_r \mu_{sat}(r)r^*r \leq M^2 \sum_r \mu_{clause}(r)r^*r,$$

where the M^2 comes from the fact that we divide by m_i^2 in the definition of μ_{clause} . Thus the identity map $(\ast_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}, \mu_{inter} + \mu_{sat}) \rightarrow (\ast_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}, \mu_{inter} + \mu_{clause})$ is an M^2 -homomorphism. \square

The following proposition shows how to construct tracial states on $\mathcal{A}_{inter}(B, \pi)$ from tracial states on $\mathcal{A}_{free}(B, B', \pi)$.

Proposition 3.4.7. *Let $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a BCS, and let π be a probability distribution on $[m] \times [m]$ which is C -diagonally dominant for some $C > 0$. Let $B' = (X, \{(V_{ij}, D_{ij})\}_{i,j})$ be a subdivision of B with m_i clauses in constraint C_i . If τ is a trace on $\mathcal{A}_{free}(B, B', \pi)$, then there is a trace $\tilde{\tau}$ on $\mathcal{A}_{inter}(B, \pi)$ such that $\text{def}(\tilde{\tau}) \leq \text{poly}(2^L, M, K) \text{def}(\tau)$, where $L = \max_{i,j} |V_{ij}|$, $K = \max_i |V_i|$, and $M = \max_i m_i$. Furthermore, if τ is finite-dimensional then so is $\tilde{\tau}$.*

Proof. Since π is C -diagonally dominant, if $\pi(i, i) = 0$ then $\pi(i, j) = \pi(j, i) = 0$ for all $j \in [m]$, and the variables in V_i do not appear in $\text{supp}(\mu_{inter})$. Thus we may assume without loss of generality that $\pi(i, i) > 0$ for all $i \in [m]$. Let τ be a trace on $\mathcal{A}_{free}(B, B', \pi)$. By the GNS construction there is a \ast -representation ρ of $\mathcal{A}_{free}(B, B', \pi)$ acting on a Hilbert space \mathcal{H}_0 with a unit cyclic vector ψ such that $\tau(a) = \langle \psi | \rho(a) | \psi \rangle$ for all $a \in \mathcal{A}_{free}(B)$. Let $\mathcal{M}_0 = \overline{\rho(\mathcal{A}_{free}(B))}$ be the weak operator closure of the image of ρ , and let τ_0 be the faithful normal tracial state on \mathcal{M}_0 corresponding to $|\psi\rangle$ (so $\tau_0 \circ \rho = \tau$).

Let $\sum_{a \in \mathbb{Z}_2^{V_i}} \mu_{comm}(a) \|a\|_\tau^2 = \varepsilon_i$ for all $i \in [m]$. The restriction of ρ to $\mathbb{Z}_2^{*V_i}$ is an ε_i -homomorphism from $\mathbb{Z}_2^{V_i}$ into (\mathcal{M}_0, τ_0) , so by Lemma 3.4.1 there is a representation $\rho_i : \mathbb{Z}_2^{V_i} \rightarrow \mathcal{U}(\mathcal{M}_0)$ such that

$$\|\rho_i(x_j) - \rho(x_j)\|_{\tau_0}^2 \leq \text{poly}(K) \varepsilon_i \quad (3.4.1)$$

for all generators $x_j \in \mathbb{Z}_2^{V_i}$. Let $\tilde{\rho} : \ast_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i} \rightarrow \mathcal{M}_0$ be the homomorphism defined by $\tilde{\rho}(x) = \rho_i(x)$ for $x \in \mathbb{Z}_2^{V_i}$. Suppose $x \in V_i \cap V_j$, then

$$\begin{aligned} \|\tilde{\rho}(\sigma_i(x) - \sigma_j(x))\|_{\tau_0}^2 &\leq 4\|\tilde{\rho}(\sigma_i(x)) - \rho(\sigma_i(x))\|_{\tau_0}^2 + 4\|\tilde{\rho}(\sigma_j(x)) - \rho(\sigma_j(x))\|_{\tau_0}^2 \\ &\quad + 4\|\rho(\sigma_i(x) - \sigma_j(x))\|_{\tau_0}^2 \\ &\leq \text{poly}(K)(\varepsilon_i + \varepsilon_j) + 4\|\sigma_i(x) - \sigma_j(x)\|_\tau^2. \end{aligned}$$

We conclude that

$$\begin{aligned}
\text{def}(\tau_0 \circ \tilde{\rho}; \mu_{inter}) &\leq \sum_{i \neq j} \sum_{x \in V_i \cap V_j} \pi(i, j) (\text{poly}(K)(\varepsilon_i + \varepsilon_j) + 4\|\sigma_i(x) - \sigma_j(x)\|_\tau^2) \\
&\leq \sum_i \sum_{x \in V_i \cap V_j} \frac{\pi(i, i)}{C} (\text{poly}(K)2\varepsilon_i + 4\|\sigma_i(x) - \sigma_j(x)\|_\tau^2) \\
&\leq O(\text{poly}(K) \text{def}(\tau; \mu_{comm}) + \text{def}(\tau; \mu_{inter})).
\end{aligned}$$

For any $S \subseteq V_i$, let $x_S := \prod_{x \in S} x \in \mathbb{Z}_2^{*V_i}$, where the order of the product is inherited from the order on X . By Equation (3.4.1),

$$\|\tilde{\rho}(x_S) - \rho(x_S)\|_{\tau_0}^2 \leq \text{poly}(K)\varepsilon_i,$$

where the degree of the polynomial $\text{poly}(K)$ has increased by one. Since

$$\Phi_{V_{ij}, \phi} = \frac{1}{2^{|V_{ij}|}} \sum_{S \subseteq V_{ij}} (-1)^{\phi(x_S)} x_S,$$

we get that

$$\|\tilde{\rho}(\Phi_{V_{ij}, \phi}) - \rho(\Phi_{V_{ij}, \phi})\|_{\tau_0}^2 \leq \frac{1}{2^{|V_{ij}|}} \sum_{S \subseteq V_{ij}} \|\tilde{\rho}(x_S) - \rho(x_S)\|_{\tau_0}^2 \leq \text{poly}(K)\varepsilon_i.$$

If $1 \leq i \leq m$, $1 \leq j \leq m_i$, and $\phi \notin D_{ij}$, then

$$\|\tilde{\rho}(\Phi_{V_{ij}, \phi})\|_{\tau_0}^2 \leq 2\|\tilde{\rho}(\Phi_{V_{ij}, \phi}) - \rho(\Phi_{V_{ij}, \phi})\|_{\tau_0}^2 + 2\|\rho(\Phi_{V_{ij}, \phi})\|_{\tau_0}^2,$$

and hence

$$\begin{aligned}
\text{def}(\tau_0 \circ \tilde{\rho}; \mu_{clause}) &= \sum_{i, j} \frac{\pi(i, i)}{m_i^2} \sum_{\phi \notin D_{ij}} \|\tilde{\rho}(\Phi_{V_{ij}, \phi})\|_{\tau_0}^2 \\
&\leq \sum_{i, j} \sum_{\phi \notin D_{ij}} \frac{\pi(i, i)}{m_i^2} (\text{poly}(K)\varepsilon_i + 2\|\Phi_{V_{ij}, \phi}\|_\tau^2) \\
&\leq \sum_i 2^L \frac{\pi(i, i)}{m_i} \text{poly}(K)\varepsilon_i + 2 \text{def}(\tau; \mu_{clause}) \\
&\leq 2^L \text{def}(\tau; \mu_{comm}) + 2 \text{def}(\tau; \mu_{clause}).
\end{aligned}$$

We conclude that $\tilde{\tau} = \tau_0 \circ \tilde{\rho}$ is a tracial state on $*_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}$ with $\text{def}(\tilde{\tau}; \mu_{inter} + \mu_{clause})$ bounded by

$$O(\text{def}(\tau; \mu_{inter}) + \text{def}(\tau; \mu_{clause}) + 2^L \text{poly}(K) \text{def}(\tau; \mu_{comm})).$$

We conclude that

$$\text{def}(\tilde{\tau}; \mu_{inter} + \mu_{clause}) \leq \text{poly}(2^L, K) \text{def}(\tau; \mu_{inter} + \mu_{clause} + \mu_{comm}).$$

By Lemma 3.4.6, there is a $O(KM^2)$ -homomorphism $\mathcal{A}_{inter}(B, \pi) \rightarrow (*_{i=1}^m \mathbb{C}\mathbb{Z}_2^{V_i}, \mu_{inter} + \mu_{clause})$, and pulling $\tilde{\tau}$ back by this homomorphism gives the proposition. \square

Finally, we can pull back tracial states from the subdivision algebra $\mathcal{A}_{inter}(B', \pi_{sub})$ to traces on $\mathcal{A}_{free}(B, B', \pi)$.

Proposition 3.4.8. *Let $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a BCS, and let $B' = (X, \{(V_{ij}, D_{ij})\}_{i,j})$ be a subdivision of B . Let π be a probability distribution on $[m] \times [m]$, and let π_{sub} be the probability distribution defined from π as above. Then there is a $\text{poly}(M, 2^L)$ -homomorphism $\mathcal{A}_{free}(B, B', \pi) \rightarrow \mathcal{A}_{inter}(B', \pi_{sub})$, where $L = \max_{i,j} |V_{ij}|$ and $M = \max_i m_i$.*

Proof. For each $1 \leq i \leq m$ and $x \in V_i$, choose an index $1 \leq r_{ix} \leq m_i$ such that $x \in V_{ir_{ix}}$. Also, for each $x, y \in V_i$, choose an index i_{xy} such that $x, y \in V_{i_{xy}}$. Define $\alpha : *_{i=1}^m \mathbb{Z}_2^{V_i} \rightarrow \mathcal{A}(B')$ by $\alpha(\sigma_i(x)) = \sigma_{ir_{ix}}(x)$. It follows immediately from the definitions that α is a $O(M^2)$ -homomorphism $(\mathcal{A}_{free}(B), \mu_{inter}) \rightarrow \mathcal{A}_{inter}(B', \pi_{sub})$. Moving on to μ_{comm} , observe that

$$\begin{aligned} \alpha(|[\sigma_i(x), \sigma_i(y)]|^2) &= |\sigma_{ir_{ix}}(x)\sigma_{ir_{iy}}(y) - \sigma_{ir_{iy}}(y)\sigma_{ir_{ix}}(x)|^2 \\ &\leq 4|(\sigma_{ir_{ix}}(x) - \sigma_{i_{xy}}(x))\sigma_{ir_{iy}}(y)|^2 + 4|\sigma_{i_{xy}}(x)(\sigma_{i_{xy}}(y) - \sigma_{ir_{iy}}(y))|^2 + \\ &\quad + 4|(\sigma_{ir_{iy}}(y) - \sigma_{i_{xy}}(y))\sigma_{ir_{ix}}(x)|^2 + 4|\sigma_{i_{xy}}(y)(\sigma_{ir_{ix}}(x) - \sigma_{i_{xy}}(x))|^2 \\ &\lesssim 8|\sigma_{ir_{ix}}(x) - \sigma_{i_{xy}}(x)|^2 + 8|\sigma_{ir_{iy}}(y) - \sigma_{i_{xy}}(y)|^2, \end{aligned}$$

where we use the fact that $[\sigma_{i_{xy}}(x), \sigma_{i_{xy}}(y)] = 0$, and that U^*a^*aU is cyclically equivalent to a^*a if $UU^* = 1$. For any given $x \in V_i$ and $1 \leq j \leq m_i$, the number of elements $y \in V_i$ with $i_{xy} = j$ is bounded by $|V_{ij}|$. Hence

$$\sum_i \sum_{x, y \in V_i} \pi(i, i) \alpha(|[\sigma_i(x), \sigma_i(y)]|^2) \lesssim O(LM^2) \sum_{i,j,j'} \sum_{x \in V_{ij} \cap V_{ij'}} \frac{\pi(i, i)}{m_i^2} |\sigma_{ij}(x) - \sigma_{ij'}(x)|^2,$$

where $\sigma_{ij} : \mathbb{C}\mathbb{Z}_2^{V_{ij}} \rightarrow \mathcal{A}(B')$ is the inclusion of the ij th factor. We conclude that there is an $O(LM^2)$ -homomorphism $(\mathcal{A}_{free}(B), \mu_{comm}) \rightarrow \mathcal{A}_{inter}(B', \pi_{sub})$.

Finally, for μ_{clause} , if $i \in [m]$, $j \in [m_i]$, and $\phi \notin D_{ij}$ then $\sigma_{ij}(\Phi_{V_{ij},\phi}) = 0$, so

$$\begin{aligned}\alpha(\Phi_{V_{ij},\phi}) &= \alpha(\Phi_{V_{ij},\phi}) - \sigma_{ij}(\Phi_{V_{ij},\phi}) \\ &= \frac{1}{2^{|V_{ij}|}} \sum_{S \subseteq V_{ij}} \prod_{x \in S} \phi(x) \sigma_{ir_{ix}}(x) - \frac{1}{2^{|V_{ij}|}} \sum_{S \subseteq V_{ij}} \prod_{x \in S} \phi(x) \sigma_{ij}(x) \\ &= \frac{1}{2^{|V_{ij}|}} \sum_{S \subseteq V_{ij}} \sum_{x \in S} u_{x,S} \phi(x) (\sigma_{ir_{ix}}(x) - \sigma_{ij}(x)) v_{x,S},\end{aligned}$$

where $u_{x,S}$ is the product of $\phi(y) \sigma_{ij}(y)$ for $y \in S$ appearing before x in the order on V_i , and $v_{x,S}$ is the product of $\phi(y) \sigma_{ir_{iy}}(y)$ for $y \in S$ appearing after x in the order on V_i . Since there are less than $|V_{ij}| \cdot 2^{|V_{ij}|}$ terms in this sum, and $\phi(x) u_{x,S}$ and $v_{x,S}$ are unitary,

$$\begin{aligned}|\alpha(\Phi_{V_{ij},\phi})|^2 &\lesssim \frac{2^{|V_{ij}|}}{2^{|V_{ij}|}} \sum_{S \subseteq V_{ij}} \sum_{x \in S} |\sigma_{ir_{ix}}(x) - \sigma_{ij}(x)|^2 \\ &= \frac{|V_{ij}|}{2^{|V_{ij}|-1}} \sum_{x \in V_{ij}} \sum_{x \in S \subseteq V_{ij}} |\sigma_{ir_{ix}}(x) - \sigma_{ij}(x)|^2 \\ &= |V_{ij}| \sum_{x \in V_{ij}} |\sigma_{ir_{ix}}(x) - \sigma_{ij}(x)|^2.\end{aligned}$$

Hence

$$\begin{aligned}\sum_{i \in [m], j \in [m_i]} \frac{\pi(i, i)}{m_i^2} \sum_{\phi \notin D_{ij}} \alpha(|\Phi_{V_{ij},\phi}|^2) &\lesssim \sum_{i,j} \frac{\pi(i, i)}{m_i^2} \sum_{\phi \notin D_{ij}} L \sum_{x \in V_{ij}} |\sigma_{ir_{ix}}(x) - \sigma_{ij}(x)|^2 \\ &\leq L 2^L \sum_{i,j} \frac{\pi(i, i)}{m_i^2} \sum_{x \in V_{ij}} |\sigma_{ir_{ix}}(x) - \sigma_{ij}(x)|^2.\end{aligned}$$

Since every term in the latter sum occurs in the sum $\sum_r \mu'(r) r^* r$ for the weight function μ' of $\mathcal{A}_{inter}(B', \pi_{sub})$, α is a $L 2^L$ -homomorphism $(\mathcal{A}_{free}(B), \mu_{clause}) \rightarrow \mathcal{A}_{inter}(B', \pi_{sub})$. We conclude that α is an $O(M^2 + LM^2 + L 2^L)$ -homomorphism $\mathcal{A}_{free}(B, \pi) \rightarrow \mathcal{A}_{inter}(B', \pi_{sub})$, and $O(M^2 + LM^2 + L 2^L) \leq \text{poly}(M, 2^L)$. \square

Proof of Theorem 3.4.4. Applying Proposition 3.4.8, Proposition 3.4.7 and Proposition 3.3.3 gives the result. \square

Chapter 4

RE-completeness of entangled CSPs

In this chapter, we prove our main result about the hardness of constraint system games. We begin with an introduction to classical and quantum constraint satisfaction problems. We then prove our main result in three parts. First, we show that NP-hard CS games that are not two-variable falsifiable are RE-hard with entanglement by using a non-TVF constraint as a commutativity gadget. Finally, we prove the RE-hardness of NP-hard boolean CS and 3-colouring games with entanglement by showing that they admit more complex commutativity gadgets.

4.1 Constraint system languages

Informally, a constraint satisfaction problem is a collection of constraint systems that are constructed by filling a fixed set of constraints with variables arbitrarily. We formalise this with the notion of a pushforward constraint.

Definition 4.1.1. Let $C \subseteq \Sigma^V$ and let $r : V \rightarrow W$. The **pushforward constraint of C by r** is $r_*C = \{\phi \in \Sigma^W \mid \phi \circ r \in C\}$.

Definition 4.1.2. Let Γ be a finite set of constraints over an alphabet Σ . The **constraint satisfaction problem (CSP)** of Γ is the set of constraint systems

$$\text{CSP}(\Gamma) = \{(X, \{(W_i, r_{i*}C_i)\}_{i=1}^m) \mid W_i \subseteq X, (V_i, C_i) \in \Gamma, r_i : V_i \rightarrow W_i\}.$$

If every element of $\text{CSP}(\Gamma)$ is a 2-CS, we say it is a **2-CSP**. We also abuse notation and say Γ is a 2-CSP.

Example 4.1.3.

- 3SAT is the CSP where $\Sigma = \{\mathbb{Z}_2^3 \setminus \{t\} \mid t \in \mathbb{Z}_2^3\}$.
- k -colouring is the CSP where $\Sigma = \mathbb{Z}_k$ and $\Gamma = \{\neq_{\mathbb{Z}_k}\}$, for $\neq_{\mathbb{Z}_k} = \{(a, b) \in \mathbb{Z}_k^2 \mid a \neq b\}$. This is a 2-CSP.

A natural complexity problem given a CSP is to decide the satisfiability of the CS instances. Much is known about the classical version of this problem, which we outline here.

Definition 4.1.4 (Constraint languages). Let Γ be a set of constraints over an alphabet Σ , and let $1 \geq c \geq s \geq 0$.

- $\text{CSP}(\Gamma)_{c,s}$ is the promise problem with instances that are CSs $S = (X, \{(V_i, C_i)\}_{i=1}^m) \in \text{CSP}(\Gamma)$, where S is a yes instance if there is an assignment $f : X \rightarrow \Sigma$ such that $f|_{V_i} \in C_i$ for at least mc values of i , and S is a no instance if, for every assignment f , $f|_{V_i} \in C_i$ for strictly less than ms values of i .
- $\text{SuccinctCSP}(\Gamma)_{c,s}$ is the promise problem with instances that are probabilistic Turing machines M that sample the constraints of a CS $S = (X, \{(V_i, C_i)\}_{i=1}^m) \in \text{CSP}(\Gamma)$ according to some probability distribution $\pi : [m] \rightarrow [0, 1]$, where M is a yes instance if there is an assignment $f : X \rightarrow \Sigma$ such that $\Pr_{i \leftarrow \pi}[f|_{V_i} \in C_i] \geq c$, and M is a no instance if for every assignment $\Pr_{i \leftarrow \pi}[f|_{V_i} \in C_i] < s$.

Definition 4.1.5. A mapping $f : \Sigma^k \rightarrow \Sigma$ induces a **polymorphism** $(\Sigma^V)^k \rightarrow \Sigma^V$ as

$$f(\phi_1, \dots, \phi_n)(v) = f(\phi_1(v), \dots, \phi_n(v)).$$

The polymorphism f **preserves** (V, C) if for all $\phi_1, \dots, \phi_n \in C$, $f(\phi_1, \dots, \phi_n) \in C$. For a set of constraints Γ , we say f is a Γ -**homomorphism** if it preserves every constraint in Γ .

Definition 4.1.6. A map $f : \Sigma^k \rightarrow \Sigma$ is a **weak near-unanimity** if for all $a, b \in \Sigma$, $f(b, a, \dots, a) = f(a, b, a, \dots, a) = \dots = f(a, \dots, a, b)$.

Theorem 4.1.7 ([70, 11]). For all finite alphabets Σ , $\text{CSP}(\Gamma)_{1,1}$ is NP-complete if there is no weak near-unanimity Γ -homomorphism; otherwise, $\text{CSP}(\Gamma)_{1,1} \in \text{P}$.

Corollary 4.1.8. If $\text{CSP}(\Gamma)_{1,1}$ is NP-complete, then there is a constant $s \in [0, 1)$ such that $\text{SuccinctCSP}(\Gamma)_{1,s}$ is NEXP-complete. Otherwise, $\text{SuccinctCSP}(\Gamma)_{1,s} \in \text{EXP}$ for all s .

Proof sketch. Using the CSP version of the complexity class $\text{PCP}[m, q] = \text{NEXP}$ for $m = \exp(n)$ and $q = O(1)$, we know that $\text{SuccinctCSP}(\Gamma_q)_{1,1/2} = \text{NEXP}$, where Γ_q is the set of all constraints on q boolean variables [5, 65]. From the proof of hardness part of the CSP dichotomy conjecture [11, 70], we see that reduction of $\text{CSP}(\Gamma_q)_{1,1}$ to $\text{CSP}(\Gamma)_{1,1}$ is done constraint by constraint, by replacing each constraint by a conjunction of constraints from Γ . As there is a finite number of constraints in Γ_q , there exists $M \in \mathbb{N}$ such that every constraint can be expressed by at most M constraints from Γ . We can apply the same reduction to any instance of $S \in \text{SuccinctCSP}(\Gamma_q)_{1,1/2}$ in polynomial time to get a succinct description of an $S' \in \text{CSP}(\Gamma)$. Note first that, if S is a yes instance, all the constraints are satisfied, and therefore all the constraints of S' are satisfied. Next, if S is a no instance, suppose $\geq 1 - \frac{1}{2M}$ constraints are satisfied by an assignment to S' . Then, $\leq \frac{1}{2M}$ constraints of S' are not satisfied. Since each constraint of S is mapped to a conjunction of at most M constraints in S' , this means there is an assignment that satisfies $\geq \frac{1}{2}$ constraints of S , a contradiction. Hence, taking $s = 1 - \frac{1}{2M}$, if S is a no instance of $\text{SuccinctCSP}(\Gamma_q)_{1,1/2}$, S' is a no instance of $\text{SuccinctCSP}(\Gamma)_{1,s}$, as wanted.

Conversely, if $\text{CSP}(\Gamma)_{1,1}$ is not NP-complete, by the CSP dichotomy theorem [11, 70] there is a polynomial-time algorithm for it. Hence, this translates to an exponential-time algorithm for $\text{SuccinctCSP}(\Gamma)_{1,1}$. As every yes (**cf.** no) instance of $\text{SuccinctCSP}(\Gamma)_{1,s}$ is a yes (**cf.** no) instance of $\text{SuccinctCSP}(\Gamma)_{1,1}$, the algorithm also decides $\text{SuccinctCSP}(\Gamma)_{1,s}$ in exponential time. Hence, $\text{SuccinctCSP}(\Gamma)_{1,s} \in \text{EXP}$. \square

To finish this section, note that in the boolean case $\Sigma = \mathbb{Z}_2$, the set of weak near-unanimity polymorphisms is well known.

Definition 4.1.9. *The following are the weak near unanimity polymorphisms for boolean constraint systems.*

- The **constant 0 polymorphism** is $0 : \mathbb{Z}_2^0 \rightarrow \mathbb{Z}_2$ with output 0.
- The **constant 1 polymorphism** is $1 : \mathbb{Z}_2^0 \rightarrow \mathbb{Z}_2$ with output 1.
- The **AND polymorphism** is $\text{AND} : \mathbb{Z}_2^2 \rightarrow \mathbb{Z}_2$, $\text{AND}(a, b) = a \wedge b$.
- The **OR polymorphism** is $\text{OR} : \mathbb{Z}_2^2 \rightarrow \mathbb{Z}_2$, $\text{OR}(a, b) = a \vee b$.
- The **majority polymorphism** is $\text{MAJ} : \mathbb{Z}_2^3 \rightarrow \mathbb{Z}_2$, $\text{MAJ}(a, b, c) = (a \wedge b) \vee (b \wedge c) \vee (c \wedge a)$.
- The **minority polymorphism** is $\text{MIN} : \mathbb{Z}_2^3 \rightarrow \mathbb{Z}_2$, $\text{MIN}(a, b, c) = a \oplus b \oplus c$.

Theorem 4.1.10 (Schaefer’s dichotomy theorem [61]). *For Γ a set of constraints over \mathbb{Z}_2 , $\text{CSP}(\Gamma)_{1,1} \in \text{P}$ if one of the polymorphisms 0, 1, AND, OR, MAJ, MIN is a Γ -homomorphism; otherwise $\text{CSP}(\Gamma)_{1,1}$ is NP-complete.*

4.2 Quantum CSPs

The quantum satisfiability problem for CSPs can be equivalently phrased in terms of algebra representations or nonlocal games.

Definition 4.2.1 (Entangled constraint systems). *Let Γ be a set of constraints over \mathbb{Z}_k , let $w \in \{c - c, c - v, a, a + \text{comm}\}$, and let $1 \geq c \geq s \geq 0$.*

- $\text{CSP}_w(\Gamma)_{c,s}^*$ is the promise problem with instances that are CSs $S = (X, \{(V_i, C_i)\}_{i=1}^m) \in \text{CSP}(\Gamma)$, where S is a yes instance if

$$\inf_{\tau} \text{def}(\tau) \leq 1 - c,$$

where the infimum is over all finite-dimensional traces τ on $\mathcal{A}_w(S, \mathbb{u}_m)$, where \mathbb{u}_m is the uniform distribution on $[m]$; and S is a no instance if

$$\inf_{\tau} \text{def}(\tau) > 1 - s.$$

- $\text{SuccinctCSP}_w(\Gamma)_{c,s}^*$ is the promise problem with instances that are probabilistic Turing machines M that sample the constraints of a CS $S = (X, \{(V_i, C_i)\}_{i=1}^m) \in \text{CSP}(\Gamma)$ according to some probability distribution $\pi : [m] \rightarrow [0, 1]$, where S is a yes instance if

$$\inf_{\tau} \text{def}(\tau) \leq 1 - c,$$

where the infimum is over all finite-dimensional traces τ on $\mathcal{A}_w(S, \pi)$; and S is a no instance if

$$\inf_{\tau} \text{def}(\tau) > 1 - s.$$

We drop the subscripts if they are not important. Note that we can always replace the infima over finite-dimensional traces by Connes-embeddable traces.

In this definition, the yes instances correspond exactly to CSs where the quantum synchronous value of the corresponding CS game, as defined in Section 3.1, is at least c ,

and the no instances correspond to CSs whose quantum synchronous value is less than s . If $w = c - c$, the corresponding game is the constraint-constraint CS game; if $w = c - v$, the corresponding game is the constraint variable CS game; and if $w = a$, there is only an associated game if Γ is a 2-CSP, in which case it is the 2-CS game.

In contrast to the classical case, where there is a complete classification of CSPs, less is known about the complexity of quantum CSPs. Prior work by Atserias, Kolaitis, and Severini [3, 58] has shown that all the classically-easy boolean CSPs except for LIN, the class of linear systems over \mathbb{Z}_2 , remain easy with entanglement; but the complexity of entangled LIN and of any classically easy CSPs over larger domains remains unknown. In this work, we are concerned with what happens to problems on the hard side of the CSP dichotomy when the provers are allowed entanglement.

We make use the following property of CSs, which precludes the construction of simple commutativity gadgets.

Definition 4.2.2. *Let $C \subseteq \Sigma^V$ be a constraint over an alphabet Σ . C is **two-variable falsifiable (TVF)** if for all $x \neq y \in V$, there exist $a, b \in \Sigma$ such that $\phi \notin C$ if $\phi(x) = a$ and $\phi(y) = b$.*

We say a CS S or a set of constraints Γ is TVF if every constraint in it is TVF.

Two-variable falsifiability can seem like a very strong restriction on the constraints. However, the example of 1-in-3-SAT shows that there are nontrivial CSPs that are TVF.

Example 4.2.3. *Consider the boolean constraint*

$$C = \{(0, 0, 1), (0, 1, 0), (1, 0, 0)\} \subseteq \mathbb{Z}_2^3.$$

This is TVF as setting any two variables to 1 makes the constraint unsatisfied. On the other hand, the language $\text{CSP}(\{C\})_{1,1}$ is NP-complete as a direct consequence of Schaefer's dichotomy theorem.

We are now ready to state the main result of this chapter.

Theorem 4.2.4. *Let Γ be a set of constraints over \mathbb{Z}_k such that $\text{CSP}(\Gamma)_{1,1}$ is NP-complete, and one of the following holds:*

1. Γ is not TVF,
2. Γ is boolean, or

3. $\Gamma = \{\neq_{\mathbb{Z}_3}\}$ is 3-colouring.

Then, there exists a constant $s \in [0, 1)$ such that $\text{SuccinctCSP}_{c-v}(\Gamma)_{1,s}^*$ is RE-complete.

It remains open whether non-boolean TVF CSPs, except for 3-colouring, are also RE-complete. As an important example, we do not know if k -colouring for $k \geq 4$ is RE-complete with entanglement. However, we do know that any 2-CSP that contains an empty constraint (where all assignments are accepted), such as the language of all 2-CSPs over \mathbb{Z}_k , is RE-complete as it is not TVF.

For RE-complete entangled CSPs, it also holds that $\text{SuccinctCSP}_{c-c}(\Gamma)_{1,s}^*$ is RE-complete, which we show using a mapping between the constraint-variable and constraint-constraint algebras given in Proposition 4.6.1. This is in contrast with Lemma 3.3.5, where the defect can scale very badly when mapping to the constraint-constraint algebra.

Corollary 4.2.5. *Let Γ be a set of constraints satisfying the conditions of Theorem 4.2.4. Then, there exists a constant $s \in [0, 1)$ such that $\text{SuccinctCSP}_{c-c}(\Gamma)_{1,s}^*$ is RE-complete.*

For some CSPs, we can also show a similar RE-hardness result for the assignment algebra via an oracularizability property. For 2-CSPs, this corresponds to the associated 2-CS game.

Corollary 4.2.6. *Let Γ be a set of constraints over \mathbb{Z}_k such that $\text{CSP}(\Gamma)_{1,1}$ is NP-complete, and Γ is boolean TVF, $\Gamma = \{\neq_{\mathbb{Z}_3}\}$, or $\Gamma = \{C \subseteq \mathbb{Z}_k^2\}$. Then, there exists a constant $s \in [0, 1)$ such that $\text{SuccinctCSP}_a(\Gamma)_{1,s}^*$ is RE-complete.*

More generally, the hardness extends to the assignment algebra with commutation.

Corollary 4.2.7. *Let Γ be a set of constraints satisfying the conditions of Theorem 4.2.4. Then, there exists a constant $s \in [0, 1)$ such that $\text{SuccinctCSP}_{a+comm}(\Gamma)_{1,s}^*$ is RE-complete.*

This follows immediately from the mappings between the constraint-variable and assignment with commutation algebras given in Lemmas 3.3.6 and 3.3.7. By these, the values of the defects for these algebras are equal up to constant factors (for constraint systems with constant-size constraints and alphabets), and therefore, the hardness of deciding one directly implies the hardness of the other.

Next, as noted in the introduction, we can extend the undecidability of entangled CSP languages to non-succinct languages, by considering computable reductions that are not polynomial-time.

