



Generalized satisfiability problems via operator assignments

Albert Atserias^{a,*}, Phokion G. Kolaitis^{b,c}, Simone Severini^{d,e}

^a Universitat Politècnica de Catalunya, Barcelona, Catalonia, Spain

^b University of California Santa Cruz, Santa Cruz, CA, USA

^c IBM Research–Almaden, San Jose, CA, USA

^d University College London, London, United Kingdom

^e Shanghai Jiao Tong University, Shanghai, China



ARTICLE INFO

Article history:

Received 9 December 2017

Received in revised form 20 March 2019

Accepted 4 May 2019

Available online 21 May 2019

Keywords:

Constraint satisfaction problem

Quantum satisfiability

Non-local games

Dichotomy theorems

Linear operators

Undecidable problems

pp-Definitions

ABSTRACT

Schaefer introduced a framework for generalized satisfiability problems on the Boolean domain and characterized the computational complexity of such problems. We investigate an algebraization of Schaefer's framework in which the Fourier transform is used to represent constraints by multilinear polynomials in a unique way. This representation of constraints gives rise to a relaxation of the notion of satisfiability in which the values to variables are linear operators on some Hilbert space. For constraints given by a system of linear equations over the two-element field, earlier work in the foundations of quantum mechanics has shown that there are systems that have no solutions in the Boolean domain, but have solutions via operator assignments on some finite-dimensional Hilbert space. Our main result is a complete characterization of the classes of Boolean relations for which there is a gap between satisfiability in the Boolean domain and the relaxation of satisfiability via operator assignments.

© 2019 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction and summary of results

In 1978, Schaefer [22] classified the computational complexity of generalized satisfiability problems. Each class A of Boolean relations gives rise to the generalized satisfiability problem $\text{SAT}(A)$. An instance of $\text{SAT}(A)$ is a conjunction of relations from A such that each conjunct has a tuple of variables as arguments; the question is whether or not there is an assignment of Boolean values to the variables, so that, for each conjunct, the resulting tuple of Boolean values belongs to the underlying relation. Schaefer's main result is a dichotomy theorem for the computational complexity of $\text{SAT}(A)$, namely, depending on A , either $\text{SAT}(A)$ is NP-complete or $\text{SAT}(A)$ is solvable in polynomial time. Schaefer's dichotomy theorem provided a unifying explanation for the NP-completeness of many well-known variants of Boolean satisfiability, such as POSITIVE 1-IN-3 SAT and MONOTONE 3SAT; moreover, it became the catalyst for numerous subsequent investigations, including the pursuit of a dichotomy theorem for constraints satisfaction problems, a pursuit that became known as the Feder-Vardi Conjecture [9].

Every Boolean relation can be identified with its characteristic function, which, via the Fourier transform, can be represented as a multilinear polynomial (i.e., a polynomial in which each variable has degree at most one) in a unique way. Moreover, in carrying out this transformation, the truth values *false* and *true* are typically represented by $+1$ and -1 , in-

* Corresponding author.

E-mail addresses: atserias@cs.upc.edu (A. Atserias), kolaitis@soe.ucsc.edu (P.G. Kolaitis), s.severini@ucl.ac.uk (S. Severini).

stead of 0 and 1. For example, it is easy to see that the multilinear polynomial representing the conjunction $x \wedge y$ of two variables x and y is $\frac{1}{2}(1 + x + y - xy)$. The multilinear polynomial representation of Boolean relations makes it possible to consider relaxations of satisfiability in which the variables take values in some suitable space, instead of the two-element Boolean algebra. Such relaxations have been considered in the foundations of physics several decades ago, where they have played a role in singling out the differences between classical theory and quantum theory. In particular, it has been shown that there is a system of linear equations over the two-element field that has no solutions over $\{+1, -1\}$, but the system of the associated multilinear polynomials has a solution in which the variables are assigned linear operators on a Hilbert space of dimension four. The Mermin-Peres magic square [16,17,20] is the most well known example of such a system. These constructions give *small proofs* of the celebrated Kochen-Specker Theorem [8] on the impossibility to explain quantum mechanics via hidden-variables [2]. More recently, systems of linear equations with this relaxed notion of solvability have been studied under the name of *binary constraint systems*, and tight connections have been established between solvability and the existence of perfect strategies in non-local games that make use of entanglement [6,7].

A Boolean relation is *affine* if it is the set of solutions of a system of linear equations over the two-element field. The collection LIN of all affine relations is prominent in Schaefer's dichotomy theorem, as it is one of the main classes A of Boolean relations for which $\text{SAT}(A)$ is solvable in polynomial time. The discussion in the preceding paragraph shows that $\text{SAT}(\text{LIN})$ has instances that are unsatisfiable in the Boolean domain, but are satisfiable when linear operators on a Hilbert space are assigned to variables (for simplicity, from now on we will use the term "operator assignments" for such assignments). Which other classes of Boolean relations exhibit such a gap between satisfiability in the Boolean domain and the relaxation of satisfiability via operator assignments? As a matter of fact, this question bifurcates into two separate questions, depending on whether the relaxation allows linear operators on Hilbert spaces of arbitrary (finite or infinite) dimension or only on Hilbert spaces of finite dimension. In a recent breakthrough paper, Slofstra [24] showed that these two questions are different for LIN by establishing the existence of systems of linear equations that are satisfiable by operator assignments on some infinite-dimensional Hilbert space, but are not satisfiable by operator assignments on any finite-dimensional Hilbert space. In a related vein, Ji [15] showed that a 2CNF-formula is satisfiable in the Boolean domain if and only if it is satisfiable by an operator assignment in some finite-dimensional Hilbert space. Moreover, Ji showed that the same holds true for Horn formulas. Note that 2SAT, HORN SAT, and DUAL HORN SAT also feature prominently in Schaefer's dichotomy theorem as, together with $\text{SAT}(\text{LIN})$, which from now on we will denote by LIN SAT, they constitute the main tractable cases of generalized satisfiability problems (the other tractable cases are the trivial cases of $\text{SAT}(A)$, where A is a class of 0-valid relations or a class of 1-valid relations, i.e., Boolean relations that contain the tuple consisting entirely of 0's or, respectively, the tuple consisting entirely of 1's).

In this paper, we completely characterize the classes A of Boolean relations for which $\text{SAT}(A)$ exhibits a gap between satisfiability in the Boolean domain and satisfiability via operator assignments. Clearly, if every relation in A is 0-valid or every relation in A is 1-valid, then there is no gap, as every constraint is satisfied by assigning to every variable the identity operator or its negation, respectively. Beyond this, we first generalize and extend Ji's results [15] by showing that if Γ is a class of Boolean relations such that every relation in A is bijunctive,¹ or every relation in A is Horn, or every relation in A is dual Horn,² then there is no gap whatsoever; this means that an instance of $\text{SAT}(A)$ is satisfiable in the Boolean domain if and only if it is satisfiable by an operator assignment on some finite-dimensional Hilbert space if and only if it is satisfiable by an operator assignment on some arbitrary Hilbert space. In contrast, we show that for all other classes A of Boolean relations, $\text{SAT}(A)$ exhibits a two-level gap: there are instances of $\text{SAT}(A)$ that are not satisfiable in the Boolean domain, but are satisfiable by an operator assignment on some finite-dimensional Hilbert space; moreover, there are instances of $\text{SAT}(A)$ that are not satisfiable by an operator assignment on any finite-dimensional Hilbert space, but are satisfiable by an operator assignment on some (infinite-dimensional) Hilbert space. For this last part of the statement we need to assume that A contains the full binary relation or, alternatively, that A admits a so-called *commutativity gadget*. This is a mild assumption that is met by all the classes of Boolean relations from the literature; we discuss this in Section 4.5.

The proof of this result uses several different ingredients. First, we use the substitution method [7] to show that there is no satisfiability gap for classes of relations that are bijunctive, Horn, and dual Horn. This gives a different proof of Ji's results [15], which were for finite-dimensional Hilbert spaces, but also shows that, for such classes of relations, there is no difference between satisfiability by linear operators on finite-dimensional Hilbert spaces and satisfiability by linear operators on arbitrary Hilbert spaces. The main tool for proving the existence of a two-level gap for the remaining classes of Boolean relations is the notion of *pp-definability*, that is, definability via primitive-positive formulas, which are existential first-order formulas having a conjunction of (positive) atoms as their quantifier-free part. In the past, primitive-positive formulas have been used to design polynomial-time reductions between decision problems; in fact, this is one of the main techniques in the proof of Schaefer's dichotomy theorem. Here, we show that primitive-positive formulas can also be used to design *gap-preserving* reductions, that is, reductions that preserve the gap between satisfiability on the Boolean domain and satisfiability by operator assignments. To prove the existence of a two-level gap for classes of Boolean relations we combine gap-preserving reductions with the two-level gap for LIN discussed earlier (i.e., the results of Mermin [16,17], Peres [16], and Slofstra [24]) and with results about Post's lattice of clones on the Boolean domain [21].