Corollary 4.2.8. *Let Γ be a set of constraints over \mathbb{Z}_k such that $\text{CSP}(\Gamma)_{1,1}$ is NP-complete, and Γ satisfies the conditions of Theorem 4.2.4. Then, there exists a constant $s \in [0, 1)$ such that $\text{CSP}_{c-v}(\Gamma)_{1,s}^*$, $\text{CSP}_{c-c}(\Gamma)_{1,s}^*$, and $\text{CSP}_{a+comm}(\Gamma)_{1,s}^*$ are RE-complete with respect to exponential-time reductions. If Γ additionally satisfies the conditions of Corollary 4.2.6, $\text{CSP}_a(\Gamma)_{1,s}^*$ is RE-complete with respect to exponential-time reductions.*

Theorem 4.2.4 and earlier corollaries give a polynomial-time mapping from instances of the halting problem to instances of $\text{SuccinctCSP}_w(\Gamma)_{1,s}^*$, preserving yes and no instances. In exponential time, the whole constraint-sampling algorithm can be described. This gives us an exponential-time mapping to $\text{CSP}_w(\Gamma)_{1,s}^*$. The only remaining issue to check is whether we can make the probability distribution uniform while preserving the constant gap. First, in the constraint-variable game, the probability can be perturbed by mixing with a uniform distribution to ensure that the probability of any constraint is at least α/m for some small constant $\alpha \in (0, 1)$ while preserving a constant gap. Next, each constraint i can be repeated with multiplicity $\lfloor \pi(i)m/\alpha \rfloor$. These constraints are sampled from a uniform distribution, which preserves the constant soundness. Note that the defect cannot be made smaller by making the trace have different values on the different copies; there exists a trace on the original algebra corresponding to the minimal value of the trace over all copies. Also, this reduction to uniform distribution requires the number of constraints to be polynomial in order to be efficient, as it may be easy to sample from a distribution while being hard to compute the probabilities. An analogous argument applies for the other algebra models.

The following sections are devoted to the proof of Theorem 4.2.4. Using the subdivision transformation from Section 3.4 allows us to ask constraints from a CSP as separate questions, rather than asking large CSP instances, while preserving a constant gap, so long as we start and end with questions of constant size. We split the proof of Theorem 4.2.4 into three parts. In Section 4.3, we prove Theorem 4.2.4 for the case of non-TVF CSPs. Next, in Section 4.4, we provide the proof of Theorem 4.2.4 for the case of boolean TVF CSPs. We also show Corollary 4.2.6 for these CSPs there. In Section 4.5, we prove Theorem 4.2.4 for 3-colouring, and Corollary 4.2.6 for 3-colouring and the language of all 2-CSPs. Finally, in Section 4.6, we prove Corollary 4.2.5.

4.3 Hardness of non-TVF CSPs

In this section, we prove the first part of Theorem 4.2.4, proving RE-hardness for noncommutative non-TVF CSPs. The main obstacle in doing so is condition (2) of Definition 3.4.2:

to subdivide a constraint, each pair of variables must appear together in one of the sub-constraints. But, in general, putting two variables in the same constraint puts nontrivial restrictions on the values that they may take. A naive way to get around this is using empty constraints.

Definition 4.3.1. We call a constraint (V, C) over an alphabet Σ **empty** if $C = \Sigma^V$.

Empty constraints impose no conditions on the variables involved, but they can be used to guarantee that variables commute. Hence, they are useful in subdivision to make sure every pair of variables appears in at least one constraint. However, if we wish to reduce to a CSP that does not contain any empty constraints, we need to find a way to replace these empty constraints by some constraint system coming from the CSP.

If a set of constraints is non-TVF (Definition 4.2.2), we can replace any empty constraint by a simple commutativity gadget coming from a non-TVF constraint.

Proposition 4.3.2. Let Γ be a non-TVF set of k -ary constraints. Then, for every CS $S = (X, \{(V_i, r_{i*} C_i)\}_{i=1}^m) \in \text{CSP}(\Gamma \cup \{\mathbb{Z}_k^2\})$, there exists a CS $S' = (X', \{(V'_i, r'_{i*} C'_i)\}_{i=1}^m) \in \text{CSP}(\Gamma)$ such that there is a $\frac{L}{2}$ -homomorphism $\alpha : \mathcal{A}_{c-v}(S, \pi) \rightarrow \mathcal{A}_{c-v}(S', \pi)$, where $L = \max_{(V, C) \in \Gamma} |V|$; and there is a 1-homomorphism $\beta : \mathcal{A}_{c-v}(S', \pi) \rightarrow \mathcal{A}_{c-v}(S, \pi)$, for every probability distribution π on $[m]$.

Proof. Without loss of generality, we may assume that there exists $1 \leq n \leq m$ such that C_i is a constraint from Γ for all $i \leq n$ and C_i is an empty constraint for all $i > n$. Since Γ is not TVF, there exists a constraint $(V, C) \in \Gamma$ such that $C \subseteq \mathbb{Z}_k^L$ and $u \neq v \in V$ such that for all $a, b \in \mathbb{Z}_k$, there exists $\phi_{a,b} \in C$ with $\phi_{a,b}(u) = a$ and $\phi_{a,b}(v) = b$. For each $i > n$, and $w \in V \setminus \{u, v\}$ let z_{iw} be a variable and take $X' = X \cup \{z_{iw} \mid i > n, w \in V \setminus \{u, v\}\}$. Also, if $i \leq n$, let $(V'_i, r'_{i*} C'_i) = (V_i, r_{i*} C_i)$; and if $i > n$, take $V'_i = V_i \cup \{z_{iw} \mid w \in V \setminus \{u, v\}\}$, $C'_i = C$, and r'_i a bijection such that $r'_i(w) = z_{iw}$ for all $w \in V \setminus \{u, v\}$.

Now, let α be the inclusion of $\mathcal{A}_{c-v}(S)$ in $\mathcal{A}_{c-v}(S')$. First, note that for every $i > n$,

$$\begin{aligned}
\alpha\left(\sum_{\substack{x \in V_i \\ \phi \in r_{i*} C_i}} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2\right) &= \sum_{\substack{x \in V_i \\ \phi \in \mathbb{Z}_k^{V_i}}} \left| \prod_{y \in V_i} \Pi_{\phi(y)}(\sigma_i(y))(1 - \Pi_{\phi(x)}(\sigma'(x))) \right|^2 \\
&= \sum_{\substack{x \in V_i \\ \phi \in \mathbb{Z}_k^{V'_i}}} \left| \prod_{y \in V'_i} \Pi_{\phi(y)}(\sigma_i(y))(1 - \Pi_{\phi(x)}(\sigma'(x))) \right|^2 \\
&= \sum_{\substack{x \in V_i \\ \phi \in r'_{i*} C'_i}} |\Phi_{V'_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \\
&\leq \sum_{\substack{x \in V'_i \\ \phi \in r'_{i*} C'_i}} |\Phi_{V'_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2.
\end{aligned}$$

Hence, we get that

$$\begin{aligned}
&\alpha\left(\sum_{i=1}^m \sum_{\substack{x \in V_i \\ \phi \in r_{i*} C_i}} \frac{\pi(i)}{|V_i|} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2\right) \\
&\leq \sum_{i=1}^n \frac{\pi(i)}{|V_i|} \sum_{\substack{x \in V'_i \\ \phi \in r'_{i*} C'_i}} |\Phi_{V'_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 + \sum_{i=n+1}^m \frac{\pi(i)}{2} \sum_{\substack{x \in V'_i \\ \phi \in r'_{i*} C'_i}} |\Phi_{V'_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \\
&\leq \frac{|V|}{2} \sum_{\substack{i, x \in V'_i \\ \phi \in r'_{i*} C'_i}} \frac{\pi(i)}{|V'_i|} |\Phi_{V'_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \leq \frac{L}{2} \sum_{\substack{i, x \in V'_i \\ \phi \in r'_{i*} C'_i}} \frac{\pi(i)}{|V'_i|} |\Phi_{V'_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2
\end{aligned}$$

For the converse, let β be defined as the map acting on $x \in X$ as $\sigma_i(x) \mapsto \sigma_i(x)$, $\sigma'(x) \mapsto \sigma'(x)$; and on z_{iw} as

$$\beta(\sigma_i(z_{iw})) = \beta(\sigma'(z_{iw})) = \sum_{a, b \in \mathbb{Z}_k} \omega_k^{\phi_{a,b}(w)} \Pi_a(\sigma_i(x)) \Pi_b(\sigma_i(y)).$$

Then, for all $i \leq n$, $\beta(\Phi_{V'_i, \phi}) = \Phi_{V_i, \phi}$, and for $i > n$, writing $V_i = \{x, y\}$,

$$\begin{aligned}
\beta(\Phi_{V'_i, \phi}) &= \Pi_{\phi(x)}(\sigma_i(x)) \Pi_{\phi(y)}(\sigma_i(y)) \prod_{w \in V \setminus \{u, v\}} \sum_{a, b, \phi_{a,b}(w) = \phi(z_{iw})} \Pi_a(\sigma_i(x)) \Pi_b(\sigma_i(y)) \\
&= \begin{cases} \Pi_{\phi(x)}(\sigma_i(x)) \Pi_{\phi(y)}(\sigma_i(y)) & \forall w \in V \setminus \{u, v\} \phi_{\phi(x), \phi(y)}(w) = \phi(z_{iw}) \\ 0 & \text{else} \end{cases}
\end{aligned}$$

That is, $\beta(\Phi_{V'_i, \phi}) = \Phi_{V_i, \phi|_{V_i}}$ if and only if $\phi = \phi_{\phi(x), \phi(y)} \circ r_i^{-1}$, and is 0 otherwise. In particular, it maps every $\Phi_{V'_i, \phi}$ for $\phi \notin r_{i*}C'_i$ to 0, meaning β is a *-homomorphism as needed. Also, in the case $\beta(\Phi_{V'_i, \phi}) \neq 0$, we have that $\beta(\Phi_{V'_i, \phi}(1 - \Pi_{\phi(z_{iw})}(\sigma'(z_{iw})))) = 0$, so

$$\begin{aligned} & \beta\left(\sum_{\substack{i, x \in V'_i \\ \phi \in r'_{i*}C'_i}} \frac{\pi(i)}{|V'_i|} |\Phi_{V'_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2\right) \\ &= \sum_{i=1}^n \frac{\pi(i)}{|V_i|} \sum_{\substack{x \in V_i \\ \phi \in r_{i*}C_i}} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 + \sum_{i=n+1}^m \frac{\pi(i)}{|V|} \sum_{\substack{x \in V_i \\ \phi \in r_{i*}C_i}} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \\ &\leq \sum_i \frac{\pi(i)}{|V_i|} \sum_{\substack{x \in V_i \\ \phi \in r_{i*}C_i}} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2. \end{aligned}$$

□

We are now ready to prove the hardness of non-TVF CSPs.

Theorem 4.3.3 (Part 1 of Theorem 4.2.4). *Let Γ be a non-TVF set of k -ary constraints such that $\text{CSP}(\Gamma)_{1,1}$ NP-complete. Then there exists a constant $s \in [0, 1)$ such that $\text{SuccinctCSP}_{c-v}(\Gamma)_{1,s}^*$ is RE-complete.*

Proof. By Corollary 3.1.2, there is a BCS – MIP* protocol $(\mathbf{G}(B_x, \pi_x), S, C)$ for the halting problem with constant soundness $0 \leq s < 1$, where $B_x = (X_x, \{(V_i^x, C_i^x)\}_{i=1}^{m_x})$, m_x is exponential in $|x|$, and $|V_i^x| = O(1)$. By the NP-completeness of $\text{CSP}(\Gamma)_{1,1}$, there is a BCS $B' = (Y_x, \{(W_i^x, D_i^x)\}_{i=1}^{n_x})$ as in Corollary 3.2.9, where $|W_i^x| = O(1)$, n_x exponential in x , and D_i^x is the boolean form of a $\text{CSP}(\Gamma)$ instance. By Lemma 3.2.6, there is a BCS – MIP* protocol $(\mathbf{G}(B', \pi_x), S, C')$ for the halting problem with the same soundness. Since $|W_i^x| = O(1)$, Theorem 3.4.4 implies that there is a BCS – MIP* protocol $(\mathbf{G}(B(S_x), \pi_{sub}^x), \tilde{S}, \tilde{C})$ for the halting problem with constant soundness, where S_x is a $\text{CSP}(\Gamma)$ instance that may contain empty constraints. By subdividing further, we may always assume that the empty constraints are on two variables. Applying Lemmas 3.3.2, 3.3.4 and 3.3.5 gives a constraint-variable CS – MIP* protocol $(\mathbf{G}(S_x, \pi_{sub}^x), \tilde{S}, \tilde{C}')$ for the halting problem with constant soundness, where S_x is a $\text{CSP}(\Gamma)$ instance that may contain empty constraints. Finally, Proposition 4.3.2 gives the result. □




Figure 4.1: The basic commutativity gadget for TVF boolean constraint systems. Exactly one variable in each triangle must be assigned value 1. These constraints bound the commutator $[x, y]$, and any assignment to x and y may be extended to an assignment to all three constraints.

4.4 Hardness of boolean TVF CSPs

In this section, we examine the hardness of boolean CSPs that are two-variable falsifiable. NP-complete classical boolean CSPs can emulate any other constraint system. What we show in this section is that this emulation can be done in a way that is sound against quantum provers. We find that boolean TVF CSPs are NP-complete if and only if they allow for a commutativity gadget similar to that constructed in [37] for 1-in-3SAT. We prove the RE-hardness of the entangled version of these CSPs by showing that the commutativity gadget is quantum sound.

4.4.1 The basic commutativity gadget

We begin with a description of the commutativity gadget, and the proof of its quantum soundness.

Lemma 4.4.1. *Let $C = \{(0, 0, 1), (0, 1, 0), (1, 0, 0)\} \in \mathbb{Z}_2^{[3]}$. Let $X = \{u, v, w, x, y, z\}$ be a set of variables; let $V_1 = \{x, u, v\}$, $V_2 = \{y, u, w\}$, and $V_3 = \{z, v, w\}$; and let $r_i : [3] \rightarrow V_i$ be bijections. Consider the BCS $B = \{X, \{(V_i, r_{i*}C)\}_{i=1}^3\}$. For all $a_x, a_y \in \mathbb{Z}_2$, there exists a classical satisfying assignment $\phi : X \rightarrow \mathbb{Z}_2$, $\phi|_{V_i} \in r_{i*}C$ such that $\phi(x) = a_x$ and $\phi(y) = a_y$. Also, in the algebra $\mathcal{A}_{c-v}(B)$,*

$$|[\sigma'(x), \sigma'(y)]|^2 \lesssim 512 \sum_{i=1}^3 \sum_{\phi \in r_{i*}C} \sum_{z \in V_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(z)}(\sigma'(z)))|^2.$$

This is a robust version of Lemma 5 in [37].

The commutativity gadget defined in this lemma is illustrated in Figure 4.1. Note also that an identical commutativity gadget can be constructed from C with any of the variables negated, one of the constraints $C' = \{(0, 0, 0), (0, 1, 1), (1, 0, 1)\}$, $C'' = \{(0, 1, 0), (0, 0, 1), (1, 1, 1)\}$, or $C''' = \{(1, 1, 0), (1, 0, 1), (0, 1, 1)\}$, where the last one, two, or three variable are negated, respectively. To achieve this, it suffices to connect three constraints in the same way, additionally taking care that negated variables are only connected to other negated variables. This is illustrated in Figure 4.2.

Proof. For the first part, we can work out all the cases: for X ordered as in the lemma statement, we get the satisfying assignments $(1, 1, 1, 0, 0, 0)$ for $(a_x, a_y) = (0, 0)$, $(0, 1, 0, 0, 1, 0)$ for $(a_x, a_y) = (0, 1)$, $(0, 0, 1, 1, 0, 0)$ for $(a_x, a_y) = (1, 0)$, and $(0, 0, 0, 1, 1, 1)$ for $(a_x, a_y) = (1, 1)$. For the second part, write $x_1^i = \Pi_1(\sigma_i(x))$ and $x_0^i = \Pi_0(\sigma_i(x))$ and similarly for the other variables. Write $x_1 = \Pi_1(\sigma'(x))$ and $x_0 = \Pi_0(\sigma'(x))$ and similarly for the other variables. First, note that $x_1^1 u_1^1 = x_1^1 v_1^1 = u_1^1 v_1^1 = 0$, so $x_1^1 + u_1^1 + v_1^1 \leq 1$, and

$$1 = x_1^1 u_0^1 v_0^1 + x_0^1 u_1^1 v_0^1 + x_0^1 u_0^1 v_1^1 \leq x_1^1 + u_1^1 + v_1^1,$$

giving equality $x_1^1 + u_1^1 + v_1^1 = 1$. In the same way, $y_1^2 + u_1^2 + w_1^2 = 1$. Therefore, the commutator

$$\begin{aligned} [x_1^1, y_1^2] &= [1 - u_1^1 - v_1^1, 1 - u_1^2 - w_1^2] = [u_1^1 + v_1^1, u_1^2 + w_1^2] \\ &= [u_1^1, u_1^2] + [u_1^1, w_1^2] + [v_1^1, u_1^2] + [v_1^1, w_1^2]. \end{aligned}$$

Noting that $[u_1^1, u_1^1] = [u_1^2, w_1^2] = [v_1^1, u_1^1] = [v_1^3, w_1^3] = 0$, we can write

$$\begin{aligned} [x_1^1, y_1^2] &= [u_1^1, u_1^2 - u_1^1] + [u_1^1 - u_1^2, w_1^2] + [v_1^1, u_1^2 - u_1^1] + [v_1^1 - v_1^3, w_1^2] + [v_1^3, w_1^2 - w_1^3] \\ &= [u_1^1 + v_1^1 + w_1^2, u_1^2 - u_1^1] - [w_1^2, v_1^1 - v_1^3] + [v_1^3, w_1^2 - w_1^3]. \end{aligned}$$

From here, we can expand

$$\begin{aligned} [x_1, y_1] &= [x_1^1, y_1^2] + [x_1 - x_1^1, y_1^2] + [x_1, y_1 - y_1^2] \\ &= [u_1^1 + v_1^1 + w_1^2, u_1^2 - u_1^1] - [w_1^2, v_1^1 - v_1^3] + [v_1^3, w_1^2 - w_1^3] + [y_1^2, x_1^1 - x_1] - [x_1, y_1^2 - y_1] \\ &= [u_1^1 + v_1^1 + w_1^2, u_1^2 - u_1^1] - [u_1^1 + v_1^1 + w_1^2, u_1^1 - u_1] - [w_1^2, v_1^1 - v_1] + [w_1^2, v_1^3 - v_1] \\ &\quad + [v_1^3, w_1^2 - w_1] - [v_1^3, w_1^3 - w_1] + [y_1^2, x_1^1 - x_1] - [x_1, y_1^2 - y_1]. \end{aligned}$$

Hence, we get the bound on the hermitian square of the commutator

$$\begin{aligned}
|[x_1, y_1]|^2 &\leq 2^{\lceil \log 8 \rceil} (|[u_1^1 + v_1^1 + w_1^2, u_1^2 - u_1]|^2 + |[u_1^1 + v_1^1 + w_1^2, u_1^1 - u_1]|^2 \\
&\quad + |[w_1^2, v_1^1 - v_1]|^2 + |[w_1^2, v_1^3 - v_1]|^2 + |[v_1^3, w_1^2 - w_1]|^2 \\
&\quad + |[v_1^3, w_1^3 - w_1]|^2 + |[y_1^2, x_1^1 - x_1]|^2 + |[x_1, y_1^2 - y_1]|^2) \\
&\lesssim 16 (|u_1^2 - u_1|^2 + |u_1^1 - u_1|^2 + |v_1^1 - v_1|^2 + |v_1^3 - v_1|^2 + |w_1^2 - w_1|^2 \\
&\quad + |w_1^3 - w_1|^2 + |x_1^1 - x_1|^2 + |y_1^2 - y_1|^2) \\
&\leq 16 \sum_{i=1}^3 \sum_{t \in V_i} |t_1^i - t_1|^2 = 4 \sum_{i=1}^3 \sum_{t \in V_i} |1 - t^i t|^2 = 16 \sum_{i=1}^3 \sum_{t \in V_i} |t_0^i t_1 + t_1^i t_0|^2 \\
&\leq 32 \sum_{i=1}^3 \sum_{t \in V_i} \sum_{a \in \mathbb{Z}_2} |t_a^i (1 - t_a)|^2 = 32 \sum_{i=1}^3 \sum_{\phi \in r_{i,*} C} \sum_{z \in V_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(z)}(\sigma'(z)))|^2.
\end{aligned}$$

Noting that $[x_1, y_1] = \frac{1}{4}[x, y]$ finishes the proof. \square

4.4.2 Compression and simulation: building the needed constraints

In this section, we analyse the structure of boolean TVF constraints. In particular, we show that, given an NP-complete set of boolean TVF constraints, we can recover the 1-in-3SAT constraint from Lemma 4.4.1, or its negation on a subset of variables. To do so, we need to study the combinatorial structure of TVF constraints, which will allow us to simplify the sets of constraints we work with and then simulate the wanted constraint.




Figure 4.2: Basic commutativity gadgets with one, two, or three variables per constraint negated. The white vertices indicate the negated variables: note that negated variables must be only connected amongst themselves to construct the gadget.

Definition 4.4.2. A **(boolean) TVF graph** is a graph with two types of undirected edges, labelled 00 and 11, and one type of directed edge, labelled 01. We denote a TVF graph $G = (V, E_{00} \sqcup E_{11} \sqcup E_{01})$. Given a vertex $v \in V$, we say that an edge e of G is 0 **on** v if there exists $u \in V$ such that $e = \{u, v\} \in E_{00}$ or $e = (v, u) \in E_{01}$. In the same way, we say that e is 1 **on** v if there exists $u \in V$ such that $e = \{u, v\} \in E_{11}$ or $e = (u, v) \in E_{01}$.

Let V be a finite set and let $C \subseteq \mathbb{Z}_2^V$ be a boolean constraint on V . The **TVF graph of C** is a graph $G_{TVF}(C) = (V, E_{00} \sqcup E_{11} \sqcup E_{01})$, where $\{u, v\} \in E_{00}$ is an edge if $\phi(u) = \phi(v) = 0$ implies $\phi \notin C$; $\{u, v\} \in E_{11}$ is an edge if $\phi(u) = \phi(v) = 1$ implies $\phi \notin C$; and $(u, v) \in E_{01}$ is a directed edge if $\phi(u) = 0$ and $\phi(v) = 1$ implies $\phi \notin C$.

A constraint C is TVF if and only if the TVF graph of C is complete. We can study TVF graphs independently of the constraints that generate them, although any TVF graph is generated by some constraint.

In a TVF graph, we sometimes denote an undirected edge $\{a, b\}$ in the same way as a directed edge, (a, b) , if we want to be more general about which edge set the edge belongs to.

Definition 4.4.3. An **assignment** to a TVF graph $G = (V, E_{00} \sqcup E_{11} \sqcup E_{01})$ is a function $\phi : V \rightarrow \mathbb{Z}_2$ such that $(\phi(u), \phi(v)) \neq (0, 0)$ for all $\{u, v\} \in E_{00}$, $(\phi(u), \phi(v)) \neq (1, 1)$ for all $\{u, v\} \in E_{11}$, and $(\phi(u), \phi(v)) \neq (0, 1)$ for all $(u, v) \in E_{01}$.

Every satisfying assignment to C induces an assignment to $G_{TVF}(C)$, but the converse is not necessarily true.

Definition 4.4.4. Let $G = (V, E_{00} \sqcup E_{11} \sqcup E_{01})$ be a TVF graph. We say that G is **compressible** to $U \subseteq V$ if for all $v \in V \setminus U$ and all assignments ϕ to G , either there is a constant $b \in \mathbb{Z}_2$ such that $\phi(v) = b$, or $u \in U$ such that $\phi(v) = \phi(u)$, or $u \in U$ such that $\phi(v) = \neg\phi(u)$.

In the first case, we say that G **compresses by a constant** to b at v . In the second case, we say that G **compresses by equality** to u at v . In the third case, we say that G **compresses by negation** to u at v . We say G is **incompressible** if G is only compressible to $U \subseteq V$ when $U = V$.

We use the same notation for a constraint $C \subseteq \mathbb{Z}_2^V$ when the properties hold for its TVF graph $G_{TVF}(C)$.

It is easy to see that if a subgraph of G is compressible, so is G . Now, we give some important examples of compressible TVF graphs.

Lemma 4.4.5. (i) A TVF graph with a 00 or 11 edge as a loop is compressible.

(ii) A TVF graph with a double edge between two distinct vertices is compressible.

(iii) The following TVF graph is compressible for $n > 1$ and all $b_1, \dots, b_{n+1} \in \mathbb{Z}_2$: $G = (\{x_1, \dots, x_n\}, E_{00} \sqcup E_{11} \sqcup E_{01})$ with $(x_i, x_{i+1}) \in E_{-b_i b_{i+1}}$ for $i = 1, \dots, n-1$ and $(x_n, x_1) \in E_{-b_n b_{n+1}}$.

Note that 01 loops are redundant as they do not affect the assignments. As such we may suppose that TVF graphs have no 01 loops. Hence, this lemma tells us in particular that incompressible TVF graphs have no loops or multiple edges, *i.e.* they are simple graphs. The cycle graph described in (iii) is illustrated in fig. 4.3.

Proof. (i) Consider the one-vertex TVF graph G with $V = \{x\}$ and $E_{bb} = \{\{x, x\}\}$. Then, we know that for every assignment ϕ to G , if $\phi(x) = b$, then $\phi(x) \neq b$. Hence, we must have that $\phi(x) = -b$. As such, G compresses by a constant at x .

(ii) Consider the two-vertex TVF graph $G = (\{x, y\}, E_{00} \sqcup E_{11} \sqcup E_{01})$. Up to relabelling of x and y , there are 4 possible double edges. In the case $E_{00} = E_{11} = \{\{x, y\}\}$ and $E_{01} = \emptyset$, we have that $\phi(x) \neq \phi(y)$ for every assignment ϕ , so G is compressible by negation. In the case $E_{00} = \{\{x, y\}\}$, $E_{11} = \emptyset$, and $E_{01} = \{(x, y)\}$, we have that $\phi(x) = 1$ for every assignment ϕ , so G is compressible by a constant. In the case $E_{00} = \emptyset$, $E_{11} = \{\{x, y\}\}$, and $E_{01} = \{(x, y)\}$, we have that $\phi(y) = 0$ for every assignment ϕ , so G is compressible by a constant. And in the case $E_{00} = E_{11} = \emptyset$ and $E_{01} = \{(x, y), (y, x)\}$, we have that $\phi(x) = \phi(y)$ for every assignment ϕ , so G is compressible by equality.

(iii) If $n = 2$, we are in the case of (ii), so we know that G is compressible. In the case $n \geq 3$, we consider two cases depending on the value of $\phi(x_2)$ for an assignment ϕ . First note that, if $\phi(x_2) = b_2$, then $\phi(x_1) = b_1$. Also, if $\phi(x_2) = -b_2$, then $\phi(x_3) = -b_3$, and by induction $\phi(x_i) = -b_i$ for $i = 2, \dots, n$. This implies $\phi(x_1) = -b_{n+1}$. If $b_1 = -b_{n+1}$, then $\phi(x_1) = b_1$ in both possible cases, so G is compressible by a constant. If $b_1 = b_{n+1}$, we get in the first case that $\phi(x_1) = b_1 = b_{n+1}$, so $\phi(x_n) = b_n$ and by induction $\phi(x_i) = b_i$ for all $i = 2, \dots, n$. As $n \geq 3$, there exist $i \neq j = 1, \dots, n$ such that $b_i = b_j$, and hence $\phi(x_i) = \phi(x_j)$ in both cases, giving that G is compressible by equality.

□




Figure 4.3: The compressible cycle TVF graph from Lemma 4.4.5.iii.

If a given set of constraints has compressible elements, there may be significant redundancy in the variables. To remedy this, we show that we can always reduce to the case of incompressible constraints while preserving the NP-hardness. We will then show that all compressed NP-hard constraint systems allow for the construction of commutativity gadgets.

Lemma 4.4.6. *Let $C \subseteq \mathbb{Z}_2^V$ be a constraint. Suppose C is compressible to $U \subseteq V$, and that $C|_U$ is compressible to W . Then, C is compressible to W .*

Proof. Let $v \in V \setminus W$. If $v \in U$, we know, since $C|_U$ is compressible to W , that either $\phi(v) = b$ for some $b \in \mathbb{Z}_2$, $\phi(v) = \phi(u)$ for some $u \in W$, or $\phi(v) = \neg\phi(u)$ for some $u \in W$, for all $\phi \in C$. If $v \notin U$, since C is compressible to U , either $\phi(v) = b$ for some $b \in \mathbb{Z}_2$, $\phi(v) = \phi(u)$ for some $u \in U$, or $\phi(v) = \neg\phi(u)$ for some $u \in U$, for all $\phi \in C$. In the latter two cases, if $u \in W$ we are done. If not, knowing that $C|_U$ compresses at u , we have that either $\phi(u) = b$, $\phi(u) = \phi(w)$, or $\phi(u) = \neg\phi(w)$, for some $w \in W$. As such, in the two cases, we get either $\phi(v) = b$, $\phi(v) = \phi(w)$, or $\phi(v) = \neg\phi(w)$; or $\phi(v) = \neg b$, $\phi(v) = \neg\phi(w)$, or $\phi(v) = \phi(w)$. \square

Proposition 4.4.7. *Every constraint $C \subseteq \mathbb{Z}_2^V$ is compressible to an incompressible constraint.*

Proof. We proceed recursively. If for every nonempty $U \subset V$, C does not compress to U , then C is incompressible. Otherwise, C compresses to some U_1 . Now, consider $C|_{U_1}$, and continue recursively. We get a descending sequence of subsets $V \supset U_1 \supset U_2 \supset \dots$ such that C compresses to U_1 , $C|_{U_1}$ compresses to U_2 , and so on. As these inclusions are strict, there will be some $k \in \mathbb{N}$ such that $C|_{U_k}$ is incompressible. But by Lemma 4.4.6, C compresses to U_k . \square

Definition 4.4.8. Let Γ be a set of constraints. For each $(V, C) \in \Gamma$, let $U_C \subseteq V$ be a set of variables such that C is compressible to U_C and $C|_{U_C}$ is incompressible. Then, the **maximal compression** of Γ is $\Gamma_{\max} = \Gamma_{\text{comp}} \cup \Gamma_{\text{aux}}$, where $\Gamma_{\text{comp}} = \{(U_C, C|_{U_C}) | C \in \Gamma\}$ and

- $\{b\} \in \Gamma_{\text{aux}}$ (constant constraint) iff, for some $(V, C) \in \Gamma$, there exists $v \in V \setminus U_C$ that compresses by a constant to $b \in \mathbb{Z}_2$;
- $C_{=} = \{(0, 0), (1, 1)\} \in \Gamma_{\text{aux}}$ (equality constraint) iff, for some $(V, C) \in \Gamma$, there exists $v \in V \setminus U_C$ that compresses by equality; and
- $C_{\neq} = \{(0, 1), (1, 0)\} \in \Gamma_{\text{aux}}$ (negation constraint) iff, for some $(V, C) \in \Gamma$, there exists $v \in V \setminus U_C$ that compresses by negation.

Note that all the constraints in Γ_{comp} are incompressible, but the constraints in Γ_{aux} are compressible. However, they are important in making sure the hardness of the CSP is preserved, both in the quantum and classical cases.

Lemma 4.4.9. Let Γ be a set of constraints and let Γ_{\max} be its maximal compression. If Γ_{\max} satisfies one of the polymorphisms 0, 1, AND, OR, MAJ, or MIN, so does Γ .

By contrapositive and Schaefer's dichotomy theorem, we get that if $\text{CSP}(\Gamma)_{1,1}$ is NP-complete, so is $\text{CSP}(\Gamma_{\max})_{1,1}$.

Proof. Suppose first that Γ_{\max} satisfies the constant polymorphism 0. Then, by construction we know that if Γ compresses by a constant, it must compress to 0, and it cannot compress by negation, as the negation constraint does not satisfy a constant polymorphism. Thus, all the compressed variables compress by a constant to 0 or by equality. As such, we know that $0 \in C|_{U_C}$ implies that $0 \in C$, so Γ satisfies the constant polymorphism 0. The same argument holds for the constant polymorphism 1.

Next, suppose that Γ_{\max} satisfies the polymorphism AND. By construction, we know that Γ does not compress by negation, as the negation constraint does not satisfy AND.

As such, every variable must compress by equality or a constant. Thus, if $C|_{U_C}$ compressed in this way satisfies AND, so does C . The same argument holds for OR.

Suppose now that Γ_{\max} satisfies the majority polymorphism MAJ. If $C|_{U_C}$ satisfies MAJ, then so does C , as every assignment to the variables in U_C is constant, or equal to, or the negation of one of the remaining variables, and majority commutes with negation. So Γ must also satisfy MAJ. The same argument holds for MIN. \square

In what follows, we can always assume that we are working with a maximally compressed set of constraints, hence with either an incompressible constraint, a constant constraint, an equality constraint, or a negation constraint. Note that, of these constraints, only the incompressible constraints can have more than three satisfying assignments, which is a necessary condition to avoid satisfying the majority polymorphism. Now, we study the structure of incompressible TVF constraints and use it to construct commutativity gadgets.

Definition 4.4.10. *Let $G = (V, E_{00} \sqcup E_{11} \sqcup E_{01})$ be a TVF graph. The **constraint generated by G** , $C_{TVF}(G)$, is the set of all assignments to G . The **TVF completion of a constraint C** is $C_{TVF}(G_{TVF}(C))$.*

Lemma 4.4.11. *For any TVF graph G , $G_{TVF}(C_{TVF}(G)) = G$.*

Proof. Let $G' = G_{TVF}(C_{TVF}(G))$. By definition, G and G' have the same vertex sets. Let (u, v) be an ab edge of G . Equivalently, $(\phi(u), \phi(v)) \neq (a, b)$ for all $\phi \in C_{TVF}(G)$. By definition, this is equivalent to (u, v) being an ab edge of G' . \square

Lemma 4.4.12. *Let $G = (V, E_{00} \sqcup E_{11} \sqcup E_{01})$ be a complete TVF graph. Then, $|C_{TVF}(G)| \leq |V| + 1$.*

Proof. We proceed by induction on $|V|$. If $|V| = 0$, there are no edges, so no constraints, and only one assignment, the vacuous one. Hence $|C_{TVF}(G)| = 1 = |V| + 1$.

Now, suppose the induction hypothesis holds for $|V| \leq k$. Now let $|V| = k + 1$ and let $v \in V$. The vertex v is connected to every element of $V \setminus \{v\}$. Let $\ell \leq k$ be the number of edges that are 0 on v . Hence, given an assignment ϕ to G , if $\phi(v) = 0$, there are ℓ vertices v_1, \dots, v_ℓ whose value of ϕ is fixed; and if $\phi(v) = 1$, the value on the remaining $k - \ell$ vertices $v_{\ell+1}, \dots, v_k$ is fixed. As such,

$$\begin{aligned} |C_{TVF}(G)| &= |\{\phi \in C_{TVF}(G) | \phi(v) = 0\}| + |\{\phi \in C_{TVF}(G) | \phi(v) = 1\}| \\ &\leq |C_{TVF}(G \setminus \{v, v_1, \dots, v_\ell\})| + |C_{TVF}(G \setminus \{v, v_{\ell+1}, \dots, v_k\})| \\ &\leq (k - \ell + 1) + (\ell + 1) = k + 2 = |V| + 1, \end{aligned}$$

by induction hypothesis. □

Definition 4.4.13. Let $C \subseteq \mathbb{Z}_2^V$ be a constraint. The **tableau form** of C with respect to orderings $V = \{v_1, \dots, v_k\}$ and $C = \{\phi_1, \dots, \phi_n\}$ is the matrix

$$\begin{bmatrix} \phi_1(v_1) & \phi_1(v_2) & \cdots & \phi_1(v_k) \\ \phi_2(v_1) & \phi_2(v_2) & \cdots & \phi_2(v_k) \\ \vdots & & \ddots & \vdots \\ \phi_n(v_1) & \phi_n(v_2) & \cdots & \phi_n(v_k) \end{bmatrix}.$$

We say C has a tableau form M if M is the tableau form for some ordering of the variables and satisfying assignment.

We say that M is **upper triangular** if it is upper triangular as a matrix.

Lemma 4.4.14. Let $G = (V, E)$ be a directed complete graph with no multiple edges or loops, i.e. for all $x, y \in V$ distinct, either $(x, y) \in E$ or $(y, x) \in E$, but not both. Suppose that every vertex of G has an incoming edge, and $|V| \geq 3$. Then G has a cycle.

Proof. We proceed by induction on $|V|$. For the base case, we have $|V| = 3$. If there is a vertex $x \in V$ with two incoming edges $(y, x), (z, x) \in E$. Then, there is an edge between y and z , which we may assume without loss of generality is $(y, z) \in E$. Then, y has no incoming edges, which contradicts the hypothesis. As such, every vertex has must have one incoming edge and one outgoing edge. So G is a directed 3-cycle, and thus contains a cycle.

Now let $|V| > 3$. If every vertex of G has both an incoming and an outgoing edge, then it has a subgraph that is a cycle. Otherwise, there exists a vertex with only incoming edges. Consider the subgraph H of G with that vertex removed. Then, every vertex of H has at least one incoming edge and H has $|V| - 1$ vertices. By induction hypothesis, H has a cycle, and therefore so does G . □

Proposition 4.4.15. Let G be an incompressible complete TVF graph with no 00 or 11 edges. Then the constraint $C_{TVF}(G)$ has a tableau form

$$\begin{bmatrix} 1 & 1 & 1 & \cdots & 1 & 1 \\ 0 & 1 & 1 & \cdots & 1 & 1 \\ 0 & 0 & 1 & \cdots & 1 & 1 \\ \vdots & & & \ddots & & \vdots \\ 0 & 0 & 0 & \cdots & 1 & 1 \\ 0 & 0 & 0 & \cdots & 0 & 1 \\ 0 & 0 & 0 & \cdots & 0 & 0 \end{bmatrix}$$

As a consequence, for any TVF constraint $C \subseteq \mathbb{Z}_2^V$ whose TVF graph has no 00 or 11 edges, there is a tableau form of C that is a subset of the rows of the above matrix.

Proof. We construct an ordering on V and $C = C_{TVF}(G)$ recursively. We say that $G_{TVF}(C) = G = (V, E_{01})$ has a directed edge from x to y if $(y, x) \in E_{01}$. Then, note that there must always be a vertex $v_1 \in V$ such that all the edges incident to v_1 must be pointing outwards. Otherwise, every vertex has at least one edge pointing inwards. Then, due to Lemma 4.4.14, we know that G has a cycle, and due to Lemma 4.4.5, this cycle is a compressible subgraph. Then, G is compressible, a contradiction. As such, for each $v \neq v_1$, $(v_1, v) \in E_{01}$. Consider the assignment ϕ_1 to G such that $\phi_1(v_1) = 1$, then $\phi_1(v) = 1$ for all $v \in V$. As ϕ_1 is an assignment to G , it is in C by definition. For every other $\phi \in C$, we have $\phi(v_1) = 0$. Now, we can continue this recursively with the subgraph of G constructed by removing vertex v_1 . Suppose we have v_1, \dots, v_k and some ϕ_1, \dots, ϕ_k such that $\phi_j(v_i) = 0$ for $i < j$ and $\phi_j(v) = 1$ for all other $v \in V$, and for all other $\phi \in C$, $\phi(v_i) = 0$. Then, take $G_k = G \setminus \{v_1, \dots, v_k\}$. As above, there must be a vertex v_{k+1} with only outgoing edges. By maximality of C , there exists a unique $\phi_{k+1} \in C$ with $\phi_{k+1}(v_{k+1}) = 1$ and $\phi_{k+1}(v_i) = 0$ for $i < k+1$, and we must have $\phi_{k+1}(v) = 1$ for all $v \notin \{v_1, \dots, v_k\}$. We have that $\phi(v_{k+1}) = 0$ for all $\phi \in C \setminus \{\phi_1, \dots, \phi_{k+1}\}$. Hence, we can continue recursively until we have exhausted all the elements of V .

At the end, we have orderings $v_1, \dots, v_{|V|}$ and $\phi_1, \dots, \phi_{|V|}$ such that $\phi_j(v_i) = 1$ if $i \geq j$ and 0 otherwise. We can then construct the assignment $\phi_{|V|+1}(v) = 0$. This brings the number of elements of C to $|V| + 1$, the maximum possible by Lemma 4.4.12, finishing the proof. \square

Corollary 4.4.16. *Let Γ be a set of boolean TVF constraints. If $\text{CSP}(\Gamma)_{1,1}$ is NP-complete, then there exists at least one $C \in \Gamma_{\text{comp}}$ whose TVF graph has a 00 or 11 edge.*

Proof. Suppose otherwise that every edge of the TVF graph of every $C \in \Gamma_{\text{comp}}$ is a 01 edge. Then, the TVF completion of C has a tableau form as given in Proposition 4.4.15. Such constraints satisfy the majority polymorphism. In fact, if $i \geq j \geq k$ with $\phi_i, \phi_j, \phi_k \in C$, $\text{MAJ}(\phi_i, \phi_j, \phi_k) = \phi_j \in C$. Therefore, by Lemma 4.4.9, every constraint in Γ must also satisfy the majority polymorphism. Hence, $\text{CSP}(\Gamma)$ is in P by Schaefer's dichotomy theorem. By contrapositive, we get the desired result. \square

Lemma 4.4.17. *Suppose that $G = (V, E_{00} \sqcup E_{11} \sqcup E_{01})$ is an incompressible complete TVF graph. If G has no 00 edges, then $C_{TVF}(G)$ has an upper triangular tableau form.*

Proof. The proof follows the structure of Proposition 4.4.15; that is, we recursively construct an order on the variables and the constraints of $C = C_{TVF}(G)$ to get the wanted

form. For the base case, we first show that there is a vertex $v_1 \in V$ such that every edge of G incident to v_1 is 1 on it. Suppose otherwise that for every vertex of v , there is a 01 edge that is 0 on v . Start with an arbitrary vertex v and follow one of the 01 edges which is 0 on v to the next vertex. Then repeat this procedure. Because the graph is finite, eventually we visit a vertex twice. But this induces a cycle satisfying the conditions of Lemma 4.4.5, so G is compressible, a contradiction. Looking at the subgraph on the vertex set $V \setminus \{v_1\}$ we can apply the same reasoning to find a vertex v_2 . Then, continuing recursively, we find vertices $v_1, \dots, v_{|V|}$ such that the (unique) edge from v_i to v_j is 1 on v_i iff $i < j$. Thus, for any assignment $\phi \in C$, if $\phi(v_i) = 1$, then the value $\phi(v_j)$ is fixed for all $j > i$. As such, for each i , there is exactly one $\phi_i \in C$ such that $\phi_i(v_k) = 0$ for all $k < i$, and $\phi_i(v_i) = 1$. Finally, the zero assignment $\phi_{|V|+1}(v) = 0$ gives the remaining element of C . \square

Definition 4.4.18. Let $r : V \rightarrow W \cup \Sigma$. Given a map $\phi \in \Sigma^W$, the **augmented composition** of ϕ with r is $\phi \circ r \in \Sigma^V$ defined as

$$(\phi \circ r)(v) = \begin{cases} \phi(r(v)) & r(v) \in W \\ r(v) & r(v) \in \Sigma \end{cases}.$$

The **augmented pushforward** of a constraint $C \subseteq \Sigma^V$ by r is $r_*C = \{\phi \in \Sigma^W \mid \phi \circ r \in C\}$. We say that C **simulates** a constraint $C' \subseteq \Sigma^W$ if there exists $W' \supseteq W$ and $r : V \rightarrow W' \cup \Sigma$ such that $C' = r_*C|_W$.

Let $r : V \rightarrow W \cup \neg W \cup \mathbb{Z}_2$. Given a map $\phi \in \mathbb{Z}_2^W$, the **augmented composition (with negation)** of ϕ with r is $\phi \circ r \in \mathbb{Z}_2^V$ defined as

$$(\phi \circ r)(v) = \begin{cases} \phi(r(v)) & r(v) \in W \\ \neg\phi(w) & r(v) = \neg w \in \neg W \\ r(v) & r(v) \in \mathbb{Z}_2 \end{cases}.$$

The **augmented pushforward (with negation)** of a constraint $C \in \mathbb{Z}_2^V$ by r is $r_*C = \{\phi \in \mathbb{Z}_2^W \mid \phi \circ r \in C\}$. We say that C **simulates** $C' \subseteq \mathbb{Z}_2^W$ **with negation** if there exists $W' \supseteq W$ and $r : V \rightarrow W' \cup \neg W' \cup \mathbb{Z}_2$ such that $C' = r_*C|_W$.

Proposition 4.4.19. Suppose that $C \subseteq \mathbb{Z}_2^V$ is a TVF constraint that does not satisfy the majority polymorphism and whose TVF graph does not have a 00 edge. Then, C either simulates $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$ or $\{(1, 0, 1), (0, 1, 1), (0, 0, 0)\}$.

Proof. By Lemma 4.4.17, the TVF completion of C has an upper triangular tableau form, induced by orderings $V = \{v_1, \dots, v_{|V|}\}$ and $C_{TVF}(G_{TVF}(C)) = \{\phi_1, \dots, \phi_{|V|+1}\}$. Let

$I \subseteq [|V| + 1]$ be the set of indices of elements of $C_{TVF}(G_{TVF}(C))$ that are in C . By construction of the upper triangular tableau, for all rows $i, j, k \in [|V| + 1]$ there exists $r_{ijk} \in [|V| + 1]$ such that $\text{MAJ}(\phi_i, \phi_j, \phi_k) = \phi_{r_{ijk}}$. By hypothesis, there exist three rows $i < j < k$ in I such that $r_{ijk} \notin I$. First, we know that for all $l < j$, $\phi_j(v_l) = \phi_k(v_l) = 0$, so $\text{MAJ}(\phi_i, \phi_j, \phi_k)(v_l) = 0$. We know that $\phi_j(v_j) = 1$ and $\phi_k(v_j) = 0$. We claim that $\phi_i(v_j) = 0$ as well. In fact, if $\phi_i(v_j) = 1$, then by construction we know that $\phi_i(v_l) = \phi_j(v_l)$ for all $l \geq j$. Hence, the majority $\text{MAJ}(\phi_i, \phi_j, \phi_k) = \phi_j$, a contradiction.

Now consider two cases. Suppose first that there is some $j < h < k$ such that $\phi_i(v_h) = \phi_j(v_h) = 1$. We can also suppose that h is the smallest index satisfying this property. Then, we know that $\phi_i(v_l) = \phi_j(v_l)$ for $l \geq h$. As such, $r_{ijk} = h$, so $h \notin I$. Now, define a map $r : V \rightarrow \{x, y, z, z'\} \cup \mathbb{Z}_2$ as follows:

$$r(v_l) = \begin{cases} \phi_i(v_l) & \text{if } \phi_i(v_l) = \phi_j(v_l) = \phi_k(v_l) \\ x & \text{if } \phi_i(v_l) = 1 \text{ and } \phi_j(v_l) = \phi_k(v_l) = 0 \\ y & \text{if } \phi_j(v_l) = 1 \text{ and } \phi_i(v_l) = \phi_k(v_l) = 0 \\ z' & \text{if } \phi_k(v_l) = 1 \text{ and } \phi_i(v_l) = \phi_j(v_l) = 0 \\ z & \text{if } \phi_i(v_l) = \phi_j(v_l) = 1 \text{ and } \phi_k(v_l) = 0 \end{cases}.$$

Note that no other cases are possible as $\phi_i(v_l) = \phi_j(v_l)$ for all $l \geq h$, while $\phi_k(v_l) = 0$ for all $l < h$. Let $\phi \in \mathbb{Z}_2^{\{x, y, z, z'\}}$ be such that $\psi = \phi \circ r \in C$. Since $\psi(v_l) = 0$ for all $l < i$, $\psi = \phi_t$ for some $t \geq i$. First, we claim that if $\phi(x) = 1$, then $\phi(y) = 0$ and $\phi(z) = 1$. Since if $\phi(x) = 1$, then $\psi(v_i) = 1$, so $\psi = \phi_i$. As such, $\psi(v_j) = 0$ and $\psi(v_h) = 1$ as wanted. Next, suppose that $\phi(x) = 0$. Then, if $\phi(y) = 1$, we have that $\psi(v_j) = 1$ and $\psi(v_l) = 0$ for all $l < j$. As such, $\psi = \phi_j$ and $\psi(v_h) = 1$, giving that $\phi(z) = 1$. On the other hand, if $\phi(x) = \phi(y) = 0$ and $\phi(z) = 1$ we would have that $\psi(v_h) = 1$ and $\psi(v_l) = 0$ for all $l < h$. This does not correspond to any ϕ_t for $t \in I$ and therefore we must have $\phi(z) = 0$. Hence, restricting to the variables $\{x, y, z\}$, we find that C simulates $\{(1, 0, 1), (0, 1, 1), (0, 0, 0)\}$.

Now, suppose that $\phi_i(v_l)$ and $\phi_j(v_l)$ are not both 1 for all $j < l < k$. Note that we must have $\phi_i(v_k) = \phi_j(v_k) = 0$, as otherwise $\text{MAJ}(\phi_i, \phi_j, \phi_k) = \phi_k$. Let h be the minimal l such that two of $\phi_i(v_l), \phi_j(v_l), \phi_k(v_l)$ are equal to 1. Then, $r_{ijk} = h$ and $h \notin I$. Now, define $r : V \rightarrow \{x, y, z, z'\} \cup \mathbb{Z}_2$ as follows:

$$r(v_l) = \begin{cases} \phi_i(v_l) & \text{if } \phi_i(v_l) = \phi_j(v_l) = \phi_k(v_l) \\ x & \text{if } \phi_i(v_l) = 1 \text{ and } \phi_j(v_l) = \phi_k(v_l) = 0 \\ y & \text{if } \phi_j(v_l) = 1 \text{ and } \phi_i(v_l) = \phi_k(v_l) = 0 \\ z & \text{if } \phi_k(v_l) = 1 \text{ and } \phi_i(v_l) = \phi_j(v_l) = 0 \\ z' & \text{if only two of } \phi_i(v_l), \phi_j(v_l) \text{ and } \phi_k(v_l) \text{ are equal to 1} \end{cases}.$$

Let $\phi \in \mathbb{Z}_2^{\{x,y,z,z'\}}$ be such that $\psi = \phi \circ r \in C$. We claim that at most one of $\phi(x), \phi(y), \phi(z)$ can be 1. If $\phi(x) = 1$, then $\psi = \phi_i$, so $\phi(y) = \phi(z) = 0$. Next, if $\phi(y) = 1$ we must have $\phi(x) = 0$ to not contradict the above, and therefore $\phi = \phi_j$ so $\phi(z) = 0$. Finally, if $\phi(z) = 1$, we have by the above that $\phi(x) = \phi(y) = 0$. To complete the argument consider two cases. Suppose that there is no $\phi \in r_*C$ such that $\phi(x) = \phi(y) = \phi(z) = 0$. Then, via r , C simulates $r_*C|_{\{x,y,z\}} = \{(1,0,0), (0,1,0), (0,0,1)\}$. For the second case, suppose that there is such an element $\phi \in r_*C$. For this element, note that if $\phi(z') = 1$, $\psi = \phi_h$. As $h \notin I$, this implies that we must have $\phi(z') = 0$. Therefore, we have that r_*C is one of the following constraints: $\{(1,0,0,0), (0,1,0,1), (0,0,1,1), (0,0,0,0)\}$, $\{(1,0,0,1), (0,1,0,0), (0,0,1,1), (0,0,0,0)\}$, $\{(1,0,0,1), (0,1,0,1), (0,0,1,0), (0,0,0,0)\}$. Take $s : \{x,y,z,z'\} \cup \mathbb{Z}_2 \rightarrow \{x,y,z\} \cup \mathbb{Z}_2$, where $s(a) = a$ for $a \in \mathbb{Z}_2$ and $s(z') = z$. In the first case, $s(x) = 0, s(y) = x, s(z) = y$; in the second case $s(x) = x, s(y) = 0, s(z) = y$; and in the third case, $s(x) = x, s(y) = y, s(z) = 0$. In all three cases, via $s \circ r$, C simulates $\{(1,0,1), (0,1,1), (0,0,0)\}$. \square

Definition 4.4.20. Let $\phi \in \mathbb{Z}_2^V$ and $U \subseteq V$. The **negation** of ϕ at U is the map $\phi_{-U} \in \mathbb{Z}_2^V$ defined as

$$\phi_{-U}(v) = \begin{cases} -\phi(v) & v \in U \\ \phi(v) & \text{otherwise.} \end{cases}$$

The **negation** of a constraint $C \subseteq \mathbb{Z}_2^V$ at U is the constraint $C_{-U} = \{\phi_{-U} | \phi \in C\}$.

Lemma 4.4.21. Let $G = (V, E_{00} \sqcup E_{11} \sqcup E_{01})$ be an incompressible complete TVF graph. There exists $U \subseteq V$ such that $C_{TVF}(G)_{-U}$ has an upper triangular tableau form.

Proof. The proof again follows the structure of Proposition 4.4.15; that is, we recursively construct an order on the variables and the constraints of $C = C_{TVF}(G)$ to get the wanted form. Here, we also construct the subset $U \subseteq V$ at the same time. For the base case, we first show that there is a vertex $v_1 \in V$ such that all edges of G incident to v_1 are 0 on v or 1 on v . Suppose otherwise that for every vertex v of G , there is an edge that is 0 on v and an edge that is 1 on v . Start with an arbitrary vertex u_1 and follow any edge to the next vertex u_2 . This edge is b on u_2 for some $b \in \mathbb{Z}_2$, so we can pick an edge that is $-b$ on u_2 , connecting to the next vertex u_3 . Then we repeat this procedure. Because the graph is finite, eventually we visit a vertex twice. But this induces a cycle satisfying the conditions of Lemma 4.4.5, so G is compressible, a contradiction. If every edge incident to v_1 is 0 on v_1 , we pass to the negated constraint $C_{-\{v_1\}}$, so that now the edges are 1 on v_1 . Looking at the subgraph on the vertex set $V \setminus \{v_1\}$ we can apply the same reasoning to find a vertex v_2 . Then, continuing recursively, we find vertices $v_1, \dots, v_{|V|}$ such that the

edge from v_i to v_j in the TVF graph of C_{-U} is 1 on v_i if $i < j$. Thus, for any assignment $\phi \in C_{-U}$, if $\phi(v_i) = 1$, then $\phi(v_j)$ can only take one value for all $j > i$. As such, for each i , there is exactly one $\phi_i \in C_{-U}$ such that $\phi_i(v_k) = 0$ for all $k < i$ and $\phi_i(v_i) = 1$. Finally, the zero assignment $\phi_{|V|+1}(v) = 0$ gives the remaining element of C_{-U} . \square

Proposition 4.4.22. *Suppose that $C \subseteq \mathbb{Z}_2^V$ is a TVF constraint that does not satisfy the majority polymorphism. Then, C simulates $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$ with negation.*

Proof. Let $U \subseteq V$ be such that the TVF completion of C_{-U} has an upper triangular tableau form induced by orderings $V = \{v_1, \dots, v_{|V|}\}$ and $C_{TVF}(G_{TVF}(C))_{-U} = \{\phi_1, \dots, \phi_{|V|+1}\}$, guaranteed by Lemma 4.4.21. Let $I \subseteq [|V| + 1]$ be the set of indices of elements of $C_{TVF}(G_{TVF}(C))_{-U}$ that are in C_{-U} . By construction, we know that for all $i, j, k \in [|V| + 1]$, there exists r_{ijk} such that $\text{MAJ}(\phi_i, \phi_j, \phi_k) = \phi_{r_{ijk}}$. Since C does not satisfy the majority polymorphism, neither does C_{-U} , and hence there exist $i < j < k$ in I such that $r_{ijk} \notin I$. Define $r : V \rightarrow \{x, y, z, \neg x, \neg y, \neg z\} \cup \mathbb{Z}_2$ as follows:

$$r(v_l) = \begin{cases} \phi_i(v_l) & \text{if } \phi_i(v_l) = \phi_j(v_l) = \phi_k(v_l) \text{ and } v_l \notin U \\ \neg \phi_i(v_l) & \text{if } \phi_i(v_l) = \phi_j(v_l) = \phi_k(v_l) \text{ and } v_l \in U \\ x & \text{if } (\phi_i(v_l) = 1, \phi_j(v_l) = \phi_k(v_l) = 0, v_l \notin U) \\ & \text{or } (\phi_i(v_l) = 0, \phi_j(v_l) = \phi_k(v_l) = 1, v_l \in U) \\ \neg x & \text{if } (\phi_i(v_l) = 0, \phi_j(v_l) = \phi_k(v_l) = 1, v_l \notin U) \\ & \text{or } (\phi_i(v_l) = 1, \phi_j(v_l) = \phi_k(v_l) = 0, v_l \in U) \\ y & \text{if } (\phi_j(v_l) = 1, \phi_k(v_l) = \phi_i(v_l) = 0, v_l \notin U) \\ & \text{or } (\phi_j(v_l) = 0, \phi_k(v_l) = \phi_i(v_l) = 1, v_l \in U) \\ \neg y & \text{if } (\phi_j(v_l) = 0, \phi_k(v_l) = \phi_i(v_l) = 1, v_l \notin U) \\ & \text{or } (\phi_j(v_l) = 1, \phi_k(v_l) = \phi_i(v_l) = 0, v_l \in U) \\ z & \text{if } (\phi_k(v_l) = 1, \phi_i(v_l) = \phi_j(v_l) = 0, v_l \notin U) \\ & \text{or } (\phi_k(v_l) = 0, \phi_i(v_l) = \phi_j(v_l) = 1, v_l \in U) \\ \neg z & \text{if } (\phi_k(v_l) = 0, \phi_i(v_l) = \phi_j(v_l) = 1, v_l \notin U) \\ & \text{or } (\phi_k(v_l) = 1, \phi_i(v_l) = \phi_j(v_l) = 0, v_l \in U) \end{cases}.$$

First, let $\phi = (0, 0, 0) \in \mathbb{Z}_2^{\{x, y, z\}}$. We claim that $\phi \notin r_*C$. Let $\psi = \phi \circ r$. Then, for $v_l \notin U$, $\psi(v_l) = 1$ if two of $\phi_i(v_l), \phi_j(v_l), \phi_k(v_l)$ are 1; and for $v_l \in U$, $\psi(v_l) = 1$ if two of $\phi_i(v_l), \phi_j(v_l), \phi_k(v_l)$ are 0. Hence, $\psi_{-U} = \phi_{r_{ijk}}$, and thus $\psi \notin C$. So $(0, 0, 0) \notin r_*C$. Now suppose that $\phi \in r_*C$ with $\phi(x) = 1$. We claim that $\phi(y) = \phi(z) = 0$. Since $\phi(x) = 1$,

$\phi \circ r = \phi_i$. Since $r_{ijk} \neq j$, we know that $\phi_i(v_j) = 0$ so $\phi(y) = 0$. If there exists $j < h < k$ such that $\phi_i(v_h) = \phi_j(v_h) = 1$, then $\neg\phi(z) = 1$ so $\phi(z) = 0$. Otherwise, we know that $r_{ijk} \neq k$, so $\phi_i(v_k) = \phi_j(v_k) = 0$, giving $\phi(z) = 0$ as well. Next, suppose $\phi(y) = 1$. By the above, we must have $\phi(x) = 0$. Then, we have that $\phi \circ r = \phi_j$, so by the same two-case argument as above $\phi(z) = 0$. Finally, note that it is possible to have $\phi(z) = 1$ and $\phi(x) = \phi(y) = 0$, as $\phi \circ r = \phi_k$ in that case. By the above, this is the only possible $\phi \in r_*C$ with $\phi(z) = 1$. As such, we have that $r_*C = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$, as wanted. \square

4.4.3 The general commutativity gadget

In this section, we show that the simplification and simulation arguments of the previous section are quantum-sound, and use this to construct a general commutativity gadget modelled on that of Lemma 4.4.1.

First, we want to show the constraints in the maximal compression can be expressed in terms of the original set of constraints, in a quantum-sound way. Then, we can work only with the maximal compression to construct gadgets.

Lemma 4.4.23. *Let Γ be a set of constraints and let Γ_{\max} be its maximal compression. For each CS $S = (X, \{(V_i, r_{i*}C_i)_{i=1}^m\}) \in \text{CSP}(\Gamma_{\max})$ and probability distribution π on $[m]$, there exists a CS $S' = (X', \{(V'_i, r'_{i*}C'_i)_{i=1}^m\}) \in \text{CSP}(\Gamma)$ such that there exists a L -homomorphism $\alpha : \mathcal{A}_{c-v}(S, \pi) \rightarrow \mathcal{A}_{c-v}(S', \pi)$, where $L = \max_{(V,C) \in \Gamma} |V|$.*

Proof. By definition of Γ_{\max} , for each C_i , there exists a $C'_i \in \Gamma$ such that $C'_i|_{V_{C_i}} = C_i$. Then, let $X' = X \cup \{x_{i,v} | v \in V_{C'_i} \setminus V_{C_i}\}$, let $V'_i = V_i \cup \{x_{i,v} | v \in V_{C'_i} \setminus V_{C_i}\}$, and let $r'_i|_{V_{C_i}} = r_i$ and for $v \in V_{C'_i} \setminus V_{C_i}$, $r'_i(v) = x_{i,v}$. Now, define α on $\mathcal{A}(V_i, r_{i*}C_i)$ via $\alpha(\Phi_{V_i, \phi}) = \sum_{\psi \in r'_{i*}C'_i, \psi|_{V_i} = \phi} \Phi_{V'_i, \psi}$

and as identity on $\mathbb{C}\mathbb{Z}_2^{*X}$. Then,

$$\begin{aligned}
& \alpha \left(\sum_{i=1}^m \frac{\pi(i)}{|V_i|} \sum_{x \in V_i, \phi \in r_{i*} C_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \right) \\
&= \sum_{i=1}^m \frac{\pi(i)}{|V_i|} \sum_{\substack{x \in V_i, \phi \in r_{i*} C_i \\ \psi \in r_{i*} C'_i, \psi|_{V_i} = \phi}} |\Phi_{V'_i, \psi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \\
&= \sum_{i=1}^m \frac{\pi(i)}{|V_i|} \sum_{x \in V_i, \phi \in r'_{i*} C'_i} |\Phi_{V'_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \\
&\leq L \sum_{i=1}^m \frac{\pi(i)}{|V'_i|} \sum_{x \in V'_i, \phi \in r'_{i*} C'_i} |\Phi_{V'_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2.
\end{aligned}$$

□

Next, we show that if a constraint can be simulated, it can also be done in a quantum-sound way, assuming that there are constraints that set variables to constants. Similarly, we also show that using a negation constraint, any negation of a constraint can be simulated in a quantum-sound way.