¹ A Boolean relation is *bijunctive* if it is the set of satisfying assignments of a 2CNF-formula.

² A Boolean relation is *Horn (dual Horn)* if it is the set of satisfying assignments of a Horn (dual Horn) formula.

We also give two additional applications of pp-definability. First, we consider an extension of pp-definability in which the existential quantifiers may range over linear operators on some finite-dimensional Hilbert space. At first sight, it appears that new Boolean relations may be pp-definable in the extended sense from a given set of Boolean relations. We show, however, that this is not the case. Specifically, by analyzing closure operations on sets of linear operators, we show that if a Boolean relation is pp-definable in the extended sense from other Boolean relations, then it is also pp-definable from the same relations. In other words, for Boolean relations, this extension of pp-definability is not more powerful than standard pp-definability. Second, we apply pp-definability to the problem of quantum realizability of contextuality scenarios. Recently, Fritz [12] used Slofstra’s results [24] to resolve two problems raised by Acin et al. in [1]. Using pp-definability and Slofstra’s results, we obtain new proofs of Fritz’s results that have the additional feature that the parameters involved are optimal.

2. Definitions and technical background

2.1. Notation

For an integer n , we write $[n]$ for the set $\{1, \dots, n\}$. We use mainly the $+1, -1$ representation of the Boolean domain ($+1$ for “false” and -1 for “true”). We write $\{\pm 1\}$ for the set $\{+1, -1\}$. If a denotes a tuple of length r we write a_1, \dots, a_r to denote its r components. If a is such a tuple and f is a function that has a_1, \dots, a_r in its domain, we write $f(a)$ to denote the tuple $(f(a_1), \dots, f(a_r))$. We write T and F for the full and empty Boolean relations, respectively. The letters stand for *true* and *false*. Their arity is unspecified by the notation and will be made clear by the context.

2.2. Linear operators and polynomials thereof

Let V be a complex vector space. A linear operator on V is a linear map from V to V . The linear operator that is the identity on V is denoted by I , and the linear operator that is identically 0 is denoted by 0 . The pointwise addition of two linear operators A and B is denoted by $A + B$, the composition of two linear operators A and B is denoted by AB , and the pointwise scaling of a linear operator A by a scalar $c \in \mathbb{C}$ is denoted by cA . All these are linear operators. As a result, if $\mathbb{C}\langle X_1, \dots, X_n \rangle$ denotes the ring of polynomials with complex coefficients and non-commuting variables in X_1, \dots, X_n , then for a polynomial $P(X_1, \dots, X_n)$ in $\mathbb{C}\langle X_1, \dots, X_n \rangle$ and linear operators A_1, \dots, A_n on V , the notation $P(A_1, \dots, A_n)$ makes sense. If A_1, \dots, A_n pairwise commute, i.e., $A_i A_j = A_j A_i$ for all $i, j \in \{1, \dots, n\}$, then the notation makes sense even for a polynomial in $\mathbb{C}[X_1, \dots, X_n]$, the ring of polynomials with commuting variables in X_1, \dots, X_n .

Let V and W be complex vector spaces. Let A be a linear operator on V and let B be a linear operator on W . We say that A and B are similar if there exists an invertible linear map $C : V \rightarrow W$ such that $A = CBC^{-1}$. Let A_1, \dots, A_n and B_1, \dots, B_n be linear operators on V and W , respectively. We say that A_1, \dots, A_n and B_1, \dots, B_n are simultaneously similar if there exists an invertible linear map $C : V \rightarrow W$ such that $A_i = CB_i C^{-1}$ holds for all $i \in [n]$. The following simple fact with an equally simple proof will be used multiple times.