Lemma 4.4.24. *Suppose that $C \subseteq \mathbb{Z}_2^V$ simulates $C' \subseteq \mathbb{Z}_2^W$ via $r : V \rightarrow W' \cup \mathbb{Z}_2$. Consider the BCS $S = (W' \cup \{x_0, x_1\}, \{(V_i, C_i)\}_{i=1}^3)$, where $V_1 = W' \cup \{x_0, x_1\}$, $C_1 = s_* C$ with $s(v) = r(v)$ for $r(v) \in W'$ and $s(v) = x_{r(v)}$ for $r(v) \notin W'$, $V_2 = \{x_0\}$, $C_2 = \{0\}$, $V_3 = \{x_1\}$, and $C_3 = \{1\}$. There exists a $24(|W'| + 2)$ -homomorphism $\alpha : \mathcal{A}_{c-v}((W, \{(W, C')\}), \mathfrak{u}_1) \rightarrow \mathcal{A}_{c-v}(S, \mathfrak{u}_3)$.*

Proof. Let α be the natural embedding $\mathcal{A}_{c-v}((W, \{(W, C')\})) \hookrightarrow \mathcal{A}_{c-v}(S)$. First, note that as $C' = r_* C|_W$, $\Phi_{W, \phi} = \sum_{\psi \in r_* C, \psi|_W = \phi} \Phi_{W', \psi}$. Therefore,

$$\sum_{x \in W, \phi \in C'} |\Phi_{W, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \leq \sum_{x \in W, \phi \in r_* C} |\Phi_{W', \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2.$$

Next, note that any $\phi \in r_* C$ admits an extension to $\phi' \in s_* C$ by setting $\phi'(x_0) = 0$, $\phi'(x_1) = 1$, and $\phi'|_{W'} = \phi$. Then, $\Phi_{V_1, \phi'} = \Phi_{W', \phi} \Pi_0(\sigma_1(x_0)) \Pi_1(\sigma_1(x_1))$. Noticing that

$\Pi_0(\sigma_2(x_0)) = 1$, we get

$$\begin{aligned}
|\Pi_0(\sigma_1(x_0)) - 1|^2 &\leq 2 |\Pi_0(\sigma_1(x_0)) - \Pi_0(\sigma'(x_0))|^2 + 2 |\Pi_0(\sigma'(x_0)) - 1|^2 \\
&\lesssim 2 \sum_b |\Pi_b(\sigma_1(x_0))(1 - \Pi_b(\sigma'(x_0)))|^2 + 2 |\Pi_0(\sigma_2(x_0))(1 - \Pi_0(\sigma'(x_0)))|^2 \\
&= 2 \sum_{\phi \in C_1} |\Phi_{V_1, \phi}(1 - \Pi_{\phi(x_0)}(\sigma'(x_0)))|^2 + 2 \sum_{x \in V_2, \phi \in C_2} |\Phi_{V_2, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2.
\end{aligned}$$

In the same way,

$$|\Pi_1(\sigma_1(x_1)) - 1|^2 \lesssim 2 \sum_{\phi \in C_1} |\Phi_{V_1, \phi}(1 - \Pi_{\phi(x_1)}(\sigma'(x_1)))|^2 + 2 \sum_{x \in V_3, \phi \in C_3} |\Phi_{V_3, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2.$$

Putting these together,

$$\begin{aligned}
|\Pi_0(\sigma_1(x_0))\Pi_1(\sigma_1(x_1)) - 1|^2 &\leq 2 |\Pi_0(\sigma_1(x_0))\Pi_1(\sigma_1(x_1)) - \Pi_0(\sigma_1(x_0))|^2 + 2 |\Pi_0(\sigma_1(x_0)) - 1|^2 \\
&\leq 2 |\Pi_1(\sigma_1(x_1)) - 1|^2 + 2 |\Pi_0(\sigma_1(x_0)) - 1|^2 \\
&\lesssim 4 \sum_{i=1}^3 \sum_{\substack{\phi \in C_i \\ x=x_0, x_1}} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2.
\end{aligned}$$

As such, we get the result

$$\begin{aligned}
\sum_{x \in W, \phi \in C'} |\Phi_{W, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 &\leq \sum_{x \in W, \phi \in r_* C} |\Phi_{W', \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \\
&\leq \sum_{x \in W, \phi \in r_* C} |(\Phi_{V_1, \phi'} - \Phi_{W', \phi}(\Pi_0(\sigma_1(x_0))\Pi_1(\sigma_1(x_1)) - 1))(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \\
&\lesssim 2 \sum_{x \in W, \phi \in r_* C} |\Phi_{V_1, \phi'}(1 - \Pi_{\phi'(x)}(\sigma'(x)))|^2 + 2|W| |\Pi_0(\sigma_1(x_0))\Pi_1(\sigma_1(x_1)) - 1|^2 \\
&\lesssim 2 \sum_{x \in W', \phi \in C_1} |\Phi_{V_1, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 + 8|W| \sum_{i=1}^3 \sum_{\substack{\phi \in C_i \\ x=x_0, x_1}} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \\
&\leq 8|W| \sum_{i=1}^3 \sum_{x \in V_i, \phi \in C_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \\
&\leq 8|W||V_1| \sum_{i=1}^3 \frac{1}{|V_i|} \sum_{x \in V_i, \phi \in C_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2.
\end{aligned}$$

□

Lemma 4.4.25. *Let $C \subseteq \mathbb{Z}_2^V$ be a constraint and $U \subseteq V$, and suppose V is ordered as $V = \{v_1, \dots, v_{|V|}\}$. Let $I = \{i | v_i \in U\}$. Consider the constraint system*

$$S = (\{v_1, w_1, \dots, v_{|V|}, w_{|V|}\}, \{(V_i, C_i)\}_{i=0}^{|V|})$$

where $V_0 = V$, $C_0 = C$, and $V_i = \{v_i, w_i\}$ and $C_i = C_{\neq}$ for $i > 0$. There exists a $4(|V| + 1)$ -homomorphism $\alpha : \mathcal{A}_{c-v}((V, \{(V, C_{-U})\}), \mathbf{u}_1) \rightarrow \mathcal{A}_{c-v}(S, \mathbf{u}_{|V|+1})$.

Proof. Define α on $\mathcal{A}(V, C)$ as $\alpha(v_i) = \sigma_0(v_i)$ if $i \notin I$ and $\alpha(v_i) = -\sigma_0(v_i)$ if $i \in I$; let $\alpha(\sigma'(v_i)) = \sigma'(v_i)$ if $i \notin I$ and $\alpha(\sigma'(v_i)) = \sigma'(w_i)$ if $i \in I$. First, we get that

$$\begin{aligned} \alpha \left(\sum_{x \in V, \phi \in C_{-U}} |\Phi_{V, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \right) \\ &= \sum_{i \in I, \phi \in C_{-U}} |\Phi_{V_0, \phi_{-U}}(1 - \Pi_{\phi(v_i)}(\sigma'(w_i)))|^2 + \sum_{i \notin I, \phi \in C_{-U}} |\Phi_{V_0, \phi_{-U}}(1 - \Pi_{\phi(v_i)}(\sigma'(v_i)))|^2 \\ &= \sum_{i \in I, \phi \in C} |\Phi_{V_0, \phi}(1 + \Pi_{\phi(v_i)}(\sigma'(w_i)))|^2 + \sum_{i \notin I, \phi \in C} |\Phi_{V_0, \phi}(1 - \Pi_{\phi(v_i)}(\sigma'(v_i)))|^2 \end{aligned}$$

Next, in $\mathcal{A}_{c-v}(S)$, we have that for $i > 0$, $\sigma_i(v_i) = -\sigma_i(w_i)$, so

$$\begin{aligned} |\sigma'(v_i) + \sigma'(w_i)|^2 &\leq 4|\sigma_i(v_i) - \sigma'(v_i)|^2 + 4|\sigma_i(w_i) - \sigma'(w_i)|^2 \\ &\leq 16 \sum_{x \in V_i, \phi \in C_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2. \end{aligned}$$

Putting these together,

$$\begin{aligned} &\sum_{i \in I, \phi \in C} |\Phi_{V_0, \phi}(1 + \Pi_{\phi(v_i)}(\sigma'(w_i)))|^2 \\ &\leq 2 \sum_{i \in I, \phi \in C} \left(|\Phi_{V_0, \phi}(1 - \Pi_{\phi(v_i)}(\sigma'(v_i)))|^2 + \frac{1}{4} |\Phi_{V_0, \phi}(\sigma'(v_i) + \sigma'(w_i))|^2 \right) \\ &\leq 2 \sum_{i \in I, \phi \in C} |\Phi_{V_0, \phi}(1 - \Pi_{\phi(v_i)}(\sigma'(v_i)))|^2 + 4 \sum_{i \in I} \sum_{x \in V_i, \phi \in C_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2. \end{aligned}$$

□

Now, in order to construct the constant and negation constraints, we appeal to the structure of TVF graphs of incompressible constraints.

Lemma 4.4.26. *Let $C \subseteq \mathbb{Z}_2^V$ be an incompressible TVF constraint whose TVF graph has a 11 edge between vertices $u, v \in V$. Let $r : V \rightarrow V \setminus \{v\}$ be defined $r(v) = u$ and $r(w) = w$ for all $w \neq v$. Then, there is a $(|V| - 1)$ -homomorphism $\alpha : \mathcal{A}_{c-v}(\{\{x\}, \{(\{x\}, \{0\})\}\}, \mathfrak{u}_1) \rightarrow \mathcal{A}_{c-v}(\{V \setminus \{v\}, \{(V \setminus \{v\}, r_*C)\}\}, \mathfrak{u}_1)$.*

Proof. Noting that $\sigma_1(x) = 1$, define α by $\alpha(\sigma'(x)) = \sigma'(u)$. Then, we have that

$$\alpha\left(|1 - \Pi_0(\sigma'(x))|^2\right) = |1 - \Pi_0(\sigma'(u))|^2.$$

On the other hand, noting that $\phi \in r_*C$ implies that $\phi(u) \neq 1$,

$$\begin{aligned} \sum_{w \in V \setminus \{v\}, \phi \in r_*C} |\Phi_{V \setminus \{v\}, \phi}(1 - \Pi_{\phi(w)}(\sigma'(w)))|^2 &\geq \sum_{\phi \in r_*C} |\Phi_{V \setminus \{v\}, \phi}(1 - \Pi_{\phi(w)}(\sigma'(u)))|^2 \\ &= |1 - \Pi_0(\sigma'(u))|^2, \end{aligned}$$

giving the wanted result. \square

Lemma 4.4.27. *Let $C \subseteq \mathbb{Z}_2^V$ and r be as in the previous lemma, and let $C' \subseteq \mathbb{Z}_2^W$ be an incompressible nonempty constraint that does not contain the all-0 assignment. In particular, there exists $W_0 \subset W$ and $\phi_0 \in C'$ such that $\phi_0(w) = 0$ if $w \in W_0$ and $\phi_0(w) = 1$ otherwise. Consider the constraint system $S = (V \setminus \{v\} \cup \{u'\}, \{(V \setminus \{v\}, r_*C), (\{u, u'\}, s_*C')\})$ where $s(w) = u$ if $w \in W_0$ and $s(w) = u'$ otherwise. Then, there exists a $8(|V| - 1)$ -homomorphism $\alpha : \mathcal{A}_{c-v}(\{\{y\}, \{(\{y\}, \{1\})\}\}, \mathfrak{u}_1) \rightarrow \mathcal{A}_{c-v}(S, \mathfrak{u}_2)$.*

Proof. As in the previous lemma, $\sigma_1(y) = -1$, so we define α by $\alpha(\sigma'(y)) = \sigma'(u')$, giving that $\alpha\left(|1 - \Pi_1(\sigma'(y))|^2\right) = |1 - \Pi_1(\sigma'(u'))|^2$. As before, we have

$$\sum_{\substack{w \in V \setminus \{v\} \\ \phi \in r_*C}} |\Phi_{V \setminus \{v\}, \phi}(1 - \Pi_{\phi(w)}(\sigma'(w)))|^2 \geq |1 - \Pi_0(\sigma'(u))|^2.$$

Next, note that $\Phi_{\{u, u'\}, (0,0)} = 0$ by hypothesis so $\Pi_1(\sigma_2(u')) \geq \Pi_0(\sigma_2(u))$. Then,

$$\begin{aligned} |1 - \Pi_1(\sigma'(u'))|^2 &\leq 2|1 - \Pi_1(\sigma_2(u'))|^2 + 2|\Pi_1(\sigma_2(u')) - \Pi_1(\sigma'(u'))|^2 \\ &\leq 2|1 - \Pi_0(\sigma_2(u))|^2 + \frac{1}{2}|\sigma_2(u') - \sigma'(u')|^2 \\ &\leq 4|1 - \Pi_0(\sigma'(u))|^2 + |\sigma_2(u) - \sigma'(u)|^2 + |\sigma_2(u') - \sigma'(u')|^2 \\ &\leq 4 \sum_{w \in V \setminus \{v\}, \phi \in r_*C} |\Phi_{V \setminus \{v\}, \phi}(1 - \Pi_{\phi(w)}(\sigma'(w)))|^2 \\ &\quad + 4 \sum_{w \in \{u, u'\}, \phi \in s_*C'} |\Phi_{\{u, u'\}, \phi}(1 - \Pi_{\phi(2)}(\sigma'(w)))|^2. \end{aligned} \quad \square$$

Lemma 4.4.28. *Let $C \subseteq \mathbb{Z}_2^V$ and $C' \subseteq \mathbb{Z}_2^{V'}$ be incompressible TVF constraints whose TVF graphs have a 00 edge between $u, v \in C$ and a 11 edge between $u', v' \in C'$, respectively. Consider the constraint systems $S = (\{x, y\}, \{(\{x, y\}, C_{\neq})\})$ and $S' = (X, \{(V_1, C), (V_2, r_*C')\})$, where $X = V \cup V' \setminus \{u', v'\}$, $V_1 = V$, $V_2 = V' \setminus \{u', v'\} \cup \{u, v\}$, and $r : V' \rightarrow V_2$ is a bijection such that $r_2(u') = u$ and $r_2(v') = v$. Then, there exists a $4 \max\{|V|, |V'|\}$ -homomorphism $\alpha : \mathcal{A}_{c-v}(S, \mathbb{u}_1) \rightarrow \mathcal{A}_{c-v}(S', \mathbb{u}_2)$.*

Proof. Define α as $\alpha(\sigma'(x)) = \alpha(\sigma_1(x)) = -\alpha(\sigma_1(y)) = \sigma'(u)$ and $\alpha(\sigma'(y)) = \sigma'(v)$. Then, $\alpha(\Phi_{\{x,y\},(0,1)}) = \Pi_0(\sigma'(u))$ and $\alpha(\Phi_{\{x,y\},(1,0)}) = \Pi_1(\sigma'(u))$, so

$$\alpha \left(\sum_{z \in \{x,y\}, \phi \in C_{\neq}} |\Phi_{\{x,y\},\phi}(1 - \Pi_{\phi(z)}(\sigma'(z)))|^2 \right) = |\Pi_0(\sigma'(u))\Pi_0(\sigma'(v))|^2 + |\Pi_1(\sigma'(u))\Pi_1(\sigma'(v))|^2.$$

We have that $\Pi_0(\sigma_1(u))\Pi_0(\sigma_1(v)) = \Pi_1(\sigma_2(u))\Pi_1(\sigma_2(v)) = 0$, so

$$\begin{aligned} |\Pi_0(\sigma'(u))\Pi_0(\sigma'(v))|^2 &= |\Pi_0(\sigma_1(u))\Pi_0(\sigma_1(v)) - \Pi_0(\sigma'(u))\Pi_0(\sigma'(v))|^2 \\ &\leq 2 |\Pi_0(\sigma_1(u)) - \Pi_0(\sigma'(u))|^2 + 2 |\Pi_0(\sigma_1(v)) - \Pi_0(\sigma'(v))|^2 \\ &\leq 2 \sum_{w \in \{u,v\}, \phi \in C} |\Phi_{V_1,\phi}(1 - \Pi_{\phi(w)}(\sigma'(w)))|^2 \\ &\leq 2 \sum_{w \in V_1, \phi \in C} |\Phi_{V_1,\phi}(1 - \Pi_{\phi(w)}(\sigma'(w)))|^2. \end{aligned}$$

By a similar argument, $|\Pi_1(\sigma'(u))\Pi_1(\sigma'(v))|^2 \leq 2 \sum_{w \in V_2, \phi \in C} |\Phi_{V_2,\phi}(1 - \Pi_{\phi(w)}(\sigma'(w)))|^2$, giving the wanted result. \square

Now, we show that at least one of the constraints necessary to construct a version of the basic commutativity gadget can be simulated.

Theorem 4.4.29. *Let Γ be a set of boolean TVF constraints such that $\text{CSP}(\Gamma)_{1,1}$ is NP-complete. Then, there exists a CS $S = (X, \{(V_i, r_{i*}C_i)\}_{i=1}^m) \in \text{CSP}(\Gamma)$ and a poly(L)-homomorphism*

$$\alpha : \mathcal{A}_{c-v}(S_0, \mathbb{u}_1) \rightarrow \mathcal{A}_{c-v}(S, \mathbb{u}_m),$$

where $L = \max_{(V,C) \in \Gamma} |V|$ and $S_0 = (\{x, y, z\}, \{(\{x, y, z\}, r_*C)\})$ for C being one of $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$, $\{(1, 0, 1), (0, 1, 1), (0, 0, 0)\}$, $\{(1, 1, 1), (0, 0, 1), (0, 1, 0)\}$, $\{(0, 1, 1), (1, 0, 1), (1, 1, 0)\}$, and $r : [3] \rightarrow \{x, y, z\}$ a bijection.

Proof. Since $\text{CSP}(\Gamma)_{1,1}$ is NP-complete, we know that $\text{CSP}(\Gamma_{\max})_{1,1}$ is NP-complete by Lemma 4.4.9. In particular, by Schaefer's dichotomy theorem, we know that there exist $C_1, C_2 \in \text{CSP}(\Gamma_{\max})$ such that C_1 does not satisfy the majority polymorphism, and C_2 does not satisfy the constant 0 polymorphism ($0 \notin C_2$ and $C_2 \neq \emptyset$). Further, by Corollary 4.4.16, there exists an incompressible $C_3 \in \Gamma_{\max}$ such that the TVF graph of C_3 has a 00 or a 11 edge. Suppose we are in the second case. Suppose also that no constraint in Γ_{\max} has a TVF graph with a 00 edge. Then, we know by Proposition 4.4.19 that C_1 simulates either $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$ or $\{(1, 0, 1), (0, 1, 1), (0, 0, 0)\}$. Let this be C . First, using Lemma 4.4.24, there is a constraint system $S_1 \in \text{CSP}(C_1, \{0\}, \{1\})$ and a poly(L)-homomorphism $\alpha_1 : \mathcal{A}_{c-v}(S_0) \rightarrow \mathcal{A}_{c-v}(S_1)$, where uniform probability distributions on the constraints are implied. Now, using Lemma 4.4.26 and Lemma 4.4.27, we can express the constraints $\{0\}$ and $\{1\}$, respectively, with constraint systems in terms of C_2 and C_3 . This induces a poly(L)-homomorphism $\alpha_2 : \mathcal{A}_{c-v}(S_1) \rightarrow \mathcal{A}_{c-v}(S_2)$ where $S_2 \in \text{CSP}(C_1, C_2, C_3) \subseteq \text{CSP}(\Gamma_{\max})$. To finish, note that using Lemma 4.4.23, we can construct a CS $S \in \text{CSP}(\Gamma)$ and a poly(L)-homomorphism $\alpha_3 : \mathcal{A}_{c-v}(S_2) \rightarrow \mathcal{A}_{c-v}(S)$. Taking $\alpha = \alpha_3 \circ \alpha_2 \circ \alpha_1$ finishes the proof in this case.

In the case that there are no 11 edges in the constraints of Γ_{\max} , we can do an identical argument with the labels 0 and 1 reversed. Then, we have that C may be taken to be $\{(1, 1, 1), (0, 0, 1), (0, 1, 0)\}$ or $\{(0, 1, 1), (1, 0, 1), (1, 1, 0)\}$, the negation of the possible C s from the previous case.

Now, suppose that there are incompressible constraints $C_3, C_4 \in \Gamma_{\max}$ such that the TVF graph of C_3 has a 11 edge and the TVF graph of C_4 has a 00 edge. Take $C = \{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$. Then, by Proposition 4.4.22, C_1 simulates C with negation. In particular, there exists $U \subseteq V_{C_1}$ such that $(C_1)_{-U}$ simulates C . As before, using Lemma 4.4.24, Lemma 4.4.26, and Lemma 4.4.27, there is a constraint system $S_1 \in \text{CSP}((C_1)_{-U}, C_2, C_3)$ and a poly(L)-homomorphism $\alpha_1 : \mathcal{A}_{c-v}(S_0) \rightarrow \mathcal{A}_{c-v}(S_1)$. Now, using Lemma 4.4.25, there exists a CS $S_2 \in \text{CSP}(C_1, C_{\neq}, C_2, C_3)$ and a poly(L)-homomorphism $\alpha_2 : \mathcal{A}_{c-v}(S_1) \rightarrow \mathcal{A}_{c-v}(S_2)$. Next, by applying Lemma 4.4.28, we see there exists a CS $S_3 \in \text{CSP}(C_1, C_2, C_3, C_4) \subseteq \text{CSP}(\Gamma_{\max})$ and a poly(L)-homomorphism $\alpha_3 : \mathcal{A}_{c-v}(S_2) \rightarrow \mathcal{A}_{c-v}(S_3)$. Finally, as before, we use Lemma 4.4.23, to construct a CS $S \in \text{CSP}(\Gamma)$ and a poly(L)-homomorphism $\alpha_4 : \mathcal{A}_{c-v}(S_3) \rightarrow \mathcal{A}_{c-v}(S)$, and take $\alpha = \alpha_4 \circ \alpha_3 \circ \alpha_2 \circ \alpha_1$, finishing the proof. \square

To finish this section, we construct the commutativity gadget.

Corollary 4.4.30. *Let Γ be a set of Boolean TVF constraints such that $\text{CSP}(\Gamma)_{1,1}$ is NP-complete. Then, there exists a BCS $S = (X, \{(V_i, r_{i*} C_i)\}_{i=1}^m) \in \text{CSP}(\Gamma)$ and variables*

$x, y \in X$ such that S is satisfiable for any assignments to x, y in \mathbb{Z}_2 and

$$\|[\sigma'(x), \sigma'(y)]\|^2 \leq \text{poly}(L) \frac{1}{m} \sum_{i=1}^m \frac{1}{|V_i|} \sum_{\phi \in C_i, z \in V_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(z)}(\sigma'(z)))|^2$$

in $\mathcal{A}_{c-v}(S)$, where $L = \max_{(V, C) \in \Gamma} |V|$.

Proof. We begin with the the BCS from Lemma 4.4.1, $B = \{X, \{(V_i, r_{i*} C)\}_{i=1}^3\}$, where C is one of $\{(1, 0, 0), (0, 1, 0), (0, 0, 1)\}$, $\{(1, 0, 1), (0, 1, 1), (0, 0, 0)\}$, $\{(1, 1, 1), (0, 0, 1), (0, 1, 0)\}$, $\{(0, 1, 1), (1, 0, 1), (1, 1, 0)\}$. Now, we can apply the $\text{poly}(L)$ -homomorphism α from Theorem 4.4.29 to each of the constraints in B to get $S \in \text{CSP}(\Gamma)$. Letting $x_0, y_0 \in X$ be the approximately-commuting variables in B , we take $\sigma'(x) = \alpha(\sigma'(x_0))$ and $\sigma'(y) = \alpha(\sigma'(y_0))$ to get the wanted result. \square

Theorem 4.4.31 (Part 2 of Theorem 4.2.4). *Let Γ be a set of boolean TVF constraints such that $\text{CSP}(\Gamma)_{1,1}$ is NP-complete. Then, there exists $s \in [0, 1)$ such that $\text{SuccinctCSP}_{c-v}(\Gamma)_{1,s}^*$ is RE-complete.*

Proof. Noting that the set of constraints $\Gamma \cup \{\mathbb{Z}_2^2\}$ is non-TVF and NP-complete, we have by Theorem 4.3.3 that $\text{SuccinctCSP}_{c-v}(\Gamma \cup \{\mathbb{Z}_2^2\})_{1,s}^*$ is RE-complete for some $s < 1$. To complete the proof, note that we can use Corollary 4.4.30 to replace any instance of an empty constraint \mathbb{Z}_2^2 by a gadget composed of the constraints in Γ while preserving completeness and constant soundness. \square

4.4.4 Oracularizability of boolean TVF CSPs

Lemma 4.4.32. *Suppose $S = (X, \{(V_i, C_i)\}_{i=1}^m)$ is a TVF BCS. The identity map is a $(16L^2 + 1)$ -homomorphism $\mathcal{A}_{a+comm}(S, \pi) \rightarrow \mathcal{A}_a(S, \pi)$, where $L = \max_i |V_i|$.*

Proof. By the TVF property, for every constraint i and pair of variables $x, y \in V_i$, there exist $a, b \in \mathbb{Z}_2$ such that $\phi \notin C_i$ if $\phi(x) = a$ and $\phi(y) = b$. Then,

$$\begin{aligned} \|[\Pi_a(x), \Pi_b(y)]\|_\tau^2 &\leq 4\|\Pi_a(x)\Pi_b(y)\|_\tau^2 = 4\tau(\Pi_a(x)\Pi_b(y)) \\ &= 4 \sum_{\substack{\phi \in \mathbb{Z}_2^{V_i} \text{ s.t.} \\ \phi(x)=a, \phi(y)=b}} \|\Phi_{V_i, \phi}\|_\tau^2 \leq 4 \sum_{\phi \notin C_i} \|\Phi_{V_i, \phi}\|_\tau^2. \end{aligned}$$

Since $\Pi_{-a}(x) = 1 - \Pi_a(x)$, and similarly for y , this upper bound holds for all $a, b \in \mathbb{Z}_2$. Then, we get that

$$\begin{aligned}
\sum_{r \in \mathcal{A}_a(S)} (\mu_{a,\pi}(r) + \mu_{\text{comm},\pi}(r)) r^* r &= \sum_{i=1}^m \pi(i) \left(\sum_{\phi \notin C_i} \|\Phi_{V_i, \phi}\|_\tau^2 + \sum_{\substack{x, y \in V_i \\ a, b \in \mathbb{Z}_2}} \|[\Pi_a(x), \Pi_b(y)]\|_\tau^2 \right) \\
&\leq \sum_{i=1}^m \pi(i) (1 + 16|V_i|^2) \sum_{\phi \notin C_i} \|\Phi_{V_i, \phi}\|_\tau^2 \\
&\leq (1 + 16L^2) \sum_{i=1}^m \pi(i) \sum_{\phi \notin C_i} \|\Phi_{V_i, \phi}\|_\tau^2. \quad \square
\end{aligned}$$

Theorem 4.4.33 (Part 1 of Corollary 4.2.6). *Let Γ be a set of boolean TVF constraints such that $\text{CSP}(\Gamma)_{1,1}$ is NP-complete. Then, there exists $s \in [0, 1)$ such that $\text{SuccinctCSP}_a(\Gamma)_{1,s}^*$ is RE-complete.*

Proof. Due to Lemmas 3.3.7 and 4.4.32, there is a mapping $\mathcal{A}_{c-v}(S, \pi) \rightarrow \mathcal{A}_a(S, \pi)$ for all $S \in \text{CSP}(\Gamma)$ that preserves the constant soundness; and due to Lemmas 3.3.6 and 3.3.8 there is a C -homomorphism in the other direction, preserving completeness. Hence, the result follows from Theorem 4.4.31. \square

4.5 Hardness of 2-CSPs

4.5.1 The case of 3-colouring

In this section, we show the RE-completeness of 3-colouring in the assignment and constraint-variable settings. Our arguments are exactly those of [37], but adapted to the context of imperfect completeness, where we can phrase them in the language of weighted algebras.

First, we show the oracularizability of 3-colouring, and use this to construct a mapping from the assignment algebra to the constraint-variable algebra for 3-colouring instances.

Lemma 4.5.1. *Let \mathcal{A} be a $*$ -algebra. Let $x, y \in \mathcal{A}$ be order-3 unitaries. Then, for any $a, b \in \mathbb{Z}_3$,*

$$|[\Pi_a(x), \Pi_b(y)]|^2 \lesssim 16 \sum_{c \in \mathbb{Z}_3} |\Pi_c(x) \Pi_c(y)|^2.$$

This can be seen as a robust version of Lemma 2 from [37].

Proof. If $a = b$, then $|\Pi_a(x), \Pi_b(y)|^2 \lesssim 4 |\Pi_a(x)\Pi_a(y)|^2$, giving the result. Else, without loss of generality, suppose $a = 0$ and $b = 1$, and write $x_i = \Pi_i(x)$ and $y_i = \Pi_i(y)$. We have that

$$\begin{aligned} [x_0, y_1] &= x_0y_1 - y_1x_0 = (y_0 + y_1 + y_2)x_0y_1 - y_1x_0(y_0 + y_1 + y_2) \\ &= y_0x_0y_1 - y_1x_0y_0 + y_2x_0y_1 - y_1x_0y_2 \\ &= y_0x_0y_1 - y_1x_0y_0 + y_2(1 - x_1 - x_2)y_1 - y_1(1 - x_1 - x_2)y_2 \\ &= y_0x_0y_1 - y_1x_0y_0 - y_2x_1y_1 + y_1x_1y_2 - y_2x_2y_1 + y_1x_2y_2. \end{aligned}$$

Thus, by triangle inequality,

$$\begin{aligned} |[x_0, y_1]|^2 &\leq 8(|y_0x_0y_1|^2 + |y_1x_0y_0|^2 + |y_2x_1y_1|^2 + |y_1x_1y_2|^2 + |y_2x_2y_1|^2 + |y_1x_2y_2|^2) \\ &\lesssim 16(|x_0y_0|^2 + |x_1y_1|^2 + |x_2y_2|^2). \quad \square \end{aligned}$$

Lemma 4.5.2. *Suppose $S = (X, \{(V_i, r_{i*} \neq \mathbb{Z}_3)\}_{i=1}^m)$ is a 3-colouring instance. Then, the identity map on $\mathcal{A}_a(S)$ is a 145-homomorphism $\mathcal{A}_{a+comm}(S, \pi) \rightarrow \mathcal{A}_a(S, \pi)$.*

Proof. Using lemma 4.5.1, we get that

$$\begin{aligned} \sum_{r \in \mathcal{A}_a(S)} \mu_{comm, \pi}(r) |r|^2 &= \sum_{\substack{x, y \in X \\ a, b \in \mathbb{Z}_3}} \sum_{\substack{i \\ x, y \in V_i}} \pi(i) |\Pi_a(x), \Pi_b(y)|^2 = \sum_{\substack{V_i \\ x, y \in V_i}} \pi(i) \sum_{a, b \in \mathbb{Z}_3} |\Pi_a(x_i), \Pi_b(y_i)|^2 \\ &\lesssim 144 \sum_i \pi(i) \sum_{c \in \mathbb{Z}_3} |\Pi_c(x_i)\Pi_c(y_i)|^2 \\ &= 144 \sum_{r \in \mathcal{A}_a(S)} \mu_{a, \pi} |r|^2. \end{aligned}$$

Therefore, we see that $\sum_{r \in \mathcal{A}_a(S)} (\mu_{a, \pi}(r) + \mu_{comm, \pi}(r)) |r|^2 \leq 145 \sum_{r \in \mathcal{A}_a(S)} \mu_{a, \pi} |r|^2$ \square

Now, we show the soundness of the prism graph construction of [37] in the assignment algebra, which we will use as a commutativity gadget.

Lemma 4.5.3. *Let \mathcal{A} be a $*$ -algebra with tracial state τ , and let $x, y, z \in \mathcal{A}$ be order-3 unitaries. Then,*

$$\sum_{a \in \mathbb{Z}_3} \|\Pi_a(x) + \Pi_a(y) + \Pi_a(z) - 1\|_\tau^2 = 2 \sum_{a \in \mathbb{Z}_3} \|\Pi_a(x)\Pi_a(y)\|_\tau^2 + \|\Pi_a(y)\Pi_a(z)\|_\tau^2 + \|\Pi_a(z)\Pi_a(x)\|_\tau^2.$$




Figure 4.4: The triangular constraint system in Lemma 4.5.3. Each vertex corresponds to a variable and each edge corresponds to a 3-colouring constraint.

This is a robust version of Lemma 3 from [37]. The right-hand side corresponds to the defect of a constraint system where the three variables x, y, z are connected by 3-coloring constraints in a triangular arrangement. See Figure 4.4 for a graphical representation.

Proof. As in the proof of Lemma 4.5.1, write $x_i = \Pi_i(x)$, $y_i = \Pi_i(y)$, and $z_i = \Pi_i(z)$. Then, for $i = 0, 1, 2$, we can expand

$$\begin{aligned}
\sum_i \|x_i + y_i + z_i - 1\|_\tau^2 &= \sum_i \tau((x_i + y_i + z_i - 1)^2) \\
&= \sum_i \tau(x_i + x_i y_i + x_i z_i + y_i x_i + y_i + y_i z_i + z_i x_i \\
&\quad + z_i y_i + z_i - 2(x_i + y_i + z_i) + 1) \\
&= \sum_i (2\tau(x_i y_i + y_i z_i + z_i x_i) + \tau(1 - (x_i + y_i + z_i))) \\
&= 2 \sum_i (\tau(x_i y_i) + \tau(y_i z_i) + \tau(z_i x_i)) + 3 - \sum_i x_i - \sum_i y_i - \sum_i z_i \\
&= 2 \sum_i \|x_i y_i\|_\tau^2 + \|y_i z_i\|_\tau^2 + \|z_i x_i\|_\tau^2. \quad \square
\end{aligned}$$

Lemma 4.5.4. *Let \mathcal{A} be a $*$ -algebra with tracial state τ , and let $x, y, z, x', y', z' \in \mathcal{A}$ be order-3 unitaries. Writing as previously $x_i = \Pi_i(x)$ and similarly for the other variables, we have*

$$\begin{aligned}
\sum_{i,j} \|[x_i, y'_j]\|_\tau^2 &\leq 6240 \sum_i (\|x_i y_i\|_\tau^2 + \|y_i z_i\|_\tau^2 + \|z_i x_i\|_\tau^2 + \|x'_i y'_i\|_\tau^2 + \|y'_i z'_i\|_\tau^2 + \|x'_i y'_i\|_\tau^2 \\
&\quad + \|x_i x'_i\|_\tau^2 + \|y_i y'_i\|_\tau^2 + \|z_i z'_i\|_\tau^2).
\end{aligned}$$

This is a version of Lemma 4 from [37] with imperfect completeness. As in the previous lemma, the right-hand side corresponds to the defect of a 3-colouring constraint system. Here, the variables are arranged as the vertices of a triangular prism, as illustrated in Figure 4.5.

Proof. Consider first the case $i = j = 0$. Using [37, Lemma 4], we see that

$$\begin{aligned}
z_0 z'_0 x_0 &= z_0(1 - x'_0 - y'_0)x_0 + z_0(x'_0 + y'_0 + z'_0 - 1)x_0 \\
&= z_0(1 - y'_0)x_0 - z_0 x'_0 x_0 + z_0(x'_0 + y'_0 + z'_0 - 1)x_0 \\
&= (1 - x_0 - y_0)(x_0 - y'_0 x_0) - z_0 x'_0 x_0 + z_0(x'_0 + y'_0 + z'_0 - 1)x_0 \\
&\quad + (x_0 + y_0 + z_0 - 1)(1 - y'_0)x_0 \\
&= -y_0 x_0 - y'_0 x_0 + x_0 y'_0 x_0 + y_0 y'_0 x_0 - z_0 x'_0 x_0 + z_0(x'_0 + y'_0 + z'_0 - 1)x_0 \\
&\quad + (x_0 + y_0 + z_0 - 1)(1 - y'_0)x_0,
\end{aligned}$$

and taking the adjoint

$$x_0 z'_0 z_0 = -x_0 y_0 - x_0 y'_0 + x_0 y'_0 x_0 + x_0 y'_0 y_0 - x_0 x'_0 z_0 + x_0(x'_0 + y'_0 + z'_0 - 1)z_0 + x_0(1 - y'_0)(x_0 + y_0 + z_0 - 1).$$

Hence, the commutator

$$\begin{aligned}
[x_0, y'_0] &= x_0 y'_0 - y'_0 x_0 \\
&= -x_0 y_0 - x_0 z'_0 z_0 + x_0 y'_0 y_0 - x_0 x'_0 z_0 + x_0(x'_0 + y'_0 + z'_0 - 1)z_0 \\
&\quad + x_0(1 - y'_0)(x_0 + y_0 + z_0 - 1) + y_0 x_0 + z_0 z'_0 x_0 - y_0 y'_0 x_0 + z_0 x'_0 x_0 \\
&\quad - z_0(x'_0 + y'_0 + z'_0 - 1)x_0 - (x_0 + y_0 + z_0 - 1)(1 - y'_0)x_0,
\end{aligned}$$

so the norm is bounded

$$\begin{aligned}
\|[x_0, y'_0]\|_\tau^2 &\leq 32(\|x_0 y_0\|_\tau^2 + \|z_0 z'_0\|_\tau^2 + \|y_0 y'_0\|_\tau^2 + \|x_0 x'_0\|_\tau^2 \\
&\quad + \|x'_0 + y'_0 + z'_0 - 1\|_\tau^2 + \|x_0 + y_0 + z_0 - 1\|_\tau^2).
\end{aligned}$$

Now, by symmetry and using Lemma 4.5.3,

$$\begin{aligned}
\sum_i \|[x_i, y'_i]\|_\tau^2 &\leq 32 \sum_i (\|x_i y_i\|_\tau^2 + \|z_i z'_i\|_\tau^2 + \|y_i y'_i\|_\tau^2 + \|x_i x'_i\|_\tau^2 + \|x'_i + y'_i + z'_i - 1\|_\tau^2 \\
&\quad + \|x_i + y_i + z_i - 1\|_\tau^2) \\
&= 32 \sum_i (\|z_i z'_i\|_\tau^2 + \|y_i y'_i\|_\tau^2 + \|x_i x'_i\|_\tau^2 + 2\|x'_i y'_i\|_\tau^2 + 2\|y'_i z'_i\|_\tau^2 + 2\|z'_i x'_i\|_\tau^2 \\
&\quad + 3\|x_i y_i\|_\tau^2 + 2\|y_i z_i\|_\tau^2 + 2\|z_i x_i\|_\tau^2).
\end{aligned}$$

Now, consider the case $i = 0, j = 1$. Again using [37, Lemma 4],

$$\begin{aligned} y_2'x_2'x_0 &= y_2'(1 - x_0' - x_1')x_0 = (1 - y_0' - y_1')(x_0 - x_1'x_0) - y_2'x_0'x_0 \\ &= x_0 - y_0'x_0 - y_1'x_0 - x_1'x_0 + y_0'x_1'x_0 + y_1'x_1'x_0 - y_2'x_0'x_0, \end{aligned}$$

and taking the adjoint $x_0x_2'y_2' = x_0 - x_0y_0' - x_0y_1' - x_0x_1' + x_0x_1'y_0' + x_0x_1'y_1' - x_0x_0'y_2'$. Hence,

$$\begin{aligned} [x_0, y_1'] &= x_0y_1' - y_1'x_0 \\ &= -x_0x_2'y_2' - x_0y_0' - x_0x_1' + x_0x_1'y_0' + x_0x_1'y_1' - x_0x_0'y_2' \\ &\quad + y_2'x_2'x_0 + y_0'x_0 + x_1'x_0 - y_0'x_1'x_0 - y_1'x_1'x_0 + y_2'x_0'x_0 \\ &= -[x_0, y_0'] - [x_0, x_1'] - x_0x_2'y_2' + y_2'x_2'x_0 + x_0x_1'y_1' - y_1'x_1'x_0 \\ &\quad - x_0x_0'y_2' + y_2'x_0'x_0 + x_0[x_1', y_0'] + [x_0, y_0']x_1' - y_0'[x_1', x_0]. \end{aligned}$$

Taking the norm, and then using Lemma 4.5.1,

$$\begin{aligned} \|[x_0, y_1']\|_\tau^2 &\leq 16(2\|[x_0, y_0']\|_\tau^2 + \|[x_0, x_1']\|_\tau^2 + 2\|x_2'y_2'\|_\tau^2 + 2\|x_1'y_1'\|_\tau^2 + 2\|x_0x_0'\|_\tau^2 \\ &\quad + \|[x_1', y_0']\|_\tau^2 + \|[x_1', x_0]\|_\tau^2) \\ &\leq 16(2\|[x_0, y_0']\|_\tau^2 + 2\|x_2'y_2'\|_\tau^2 + 2\|x_1'y_1'\|_\tau^2 + 2\|x_0x_0'\|_\tau^2 \\ &\quad + 16 \sum_i \|x_i'y_i'\|_\tau^2 + 32 \sum_i \|x_ix_i'\|_\tau^2). \end{aligned}$$

By symmetry,

$$\begin{aligned} \sum_{i,j} \|[x_i, y_j']\|_\tau^2 &\leq \sum_i (65\|[x_i, y_i']\|_\tau^2 + 1664\|x_i'y_i'\|_\tau^2 + 3136\|x_ix_i'\|_\tau^2) \\ &\leq \sum_i (5824\|x_i'y_i'\|_\tau^2 + 4160\|y_i'z_i'\|_\tau^2 + 4160\|x_i'y_i'\|_\tau^2 \\ &\quad + 6240\|x_iy_i\|_\tau^2 + 4160\|y_iz_i\|_\tau^2 + 4160\|z_ix_i\|_\tau^2 \\ &\quad + 5216\|x_ix_i'\|_\tau^2 + 2080\|y_iy_i'\|_\tau^2 + 2080\|z_iz_i'\|_\tau^2). \quad \square \end{aligned}$$

Definition 4.5.5. The *triangular prism graph* is the graph illustrated in Figure 4.5, i.e. $G_{prism} = (V_{prism}, E_{prism})$ where $V_{prism} = \{x, y, z, x', y', z'\}$ and

$$E_{prism} = \{\{x, y\}, \{y, z\}, \{z, x\}, \{x', y'\}, \{y', z'\}, \{z', x'\}, \{x, x'\}, \{y, y'\}, \{z, z'\}\}.$$

Corollary 4.5.6. Let $S = (X, \{(V_i, r_{i*} \neq z_3)\}_{i=1}^{m_0} \cup \{(V_i, \mathbb{Z}_3^V)\}_{i=m_0+1}^m)$ be a 3-ary 2-CS and let π be a probability distribution on $[m]$. Write $V_i = \{x_i, y_i'\}$. Then, define $S' =$




Figure 4.5: The triangular prism constraint system in Lemma 4.5.4. Each vertex corresponds to a variable and each edge correspond to a 3-colouring constraint.

$(X', \{(V_i, r_{i*} \neq \mathbb{Z}_3)\}_{i=1}^{m_0} \cup \{(V_{ie}, r_{ie*} \neq \mathbb{Z}_3)\}_{i \in [m] \setminus [m_0], e \in E_{prism}})$, where the new set of variables $X' = X \cup \{y_i, y'_i, z_i, z'_i | i = m_0 + 1, \dots, m\}$ and $V_{ie} = \{\alpha_i, \beta_i\}$, with $\{\alpha, \beta\} = e$, and $r_{ie} : [2] \rightarrow V_{ie}$ is a bijection. Let π' be the probability distribution $\pi'(i) = \pi(i)$ if $i \in [m_0]$ and $\pi'(ie) = \frac{\pi(i)}{9}$ otherwise. Then, for any trace τ on $\mathcal{A}_a(S', \pi')$, there exists a trace τ' on $\mathcal{A}_{c-v}(S, \pi)$ such that $\text{def}(\tau') \leq C \text{def}(\tau)$ for some universal constant $C > 0$.

The constraint system S' is constructed by replacing every empty constraint in S by the triangular prism gadget from Lemma 4.5.4.

We can also use Lemmas 3.3.6 and 3.3.8 and Lemmas 3.3.7 and 4.5.2 to relate traces on $\mathcal{A}_a(S', \pi')$ and $\mathcal{A}_{c-v}(S', \pi')$.

Proof. First, via Lemma 4.5.1, for each $i = 1, \dots, m_0$, we have that

$$\sum_{a, b \in \mathbb{Z}_3} \|\Pi_a(x_i), \Pi_b(y'_i)\|_\tau^2 \leq 144 \sum_{c \in \mathbb{Z}_3} \|\Pi_c(x_i) \Pi_c(y'_i)\|_\tau^2 = 144 \sum_{\phi \notin r_{i*} \neq \mathbb{Z}_3} \|\Phi_{V_i, \phi}\|_\tau^2$$

Also, using Lemma 4.5.4, we have that for each $i = m_0 + 1, \dots, m$,

$$\begin{aligned} \sum_{a, b \in \mathbb{Z}_3} \|\Pi_a(x_i), \Pi_b(y'_i)\|_\tau^2 &\leq 6240 \sum_{c \in \mathbb{Z}_3} (\|\Pi_c(x_i) \Pi_c(y_i)\|_\tau^2 + \|\Pi_c(y_i) \Pi_c(z_i)\|_\tau^2 + \|\Pi_c(z_i) \Pi_c(x_i)\|_\tau^2 \\ &\quad + \|\Pi_c(x'_i) \Pi_c(y'_i)\|_\tau^2 + \|\Pi_c(y'_i) \Pi_c(z'_i)\|_\tau^2 + \|\Pi_c(x'_i) \Pi_c(y'_i)\|_\tau^2 \\ &\quad + \|\Pi_c(x_i) \Pi_c(x'_i)\|_\tau^2 + \|\Pi_c(y_i) \Pi_c(y'_i)\|_\tau^2 + \|\Pi_c(z_i) \Pi_c(z'_i)\|_\tau^2) \\ &= 6240 \sum_{\substack{e \in E_{prism} \\ \phi \notin r_{ie*} \neq \mathbb{Z}_3}} \|\Phi_{V_{ie}, \phi}\|_\tau^2. \end{aligned}$$


Therefore, as $\mathcal{A}_a(S)$ is a subalgebra of $\mathcal{A}_a(S')$, the defect

$$\begin{aligned} \text{def}(\tau|_{\mathcal{A}_a(S)}; \mu_{a,\pi} + \mu_{a,comm}) &= \sum_{i=1}^{m_0} \pi(i) \sum_{\varphi \notin r_{i*} \neq \mathbb{Z}_3} \|\Phi_{V_i, \phi}\|_{\tau}^2 + \sum_{i=1}^m \pi(i) \sum_{x, y \in V_i, a, b \in \mathbb{Z}_3} \|[\Pi_a(x), \Pi_b(y)]\|_{\tau}^2 \\ &\leq 145 \sum_{i=1}^{m_0} \pi(i) \sum_{\varphi \notin r_{i*} \neq \mathbb{Z}_3} \|\Phi_{V_i, \phi}\|_{\tau}^2 + 6240 \sum_{i=m_0+1}^m \pi(i) \sum_{\substack{e \in E_{prism} \\ \phi \notin r_{ie*} \neq \mathbb{Z}_3}} \|\Phi_{V_{ie}, \phi}\|_{\tau}^2 \\ &\leq 56160 \text{def}(\tau). \end{aligned}$$

Now, due to Lemma 3.3.7, there exists a trace τ' on $\mathcal{A}_{c-v}(S, \pi)$ such that $\text{def}(\tau) \leq 56160 \text{poly}(k^L) \text{def}(\tau)$, which is a constant as $k = 3, L = 2$. \square

Theorem 4.5.7 (Part 3 of Theorem 4.2.4 and Part 2 of Corollary 4.2.6). *There exists a constant $s \in [0, 1)$ such that $\text{SuccinctCSP}_a(\{\neq_{\mathbb{Z}_3}\})_{1,s}^*$ and $\text{SuccinctCSP}_{c-v}(\{\neq_{\mathbb{Z}_3}\})_{1,s}^*$ are RE-complete.*

Proof. We proceed similarly to Theorem 4.4.31. By Theorem 4.3.3, we know that the promise problem $\text{SuccinctCSP}_{c-v}(\{\neq_{\mathbb{Z}_3}, \mathbb{Z}_3^2\})_{1,s_0}^*$ is RE-complete. Next, using Corollary 4.5.6, we can replace the commutation constraints by gadgets over the assignment algebra, and find that $\text{SuccinctCSP}_a(\{\neq_{\mathbb{Z}_3}\})_{1,s_1}^*$ is RE-complete. Finally, using Lemmas 3.3.7 and 4.5.2 and Lemmas 3.3.6 and 3.3.8, we find that $\text{SuccinctCSP}_{c-v}(\{\neq_{\mathbb{Z}_3}\})_{1,s_2}^*$ is also RE-complete. \square



(a) Succinct entangled 3-colouring

(b) Entangled 3-colouring

Figure 4.6: Transitions in complexity based on soundness parameter for entangled 3-colouring

An interesting consequence of part 3 of Theorem 4.2.4 is that there exists a constant soundness parameter $s \in [0, 1)$ such that the class of succinct graph 3-colouring games is RE-complete with entanglement. On the other hand, Culf, Mousavi, and Spirig show that entangled graph 3-colouring with $s = 0.864$ is in P, and thus the succinct version with this soundness is in EXP [19]. This situation is illustrated in Figure 4.6. The question of what the complexity is when s is between these two values is an interesting open problem.

4.5.2 The case of 2-CSP(k)

Definition 4.5.8. Define the CSP 2-CSP(k) = CSP(Γ), where $\Gamma = \{C \subseteq \mathbb{Z}_k^2\}$. This corresponds to the set of all 2-CSPs over an alphabet of size k . Define the corresponding promise problems as 2-CSP(k) $_{c,s}$ = CSP(Γ) $_{c,s}$, Succinct-2-CSP(k) $_{c,s}$ = SuccinctCSP(Γ) $_{c,s}$, 2-CSP $_w(k)$ $_{c,s}^*$ = CSP $_w(\Gamma)$ $_{c,s}^*$, Succinct-2-CSP $_w(k)$ $_{c,s}^*$ = SuccinctCSP $_w(\Gamma)$ $_{c,s}^*$, where $w \in \{c - c, c - v, a, a + comm\}$.

Next, we relate the assignment algebra for the language of all 2-CSPs to a language where the c-v algebra is hard.

Proposition 4.5.9. Let $k \geq 3$ and set $C = \{(x_1, \dots, x_k) \in \mathbb{Z}_2^k \mid \exists! i. x_i = 1\}$. Note that $|C| = k$, so for $a \in \mathbb{Z}_k$, let $c_a \in C$ be the element with 1 in the a -th position. Consider some BCS $S = (X, \{(V_i, r_{i*}C)\}_{i=1}^m) \in \text{CSP}(\{C\})$ and a probability distribution π on $[m]$. Define a k -ary 2-CS $S' = (Y, \{(W_{ix}, D_x)\}_{i \in [m], x \in V_i})$ by $Y = X \cup \{y_i \mid i \in [m]\}$, $W_{ix} = \{y_i, x\}$, and $D_x = \{(a, 0) \mid f(x) = 0, f \circ r_i = c_a\} \cup \{(a, b) \mid f(x) = 1, f \circ r_i = c_a, b \neq 0\} \subseteq \mathbb{Z}_k^2$. Let $\pi'(i, x) = \frac{\pi(i)}{|V_i|}$. Then, there is a trace τ on $\mathcal{A}_{c-v}(S, \pi)$ with $\text{def}(\tau) = \varepsilon$ if and only if there is a trace τ' on $\mathcal{A}_a(S', \pi')$ with $\text{def}(\tau') = \varepsilon$.

Proof. Let τ be a trace on $\mathcal{A}_{c-v}(S)$, and $\tau = \rho \circ \varphi$ be its GNS representation. Let χ be the representation of $\mathcal{A}_a(S')$ defined by

$$\chi(\Pi_a(y_i)) = \begin{cases} \varphi(\Phi_{V_i, f}) & \text{if there is } f \in r_{i*}C. f \circ r_i = c_a \\ 0 & \text{otherwise,} \end{cases}$$

$\chi(\Pi_0(x)) = \varphi(\Pi_0(\sigma'(x)))$, $\chi(\Pi_1(x)) = \varphi(\Pi_1(\sigma'(x)))$, and $\chi(\Pi_a(x)) = 0$ for $a \neq 0, 1$. Take

$\tau' = \rho \circ \chi$. Then, the defect

$$\begin{aligned}
\text{def}(\tau') &= \sum_{i \in [m], x \in V_i} \pi'(i, x) \sum_{(a,b) \notin D_x} \|\Pi_a(y_i) \Pi_b(x)\|_{\tau'}^2 = \sum_{i \in [m], x \in V_i} \frac{\pi(i)}{|V_i|} \sum_{(a,b) \notin D_x} \tau'(\Pi_a(y_i) \Pi_b(x)) \\
&= \sum_{i \in [m], x \in V_i} \frac{\pi(i)}{|V_i|} \left(\sum_{\substack{a. (a,0) \notin D_x \\ f \in r_{i*} C. \text{for } i=c_a}} \tau(\Phi_{V_i, f} \Pi_0(\sigma'(x))) + \sum_{\substack{a. (a,1) \notin D_x, \\ f \in r_{i*} C. \text{for } i=c_a}} \tau(\Phi_{V_i, f} \Pi_1(\sigma'(x))) \right) \\
&= \sum_{i \in [m], x \in V_i} \frac{\pi(i)}{|V_i|} \sum_{\phi \in r_{i*} C} \tau(\Phi_{V_i, \phi} (1 - \Pi_{\phi(x)}(\sigma'(x)))) = \text{def}(\tau).
\end{aligned}$$

Conversely, suppose that τ' is a trace on $\mathcal{A}_a(S')$ with GNS representation $\tau' = \rho' \circ \varphi'$. Then, define the representation χ' of $\mathcal{A}_{c-v}(S)$ by $\chi'(\Phi_{V_i, f}) = \sum_{a. f \circ r_i = c_a} \varphi'(\Pi_a(y_i))$, $\chi'(\Pi_0(\sigma'(x))) = \varphi'(\Pi_0(x))$, and $\chi'(\Pi_1(\sigma'(x))) = \sum_{a \neq 0} \varphi'(\Pi_a(x))$. Let $\tau = \rho' \circ \chi'$. By a similar calculation as above, we have that the defect

$$\begin{aligned}
\text{def}(\tau) &= \sum_{i \in [m], x \in V_i} \frac{\pi(i)}{|V_i|} \sum_{\phi \in r_{i*} C} \tau(\Phi_{V_i, \phi} (1 - \Pi_{\phi(x)}(\sigma'(x)))) \\
&= \sum_{i \in [m]} \pi'(i, x) \left(\sum_{\substack{x \in V_i, f \in r_{i*} C \\ f(x)=0}} \tau(\Phi_{V_i, f} \Pi_1(\sigma'(x))) + \sum_{\substack{x \in V_i, f \in r_{i*} C \\ f(x)=1}} \tau(\Phi_{V_i, f} \Pi_0(\sigma'(x))) \right) \\
&= \sum_{i \in [m], a \in \mathbb{Z}_k} \pi'(i, x) \left(\sum_{\substack{x \in V_i, \\ f \circ r_i = c_a, f(x)=0}} \sum_{b \neq 0} \tau'(\Pi_a(y_i) \Pi_b(x)) + \sum_{\substack{x \in V_i, \\ f \circ r_i = c_a, f(x)=1}} \tau'(\Pi_a(y_i) \Pi_0(x)) \right) \\
&= \sum_{i \in [m], x \in V_i} \pi'(i, x) \sum_{(a,b) \notin D_x} \tau'(\Pi_a(y_i) \Pi_b(x)) = \text{def}(\tau'). \quad \square
\end{aligned}$$

Theorem 4.5.10 (Part 3 of Corollary 4.2.6). *There exists $s \in [0, 1)$ such that the promise problem $\text{Succinct-2-CSP}_a(k)_{1,s}^*$ is RE-complete.*

Proof. Let C be as in Proposition 4.5.9. By Theorem 4.4.31, we know that the promise problem $\text{SuccinctCSP}_{c-v}(\{C\})_{1,s}^*$ is RE-complete. But, by Proposition 4.5.9, there is a value-preserving mapping from instances of the $\text{SuccinctCSP}_{c-v}(\{C\})_{1,s}^*$ to instances of the assignment problem $\text{Succinct-2-CSP}_a(k)_{1,s}^*$, completing the proof. \square

4.6 Constraint-variable to constraint-constraint for CSPs

Proposition 4.6.1. *Let Γ be a set of k -ary constraints with $(V_0, C_0) \in \Gamma$ and $v \in V_0$ such that for all $a \in \mathbb{Z}_k$ there exists $\phi \in C_0$ such that $\phi(v) = a$. Let $S = (X, \{(V_i, r_{i*}C_i)\}_{i=1}^m) \in \text{CSP}(\Gamma)$, and let π be a probability distribution on $[m]$. Then, there exists a CS $S' = (X', \{(V'_i, r'_{i*}C'_i)\}_{i=1}^{m'}) \in \text{CSP}(\Gamma)$ and a probability distribution on $[m']$, such that there is a 2-homomorphism $\alpha : \mathcal{A}_{c-v}(S, \pi) \rightarrow \mathcal{A}_{c-v}(S', \pi')$ and a $\frac{1}{2}$ -homomorphism $\beta : \mathcal{A}_{c-v}(S', \pi') \rightarrow \mathcal{A}_{c-v}(S, \pi)$. Also, there exists a probability distribution π'' on $[m] \times [m]$ and a $\frac{1}{2}$ -homomorphism $\gamma : \mathcal{A}_{c-v}(S', \pi') \rightarrow \mathcal{A}_{c-c}(S', \pi'')$.*

As will be shown below in the proof of Corollary 4.2.5, note that all the sets of constraints considered in Theorem 4.2.4 satisfy the necessary condition of this proposition.

Proof. Define $S' = (X', \{(V_i, r_{i*}C_i)\}_{i=1}^m \cup \{(V_x, r_{x*}C_0)\}_{x \in X})$, where $X' = X \cup \bigcup_{x \in X} V_x$, $V_x = x \cup \{u_x | u \in V_0 \setminus \{v\}\}$, and

$$r_x(u) = \begin{cases} x & u = v \\ u_x & \text{otherwise.} \end{cases}$$

Define $\pi'(i) = \frac{\pi(i)}{2}$ and $\pi'(x) = \sum_{i \in [m]. x \in V_i} \frac{\pi(i)}{2|V_i|}$. Taking α to be the natural embedding $\mathcal{A}_{c-v}(S, \pi) \hookrightarrow \mathcal{A}_{c-v}(S', \pi')$ gives a 2-homomorphism. Next, we can take β to be the identity on $\mathcal{A}_{c-v}(S, \pi) \subseteq \mathcal{A}_{c-v}(S', \pi')$ and then take $\beta(\sigma_x(x)) = \sigma'(x)$, and $\beta(\sigma'(u_x)) = \beta(\sigma_x(u_x)) \in \langle \sigma'(x) \rangle$ such that they give a satisfying assignment to $r_{x*}C_0$. Therefore, the terms corresponding to $(V_x, r_{x*}C_0)$ are sent to 0 and we have that

$$\begin{aligned} & \beta \left(\sum_{i=1}^{m'} \frac{\pi'(i)}{|V_i|} \sum_{x \in V_i, \phi \in r_{i*}C_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2 \right) \\ &= \frac{1}{2} \sum_{i=1}^m \frac{\pi(i)}{|V_i|} \sum_{x \in V_i, \phi \in r_{i*}C_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2, \end{aligned}$$

giving the wanted $\frac{1}{2}$ -homomorphism.

Now, take $\pi''(i, i) = \pi''(x, x) = \pi''(x, i) = 0$ and

$$\pi''(i, x) = \begin{cases} \frac{\pi(i)}{|V_i|} & x \in V_i \\ 0 & \text{else} \end{cases}.$$

Let $\gamma(\Phi_{V_i, \phi}) = \Phi_{V_i, \phi}$, $\gamma(\sigma'(x)) = \gamma(\sigma_x(x)) = \sigma_x(x)$, and $\gamma(\sigma'(u_x)) = \gamma(\sigma_x(u_x)) = \sigma_x(u_x)$. Then, we have that

$$\begin{aligned}
& \gamma\left(\sum_{i=1}^{m'} \frac{\pi'(i)}{|V_i|} \sum_{x \in V_i, \phi \in r_{i*} C_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma'(x)))|^2\right) \\
&= \sum_{i=1}^m \frac{\pi(i)}{2|V_i|} \sum_{x \in V_i, \phi \in r_{i*} C_i} |\Phi_{V_i, \phi}(1 - \Pi_{\phi(x)}(\sigma_x(x)))|^2 \\
&\lesssim \frac{1}{2} \sum_{i \in [m], x \in V_i} \pi''(i, x) \sum_{\phi \in r_{i*} C_i, a \neq \phi(x)} \Phi_{V_i, \phi} \Pi_a(\sigma_x(x)) \\
&= \frac{1}{2} \sum_{i \in [m], x \in V_i} \pi''(i, x) \sum_{\substack{\phi \in r_{i*} C_i, \psi \in r_{x*} C_0 \\ \phi|_{V_i \cap V_x} \neq \psi|_{V_i \cap V_x}}} \Phi_{V_i, \phi} \Phi_{V_x, \psi} \\
&\lesssim \frac{1}{2} \sum_{i, j=1}^{m'} \pi''(i, j) \sum_{\substack{\phi \in r_{i*} C_i, \psi \in r_{j*} C_j \\ \phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}}} |\Phi_{V_i, \phi} \Phi_{V_j, \psi}|^2,
\end{aligned}$$

giving the wanted $\frac{1}{2}$ -homomorphism. \square

Proof of Corollary 4.2.5. Let Γ be a set constraints satisfying the conditions of Theorem 4.2.4. In the case that Γ is boolean, there must be a variable that takes both assignment 0 and 1, as there must be at least one constraint that has two distinct satisfying assignments. Next, in the case that $\Gamma = \{\neq_{\mathbb{Z}_3}\}$, either of the two variables can take all three values in \mathbb{Z}_3 . Finally, in the case that Γ is non-TVF, there is a pair of variables that can take any pair of values, so in particular either one of them can take any value. In all three cases, Γ satisfies the condition of Proposition 4.6.1. Further, we know by Theorem 4.2.4 that $\text{SuccinctCSP}_{c-v}(\Gamma)_{1,s}^*$ is RE-complete. Then, by the first part of Proposition 4.6.1, the instances of this problem can be mapped to a subset of the instances in such a way that the constant gap is preserved. Then, using the second part of Proposition 4.6.1 and Lemma 3.3.5, we find a C -homomorphism between these instances and the corresponding constraint-constraint algebras, thus $\text{SuccinctCSP}_{c-c}(\Gamma)_{1,s'}^*$ is also RE-complete. \square

Chapter 5

Two prover perfect zero knowledge for MIP*

This chapter shows that MIP* admits two prover perfect zero knowledge proofs. Our proof uses weighted algebras to show that a modified version of a two prover perfect zero knowledge protocol for NP due to [21] is sound against entangled provers. We use an algebraic argument to show that the players' strategy in an accept instance employs a correlation that the verifier can simulate efficiently.