Lemma 1. *Let V and W be complex vector spaces, and let $P(X_1, \dots, X_n)$ be a polynomial in $\mathbb{C}\langle X_1, \dots, X_n \rangle$. If A_1, \dots, A_n and B_1, \dots, B_n are simultaneously similar linear operators on V and W , respectively, then so are $P(A_1, \dots, A_n)$ and $P(B_1, \dots, B_n)$.*

Proof. We write $[n]^*$ for the set of finite sequences with components in $[n]$, and $|\alpha|$ for the length of the sequence α . Let $P(X_1, \dots, X_n) = \sum_{\alpha \in [n]^*} c_\alpha \prod_{i=1}^{|\alpha|} X_{\alpha_i}$, where only finitely many of the coefficients c_α are non-zero. Let $C : V \rightarrow W$ be an invertible linear map witnessing that A_1, \dots, A_n and B_1, \dots, B_n are simultaneously similar; thus $A_j = CB_j C^{-1}$ holds for every $j \in [n]$. Note that for every $\alpha \in [n]^*$ of length ℓ we have $\prod_{i=1}^\ell (CB_{\alpha_i} C^{-1}) = C(\prod_{i=1}^\ell B_{\alpha_i})C^{-1}$. It follows that $P(A_1, \dots, A_n) = \sum_{\alpha \in [n]^*} c_\alpha C(\prod_{i=1}^{|\alpha|} B_{\alpha_i})C^{-1}$, and linearity gives $P(A_1, \dots, A_n) = CP(B_1, \dots, B_n)C^{-1}$. \square

2.3. Unique multilinear polynomial representations

A polynomial $P(X_1, \dots, X_n)$ is called multilinear if it has individual degree at most one on each variable. Each function $f : \{\pm 1\}^n \rightarrow \mathbb{C}$ has a unique representation as a multilinear polynomial in $\mathbb{C}\langle X_1, \dots, X_n \rangle$ given by the Fourier or Walsh-Hadamard transform [18]. Explicitly:

$$P_f(X_1, \dots, X_n) = \sum_{S \subseteq [n]} \hat{f}(S) \prod_{i \in S} X_i, \tag{1}$$

where

$$\hat{f}(S) = \frac{1}{2^n} \sum_{a \in \{\pm 1\}^n} f(a) \prod_{i \in S} a_i. \tag{2}$$

The polynomial represents f in the sense that $P_f(a) = f(a)$ holds for every $a \in \{\pm 1\}^n$. If the range of f is a subset of \mathbb{R} , then each $\hat{f}(S)$ is indeed a real number. The Convolution Formula describes the Fourier coefficients of pointwise products $f g$ of functions $f, g : \{\pm 1\}^n \rightarrow \mathbb{C}$. It states that

$$\widehat{fg}(S) = \sum_{T \subseteq [n]} \widehat{f}(S) \widehat{g}(S \Delta T) \quad (3)$$

for every $S \subseteq [n]$, where $S \Delta T$ denotes symmetric difference; i.e. $S \Delta T = (S \setminus T) \cup (T \setminus S)$.

We give an example of use of the uniqueness of the Fourier transform that will be useful later on. We begin by recalling some notation and terminology. A *literal* is a Boolean variable x or its negation $\neg x$. The literals x and $\neg x$ are said to be *complementary* of each other, and x is their *underlying variable*. If ℓ is a literal, then $\bar{\ell}$ denotes its complementary literal. The *sign* $\text{sg}(\ell)$ of ℓ is defined as follows: $\text{sg}(\ell) = 1$ if $\ell = x$, and $\text{sg}(\ell) = -1$ if $\ell = \neg x$, where x is its underlying variable. Clearly, $\text{sg}(\bar{\ell}) = -\text{sg}(\ell)$.

A *clause* is a disjunction of literals. Let $C = (\ell_1 \vee \dots \vee \ell_r)$ be a clause. In the ± 1 representation of Boolean values, the clause C represents the relation $\{\pm 1\}^r \setminus \{(\text{sg}(\ell_1), \dots, \text{sg}(\ell_r))\}$, which will be denoted by R_C . The *indicator* function of the clause $C = (\ell_1 \vee \dots \vee \ell_r)$ is the Boolean function from $\{\pm 1\}^r \rightarrow \{\pm 1\}$ that maps the tuple $(\text{sg}(\ell_1), \dots, \text{sg}(\ell_r))$ to $+1$ and every other tuple to -1 . We write $P_C(X_1, \dots, X_r)$ to denote the unique multilinear polynomial representation of the indicator function of the clause C .