5.1 Definitions

An MIP protocol is perfect zero knowledge if the verifier gains no new information from interacting with the provers. If the players' behaviour in a game $\mathbf{G} = (I, \{O_i\}_{i \in I}, \pi, V)$ is given by the correlation p , then what the verifier (or any outside observer) sees is the distribution $\{\pi(i, j)p(a, b|i, j)\}$ over tuples $(a, b|i, j)$. Consequently an MIP* protocol $(\{\mathbf{G}_x\}, Q, V)$ is said to be perfect zero-knowledge against an honest verifier if the players can use correlations p_x for \mathbf{G}_x such that the distribution $\{\pi(i, j)p(a, b|i, j)\}$ can be sampled in polynomial time in $|x|$. However, a dishonest verifier seeking to get more information from the players might sample the questions from a different distribution π' from π . To be perfect zero-knowledge against a dishonest verifier, it must be possible to efficiently sample $\{\pi'(i, j)p_x(a, b|i, j)\}$ for any efficiently sampleable distribution π' , and this is equivalent to being able to efficiently sample from $\{p_x(a, b|i, j)\}_{(a, b) \in O_i \times O_j}$ for any i, j . This leads to the definition (following [17, Definition 6.3]):

Definition 5.1.1. Let $\mathcal{P} = (\{\mathbf{G}_x\}, Q, V)$ be a two-prover one-round MIP* protocol for a language \mathcal{L} with completeness c and soundness s , where $\mathbf{G}_x = (I_x, \{O_{xi}\}, \pi_x, V_x)$. The protocol \mathcal{P} is **perfect zero knowledge** if for every string x , there is a correlation p_x for \mathbf{G}_x such that

1. for all $i, j \in I_x$, the distribution $\{p_x(a, b|i, j)\}$ can be sampled in polynomial time in $|x|$, and
2. if $x \in \mathcal{L}$ then $p_x \in C_{qa}$ and $\mathfrak{w}(\mathbf{G}_x, p_x) = 1$.

The class $\text{PZK-MIP}^*(2, 1, c, s)$ is the class of languages with a perfect zero knowledge two-prover one round MIP* protocol with completeness c and soundness s .

By replacing C_{qa} with C_{qc} , we get another class PZK-MIP^{co} . If we replace MIP* protocols with BCS-MIP* (resp. BCS-MIP^{co}) protocols and C_{qa} with C_{qa}^s (resp. C_{qc}^s) we get the class PZK-BCS-MIP^* (resp. PZK-BCS-MIP^{co}).

5.2 Parallel repetition

Let $\mathbf{G} = (I, \{O_i\}_{i \in I}, \pi, V)$ be a nonlocal game. The **n -fold parallel repetition** of \mathbf{G} is the game

$$\mathbf{G}^{\otimes n} = (I^n, \{O_{\underline{i}}\}_{\underline{i} \in I^n}, \pi^{\otimes n}, V^{\otimes n}),$$

where

1. I^n is the n -fold product of I ,
2. if $\underline{i} \in I^n$, then $O_{\underline{i}} := O_{i_1} \times O_{i_2} \times \cdots \times O_{i_n}$,
3. if $\underline{i}, \underline{j} \in I^n$, then $\pi^{\otimes n}(\underline{i}, \underline{j}) = \prod_{k=1}^n \pi(i_k, j_k)$, and
4. if $\underline{i}, \underline{j} \in I^n$, $\underline{a} \in O_{\underline{i}}$, $\underline{b} \in O_{\underline{j}}$, then $V^{\otimes n}(\underline{a}, \underline{b}|\underline{i}, \underline{j}) = \prod_{k=1}^n V(a_k, b_k|i_k, j_k)$.

In other words, the players each receive a vector of questions $\underline{i} = (i_1, \dots, i_n)$ and $\underline{j} = (j_1, \dots, j_n)$ from \mathbf{G} , and must reply with vectors of answers (a_1, \dots, a_n) and (b_1, \dots, b_n) to each question. Each pair of questions (i_k, j_k) , $1 \leq k \leq n$ is sampled independently from π , and the players win if and only if (a_k, b_k) is a winning answer to questions (i_k, j_k) for all

$1 \leq k \leq n$. If \mathbf{G} has questions of length q and answers of length a , then $\mathbf{G}^{\otimes n}$ has questions of length nq and answers of length na .

If p is a correlation for \mathbf{G} , let $p^{\otimes n}$ be the correlation for $\mathbf{G}^{\otimes n}$ defined by

$$p^{\otimes n}(\underline{a}, \underline{b} | \underline{i}, \underline{j}) = \prod_{k=1}^n p(a_k, b_k | i_k, j_k).$$

It is easy to see that $p^{\otimes n}$ is a quantum (resp. commuting operator) correlation if and only if p is a quantum (resp. commuting operator) correlation, and that $\mathfrak{w}(\mathbf{G}^{\otimes n}; p^{\otimes n}) = \mathfrak{w}(\mathbf{G}, p)^n$. Hence if $\mathfrak{w}_q(\mathbf{G}) = 1$ (resp. $\mathfrak{w}_{qc}(\mathbf{G}) = 1$) then $\mathfrak{w}_q(\mathbf{G}^{\otimes n}) = 1$ (resp. $\mathfrak{w}_{qc}(\mathbf{G}^{\otimes n}) = 1$) as well. If $\mathfrak{w}_q(\mathbf{G}) < 1$, then $\mathfrak{w}_q(\mathbf{G}^{\otimes n}) \geq \mathfrak{w}_q(\mathbf{G})^n$ (and the same for the commuting operator value), but this inequality is not always tight. However, Yuen's parallel repetition theorem states that the game value goes down at least polynomially in n :

Theorem 5.2.1 ([69]). *For any nonlocal game \mathbf{G} , if $\delta = 1 - \mathfrak{w}_q(\mathbf{G}) > 0$, then $\mathfrak{w}_q(\mathbf{G}^{\otimes n}) \leq b / \text{poly}(\delta, n)$, where b is the length of the answers of \mathbf{G} .*

Suppose $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ is a BCS and that π is a probability distribution on $[m] \times [m]$. For any $n \geq 1$, let $X^{(n)} := X \times [n]$, and $V_i^{(k)} = V_i \times \{k\} \subseteq X^{(n)}$. We can think of $X^{(n)}$ as the disjoint union of n copies of X , and $V_i^{(k)}$ as the copy of V_i from the k^{th} copy of X . Since $V_i^{(k)}$ is a copy of V_i , we can identify $\mathbb{Z}_2^{V_i^{(k)}}$ with $\mathbb{Z}_2^{V_i}$ in the natural way. If $\underline{i} \in [m]^n$, let $V_{\underline{i}} = \bigcup_{j=1}^n V_{i_j}^{(j)}$ and $C_{\underline{i}} = C_{i_1} \times \cdots \times C_{i_n} \subseteq \mathbb{Z}_2^{V_{\underline{i}}} = \mathbb{Z}_2^{V_{i_1}^{(1)}} \times \cdots \times \mathbb{Z}_2^{V_{i_n}^{(n)}}$. Let $B^{(n)} := (X^{(n)}, \{(V_{\underline{i}}, C_{\underline{i}})_{\underline{i} \in [m]^n}\})$. Given a distribution π on $[m] \times [m]$, consider the game $\mathbf{G}(B^{(n)}, \pi^{\otimes n})$, where $\pi^{\otimes n}$ is the product distribution as above. In this game, the players are given questions \underline{i} and \underline{j} from $[m]^n$ respectively, and must reply with elements $\underline{\phi} \in C_{\underline{i}}$ and $\underline{\psi} \in C_{\underline{j}}$ respectively. They win if and only if $\underline{\phi}$ and $\underline{\psi}$ agree on $V_{\underline{i}} \cap V_{\underline{j}} = \bigcup_{k=1}^n V_{i_k}^{(k)} \cap V_{j_k}^{(k)}$. But this happens if and only if ϕ_k and ψ_k agree on $V_{i_k} \cap V_{j_k}$. Thus $\mathbf{G}(B^{(n)}, \pi^{\otimes n})$ is the parallel repetition $\mathbf{G}(B, \pi)^{\otimes n}$. We record this in the following lemma:

Lemma 5.2.2. *If \mathbf{G} is a BCS game, then so is the parallel repetition $\mathbf{G}^{\otimes n}$.*

To illustrate the purpose of parallel repetition, suppose that $(\{\mathbf{G}_x\}, Q, V)$ is an $\text{MIP}^*(2, 1, 1, s)$ protocol for a language \mathcal{L} , where $\mathbf{G}_x = (I_x, \{O_{xi}\}, \pi_x, V_x)$ and has answer length a_x . If n_x is a polynomial in $|x|$, then $\pi_x^{\otimes n_x}$ can be sampled in polynomial time by running Q independently n_x times, and $V_x^{\otimes n_x}$ can also be computed in polynomial time by running V repeatedly. If $Q^{\otimes n_x}$ and $V^{\otimes n_x}$ are these Turing machines for sampling $\pi_x^{\otimes n_x}$ and computing $V_x^{\otimes n_x}$ respectively, then $(\{\mathbf{G}_x^{\otimes n_x}\}, Q^{\otimes n_x}, V^{\otimes n_x})$ is an $\text{MIP}^*(2, 1, 1, s')$ protocol for \mathcal{L} , where

$s' = a_x / \text{poly}(1 - s) \cdot \text{poly}(n_x)$. Since a_x is polynomial in $|x|$, if $1 - s = 1 / \text{poly}(|x|)$, then we can choose n_x such that s' is any constant < 1 . By Lemma 5.2.2 the same can be done for BCS-MIP*.

For the one-round protocols that we are considering, parallel repetition preserves the property of being perfect zero knowledge.

Proposition 5.2.3. *Let $(\{\mathbf{G}_x\}, S, V)$ be a PZK-MIP*(2, 1, 1, s) protocol, and let n_x be a polynomial function of $|x|$. Then the parallel repeated protocol $(\{\mathbf{G}_x^{\otimes n_x}\}, S^{\otimes n_x}, V^{\otimes n_x})$ is also perfect zero knowledge.*

Proof. Let p_x be a correlation for the game \mathbf{G} that satisfies the two requirements of Definition 5.1.1. Then $\{p_x^{\otimes n_x}(\underline{a}, \underline{b}|\underline{i}, \underline{j})\}_{\underline{a}, \underline{b}}$ can be sampled in polynomial time in $|x|$ for all $\underline{i}, \underline{j}$ by independently sampling from $\{p_x(a, b, i_\ell, j_\ell)\}_{a, b}$ for each pair (i_ℓ, j_ℓ) from $\underline{i} = (i_1, \dots, i_{n_x})$ and $\underline{j} = (j_1, \dots, j_{n_x})$. If $x \in \mathcal{L}$, then $\mathfrak{w}(\mathbf{G}_x^{\otimes n_x}; p_x^{\otimes n_x}) = 1$, and it is not hard to see that $p_x^{\otimes n_x} \in C_{qa}$. \square

5.3 The tableau construction

We will now prove the main result of this chapter that any proof system in BCS-MIP* or BCS-MIP^{co} can be turned into a perfect zero knowledge BCS-MIP* or BCS-MIP^{co} protocol. For this purpose, we use the perfect zero knowledge proof system for 3SAT due to Dwork, Feige, Kilian, Naor, and Safra [21], slightly modified to prove quantum soundness. For the construction, we assume that we start with a BCS-MIP* protocol (and in the proof of the main result, this will be a 3SAT-MIP* protocol). Following [21], the new proof system is constructed in three steps. First, we apply a transformation called obliviation, then turn the resulting system into a permutation branching program via Barrington's theorem [6], and finally rewrite the permutation branching programs using the randomizing tableaux of Kilian [44]. We start by describing obliviation.

Definition 5.3.1. *Given a BCS $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ and $n \geq 1$, let $Z = X \times [n]$, and $U_i = V_i \times [n]$ for any $1 \leq i \leq m$. To make the elements of Z look more like variables, we denote (x, i) by $x(i)$. Let $E_i \subseteq \mathbb{Z}_2^{U_i}$ be the set of assignments ϕ to U_i such that the assignment ψ to V_i defined by $\psi(x) = \psi(x(1)) \cdots \psi(x(n))$ is in C_i . The **obliviation of B of degree n** is the constraint system $\text{Obl}_n(B) = (Z, \{(U_i, E_i)\}_{i=1}^m)$. We call the variables of $\text{Obl}_n(B)$ **oblivious variables**.*

The point of obliviation is the following:

Lemma 5.3.2. *Suppose $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ is a BCS, and let $B' = \text{Obl}_n(B)$ for some $n \geq 1$, using the notation from Definition 5.3.1. Then:*

- (a) *There is a classical homomorphism $\alpha : \mathcal{A}_{c-c}(B) \rightarrow \mathcal{A}_{c-c}(B')$ such that $\alpha(\sigma_i(x)) = \sigma_i(x(1) \cdots x(n))$ for all $i \in [m]$ and $x \in V_i$, where σ_i is the inclusion of the i th factor for $\mathcal{A}_{c-c}(B)$ and $\mathcal{A}_{c-c}(B')$.*
- (b) *Let Γ be the set of sequences x_1, \dots, x_k in Z of length $1 \leq k \leq n-1$, such that there is some $i \in [k]$ with $x_i \neq x_j$ for all $j \in [k] \setminus \{i\}$. If π is a probability distribution on $[m] \times [m]$, and τ is a tracial state on $\mathcal{A}_{c-c}(B)$, then there is a tracial state $\tilde{\tau}$ on $\mathcal{A}_{c-c}(B')$ such that $\tau = \tilde{\tau} \circ \alpha$, $\text{def}(\tilde{\tau}; \mu_\pi) = \text{def}(\tau; \mu_\pi)$, and $\tilde{\tau}(\sigma_{i_1}(x_1) \cdots \sigma_{i_k}(x_k)) = 0$ for all sequences x_1, \dots, x_k in Γ and indices $i_1, \dots, i_k \in [m]$ such that $x_j \in U_{i_j}$ for all $1 \leq j \leq k$. If τ is finite-dimensional (resp. Connes-embeddable), then $\tilde{\tau}$ is also finite-dimensional (resp. Connes-embeddable).*
- (c) *For any $1 \leq i \leq m$, the set $\{\prod_{x \in S} x : S \subseteq U_i, |S| < n/2\}$ of monomials in U_i of degree less than $n/2$ is linearly independent in $\mathcal{A}(U_i, E_i)$.*

In particular, if τ is perfect then $\tilde{\tau}$ is perfect.

Proof. Define $f_i : \mathbb{Z}_2^{U_i} \rightarrow \mathbb{Z}_2^{V_i}$ for each $i \in [m]$ by $f_i(\phi)(x) = \phi(x(1)) \cdots \phi(x(n))$ for $\phi \in \mathbb{Z}_2^{U_i}$ and $x \in V_i$. By definition, $\phi \in E_i$ if and only if $f_i(\phi) \in C_i$, so $f_i(E_i) = C_i$. If $f_i(\phi)(x) \neq f_j(\psi)(x)$ for some $\phi \in \mathbb{Z}_2^{U_i}$, $\psi \in \mathbb{Z}_2^{U_j}$, and $x \in V_i \cap V_j$, then we must have $\phi(x(i)) \neq \psi(x(i))$ for some i . Since

$$\sigma_i(x(1) \cdots x(n)) = \sum_{\phi \in \mathbb{Z}_2^{U_i}} f_i(\phi)(x) \Phi_{U_i, \phi}$$

for all $x \in V_i$, $i \in [m]$, the functions f_i correspond to a classical homomorphism $\alpha : \mathcal{A}_{c-c}(B) \rightarrow \mathcal{A}_{c-c}(B')$ with $\alpha(\sigma_i(x)) = \sigma_i(x(1) \cdots x(n))$ for all $i \in [m]$ and $x \in V_i$. This proves part (a).

Conversely, given $y \in \mathbb{Z}_2^{X \times [n-1]}$ and $\phi \in \mathbb{Z}_2^{V_i}$, define $\phi_y \in \mathbb{Z}_2^{U_i}$ by $\phi_y(x(1)) = \phi(x)y(x, 1)$, $\phi_y(x(j)) = y(x, j-1)y(x, j)$ for $2 \leq j \leq n-1$, and $\phi_y(x(n)) = y(x, n-1)$. Since $f_i(\phi_y) = \phi$, the function $\phi \mapsto \phi_y$ sends C_i to E_i . Also if $\phi \in \mathbb{Z}_2^{V_i}$ and $\psi \in \mathbb{Z}_2^{V_j}$, then $\phi_y|_{U_i \cap U_j} \neq \psi_y|_{U_i \cap U_j}$ if and only if $\phi|_{V_i \cap V_j} \neq \psi|_{V_i \cap V_j}$, so the functions $\phi \mapsto \phi_y$ determine a classical homomorphism $\beta_y : \mathcal{A}_{c-c}(B') \rightarrow \mathcal{A}_{c-c}(B)$ with $\beta_y(\sigma_i(x(1))) = \sigma_i(x)y(x, 1)$, $\beta_y(\sigma_i(x(j))) = y(x, j-1)y(x, j)$ for $2 \leq j \leq n-1$, and $\beta_y(\sigma_i(x(n))) = y(x, n-1)$ for all $i \in [m]$ and $x \in V_i$.

Given a tracial state τ on $\mathcal{A}_{c-c}(B)$, define a tracial state $\tilde{\tau}$ on $\mathcal{A}_{c-c}(B')$ by $\tilde{\tau} = 2^{-|X|(n-1)} \sum_y \tau \circ \beta_y$, where the sum is over all $y \in \mathbb{Z}_2^{X \times [n-1]}$. Notice that if τ is finite-dimensional (resp. Connes-embeddable), then $\tilde{\tau}$ is also finite-dimensional (resp. Connes-embeddable). Since $\beta_y \circ \alpha$ is the identity on $\mathcal{A}_{c-c}(B)$, $\tilde{\tau} \circ \alpha = \tau$. Since β_y and α are 1-homomorphisms,

$$\text{def}(\tau \circ \beta_y; \mu_\pi) \leq \text{def}(\tau; \mu_\pi) = \text{def}(\tau \circ \beta_y \circ \alpha; \mu_\pi) \leq \text{def}(\tau \circ \beta_y; \mu_\pi)$$

for any y , so $\text{def}(\tau \circ \beta_y; \mu_\pi) = \text{def}(\tau; \mu_\pi)$ and hence $\text{def}(\tilde{\tau}; \mu_\pi) = \text{def}(\tau; \mu_\pi)$.

Finally, if x_1, \dots, x_k is a sequence in Z , and i_1, \dots, i_k is a sequence in $[m]$ such that $x_j \in U_{i_j}$, then there is an element $a \in \mathcal{A}_{c-c}(B)$ and set $S \subseteq X \times [n-1]$ such that

$$\beta_y(\sigma_{i_1}(x_1) \cdots \sigma_{i_k}(x_k)) = m_y \tau(a)$$

for all $y \in \mathbb{Z}_2^{X \times [n-1]}$, where $m_y := \prod_{(x,j) \in S} y(x, j)$. If x_1, \dots, x_k is in Γ , then S is non-empty, and $\sum_y m_y = 0$. Hence

$$\tilde{\tau}(\sigma_{i_1}(x_1) \cdots \sigma_{i_k}(x_k)) = 2^{-|X|(n-1)} \sum_y m_y \tau(a) = 0.$$

This proves part (b).

For part (c), pick a tracial state τ on the finite-dimensional C^* -algebra $\mathcal{A}(V_i, C_i)$ (since C_i is non-empty, this algebra is non-trivial). As in the proof of part (b), we can define a tracial state $\tilde{\tau} = 2^{|X|(n-1)} \sum_y \tau \circ \beta_y$ on $\mathcal{A}(U_i, E_i)$ with the property that $\tilde{\tau}(x_1 \cdots x_k) = 0$ if $1 \leq k \leq n-1$ and $x_1, \dots, x_k \in U_i$ are distinct. If $S, T \subseteq U_i$, then

$$\prod_{x \in S} x \cdot \prod_{x \in T} x = \prod_{x \in S \Delta T} x,$$

where $S \Delta T := (S \cup T) \setminus (S \cap T)$. If $|S|, |T| < n/2$, then $|S \Delta T| < n$, and $S \Delta T = \emptyset$ if and only if $S = T$. Hence by part (b),

$$\tilde{\tau}\left(\prod_{x \in S} x \cdot \prod_{x \in T} x\right) = \begin{cases} 1 & S = T \\ 0 & S \neq T \end{cases}.$$

It follows that the monomials $\{\prod_{x \in S} x : S \subseteq U_i, |S| < n/2\}$ are linearly independent. \square

A **permutation branching program** of width 5 and depth d on a set of variables X is a tuple $P = (X, \{(x_i, \pi_1^{(i)}, \pi_{-1}^{(i)})\}_{i=1}^d, \sigma)$ where $x_i \in X$ and $\pi_1^{(i)}, \pi_{-1}^{(i)}$ are elements of the

permutation group S_5 for all $1 \leq i \leq d$, and $\sigma \in S_5$ is a 5-cycle. A permutation branching program P defines a map $P : \mathbb{Z}_2^X \rightarrow S_5$ via $P(\phi) = \prod_{i=1}^d \pi_{\phi(x_i)}^{(i)}$. A program P **recognizes a constraint** $C \subseteq \mathbb{Z}_2^X$ if $P(\phi) = \sigma$ for all $\phi \in C$, and $P(\phi) = e$ for all $\phi \notin C$, where e is the identity in S_5 .

Theorem 5.3.3 (Barrington [6]). *Suppose a constraint $C \subseteq \mathbb{Z}_2^X$ is recognized by a depth d fan-in 2 boolean circuit. Then C is recognized by a permutation branching program of width 5 and depth 4^d on the variables X .*

For the rest of the chapter, we assume that we have a canonical way of turning constraints described by fan-in 2 boolean circuits into permutation branching programs using Barrington's theorem.

The final ingredient is randomizing tableaux, which are described using constraints of the form $x_1 \cdots x_n = \gamma$, where the variables x_1, \dots, x_n take values in S_5 , γ is a constant in S_5 , and the product is the group multiplication. Since $|S_5| = 120 < 2^7$, we can encode permutations as bit strings of length 7 by choosing an enumeration $S_5 = \{e = \gamma_0, \dots, \gamma_{119}\}$, and identifying γ_j by its index j in binary. This means that any permutation-valued variable can be represented by 7 boolean variables. Similarly, a permutation-valued constraint $x_1 \cdots x_n = \gamma$ can be rewritten as the constraint on $7n$ boolean variables which requires the boolean variables corresponding to x_i to encode a permutation value, and the product of all the permutations to be equal to γ . Since we want our final output to be a boolean constraint system, we use permutation-valued variables and permutation-valued constraints as short-hand for boolean constraint systems constructed in this way. We can now define randomizing tableaux, still following [21] with small modifications.

Definition 5.3.4. *Let $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ be a BCS, where each C_i is described by a fan-in 2 boolean circuit. Let $P_i = (V_i, \{(x_{ij}, \pi_1^{(ij)}, \pi_{-1}^{(ij)})\}_{j=1}^{d_i}, \sigma_i)$ be the permutation branching program recognizing C_i . For each $i \in [m]$, let*

$$W_i = V_i \sqcup \{T_i(p, q) : (p, q) \in [4] \times [d_i]\} \sqcup \{r_i(j, k) : (j, k) \in [3] \times [d_i - 1]\},$$

where $T_i(p, q)$ and $r_i(j, k)$ are new permutation-valued variables (and thus represent 7 boolean variables each), and let

$$Y = X \sqcup \{T_i(p, q), r_i(j, k) : (i, p, q, j, k) \in [m] \times [4] \times [d_i] \times [3] \times [d_i - 1]\}$$

be the union of all the original and new variables. The variables $T_i(p, q)$ are called *tableau elements*, and the variables $r_i(j, k)$ are called *randomizers*.

Let D_i be the constraint on variables W_i which is the conjunction of the following clauses:

1. $T_i(1, q) = \pi_{x_q}^{(iq)}$ for all $q \in [d_i]$,
2. $T_i(p+1, q) = r_i(p, q-1)^{-1}T_i(p, q)r_i(p, q)$ for $q \in [d_i]$ and $p \in [3]$, where we use the notation $r_i(p, 0) = r_i(p, d_i) = e$,
3. $\prod_{1 \leq q \leq d_i} T_i(4, q) = \sigma_i$, and
4. a trivial constraint (meaning that all assignment are allowed) on any pair x, y of original or permutation-valued variables which do not appear in one of the above constraints.

The **tableau** of B is $\text{Tab}(B) = (Y, \{(W_i, D_i)\}_{i=1}^m)$, interpreted as a boolean constraint system. For a fixed i and p , we call the set of elements $\{T_i(p, q) : 1 \leq q \leq d_i\}$ a **row** of the tableau for constraint i . We further let $\{(W_{ij}, D_{ij})\}_{j=1}^{m_i}$ be a list of the clauses in (1)-(4) making up D_i . The **subdivided tableau** of B is $\text{Tab}_{\text{sub}}(B) = (Y, \{(W_{ij}, D_{ij})\}_{i \in [m], j \in [m_i]})$.

Compared to [21], we've added the trivial constraints (4), as well as an extra row of the tableau. As mentioned above, the product in the constraints on the permutation-valued variables in parts (1)-(4) of the definition is the group product in S_5 . The constraints in part (1) involve both original variables x_q and permutation-valued variables $T_i(1, q)$, and say that the value of $T_i(1, q)$ is either $\pi_1^{(iq)}$ or $\pi_{-1}^{(iq)}$ depending on the value of x_q . In part (4), x and y can be either an original or a permutation-valued variable. If one of them is a permutation-valued variable, then all the corresponding boolean variables encoding the permutation-valued variable are included in the constraint (so the constraint on x and y may involve up to 14 boolean variables). Since the constraints in part (4) are trivial, they do not contribute to D_i , but they are included in the list of clauses (W_{ij}, D_{ij}) of the subdivided tableau. The point of the constraints in part (4) is that, with them, $\text{Tab}_{\text{sub}}(B)$ is a subdivision of $\text{Tab}(B)$. The extra row of the tableau is needed to compensate for the inclusion of these constraints in $\text{Tab}_{\text{sub}}(B)$ (see Remark 5.4.4). As in [21], the constraints D_i encode the constraints C_i as follows:

Lemma 5.3.5 ([21]). *Suppose $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ is a BCS with tableau $\text{Tab}(B) = (Y, \{(W_i, D_i)\}_{i=1}^m)$. If $\psi \in D_i$, then $\psi|_{V_i} \in C_i$. Conversely, if $r \in S_5^{R_i}$, where $R_i = \{r_i(j, k) : (j, k) \in [3] \times [d_i]\}$ is the set of randomizers in W_i , and $\phi \in C_i$, then there is a unique element $\phi_r \in D_i$ such that $\phi_r|_{V_i} = \phi$ and $\phi_r|_{R_i} = r$.*

In this lemma, the statement that $\phi_r|_{R_i} = r$ means that for every randomizer $r_i(j, k) \in R_i$, the restriction of ϕ to the boolean variables corresponding to $r_i(j, k)$ is the encoding of the permutation $r(r_i(j, k))$.

Proof. If $\psi \in D_i$, then by constraint (2), $\prod_q T_i(p+1, q) = \prod_q T_i(p, q)$. Since $\prod_q T_i(4, q) = \sigma_i$ by constraint (3), $\prod_q \pi_{x_q}^{(iq)} = \sigma_i$. Since the permutation branching program P_i recognizes C_i , we conclude that $\psi|_{V_i} \in C_i$.

Conversely, given an assignment $r \in S_5^{R_i}$ to the variables R_i and $\phi \in C_i$, we can set $T_i(1, q) = \pi_{\phi(x_q)}^{(iq)}$ and $T_i(p+1, q) = r_i(p, q-1)^{-1} T_i(p, q) r_i(p, q)$ to get an assignment where $\prod_q T_i(4, q) = \sigma_i$. \square

5.4 Perfect zero knowledge

Although the permutation-valued variables in $\text{Tab}(B)$ are shorthand for boolean variables, it is helpful to be able to work with the permutation-valued variables directly in $\mathcal{A}_{c-c}(\text{Tab}(B))$. Suppose for a moment that x_1, \dots, x_7 are variables in a set V , and C is a constraint on V which includes the requirement that x_1, \dots, x_7 encode a permutation-valued variable x . Let $S = \{x_1, \dots, x_7\}$. If $\phi \in \mathbb{Z}_2^S$, then $\Phi_{S, \phi} = 0$ in $\mathcal{A}(V, C)$ unless ϕ is the binary representation of an index $0 \leq j < 120$, in which case we also write $\Phi_{S, \phi}$ as $\Phi_{S, j}$. Hence the subalgebra of $\mathcal{A}(V, C)$ is generated by the single unitary $\sum_{j=0}^{119} e^{2\pi i j / 120} \Phi_{S, j}$, which we denote by the same symbol as the permutation-valued variable x . In particular, if $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ and $\text{Tab}(B) = (Y, \{(W_i, D_i)\}_{i=1}^m)$ as in Definition 5.3.4, then we can refer to $T_i(p, q)$ and $r_i(j, k)$ as unitary elements of $\mathcal{A}(W_i, D_i)$ of order 120, and they generate the same subalgebra as the boolean variables encoding them. Since these variables do not occur in any other context W_j for $j \neq i$, we also use $T_i(p, q)$ and $r_i(j, k)$ to refer to $\sigma_i(T_i(p, q))$ and $\sigma_i(r_i(j, k))$ in $\mathcal{A}_{c-c}(\text{Tab}(B))$. We use the same convention for $\mathcal{A}(W_{i\ell}, D_{i\ell})$, although since the variables $T_i(p, q)$ and $r_i(j, k)$ occur in more than one constraint of $\text{Tab}_{\text{sub}}(B)$, we are stuck with the notation $\sigma_{i\ell}(T_i(p, q))$ and $\sigma_{i\ell}(r_i(j, k))$ when referring to these variables in $\mathcal{A}_{c-c}(\text{Tab}_{\text{sub}}(B))$. With these conventions, we can state the following noncommutative version of Lemma 5.3.5.

Lemma 5.4.1. *Suppose that $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ is a BCS with tableau $\text{Tab}(B) = (Y, \{(W_i, D_i)\}_{i=1}^m)$. Let $R_i = \{r_i(j, k) : (j, k) \in [3] \times [d_i - 1]\}$ be the set of randomizers in W_i , and let $R = \bigcup_i R_i$.*

(a) *The natural map*

$$\mathcal{A}(V_i, C_i) \otimes \mathbb{C}\mathbb{Z}_{120}^{R_i} \rightarrow \mathcal{A}(W_i, D_i) : x_i \mapsto x_i, r_i(j, k) \mapsto r_i(j, k)$$

is an isomorphism. In particular, $\mathcal{A}(W_i, D_i)$ is generated as an algebra by $V_i \cup R_i$, and $\mathcal{A}_{c-c}(\text{Tab}(B))$ is generated by $\bigcup_i \{\sigma_i(x) : x \in V_i\} \cup R$.

- (b) The natural inclusion $\alpha : \mathcal{A}_{c-c}(B) \rightarrow \mathcal{A}_{c-c}(\text{Tab}(B))$ defined by $\alpha(\sigma_i(x)) = \sigma_i(x)$ for $i \in [m]$ and $x \in V_i$ is a classical homomorphism.
- (c) If $r \in S_5^R$, then there is a classical homomorphism $\beta_r : \mathcal{A}_{c-c}(\text{Tab}(B)) \rightarrow \mathcal{A}_{c-c}(B)$ such that for all $i \in [m]$, if $x \in V_i$ then $\beta_r(\sigma_i(x)) = \sigma_i(x)$, and if $x \in R_i$ then $\beta_r(x) = e^{2\pi ij/120}$ where $r(x) = \gamma_j$ in the enumeration of S_5 fixed above.
- (d) Let \mathcal{M} be the set of monomials in $\mathcal{A}_{c-c}(B)$ of the form $u\sigma_i(z)^av$, where $z \in R_i$ for some $i \in [m]$, $1 \leq a < 120$, and u and v are monomials in $\{\sigma_j(x) : j \in [m], x \in V_j \cup R_j\}$ which do not contain z . If π is a probability distribution on $[m] \times [m]$, and τ is a tracial state on $\mathcal{A}_{c-c}(B)$, then there is a tracial state $\tilde{\tau}$ on $\mathcal{A}_{c-c}(\text{Tab}(B))$ such that $\tau = \tilde{\tau} \circ \alpha$, where α is the classical homomorphism from part (b), $\text{def}(\tilde{\tau}; \mu_\pi) = \text{def}(\tau; \mu_\pi)$, and $\tilde{\tau}(y) = 0$ for all $y \in \mathcal{M}$. Furthermore, if τ is finite-dimensional (resp. Connes-embeddable), then $\tilde{\tau}$ is also finite-dimensional (resp. Connes-embeddable).

Proof. For part (a), the algebra $\mathbb{C}\mathbb{Z}_{120}^{R_i}$ has a basis consisting of the joint spectral projections

$$\Phi_{R_i, r} = \prod_{x \in R_i} \lambda(r(x))^{-1} \prod_{\gamma_k \neq r(x)} (x - e^{2\pi ik/120}), \quad r \in \mathbb{Z}_{120}^{R_i},$$

where $\lambda(\gamma_j) = \prod_{k \neq j} (e^{2\pi ij/120} - e^{2\pi ik/120})$. Hence $\mathcal{A}(V_i, C_i) \otimes \mathbb{C}\mathbb{Z}_{120}^{R_i}$ has a basis consisting of the elements $\Phi_{V_i, \phi} \otimes \Phi_{R_i, r}$ for $\phi \in C_i$ and $r \in \mathbb{Z}_{120}^{R_i}$. Using the enumeration of S_5 fixed earlier, we can interpret $\mathbb{Z}_{120}^{R_i}$ as the set $S_5^{R_i}$ of permutation-valued assignments to R_i . The natural homomorphism $\mathcal{A}(V_i, C_i) \otimes \mathbb{C}\mathbb{Z}_{120}^{R_i} \rightarrow \mathcal{A}(W_i, D_i)$ sends $\Phi_{V_i, \phi} \otimes \Phi_{R_i, r}$ to $\sum_{\psi} \Phi_{W_i, \psi}$, where the sum is across all $\psi \in D_i$ such that $\psi|_{V_i} = \phi$ and $\psi|_{R_i} = r$. By Lemma 5.3.5, the restriction map $\phi \mapsto \phi|_{V_i \cup R_i}$ is a bijection between D_i and $C_i \times S_5^{R_i}$, so this homomorphism is an isomorphism.

Parts (b) and part (c) follow immediately from Lemma 5.3.5 and the definition of a classical homomorphism.

The proof of part (d) is similar to the proof of Lemma 5.3.2, part (b). Given a tracial state τ on $\mathcal{A}_{c-c}(B)$, let $\tilde{\tau}$ be the tracial state on $\mathcal{A}_{c-c}(\text{Tab}(B))$ defined by $\tilde{\tau} = \frac{1}{120^{|R|}} \sum_r \tau \circ \beta_r$, where the sum is over $r \in S_5^R$. If τ is finite-dimensional (resp. Connes-embeddable), then $\tilde{\tau}$ is finite-dimensional (resp. Connes-embeddable). Since $\beta_r(\sigma_i(x)) = \sigma_i(x)$ for all $i \in [m]$ and $x \in V_i$, $\beta_r \circ \alpha$ is the identity on $\mathcal{A}_{c-c}(B)$, and $\tilde{\tau} \circ \alpha = \tau$. By parts (b) and (c), $\text{def}(\tau \circ \beta_r) \leq \text{def}(\tau) = \text{def}(\tau \circ \beta_r \circ \alpha) \leq \text{def}(\tau \circ \beta_r)$. This means that $\text{def}(\tau \circ \beta_r) = \text{def}(\tau)$, so $\text{def}(\tau) = \text{def}(\tilde{\tau})$. Finally, suppose $y \in \mathcal{M}$, so $y = u\sigma_i(z)^av$ for some $z \in R_i$, $1 \leq a < 120$, and monomials u, v which do not contain z . By part (c), there is some monomial y' in

$\{\sigma_j(x) : j \in [m], x \in V_i\}$ such that for all $r \in S_5^R$, we have $\beta_r(y) = e^{2\pi aij/120} c_{r'} y'$, where $r(z) = \gamma_j$, and $c_{r'} \in \mathbb{C}$ depends only on $r' = r|_{R \setminus \{z\}}$. Hence

$$\tilde{\tau}(y) = \frac{1}{120^{|R|}} \sum_{r \in S_5} \tau(\beta_r(y)) = \frac{1}{120^{|R|}} \sum_{j=0}^{120} e^{2\pi aij/120} \sum_{r' \in S_5^{R \setminus \{z\}}} c_{r'} \tau(y') = 0,$$

finishing the proof of part (d). □

We need one more general fact about permutation-valued variables.

Lemma 5.4.2. *Let $f : S_5^m \rightarrow S_5$ be a function, and suppose (V, C) is a boolean constraint encoding the constraint $x = f(y_1, \dots, y_m)$ on permutation-valued variables x, y_1, \dots, y_m . If $1 \leq n < 120$, then*

$$x^n = \sum_a c_a y_1^{a_1} \cdots y_m^{a_m}$$

for some coefficients $c_a \in \mathbb{C}$, where the sum is over all integer vectors $a = (a_1, \dots, a_m)$ with $0 \leq a_1, \dots, a_m < 120$. Furthermore, if for every $\pi_1, \dots, \pi_{m-1} \in S_5$, we have the equality

$$\{f(\pi_1, \dots, \pi_{k-1}, \pi, \pi_k, \dots, \pi_{m-1}) : \pi \in S_5\} = S_5,$$

then $c_a = 0$ if $a_k = 0$.

Proof. Let Y_k be the set of boolean variables representing y_k , and let X be the set of boolean variables representing x . The constraint $x = f(y_1, \dots, y_m)$ states that

$$\Phi_{X,\ell} = \sum_{(\gamma_{j_1}, \dots, \gamma_{j_m}) \in f^{-1}(\gamma_\ell)} \Phi_{Y_1, j_1} \cdots \Phi_{Y_m, j_m},$$

where $\{\gamma_0, \dots, \gamma_{119}\}$ is our chosen enumeration of S_5 . Since Φ_{Y_k, j_k} is a polynomial in y_k , and x^m is a linear combination of the projections $\Phi_{X,\ell}$ for $0 \leq \ell < 120$, we get $x^n = g(y_1, \dots, y_m)$, where $g = \sum_a c_a y_1^{a_1} \cdots y_m^{a_m}$ is a polynomial in y_1, \dots, y_m . Since $y_k^{120} = 1$, we can further assume that the sum is over vectors $a = (a_1, \dots, a_k)$ with $0 \leq a_k < 120$ for all k .

Given $0 \leq j_1, \dots, j_m < 120$, let $\phi_j : \mathcal{A}(V, C) \rightarrow \mathbb{C}$ be the homomorphism sending $\Phi_{Y_k, a} \mapsto \delta_{a, j_k}$ for all $1 \leq k \leq m$. This homomorphism sends $y_k \mapsto \omega^{j_k}$ and $x \mapsto \omega^\ell$, where $\omega = e^{2\pi i/120}$, and $\gamma_\ell = f(\gamma_{j_1}, \dots, \gamma_{j_m})$. We use the notation

$$A_1, \dots, \check{A}_k, \dots, A_m$$

to denote the list A_1, \dots, A_m with the element A_k omitted. If, for some k , we fix $0 \leq j_1, \dots, \check{j}_k, \dots, j_m < 120$, then

$$\sum_{0 \leq j_k < 120} \phi_j(g) = \sum_a c_a \prod_{t \neq k} \omega^{j_t a t} \sum_{0 \leq j_k < 120} \omega^{j_k a k} = h(\omega^{j_1}, \dots, \check{\omega}^{j_k}, \dots, \omega^{j_m}),$$

where $h = g(y_1, \dots, y_{k-1}, 0, y_{k+1}, \dots, y_m)$. If $\{f(\gamma_{j_1}, \dots, \gamma_{j_m}) : 0 \leq j_k < 120\}$ is equal to S_5 , then

$$\sum_{0 \leq j_k < 120} \phi_j(x^n) = \sum_{0 \leq \ell < 120} \omega^{n\ell} = 0$$

for $1 \leq n < 120$, and we conclude that

$$h(\omega^{j_1}, \dots, \check{\omega}^{j_k}, \dots, \omega^{j_m}) = 0.$$

If this occurs for all choices of $0 \leq j_1, \dots, \check{j}_k, \dots, j_m < 120$, then h must be the zero polynomial, so $c_a = 0$ if $a_k = 0$. \square

Although Lemma 5.4.2 is stated for general functions f , we are only going to use it for the group multiplication and inverse functions, i.e. $f(y_1, y_2) = y_1 y_2$ and $f(y) = y^{-1}$. For these functions, the additional hypothesis on f holds for all indices k . Thus the lemma states that if (V, C) encodes the constraint $x = y_1 y_2$, then x is a polynomial in y_1 and y_2 such that all monomials contain both y_1 and y_2 , and similarly for the constraint $x = y^{-1}$.

We can now prove the main algebraic lemma that we use to prove perfect zero knowledge.

Lemma 5.4.3. *Given a BCS $B = (X, \{(V_i, C_i)\}_{i=1}^m)$, let $\text{Tab}(B) = (Y, \{(W_i, D_i)\}_{i=1}^m)$, and let $\text{Tab}_{\text{sub}}(B) = (Y, \{(W_{ij}, D_{ij})\}_{i \in [m], j \in [m_i]})$. Let $R_i = \{r_i(j, k) : (j, k) \in [3] \times [d_i - 1]\}$ be the set of randomizers in W_i . Then:*

- (a) *Suppose (W_{ij}, D_{ij}) is a constraint from $\text{Tab}_{\text{sub}}(B)$ of type (1), (2), or (4) in Definition 5.3.4. If y is a polynomial in W_{ij} , then y is equal in $\mathcal{A}(W_i, D_i)$ to a polynomial in $S \cup R_i$, where $W_{ij} \cap V_i \subseteq S \subseteq V_i$ and $|S| \leq 2$.*
- (b) *Suppose (W_{ij}, D_{ij}) is a constraint from $\text{Tab}_{\text{sub}}(B)$ of type (3). If y is a polynomial in W_{ij} then y is equal in $\mathcal{A}(W_i, D_i)$ to a polynomial in $V_i \cup R_i$ where every non-scalar monomial contains a variable from R_i .*
- (c) *If y is a polynomial in W_{ij} and z is a polynomial in W_{ik} for some $i \in [m]$, $j, k \in [m_i]$, then yz is equal in $\mathcal{A}(W_i, D_i)$ to a polynomial in $V_i \cup R_i$ in which every monomial either contains a variable from R_i or has degree ≤ 4 .*

Proof. Fix $i \in [m]$, and consider the permutation-valued variables $T_i(p, q)$ in $\mathcal{A}(W_i, D_i)$. The constraints of type (1) in Definition 5.3.4 imply that $T_i(1, q)$ is a polynomial in x_q for all $q \in [d_i]$. The constraints of type (2) along with Lemma 5.4.2 imply that $T_i(p+1, q)$ is a polynomial in $\{r_i(p, q-1), r_i(p, q), T_i(p, q)\}$, and vice versa $T_i(p, q)$ is a polynomial in $\{r_i(p, q-1), r_i(p, q), T_i(p+1, q)\}$. Recall that $r_i(p, 0) = r_i(p, d_i) = 1$; for notational convenience we use the convention that they are present in every monomial, although note they aren't elements of R_i . It follows that $T_i(p, q)$ is a polynomial in $\{x_q\} \cup \{r_i(p', q-1), r_i(p', q) : 1 \leq p' < p\}$, and also a polynomial in $\{T_i(4, q)\} \cup \{r_i(p', q-1), r_i(p', q) : p \leq p' \leq 3\}$. Finally, the constraint of type (3) implies that for any $q \in [d_i]$, the variable $T_i(4, q)$ is a polynomial in $\{T_i(4, q') : q' \neq q\}$.

For part (a), suppose that y is a polynomial in W_{ij} . By the previous paragraph, if (W_{ij}, D_{ij}) is a constraint of type (1), then y can be written as a polynomial in x_q , where $\{x_q\} = W_{ij} \cap V_i$. If (W_{ij}, D_{ij}) is a constraint of type (2) then y can be written as a polynomial in $\{x_q\} \cup R_i$ for some $q \in [d_1]$ (and $W_{ij} \cap V_i = \emptyset$). If (W_{ij}, D_{ij}) is a constraint of type (4) then W_{ij} has size two, and y can be written as a polynomial in $\{x_q, x_{q'}\} \cup R_i$ for some $q, q' \in [d_i]$, where $W_{ij} \cap V_i \subseteq \{x_q, x_{q'}\}$. This finishes the proof of part (a).

For part (b), if (W_{ij}, D_{ij}) has type (3), then we can write y as a polynomial in $\{T_i(4, q) : q \in [d_i - 1]\}$. Suppose $M = T_i(4, q_1)^{a_1} \cdots T_i(4, q_k)^{a_k}$ is a monomial in this latter set of variables, where $k \geq 1$, $1 \leq q_1 < \dots < q_k < d_i$, and $0 \leq a_1, \dots, a_k < 120$. By Lemma 5.4.2, $T_i(4, q_j)^{a_j}$ is a polynomial in $\{x_{q_j}\} \cup \{r_i(p', q_j - 1), r_i(p', q_j) : p' \in [3]\}$ such that every monomial contains all the randomizers. When we multiply these polynomials together to get the monomial M , some of these randomizers may cancel out. However the randomizers $r_i(p', q_k)$ for $p' \in [3]$ appear only in the polynomial for $T_i(4, q_k)$. As a result, M is a polynomial in $V_i \cup R_i$ such that every monomial contains $r_i(p', q_k)$ for all $p' \in [3]$. We conclude that y can be written as a sum of monomials in $V_i \cup R_i$, such that each non-scalar monomial contains the randomizers $\{r_i(p', q) : p' \in [3]\}$ for some $q \in [d_i - 1]$. In particular, every non-scalar monomial contains some randomizer, finishing the proof of (b).

For part (c), suppose y and z are polynomials in W_{ij} and W_{ik} respectively. By part (a), if (W_{ij}, D_{ij}) and (W_{ik}, D_{ik}) are constraints of type (1), (2) or (4) then y and z both have V_i -degree less than or equal to two, and thus yz has V_i -degree less than or equal to four. Suppose without loss of generality that (W_{ij}, D_{ij}) is the constraint of type (3). If (W_{ik}, D_{ik}) is the same constraint, then yz is a polynomial in W_{ij} , and is covered by part (b).

Suppose (W_{ik}, D_{ik}) has type (2), so $W_{ik} = \{r_i(p, q-1), r_i(p, q), T_i(p, q), T_i(p+1, q)\}$ for some $p \in [3]$, $q \in [d_i]$. If $p \in [2]$, then z is a polynomial in $\{x_q\} \cup \{r_i(p', q-1), r_i(p', q) : 1 \leq p' \leq p\}$. Since y can be written as a polynomial in $V_i \cup R_i$ such that every non-scalar

monomial contains $r_i(3, q)$ for some $q \in [d_i - 1]$, yz can be written as a polynomial in $V_i \cup R_i$ such that every monomial either has V_i -degree at most one or contains $r_i(3, q)$ for some $q \in [d_i - 1]$. If $p = 3$, then z can be written as a polynomial in $\{T_i(4, q), r_i(3, q - 1), r_i(3, q)\}$ for some $q \in [d_i]$. For any $0 \leq a < 120$, $T_i(4, q)^a y$ can be written as a polynomial in $V_i \cup R_i$ such that every non-scalar monomial contains the randomizers $r_i(1, q')$, $r_i(2, q')$ for some $q' \in [d_i - 1]$. So yz is a polynomial in $V_i \cup R_i$ such that every monomial either has V_i -degree zero or contains $r_i(1, q')$, $r_i(2, q')$ for some $q' \in [d_i - 1]$.

Next suppose (W_{ik}, D_{ik}) has type (4), and let $F_i = \{T_i(4, q) : q \in [d_i]\}$. For $q \in [d_i]$, $T_i(1, q)$ can be written as a polynomial in x_q , $T_i(2, q)$ can be written as a polynomial in $\{x_q, r_i(1, q - 1), r_i(1, q)\}$, and $T_i(3, q)$ can be written as a polynomial in $\{T_i(4, q), r_i(3, q - 1), r_i(3, q)\}$. Hence every element W_i can be written as a polynomial in $V_i \cup R_i \cup F_i$ of V_i -degree at most one such that no monomial contains $r_i(p, q)$, $r_i(p', q)$ for some $q \in [d_i - 1]$, and $p \neq p'$. Thus z can be written as a polynomial $V_i \cup R_i \cup F_i$ with V_i -degree at most two, and such that for all $q \in [d_i - 1]$, no monomial contains all the randomizers $\{r_i(p, q) : p \in [3]\}$. If M is any monomial in F_i then My can be written as a polynomial in $V_i \cup R_i$ such that every non-scalar monomial contains $\{r_i(p, q) : p \in [3]\}$ for some $q \in [d_i - 1]$. Hence yz can be written as a polynomial in $V_i \cup R_i$ where every monomial either has V_i -degree at most two or contains a variable from R_i .

Finally if (W_{ik}, D_{ik}) has type (1), then z is a polynomial in x_q for some q , and as in the previous paragraph, yz can be written as a polynomial in $V_i \cup R_i$ where every monomial either has V_i -degree at most two or contains a variable from R_i . We conclude that part (c) holds. \square

Remark 5.4.4. *Note that the proof of Lemma 5.4.3, part (c) fails if we use a three-row tableau in Definition 5.3.4 rather than a four-row tableau. Indeed, suppose we used three-row tableaux. If (W_{ij}, D_{ij}) is the constraint of type (3), and (W_{ik}, D_{ik}) is the constraint of type (4) with $W_{ik} = \{r_i(1, q), r_i(2, q)\}$, then it is possible for yz to have monomials of degree ≥ 5 that do not contain any randomizers. For instance, when $q = 5$, we can take $y = T_i(3, 1) \cdots T_i(3, 5)$. This corresponds to the fact that, with three-row tableaux, we can recover the group product $T_i(1, 1) \cdots T_i(1, q)$ from the variables $T_i(3, q')$, $q' \in [q]$ and the randomizers $r_i(1, q)$, $r_i(2, q)$.*

Intuitively, Remark 5.4.4 prevents the three-row tableau from being perfect zero knowledge, since the verifier can gain access to a product of more than five of the permutations from the first row of the randomizing tableau. Using this information, the verifier may be able to construct one of the original non-oblivious variables from the BCS – MIP* protocol we reduce from. It is not necessarily possible to simulate the distribution of the players'

strategy for even one of the original variables in polynomial time. Adding a fourth row to the tableau solves this problem, as the verifier would need to learn three randomizers to reconstruct a product of elements on the first row of the tableau from an assignment to a type (3) constraint. No question allows the verifier to learn three randomizers. In the classical setting, constraints of type (4) are not necessary because everything already commutes, so a three-row tableau is enough.

Combining Lemma 5.4.3 with Lemma 5.3.2, for any BCS B we can define a perfect correlation p for the BCS game $\mathcal{G}(\text{Tab}_{\text{sub}}(\text{Obl}_5(B)))$ such that p is a quantum correlation if and only if $\mathcal{G}(B)$ has a perfect quantum strategy.

Proposition 5.4.5. *Suppose $B = (X, \{(V_i, C_i)\}_{i=1}^m)$ is a BCS with m constraints, and π is a probability distribution on $[m] \times [m]$ such that $\pi(i, j) > 0$ for all $i, j \in [m]$. Let $\text{Obl}_5(B) = (Z, \{(U_i, E_i)\}_{i=1}^m)$, $\text{Tab}(\text{Obl}_5(B)) = (Y, \{(W_i, D_i)\}_{i=1}^m)$, and $\text{Tab}_{\text{sub}}(\text{Obl}_5(B)) = (Y, \{(W_{ij}, D_{ij})\}_{i \in [m], j \in [m_i]})$. Let $R_i = \{r_i(j, k) : (j, k) \in [3] \times [d_i - 1]\}$ be the set of randomizers in W_i . For any $i \in [m]$ and $n \geq 1$, let $\Lambda_{i,n}$ be the set of non-scalar monomials over $U_i \cup R_i$ which either contain an element of R_i , or have degree at most n . Let Λ be the subspace of $\mathcal{A}_{c-c}(\text{Tab}(\text{Obl}_5(B)))$ defined by*

$$\Lambda = \mathbb{C}1 \oplus \text{span} \bigcup_{i \in [m]} \sigma_i(\Lambda_{i,4}) \oplus \text{span} \bigcup_{i \neq j \in [m]} \sigma_i(\Lambda_{i,2}) \sigma_j(\Lambda_{j,2}),$$

and let $f : \Lambda \rightarrow \mathbb{C}$ be the linear functional defined by $f(1) = 1$, $f(\sigma_i(x)) = 0$ for all $x \in \Lambda_{i,4}$, and $f(\sigma_i(x) \sigma_j(y)) = \delta_{xy}$ for all $x \in \Lambda_{i,2}$, $y \in \Lambda_{j,2}$, where δ_{ab} is the Kronecker delta, i.e. $\delta_{ab} = 1$ if $a = b$, and is 0 otherwise. Let $\alpha : \mathcal{A}_{c-c}(\text{Tab}_{\text{sub}}(\text{Obl}_5(B))) \rightarrow \mathcal{A}_{c-c}(\text{Tab}(\text{Obl}_5(B)))$ be the homomorphism sending $\sigma_{ij}(x) \mapsto \sigma_i(x)$ for all $x \in \mathcal{A}(W_{ij}, D_{ij})$, as in Proposition 3.4.3.

For every $i, k \in [m], j \in [m_i], l \in [m_k]$ and assignments ϕ and ψ to W_{ij} and W_{kl} respectively, let

$$p(\phi, \psi | ij, kl) = f(\alpha(\Phi_{W_{ij}, \phi} \Phi_{W_{kl}, \psi})).$$

Then p is a perfect correlation for the BCS game $\mathcal{G}(\text{Tab}_{\text{sub}}(\text{Obl}_5(B)), \pi_{\text{sub}})$, and $p \in C_q$ (resp. C_{qa}, C_{qc}) if and only if $\mathcal{G}(B, \pi)$ has a perfect quantum correlation in C_q (resp. C_{qa}, C_{qc}).

Proof. We first observe that the linear functional f is well-defined, by showing that it can be defined on a larger subspace. Indeed, for any set of variables S , let $\mathcal{M}(S)$ be the set of non-scalar monomials in S , and let $\mathcal{M}_n(S) \subseteq \mathcal{M}(S)$ be the subset of monomials of degree at most n . Since we're assuming that (V_i, C_i) has at least one satisfying assignment, $\mathcal{A}(V_i, C_i)$

has a tracial state. Applying part (b) of Lemma 5.3.2 to the constraint system containing the single constraint (V_i, C_i) , we see that $\mathcal{A}(U_i, E_i)$ has a tracial state τ_i such that $\tau_i(x) = 0$ for all $x \in \mathcal{M}_4(U_i)$. Hence $\mathbb{C}1 \cap \text{span } \mathcal{M}_4(U_i) = \{0\}$ in $\mathcal{A}(U_i, E_i)$. By Lemma 5.3.2, part (c), the set $\mathcal{M}_2(U_i)$ is linearly independent in $\mathcal{A}(U_i, E_i)$. Hence we can choose a basis Ξ_i for $\mathcal{A}(U_i, E_i)$ which contains $\{1\} \cup \mathcal{M}_2(U_i)$, and such that $\text{span } \mathcal{M}_4(U_i) \subseteq \text{span } \Xi_i \setminus \{1\}$. By Lemma 5.4.1, part (a), the set $\{ab : a \in \Xi_i, b \in \mathcal{M}(R_i)\}$ is a basis for $\mathcal{A}(W_i, D_i)$. Let Θ_i be the set of non-identity elements in this basis. Because $\mathcal{A}_{c-c}(\text{Tab}(B))$ is a free product of the algebras $\mathcal{A}(W_i, D_i)$, the set

$$\Theta := \{1\} \cup \bigcup_{i \in [m]} \sigma_i(\Theta_i) \cup \bigcup_{i \neq j \in [m]} \sigma_i(\Theta_i)\sigma_j(\Theta_j)$$

is linearly independent in $\mathcal{A}_{c-c}(\text{Tab}(B))$. Define a linear functional f on the span of Θ by setting $f(1) = 1$, $f(\sigma_i(x)) = 0$ for all $x \in \Theta_i$, and

$$f(\sigma_i(x)\sigma_j(y)) = \begin{cases} 1 & x \text{ and } y \text{ are both in } \mathcal{M}_2(U_i \cap U_j) \text{ and } x = y \\ 0 & \text{otherwise} \end{cases}$$

for all $x \in \Theta_i, y \in \Theta_j$ with $i \neq j$. The image of the set $\Lambda_{i,4}$ in $\mathcal{A}(W_i, D_i)$ is contained in the span of Θ_i , so the span of Θ contains the subspace Λ . Furthermore, if $x \in \Lambda_{i,4}$, then $f(\sigma_i(x)) = 0$. Suppose $x \in \Lambda_{i,2}$ and $y \in \Lambda_{j,2}$ with $i \neq j$. If x contains an element of R_i , then x is contained in the span of $\{ab : a \in \mathcal{M}(U_i), b \in \mathcal{M}(R_i), b \neq 1\}$, and $f(\sigma_i(x)\sigma_j(y)) = 0 = \delta_{xy}$. The same is true if y contains an element of R_j . If neither x or y contains an element of R_i or R_j respectively, then $x \in \mathcal{M}_2(U_i)$ and $y \in \mathcal{M}_2(U_j)$ are elements of Θ_i and Θ_j respectively. The only way for x and y to be equal is if both belong to $\mathcal{M}_2(U_i \cap U_j)$, so $f(\sigma_i(x)\sigma_j(y)) = \delta_{xy}$. Thus the restriction of f to Λ is the linear functional defined in the proposition.

Since f is well defined, Lemma 5.4.3 implies that p is well defined. Since $\sum_{\phi} \Phi_{W_{ij}\phi} = 1$, it follows that $\sum_{\phi, \psi} p(\phi, \psi | ij, kl) = 1$ for every $i, k \in [m], j \in [m_i]$ and $l \in [m_k]$. To show that p is a perfect correlation for $\mathbf{G}(\text{Tab}_{\text{sub}}(\text{Obl}_5(B)))$, we need to show that $p(\phi, \psi | ij, kl) \geq 0$ for all $\phi \in D_{ij}, \psi \in D_{kl}, i, k \in [m], j \in [m_i]$ and $l \in [m_k]$, and that $p(\phi, \psi | ij, kl) = 0$ if $\phi|_{W_{ij} \cap W_{kl}} \neq \psi|_{W_{ij} \cap W_{kl}}$. If $i = k$, then $\alpha(\Phi_{W_{ij}, \phi})$ and $\alpha(\Phi_{W_{kl}, \psi})$ are both projections in the commutative algebra $\mathcal{A}(W_i, D_i)$, and thus their product is also a projection. Since C_i is non-empty by assumption, $\mathcal{A}(V_i, C_i)$ has a tracial state. If $B_i = (V_i, \{(V_i, C_i)\})$ is the constraint system for the single constraint C_i , then $\text{Obl}_5(B_i) = (U_i, \{(U_i, E_i)\})$ and $\text{Tab}(\text{Obl}_5(B_i)) = (W_i, \{(W_i, D_i)\})$. By Lemma 5.3.2, part (b), there is a tracial state τ_i on $\mathcal{A}(U_i, E_i)$ such that $\tau_i(x) = \delta_{x,1}$ for all $x \in \mathcal{M}_4(U_i)$. By Lemma 5.4.1, part (d), there is a tracial state $\tilde{\tau}_i$ on $\mathcal{A}(W_i, D_i)$ such that $\tilde{\tau}_i(x) = \tau_i(x)$ for all $x \in$

$\mathcal{M}(U_i)$, and $\tilde{\tau}_i(x) = 0$ for all monomials $x \in \mathcal{M}(U_i \cup R_i)$ containing an element of R_i . Since $\tilde{\tau}_i(1) = 1$ and $\tilde{\tau}_i(x) = 0$ for all $x \in \Lambda_{i,4}$, the linear functionals f and $\tilde{\tau}_i$ agree on $\mathbb{C}1 \oplus \Lambda_{i,4}$, and $f(\alpha(\Phi_{W_{ij},\phi}\Phi_{W_{kl},\psi})) = \tilde{\tau}_i(\alpha(\Phi_{W_{ij},\phi}\Phi_{W_{kl},\psi})) \geq 0$. If $\phi|_{W_{ij} \cap W_{kl}} \neq \psi|_{W_{ij} \cap W_{kl}}$ then $\alpha(\Phi_{W_{ij},\phi})\alpha(\Phi_{W_{kl},\psi}) = 0$ in $\mathcal{A}(W_i, D_i)$, and $f(\alpha(\Phi_{W_{ij},\phi}\Phi_{W_{kl},\psi})) = 0$.

If $i \neq k$ and neither (W_{ij}, D_{ij}) or (W_{kl}, D_{kl}) are constraints of type (3) in Definition 5.3.4, then by Lemma 5.4.3 there exist $S_i \subseteq U_i$ and $S_k \subseteq U_k$ of size at most two, such that $W_{ij} \cap U_i \subseteq S_i$, $W_{kl} \cap U_k \subseteq S_k$, $\Phi_{W_{ij},\phi}$ is a polynomial in $S_i \cup R_i$, and $\Phi_{W_{kl},\psi}$ is a polynomial in $S_k \cup R_k$. Since $\mathcal{M}_2(S_i)$ is a linearly independent in $\mathcal{A}(U_i, E_i)$, part (a) of Lemma 5.4.1 implies that the subalgebra of $\mathcal{A}(W_i, D_i)$ generated by $S_i \cup R_i$ is isomorphic to $\mathbb{C}\mathbb{Z}_2^{S_i} \times \mathbb{Z}_{120}^{R_i}$, and similarly with the algebra generated by $S_k \cup R_k$ in $\mathcal{A}(W_k, D_k)$. Hence the subalgebra \mathcal{C} of $\mathcal{A}_{c-c}(\text{Tab}(\text{Obl}_5(B)))$ generated by $S_i \cup S_k \cup R_i \cup R_k$ is isomorphic to the group algebra of $(\mathbb{Z}_2^{S_i} \times \mathbb{Z}_{120}^{R_i}) * (\mathbb{Z}_2^{S_k} \times \mathbb{Z}_{120}^{R_k})$. Let H be the quotient of this free product group by the relations $\sigma_i(x) = \sigma_k(x)$ for all $x \in S_i \cap S_k$, where $\sigma_i(x)$ and $\sigma_k(x)$ are the group generators corresponding to x in the first and second factors of the free product respectively, and let

$$q : (\mathbb{Z}_2^{S_i} \times \mathbb{Z}_{120}^{R_i}) * (\mathbb{Z}_2^{S_k} \times \mathbb{Z}_{120}^{R_k}) \rightarrow H$$

be the quotient map. Observe that

$$H = \mathbb{Z}_2^{*S_i \cup S_k} * \mathbb{Z}_{120}^{*R_i \cup R_k} / \langle xy = yx \text{ for } x, y \text{ in } S_i \cup R_i \text{ or } S_k \cup R_k \rangle$$

is a graph product. By the normal form theorem for graph products [28], if $g \in \mathcal{M}(S_i \cup R_i)$ and $h \in \mathcal{M}(S_j \cup R_j)$, then $q(gh) = 1$ if and only if $g, h \in \mathcal{M}(S_i \cap S_j)$ and $g = h$. Hence if τ is the canonical trace on the group algebra $\mathbb{C}H$, then $\tau \circ q(\sigma_i(g)\sigma_k(h)) = f(\sigma_i(g)\sigma_k(h))$. We conclude that $f(\alpha(\Phi_{W_{ij},\phi}\Phi_{W_{kl},\psi})) = \tau \circ q(\alpha(\Phi_{W_{ij},\phi}\Phi_{W_{kl},\psi}))$. Since $\Phi_{W_{ij},\phi}$ and $\Phi_{W_{kl},\psi}$ are projections, $\tau \circ q(\alpha(\Phi_{W_{ij},\phi}\Phi_{W_{kl},\psi})) \geq 0$. Suppose $\phi(x) \neq \psi(x)$ for some $x \in W_{ij} \cap W_{kl}$. Then we must have $x \in U_i \cap U_k$, so $x \in S_i \cap S_k$. Since $\frac{1+\phi(x)x}{2}\Phi_{W_{ij},\phi} = \Phi_{W_{ij},\phi}$ and $\frac{1+\psi(x)x}{2}\Phi_{W_{kl},\psi} = \Phi_{W_{kl},\psi}$, we have

$$\begin{aligned} q(\alpha(\Phi_{W_{ij},\phi}\Phi_{W_{kl},\psi})) &= q\left(\alpha(\Phi_{W_{ij},\phi})\sigma_i\left(\frac{1+\phi(x)x}{2}\right)\sigma_k\left(\frac{1+\psi(x)x}{2}\right)\alpha(\Phi_{W_{kl},\psi})\right) \\ &= q(\alpha(\Phi_{W_{ij},\phi}))\left(\frac{1+\phi(x)x}{2}\right)\left(\frac{1+\psi(x)x}{2}\right)q(\alpha(\Phi_{W_{kl},\psi})) = 0. \end{aligned}$$

Thus $f(\alpha(\Phi_{W_{ij},\phi}\Phi_{W_{kl},\psi})) = 0$ if $\phi|_{W_{ij} \cap W_{kl}} \neq \psi|_{W_{ij} \cap W_{kl}}$.