Lemma 2. Let $C = (\ell_1 \vee \dots \vee \ell_r)$ be a clause on r different variables. Then, in the ring of polynomials $\mathbb{C}[X_1, \dots, X_r]$, the following identity holds.

$$P_C(X_1, \dots, X_r) = 2^{1-r} \prod_{i=1}^r (1 + \text{sg}(\ell_i) X_i) - 1. \quad (4)$$

Proof. Let $R_C = \{\pm 1\}^r \setminus \{(\text{sg}(\ell_1), \dots, \text{sg}(\ell_r))\}$ be the Boolean relation represented by C . Since the right-hand side of equation (4) is a multilinear polynomial and its left-hand side is the unique multilinear polynomial that agrees with the indicator function of R_C on $\{\pm 1\}$, it suffices to check that the right-hand side also agrees with the indicator function of R_C on $\{\pm 1\}^r$. In other words, we claim that for every $(a_1, \dots, a_r) \in \{\pm 1\}^r$, the right-hand side evaluates to -1 if the truth-assignment (a_1, \dots, a_r) satisfies the clause C , and it evaluates to 1 , otherwise.

Assume that (a_1, \dots, a_r) satisfies the clause C . Then there is some $j \in \{1, \dots, r\}$ such that $a_j = -\text{sg}(\ell_j)$. It follows that $1 + \text{sg}(\ell_j) a_j = 0$ and so $\prod_{i=1}^r (1 + \text{sg}(\ell_i) a_i) = 0$, which, in turn, implies that $P_C(a_1, \dots, a_r) = -1$. Assume that (a_1, \dots, a_r) does not satisfy the clause C . Then, for every $i \in \{1, \dots, r\}$, we have that $a_i = \text{sg}(\ell_i)$. Consequently, for every $i \in \{1, \dots, r\}$, we have that $1 + \text{sg}(\ell_i) a_i = 2$ and so $\prod_{i=1}^r (1 + \text{sg}(\ell_i) a_i) = 2^r$, which, in turn, implies that $P_C(a_1, \dots, a_r) = 1$. This completes the proof of the claim. \square

2.4. Hilbert space

A Hilbert space is a complex vector space with an inner product whose norm induces a complete metric. All Hilbert spaces of finite dimension d are isomorphic to \mathbb{C}^d with the standard complex inner product. In particular, this means that after the choice of a basis, we can identify the linear operators on a d -dimensional Hilbert space with the matrices in $\mathbb{C}^{d \times d}$. Composition of operators becomes matrix multiplication. A matrix A is Hermitian if it is equal to its conjugate transpose A^* . A diagonal matrix is one all whose off-diagonal entries are 0. A matrix A is unitary if $A^* A = A A^* = I$, where I is the identity matrix. Two matrices A and B commute if $AB = BA$, and a collection of matrices A_1, \dots, A_r pairwise commute if $A_i A_j = A_j A_i$ for all $i, j \in [r]$.

For the basics of general Hilbert spaces and their linear operators we refer the reader to Halmos' monograph [14]. We need from it the concepts of bounded linear operator and of adjoint A^* of a densely defined linear operator A . Two operators A and B commute if $AB = BA$. A sequence of operators A_1, \dots, A_r pairwise commute if $A_i A_j = A_j A_i$ for all $i, j \in [r]$. A linear operator A is called normal if it commutes with its adjoint A^* ; i.e., $AA^* = A^*A$. A linear operator is called self-adjoint if $A^* = A$. A linear map from a Hilbert space \mathcal{H}_1 to another Hilbert space \mathcal{H}_2 is called unitary if it preserves norms.

We also make elementary use of general L^2 - and L^∞ -spaces. Let $(\Omega, \mathcal{M}, \mu)$ be a measure space. Then $L^2(\Omega, \mu)$ denotes the collection of square integrable measurable functions, up to almost everywhere equality. Also $L^\infty(\Omega, \mu)$ denotes the collection of essentially bounded measurable functions, up to almost everywhere equality. All measure-theoretic terms in these definitions refer to μ . See [11] for definitions.

2.5. Constraint languages, instances, value and satisfiability

A *Boolean constraint language* A is a collection of relations over the Boolean domain $\{\pm 1\}$. Let $V = \{X_1, \dots, X_n\}$ be a set of variables. An *instance* \mathcal{I} on the variables V over the constraint language A is a finite collection of pairs

$$\mathcal{I} = ((Z_1, R_1), \dots, (Z_m, R_m)) \quad (5)$$

where each R_i is a relation from A and $Z_i = (Z_{i,1}, \dots, Z_{i,r_i})$ is a tuple of variables from V or constants from $\{\pm 1\}$, where r_i is the arity of R_i . Each pair (Z_i, R_i) is called a *constraint*, and each Z_i is called its *constraint-scope*. A *Boolean assignment* is a

mapping f assigning a Boolean value $a_i \in \{\pm 1\}$ to each variable X_i , and assigning -1 and $+1$ to the constants -1 and $+1$, respectively. We say that the assignment satisfies the i -th constraint if the tuple $f(Z_i) = (f(Z_{i,1}), \dots, f(Z_{i,r_i}))$ belongs to R_i . The value of f on \mathcal{I} is the fraction of constraints that are satisfied by f . The value of \mathcal{I} , denoted by $v(\mathcal{I})$, is the maximum value over all Boolean assignments. We say that \mathcal{I} is *satisfiable* in the Boolean domain if there is a Boolean assignment that satisfies all constraints; equivalently, if $v(\mathcal{I}) = 1$.

Three examples of Boolean constraint languages that will play an important role in the sequel are LIN, 3LIN, and 1-IN-3-SAT. The constraint language LIN consists of all affine relations, i.e., the Boolean relations whose tuples are solutions to a system of linear equations over the two-element field. In the ± 1 -representation, every such equation is a parity equation of the form $\prod_{i=1}^r x_i = b$, where $b \in \{\pm 1\}$. Note that LIN is an infinite constraint language since it contains infinitely many different relations. The constraint language 3LIN is the finite constraint language that consists of the two types of parity equations of arity three; in other words, $3LIN = \{(a_1, a_2, a_3) \in \{\pm 1\}^3 : a_1 a_2 a_3 = b\} : b \in \{\pm 1\}$. Finally, the constraint language 1-IN-3-SAT consists of a single relation, namely, the relation $R_{1/3} = \{(+1, +1, -1), (+1, -1, +1), (-1, +1, +1)\}$ of Boolean triples that have a single -1 . Observe that $3LIN \subseteq LIN$, but $1\text{-IN-3-SAT} \not\subseteq LIN$ since the cardinality of every affine subspace over the two-element field is an exact power of 2.

The instances over 3LIN are precisely the systems of linear equations over the two-element field with three variables per equation. Note that, formally, the variables in each equation need not be different. For example, the single linear equation $X_1 + X_2 + X_2 = 1 \pmod 2$ or, equivalently, the parity equation $X_1 X_2 X_2 = -1$ in the ± 1 -representation, is a valid instance over 3LIN. The set of Boolean assignments that satisfy this instance is the set of two assignments that map $X_1 \mapsto -1$ and $X_2 \mapsto b$ for any $b \in \{\pm 1\}$.

2.6. Operator assignments and satisfiability via operators

Let X_1, \dots, X_n be n variables, and let \mathcal{H} be a Hilbert space. An *operator assignment* for X_1, \dots, X_n over \mathcal{H} is an assignment of a bounded linear operator on \mathcal{H} to each variable, $f : X_1, \dots, X_n \mapsto A_1, \dots, A_n$, such that the following conditions hold:

1. A_j is self-adjoint for every $j \in [n]$,
2. $A_j^2 = I$ for every $j \in [n]$.

If S is a subset of $\{X_1, \dots, X_n\}$, we say that the operator assignment A_1, \dots, A_n *pairwise commutes* on S if in addition it satisfies $A_j A_k = A_k A_j$ for every X_j and X_k in the set S . If it pairwise commutes on the whole set $\{X_1, \dots, X_n\}$, we say that the assignment *fully commutes*.