Finally, suppose $i \neq k$ and (W_{ij}, D_{ij}) is a constraint of type (3). By Lemma 5.4.3, part (b), we can write $\alpha(\Phi_{W_{ij},\phi}) = \lambda 1 + \sum_x c_x x$ for some coefficients $\lambda, c_x \in \mathbb{C}$, where the sum is over monomials $x \in \mathcal{M}(U_i \cup R_i)$ containing an element of R_i . Let $\tilde{\tau}_i$ and $\tilde{\tau}_k$ be the

tracial states on $\mathcal{A}(W_i, D_i)$ and $\mathcal{A}(W_k, D_k)$ defined above. Since $\tilde{\tau}_i$ is equal to f on $\Lambda_{i,4}$ and $\Phi_{W_{ij},\phi}$ is a projection, $\lambda = f(\alpha(\Phi_{W_{ij},\phi})) = \tilde{\tau}_i(\alpha(\Phi_{W_{ij},\phi})) \geq 0$. Similarly, $f(\alpha(\Phi_{W_{kl},\psi})) = \tilde{\tau}_k(\alpha(\Phi_{W_{kl},\psi})) \geq 0$. If $x \in \mathcal{M}(U_i \cup R_i)$ contains an element of R_i , then $x\alpha(\Phi_{W_{kl},\psi}) \in \sigma_i(\Lambda_{i,4}) \oplus \sigma_i(\Lambda_{i,2})\sigma_k(\Lambda_{k,2})$, so $f(x\alpha(\Phi_{W_{kl},\psi})) = 0$. We conclude that $f(\alpha(\Phi_{W_{ij},\phi}\Phi_{W_{kl},\psi})) = \lambda f(\alpha(\Phi_{W_{kl},\psi})) \geq 0$. It is not possible to have $\phi|_{W_{ij} \cap W_{kl}} \neq \psi|_{W_{ij} \cap W_{kl}}$ when $i \neq k$ and (W_{ij}, D_{ij}) has type (3), since $W_{ij} \cap W_{kl} = \emptyset$.

This finishes the proof that p is a perfect correlation for $\mathfrak{G}(\text{Tab}_{sub}(\text{Obl}_5(B)), \pi_{sub})$. If $p \in C_{qc}$ (resp. C_{qa}, C_q), then $\mathfrak{G}(\text{Tab}(\text{Obl}_5(B)), \pi)$ also has a perfect strategy in C_{qc} (resp. C_{qa}, C_q) by Proposition 3.4.3. This means that there is a tracial state (resp. Connes-embeddable tracial state, finite-dimensional tracial state) $\tilde{\tau}$ on $\mathcal{A}_{c-c}(\text{Tab}(\text{Obl}_5(B)))$ with $df(\tau; \mu_\pi) = 0$. By Lemma 5.4.1, part (b), and Lemma 5.3.2, part (a), there is a 1-homomorphism $\mathcal{A}_{c-c}(B) \rightarrow \mathcal{A}_{c-c}(\text{Tab}(\text{Obl}_5(B)))$, and pulling back $\tilde{\tau}$ by this 1-homomorphism yields a perfect strategy for $\mathfrak{G}(B, \pi)$ in C_{qc} (resp. C_{qa}, C_q). Conversely, if $\mathfrak{G}(B, \pi)$ has a perfect strategy in C_{qc} , then there is a tracial state τ on $\mathcal{A}_{c-c}(B)$ with $df(\tau; \mu_\pi) = 0$. By Lemma 5.3.2, part (b), there is a tracial state τ' on $\mathcal{A}_{c-c}(\text{Obl}_5(B))$ such that $df(\tau'; \mu_\pi) = 0$, $\tau'(\sigma_i(x)) = 0$ for all $x \in \mathcal{M}_4(U_i) \setminus \{1\}$, $i \in [m]$, and $\tau'(\sigma_i(x)\sigma_j(y)) = 0$ for all $x \in \mathcal{M}_2(U_i)$, $y \in \mathcal{M}_2(U_j)$, $i \neq j \in [m]$ with $x \neq y$. By Lemma 5.4.1, part (d), there is a tracial state $\tilde{\tau}$ on $\mathcal{A}_{c-c}(\text{Tab}(\text{Obl}_5(B)))$ with $df(\tilde{\tau}; \mu_\pi) = 0$, $\tilde{\tau}(u) = \tau'(u)$ for all monomials u in $\{\sigma_i(x) : i \in [m], x \in U_i\}$, and $\tilde{\tau}(u\sigma_i(z)^a v) = 0$ for all $i \in [m]$, $z \in R_i$, $1 \leq a < 120$, and monomials u, v in $\{\sigma_j(x) : j \in [m], x \in U_j \cup R_j\}$ which do not contain z . Observe that $\tilde{\tau}(1) = 1$, and $\tilde{\tau}(\sigma_i(x)) = 0$ for all $x \in \Lambda_{i,4}$. Similarly, if $x \in \Lambda_{i,2}$ and $y \in \Lambda_{j,2}$ are not equal, then $\tilde{\tau}(\sigma_i(x)\sigma_j(y)) = 0$. By Proposition 3.3.3, $df(\tilde{\tau}; \mu_{inter}) = 0$, and since $\pi(i, j) > 0$ for all $i, j \in [m]$, $\|\sigma_i(x) - \sigma_j(x)\|_{\tilde{\tau}} = 0$ for all $x \in U_i \cap U_j$. Since $\|\cdot\|_{\tilde{\tau}}$ is unitarily bi-invariant, we get that $\|\sigma_i(x) - \sigma_j(x)\|_{\tilde{\tau}} = 0$ for all $x \in \mathcal{M}(U_i \cap U_j)$, and hence $\tilde{\tau}(\sigma_i(x)\sigma_j(x)) = 1$ for all $x \in \mathcal{M}(U_i \cap U_j)$. It follows that $\tilde{\tau}|_\Lambda = f$, so $p(\phi, \psi|_{ij, kl}) = \tilde{\tau} \circ \alpha(\Phi_{W_{ij},\phi}\Phi_{W_{kl},\psi})$ for all $\phi \in D_{ij}$, $\psi \in D_{kl}$, $i, k \in [m]$, $j \in [m_i]$, $l \in [m_k]$. We conclude that $p \in C_{qc}$. If $\mathcal{A}_{c-c}(B)$ has a perfect strategy in C_{qa} (resp. C_q), then we can take τ to be Connes-embeddable (resp. finite-dimensional), so $\tilde{\tau}$ will be Connes-embeddable (resp. finite-dimensional), and $p \in C_{qa}$ (resp. C_q). \square

Remark 5.4.6. *The correlation p in Proposition 5.4.5 is described algebraically. Alternatively, it's not hard to see that the correlation $p(\phi, \psi|_{ij, kl})$ can be simulated using the following procedure: If neither (W_{ij}, D_{ij}) or (W_{kl}, D_{kl}) has type (3), then pick an assignment to the variables $Z \cup R$ uniformly at random, and fill in the variables $T_i(p, q)$ so that the constraints (W_{ij}, D_{ij}) of types (1), (2), and (4) are satisfied, to get an assignment γ to Y . Output $\phi = \gamma|_{W_{ij}}$ and $\psi = \gamma|_{W_{kl}}$. If one of (W_{ij}, D_{ij}) or (W_{kl}, D_{kl}) has type (3), then for each $r \in [m]$, pick an assignment to R_r and a satisfying assignments to (U_r, E_r) uniformly at random, and fill in the variables $T_j(p, q)$ to get a satisfying assignment γ_r*

to (W_r, D_r) . Output $\phi = \gamma_i|_{W_{ij}}$ and $\psi = \gamma_k|_{W_{kl}}$. This procedure will output ϕ, ψ with probability $p(\phi, \psi|ij, kl)$.

The description of the correlation p in Remark 5.4.6 is simpler than the algebraic description. On the other hand, without the algebraic description, it's harder to see that the correlation generated in Remark 5.4.6 is quantum when $\mathcal{A}_{c-c}(B)$ has a perfect quantum strategy. In fact, if we use three-row tableaux rather than four-row tableaux in Definition 5.3.4, then the procedure in Remark 5.4.6 is still well-defined, and simulates a perfect strategy for the game. However, by Remark 5.4.4, the simulated correlation is not necessarily quantum even if $\mathcal{A}_{c-c}(B)$ has a perfect quantum strategy — something that is not immediately apparent from the description of the procedure.

We are now ready to prove this chapter's main result.

Theorem 5.4.7. *Let $(\{\mathbf{G}(B_x, \pi_x)\}, Q, C)$ be a BCS-MIP* protocol for a language \mathcal{L} with completeness 1 and soundness $1 - f(x)$, such that each context of B_x has constant size, and π_x is E -diagonally dominant for some constant $E > 0$. Then there is a PZK-BCS-MIP* protocol $(\{\mathbf{G}(B'_x, \pi'_x)\}, \tilde{Q}, \tilde{C})$ for \mathcal{L} with completeness 1 and soundness $1 - Df(x)$, where $D > 0$ is the a universal constant. If π_x is uniform, then π'_x is also uniform, and if $\mathbf{G}(B_x, \pi_x)$ has a perfect finite-dimensional tracial state, then so does $\mathbf{G}(B'_x, \pi'_x)$.*

Proof. Let $B'_x = \text{Tab}_{\text{sub}}(\text{Obl}_5(B_x))$, and let π'_x be the subdivision of π_x corresponding to the subdivision of $\text{Tab}(\text{Obl}_5(B_x))$ into $\text{Tab}_{\text{sub}}(\text{Obl}_5(B_x))$. If π_x is uniform, then π'_x is also uniform. Let p_x be the correlation for $\mathbf{G}(B'_x, \pi'_x)$ defined in Proposition 5.4.5. Because B_x has contexts of constant size, $\text{Obl}_5(B_x)$ and $\text{Tab}(\text{Obl}_5(B_x))$ also have contexts of constant size. As a result, the number of clauses in the constraints of $\text{Tab}(\text{Obl}_5(B_x))$ is constant, as is the size of each clause (where by clause we mean the constraints of type (1)-(4) in Definition 5.3.4). Hence the Turing machines Q and C can be turned into Turing machines \tilde{Q} and \tilde{C} such that $(\{\mathbf{G}(B'_x, \pi'_x)\}, \tilde{Q}, \tilde{C})$ is a BCS-MIP* protocol. Similarly, since all the constraints of $\text{Tab}(\text{Obl}_5(B))$ have constant size, there is a Turing machine which, given questions and answers i, j, ϕ, ψ for $\mathbf{G}(B'_x, \pi'_x)$, can produce $p_x(\phi, \psi|i, j)$ in polynomial time in i, j , and x . Since the number of answers for any question is constant, the correlation p_x can be simulated in polynomial time in x .

If $x \in \mathcal{L}$, then B_x has a perfect strategy in C_{qa} , so $p_x \in C_{qa}$, and hence $\mathbf{G}(B'_x, \pi'_x)$ has a perfect strategy in C_{qa} . Similarly, if B_x has a perfect strategy in C_q , then $\mathbf{G}(B'_x, \pi'_x)$ has a perfect strategy in C_q as well. Conversely, suppose that τ is a tracial state on $\mathcal{A}_{c-c}(B'_x)$. Since the size of contexts and number of clauses in each constraint of $\text{Tab}(\text{Obl}_5(B_x))$ are constant, the parameters L, M , and K in Theorem 3.4.4 when going from $\text{Tab}(\text{Obl}_5(B_x))$

to $\text{Tab}_{\text{sub}}(\text{Obl}_5(B_x))$ are all constant. Theorem 3.4.4 implies that there is a tracial state τ_0 on $\mathcal{A}_{c-c}(\text{Tab}(\text{Obl}_5(B_x)))$ with $\text{def}(\tau_0) \leq \text{poly}(2^L, M, K) \text{def}(\tau)$. Since there is a classical homomorphism $\mathcal{A}_{c-c}(B_x) \rightarrow \mathcal{A}_{c-c}(\text{Tab}(\text{Obl}_5(B_x)))$ by Lemmas 5.3.2 and 5.4.1, we conclude that there is a tracial state τ_1 on $\mathcal{A}_{c-c}(B_x)$ with $\text{def}(\tau_1) \leq \text{poly}(2^L, M, K) \text{def}(\tau)$. Hence if $x \notin \mathcal{L}$, then there is no synchronous strategy p for $\mathbb{G}(B'_x, \pi'_x)$ with $\mathfrak{w}_q(\mathbb{G}(B'_x, \pi'_x), p) \geq 1 - f(n)/\text{poly}(2^L, M, K)$. Hence $(\{\mathbb{G}(B'_x, \pi'_x)\}, \tilde{S}, \tilde{C})$ is a BCS-MIP* protocol for \mathcal{L} with soundness $1 - f(x)/\text{poly}(2^L, M, K)$. \square

Theorem 5.4.8. *There is a perfect zero knowledge BCS-MIP*(2, 1, 1, 1/2) protocol for the halting problem in which the questions have length $\text{poly}(n)$ and the answers have constant length. Furthermore, if a game in the protocol has a perfect strategy, then it has a perfect synchronous quantum strategy.*

Proof. By Corollary 3.1.2, there is a BCS-MIP* protocol $(\{\mathbb{G}(B_x, \pi_x)\}, Q, V)$ for the halting problem with constant soundness $s < 1$, where B_x has exponentially many contexts of constant size. Furthermore, if $(\{\mathbb{G}(B_x, \pi_x)\}, Q, V)$ has a perfect strategy, then it has a perfect synchronous strategy. Then by Theorem 5.4.7 there is a perfect zero knowledge BCS-MIP*(2, 1, 1, s') protocol for the halting problem with constant soundness $s' > 0$, polynomial length questions and constant length answers. A constant amount of parallel repetition yields the result. \square

The proof of Theorem 5.4.8 implies that the halting problem is many-one reducible to membership in C_q . In fact, there is a reduction such that if the Turing machine does not halt, then the corresponding correlation is bounded away from the closure C_{qa} of C_q :

Corollary 5.4.9. *There is a polynomial-time computable function p from Turing machines to synchronous correlations such that if M halts then $p(M) \in C_q^s$, and if M does not halt then there is a linear functional f on the space of correlations such that $f(p(M)) = 1$ and $f(p') \leq 1/2$ for all $p' \in C_{qa}$.*

Proof. Let $(\{\mathbb{G}(B_M, \pi_M)\}, S, C)$ be the BCS-MIP* protocol for the halting problem with completeness one and soundness $1/2$ constructed in the proof of Theorem 5.4.8, where the index M runs through Turing machines. Let $p(M)$ be the correlation for $\mathbb{G}(B_M, \pi_M)$ as in Definition 5.1.1. That $p(M)$ is in C_q follows from Theorem 5.4.8, and the fact that if $p \in C_q$, then $p^{\otimes n} \in C_q$. The corollary then follows with the linear functional f defined by $f(p') = \mathfrak{w}(\mathbb{G}(B_M, \pi_M), p')$. \square

Note that the number of inputs for the correlation $p(M)$ depends on the size of the Turing machine M .

If we change the input protocol of the reduction, we can get a PZK – MIP* protocol for the halting problem with a uniform question distribution at the cost of polynomial length answers. Consider the following strong version of MIP* = RE due to Natarajan and Zhang [55].

Theorem 5.4.10. *There is a two-prover one round AM*(2) protocol $(\{\mathbf{G}_x\}, Q, V)$ for the halting problem with completeness $c = 1$ and soundness $s = 1/2$, such that \mathbf{G}_x is a synchronous game with constant length questions, and $\text{polylog}(|x|)$ length answers. Furthermore, if \mathbf{G}_x has a perfect strategy, then it has a perfect oracularizable synchronous quantum strategy.*

Recall that AM*(2) is the complexity class of languages with two-prover MIP* protocols where the verifier chooses their messages to the prover uniformly at random. We can now show the following.

Theorem 5.4.11. *There is a perfect zero knowledge BCS-MIP*(2, 1, 1, 1 – 1/poly(n)) protocol for the halting problem in which the verifier selects questions according to the uniform distribution, the questions have length $\text{polylog}(n)$, and the answers have constant length. Furthermore, if a game in the protocol has a perfect strategy, then it has a perfect synchronous quantum strategy.*

Proof. By oracularizing Theorem 5.4.10, there is a BCS-MIP* protocol $(\{\mathbf{G}(B_x, \pi_x)\}, S, V)$ for the halting problem with constant soundness $s < 1$, in which B_x has a constant number of contexts and contexts of size $\text{polylog}(|x|)$, and π_x is the uniform distribution on pairs of contexts. Furthermore, if $\mathbf{G}(B_x, \pi_x)$ has a perfect strategy, then it has a perfect synchronous quantum strategy. By applying CSP reductions in each context $(\{\mathbf{G}(B_x, \pi_x)\}, S, C)$ can be turned into a BCS-MIP* protocol $(\{\mathbf{G}(B'_x, \pi_x)\}, S, C)$ where $B'_x = (X'_x, \{(W_i^x, D_i^x)\})$, D_i is a 3SAT instance with number of clauses polynomial in $|x|$, and $|W_i^x|$ is polynomial in $|x|$. Then by subdividing the B'_x into a 3SAT we obtain a 3SAT protocol $(\{\mathbf{G}(B_x^{3SAT}, \pi_x^{3SAT})\}, S, C)$ with number of clauses polynomial in $|x|$, and π_x^{3SAT} is uniform. There is a 1-homomorphism $\mathcal{A}_{c-c}(B_x^{3SAT}, \mu_{\pi_x^{3SAT}}) \rightarrow \mathcal{A}_{c-c}(B_x, \mu_{\pi_x})$, so if $\mathbf{G}(B_x, \pi_x)$ has a perfect synchronous quantum strategy, so does $\mathbf{G}(B_x^{3SAT}, \pi_x^{3SAT})$. The theorem follows from Theorem 5.4.7. \square

Theorem 5.4.12. *There is a perfect zero knowledge BCS-MIP*(2, 1, 1, 1/2) protocol for the halting problem in which the verifier chooses questions uniformly at random, and the questions and answers have length $\text{poly}(n)$. Furthermore, if a game in the protocol has a perfect strategy, then it has a perfect synchronous quantum strategy.*

Proof. Let $(\{\mathbf{G}(B_x, \pi_x)\}, S, C)$ be the BCS-MIP* protocol from Theorem 5.4.11, so in particular B_x has m_x contexts, where $m_x = \text{poly}(|x|)$, and π_x is the uniform distribution on $[m_x] \times [m_x]$. Since the uniform distribution is $1/2m_x$ -diagonally dominant, Theorem 2.5.1 implies that $(\{\mathbf{G}(B_x, \pi_x)\}, S, C)$ has soundness $1 - 1/\text{poly}(n)$ when considered as a MIP* protocol. The result follows from Theorem 5.2.1 using a polynomial amount of parallel repetition. \square

Finally, we also have:

Theorem 5.4.13. $\text{PZK-BCS-MIP}^{co}(2, 1, 1, 1 - 1/\text{poly}(n)) = \text{BCS-MIP}^{co}(2, 1, 1, 1 - 1/\text{poly}(n))$.

The proof is similar to the proof of Theorem 5.4.8, but without the parallel repetition.

References

- [1] Alex Arkhipov. Extending and characterizing quantum magic games. *arXiv preprint arXiv:1209.3819*, 2012.
- [2] Sanjeev Arora and Boaz Barak. *Computational Complexity: A Modern Approach*. Cambridge University Press, USA, 1st edition, 2009.
- [3] Albert Atserias, Phokion G Kolaitis, and Simone Severini. Generalized satisfiability problems via operator assignments. In *Fundamentals of Computation Theory: 21st International Symposium, FCT 2017, Bordeaux, France, September 11–13, 2017, Proceedings 21*, pages 56–68. Springer, 2017.
- [4] L. Babai, L. Fortnow, and C. Lund. Nondeterministic exponential time has two-prover interactive protocols. In *Proceedings [1990] 31st Annual Symposium on Foundations of Computer Science*, pages 16–25 vol.1, 1990.
- [5] László Babai, Lance Fortnow, and Carsten Lund. Non-deterministic exponential time has two-prover interactive protocols. *Computational complexity*, 1:3–40, 1991.
- [6] D A Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in NC1. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing*, STOC '86, page 1–5, New York, NY, USA, 1986. Association for Computing Machinery.
- [7] J. S. Bell. On the Einstein Podolsky Rosen paradox. *Physics Physique Fizika*, 1:195–200, Nov 1964.
- [8] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing*, STOC '88, page 113–131, New York, NY, USA, 1988. Association for Computing Machinery.

- [9] Bruce Blackadar. *Operator algebras: theory of C^* -algebras and von Neumann algebras*, volume 122. Springer Berlin, Heidelberg, 2006.
- [10] Gilles Brassard, Richard Cleve, and Alain Tapp. Cost of exactly simulating quantum entanglement with classical communication. *Phys. Rev. Lett.*, 83:1874–1877, Aug 1999.
- [11] Andrei A Bulatov. A dichotomy theorem for nonuniform CSPs. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 319–330. IEEE, 2017.
- [12] Peter J. Cameron, Ashley Montanaro, Michael W. Newman, Simone Severini, and Andreas Winter. On the quantum chromatic number of a graph. *Electronic Journal of Combinatorics*, 14(1), 2007.
- [13] Michael Chapman, Thomas Vidick, and Henry Yuen. Efficiently stable presentations from error-correcting codes. *arXiv preprint arXiv:2311.04681*, 2023.
- [14] Alessandro Chiesa, Michael A. Forbes, Tom Gur, and Nicholas Spooner. Spatial isolation implies zero knowledge even in a quantum world. *J. ACM*, 69(2), jan 2022.
- [15] R. Cleve, P. Hoyer, B. Toner, and J. Watrous. Consequences and limits of nonlocal strategies. In *Proceedings. 19th IEEE Annual Conference on Computational Complexity, 2004.*, pages 236–249, 2004.
- [16] Richard Cleve and Rajat Mittal. Characterization of binary constraint system games. In *Automata, Languages, and Programming: 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part I 41*, pages 320–331. Springer, 2014.
- [17] Matt Coudron and William Slofstra. Complexity lower bounds for computing the approximately-commuting operator value of nonlocal games to high precision. *Computational Complexity Conference (CCC)*, 2019.
- [18] Eric Culf and Kieran Mastel. RE-completeness of entangled constraint satisfaction problems. *arXiv preprint arXiv:2410.21223*, 2024.
- [19] Eric Culf, Hamoon Mousavi, and Taro Spirig. Approximation Algorithms for Non-commutative CSPs . In *2024 IEEE 65th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 920–929, Los Alamitos, CA, USA, October 2024. IEEE Computer Society.

- [20] Yangjing Dong, Honghao Fu, Anand Natarajan, Minglong Qin, Haochen Xu, and Penghui Yao. The computational advantage of MIP* vanishes in the presence of noise. *arXiv preprint arXiv:2312.04360*, 2023.
- [21] Cynthia Dwork, Uriel Feige, Joe Kilian, Moni Naor, and Shmuel Safra. Low communication 2-prover zero-knowledge proofs for NP. In *Annual International Cryptology Conference*, 1992.
- [22] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47:777–780, May 1935.
- [23] Joseph Fitzsimons, Zhengfeng Ji, Thomas Vidick, and Henry Yuen. Quantum proof systems for iterated exponential time, and beyond. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2019, page 473–480, New York, NY, USA, 2019. Association for Computing Machinery.
- [24] Honghao Fu, Carl A. Miller, and William Slofstra. The membership problem for constant-sized quantum correlations is undecidable. *Communications in Mathematical Physics*, 406(5):96, 2025.
- [25] Viktor Galliard and Stefan Wolf. Pseudo-telepathy, entanglement, and graph colorings. In *Proceedings IEEE International Symposium on Information Theory*,, page 101. IEEE, 2002.
- [26] Adina Goldberg. Synchronous linear constraint system games. *Journal of Mathematical Physics*, 62(3), mar 2021.
- [27] S Goldwasser, S Micali, and C Rackoff. The knowledge complexity of interactive proof-systems. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing*, STOC '85, page 291–304, New York, NY, USA, 1985. Association for Computing Machinery.
- [28] Elisabeth Ruth Green. Graph products of groups. PhD thesis, 1990.
- [29] Alex Bredariol Grilo, William Slofstra, and Henry Yuen. Perfect zero knowledge for quantum multiprover interactive proofs. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 611–635. IEEE Computer Society, 2019.
- [30] Samuel J. Harris. Universality of graph homomorphism games and the quantum coloring problem. *arXiv preprint arXiv:2305.18116*, 2023.

- [31] Samuel J Harris. Approximate quantum 3-colorings of graphs and the quantum max 3-cut problem. *arXiv preprint arXiv:2412.19405*, 2024.
- [32] Samuel J Harris. Universality of graph homomorphism games and the quantum coloring problem. In *Annales Henri Poincaré*, pages 1–36. Springer, 2024.
- [33] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.
- [34] J William Helton, Kyle P Meyer, Vern I Paulsen, and Matthew Satriano. Algebras, synchronous games, and chromatic numbers of graphs. *New York J. Math*, 25:328–361, 2019.
- [35] Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto. Oracularization and two-prover one-round interactive proofs against nonlocal strategies. In *2009 24th Annual IEEE Conference on Computational Complexity*, pages 217–228, 2009.
- [36] Tsuyoshi Ito and Thomas Vidick. A multi-prover interactive proof for NEXP sound against entangled provers. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 243–252, 2012.
- [37] Zhengfeng Ji. Binary constraint system games and locally commutative reductions. *arXiv preprint arXiv:1310.3794*, 2013.
- [38] Zhengfeng Ji. Classical verification of quantum proofs. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '16, page 885–898, New York, NY, USA, 2016. Association for Computing Machinery.
- [39] Zhengfeng Ji. Compression of quantum multi-prover interactive proofs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2017, page 289–302, New York, NY, USA, 2017. Association for Computing Machinery.
- [40] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. $MIP^*=RE$. *arXiv preprint arXiv:2001.04383*, 2022.
- [41] Julia Kempe, Hirotada Kobayashi, Keiji Matsumoto, Ben Toner, and Thomas Vidick. Entangled games are hard to approximate. *SIAM Journal on Computing*, 40(3):848–877, 2011.
- [42] Julia Kempe, Oded Regev, and Ben Toner. Unique games with entangled provers are easy. *SIAM Journal on Computing*, 39(7):3207–3229, 2010.

- [43] S. Khot. On the power of unique 2-prover 1-round games. In *Proceedings 17th IEEE Annual Conference on Computational Complexity*, page 25, 2002.
- [44] Joe Kilian. *Uses of randomness in algorithms and protocols*. MIT Press, 1990.
- [45] Se-Jin Kim, Vern Paulsen, and Christopher Schafhauser. A synchronous game for binary constraint systems. *Journal of Mathematical Physics*, 59(3), Mar 2018.
- [46] Junqiao Lin. Tracial embeddable strategies: Lifting MIP* tricks to MIPco. *arXiv preprint arXiv:2304.01940*, 2024.
- [47] Amine Marrakchi and Mikael de la Salle. Almost synchronous correlations and tomita-takesaki theory. *arXiv preprint arXiv:2307.08129*, 2023.
- [48] Kieran Mastel and William Slofstra. Two prover perfect zero knowledge for MIP*. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024*, page 991–1002, New York, NY, USA, 2024. Association for Computing Machinery.
- [49] N. David Mermin. Simple unified form for the major no-hidden-variables theorems. *Phys. Rev. Lett.*, 65:3373–3376, Dec 1990.
- [50] Hamoon Mousavi and Taro Spirig. A quantum unique games conjecture. *arXiv preprint arXiv:2409.20028*, 2024.
- [51] A. Natarajan and J. Wright. NEXP is contained in MIP*. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 510–518, Los Alamitos, CA, USA, nov 2019. IEEE Computer Society.
- [52] Anand Natarajan and Chinmay Nirkhe. The status of the quantum pcp conjecture (games version). *arXiv preprint arXiv:2403.13084*, 2024.
- [53] Anand Natarajan and Thomas Vidick. Low-degree testing for quantum states, and a quantum entangled games PCP for QMA. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, Oct 2018.
- [54] Anand Natarajan and Thomas Vidick. Two-player entangled games are NP-hard. In *Proceedings of the 33rd Computational Complexity Conference, CCC '18, Dagstuhl, DEU*, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

- [55] Anand Natarajan and Tina Zhang. Quantum free games. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, STOC 2023, page 1603–1616, New York, NY, USA, 2023. Association for Computing Machinery.
- [56] Narutaka Ozawa. About the Connes embedding conjecture: algebraic approaches. *Japanese Journal of Mathematics*, 8(1):147–183, 2013.
- [57] Connor Paddock. Rounding near-optimal quantum strategies for nonlocal games to strategies using maximally entangled states. *arXiv preprint arXiv:2203.02525*, 2022.
- [58] Connor Paddock and William Slofstra. Satisfiability problems and algebras of boolean constraint system games. *arXiv preprint arXiv:2310.07901*, 2023.
- [59] Asher Peres. Incompatible results of quantum measurements. *Physics Letters A*, 151(3):107–108, 1990.
- [60] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, STOC '08, page 245–254, New York, NY, USA, 2008. Association for Computing Machinery.
- [61] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing*, STOC '78, page 216–226, New York, NY, USA, 1978. Association for Computing Machinery.
- [62] Konrad Schmüdgen. *An Invitation to Unbounded Representations of *-Algebras on Hilbert Space*. Springer International Publishing, 2020.
- [63] Adi Shamir. $IP = PSPACE$. *J. ACM*, 39(4):869–877, oct 1992.
- [64] William Slofstra. The set of quantum correlations is not closed. *Forum of Mathematics, Pi*, 7(E1), 2019.
- [65] Thomas Vidick. CS286 Seminar in Computer Science: Around the quantum PCP conjecture, 2014. URL: users.cms.caltech.edu/~vidick/teaching/286_qPCP/. Last visited on 2024/10/04.
- [66] Thomas Vidick. Three-player entangled XOR games are NP-hard to approximate. *SIAM Journal on Computing*, 45(3):1007–1063, 2016.
- [67] Thomas Vidick. Erratum: Three-player entangled XOR games are NP-hard to approximate. *SIAM Journal on Computing*, 49(6):1423–1427, 2020.

- [68] Thomas Vidick. Almost synchronous quantum correlations. *Journal of Mathematical Physics*, 63(2), Feb 2022.
- [69] Henry Yuen. A Parallel Repetition Theorem for All Entangled Games. In *43rd International Colloquium on Automata, Languages, and Programming (ICALP 2016)*, volume 55 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 77:1–77:13, Dagstuhl, Germany, 2016. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [70] Dmitriy Zhuk. A proof of CSP dichotomy conjecture. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 331–342. IEEE, 2017.