Let A be a Boolean constraint language, let \mathcal{I} be an instance over A , with n variables X_1, \dots, X_n as in (5), and let \mathcal{H} be a Hilbert space. An *operator assignment for \mathcal{I} over \mathcal{H}* is an operator assignment $f : X_1, \dots, X_n \mapsto A_1, \dots, A_n$ for the variables X_1, \dots, X_n that pairwise commutes on the set of variables of each constraint scope Z_i in \mathcal{I} ; explicitly

$$A_j A_k = A_k A_j \quad \text{for every } X_j \text{ and } X_k \text{ in } Z_i, \text{ for every } i \in \{1, \dots, m\}. \tag{6}$$

We also require that f maps the constant -1 and $+1$ to $-I$ and I , respectively, where I is the identity operator on \mathcal{H} . We say that the assignment f *satisfies* the i -th constraint if

$$P_{R_i}(f(Z_i)) = P_{R_i}(f(Z_{i,1}), \dots, f(Z_{i,r_i})) = -I, \tag{7}$$

where P_{R_i} denotes the unique multilinear polynomial representation of indicator function of the relation R_i , i.e., the function that maps each tuple in R_i to -1 , and each tuple in its complement $\{\pm 1\}^{r_i} \setminus R_i$ to $+1$. Note that since $f(Z_{i,1}), \dots, f(Z_{i,r_i})$ are required to commute by definition, this notation is unambiguous despite the fact that P_{R_i} is defined as a polynomial in commuting variables. The value of f on \mathcal{I} is the fraction of constraints that are satisfied by f ; note that this quantity takes one of a finite set of values in the set $\{0, 1/m, 2/m, \dots, (m-1)/m, 1\}$. The value of \mathcal{I} over \mathcal{H} is the maximum value over all operator assignments for \mathcal{I} over \mathcal{H} . We say that f *satisfies \mathcal{I}* if it satisfies all constraints. In that case we also say that f is a *satisfying operator assignment for \mathcal{I} over \mathcal{H}* .

The *finite-dimensional value of \mathcal{I}* , denoted by $v^*(\mathcal{I})$, is the maximum of its value over all finite-dimensional Hilbert spaces. The *value of \mathcal{I}* , denoted by $v^{**}(\mathcal{I})$, is the maximum of its value over all Hilbert spaces. We say that an instance \mathcal{I} is *satisfiable via finite-dimensional operator assignments*, or *satisfiable via fd-operators* for short, if $v^*(\mathcal{I}) = 1$. We say that \mathcal{I} is *satisfiable via operator assignments*, or *satisfiable via operators* for short, if $v^{**}(\mathcal{I}) = 1$.

3. The strong Spectral Theorem

The Spectral Theorem plays an important role in linear algebra and functional analysis. It has also been used in the foundations of quantum mechanics (for some recent uses see [7,15]). We will make a similar use of it, but we will also need the version of this theorem for infinite-dimensional Hilbert spaces. In this section we discuss the statement, both for finite- and infinite-dimensional Hilbert spaces, as well as one of its important applications that we encapsulate in a lemma for later reuse.

3.1. Statement

In its most basic form, the Spectral Theorem for complex matrices states that every Hermitian matrix is unitarily equivalent to a diagonal matrix. Explicitly: if A is a $d \times d$ Hermitian matrix, then there exist a unitary matrix U and a diagonal matrix E such that $A = U^{-1}EU$. In its strong form, the Strong Spectral Theorem (SST) applies to sets of pairwise commuting Hermitian matrices and is stated as follows.

Theorem 1 (Strong Spectral Theorem; finite-dimensional case). *Let A_1, \dots, A_r be $d \times d$ Hermitian matrices, for some positive integer d . If A_1, \dots, A_r pairwise commute, then there exists a unitary matrix U and diagonal matrices E_1, \dots, E_r such that $A_i = U^{-1}E_iU$ for every $i \in [r]$.*

This form of the SST will be enough to discuss satisfiability via fd-operators. For operator assignments over arbitrary Hilbert spaces, we need to appeal to the most general form of the SST in which the role of diagonal matrices is played by multiplication operators on an $L^2(\Omega, \mu)$ -space. These are defined as follows.

Let V be a complex function space; a complex vector space of functions mapping indices from an index set X to \mathbb{C} . A multiplication operator of V is a linear operator whose value at a function $f : X \rightarrow \mathbb{C}$ in V is given by pointwise multiplication by a fixed function $a : X \rightarrow \mathbb{C}$. In symbols, the multiplication operator given by a is

$$(T_a(f))(x) = a(x)f(x) \quad \text{for each } x \in X. \quad (8)$$

In its weak form, the general Spectral Theorem states that any normal bounded linear operator on a Hilbert space is unitarily equivalent to a multiplication operator on an L^2 -space. We need the following strong version of the Spectral Theorem that states that the same is true for a collection of such operators, simultaneously through the same unitary transformation, provided they commute. The statement we use is a direct consequence of Theorem 1.47 in Folland's monograph [10].

Theorem 2 (Strong Spectral Theorem; general case). *Let A_1, \dots, A_r be normal bounded linear operators on a Hilbert space \mathcal{H} . If A_1, \dots, A_r pairwise commute, then there exist a measure space $(\Omega, \mathcal{M}, \mu)$, a unitary map $U : \mathcal{H} \rightarrow L^2(\Omega, \mu)$, and functions $a_1, \dots, a_r \in L^\infty(\Omega, \mu)$ such that $A_i = U^{-1}T_{a_i}U$ for every $i \in [r]$.*

The special case in which \mathcal{H} has finite dimension d , the measure space is actually a finite set of cardinality d with the counting measure, and thus $L^2(\Omega, \mu)$ is isomorphic to \mathbb{C}^d with the usual complex inner product.

3.2. An oft-used application

The following lemma encapsulates a frequently used application of the Strong Spectral Theorem. It states that whenever a set of polynomial equations entail another polynomial equation over the Boolean domain, then the entailment holds as well for fully commuting operator assignments.

Lemma 3. *Let X_1, \dots, X_r be variables, let Q_1, \dots, Q_m, Q be polynomials in $\mathbb{C}[X_1, \dots, X_r]$, and let \mathcal{H} be a Hilbert space. If every Boolean assignment that satisfies the equations $Q_1 = \dots = Q_m = 0$ also satisfies the equation $Q = 0$, then every fully commuting operator assignment over \mathcal{H} that satisfies the equations $Q_1 = \dots = Q_m = 0$ also satisfies the equation $Q = 0$.*

Although the same proof applies to all Hilbert spaces, the proof of the finite-dimensional case can be made more elementary. Since for certain applications only the finite-dimensional case of the lemma is relevant, we split the proof accordingly into cases.

Proof of Lemma 3; finite-dimensional case. Assume \mathcal{H} has finite dimension d . Since all Hilbert spaces of dimension d are isometrically isomorphic to \mathbb{C}^d , let us assume without loss of generality that $\mathcal{H} = \mathbb{C}^d$. In such a case, a self-adjoint bounded linear operator is just a Hermitian $d \times d$ matrix, and the composition of linear operators is matrix multiplication.

Assume the hypotheses of the lemma and let A_1, \dots, A_r be Hermitian $d \times d$ matrices. Assume that A_1, \dots, A_r make a fully commuting operator assignment for X_1, \dots, X_r such that the equations $Q_1 = \dots = Q_m = 0$ are satisfied. The matrices A_1, \dots, A_r pairwise commute, so the Strong Spectral Theorem (i.e. Theorem 1) applies to them. Thus, there exist a unitary matrix U and diagonal $d \times d$ matrices E_1, \dots, E_m such that $A_i = U^{-1}E_iU$ for every $i \in [r]$. Equivalently, $UA_iU^{-1} = E_i$. From $A_i^2 = I$ we conclude $E_i^2 = I$. Hence, if $a_i(j)$ denotes the j -th diagonal entry of E_i , then $a_i(j)^2 = 1$ for all $j \in [d]$. Thus $a_i(j) \in \{\pm 1\}$ for all $j \in [d]$. The conditions of Lemma 1 apply, so $Q_k(A_1, \dots, A_r)$ and $Q_k(E_1, \dots, E_r)$ are similar matrices for each $k \in [m]$. Since $Q_k(A_1, \dots, A_r) = 0$ and the unique matrix that is similar to the null matrix is the null matrix itself, we conclude that $Q_k(E_1, \dots, E_r) = 0$. Now, E_i is the diagonal matrix that has the vector $(a_i(1), \dots, a_i(d))$ in the diagonal, so $Q_k(a_1(j), \dots, a_r(j)) = 0$ for all $j \in [d]$. Since $a_i(j) \in \{\pm 1\}$ for each $i \in [r]$ and $j \in [d]$, the hypothesis of the lemma says that also $Q(a_1(j), \dots, a_r(j)) = 0$ for all $j \in [d]$. Thus $Q(E_1, \dots, E_r) = 0$, and another application of Lemma 1 shows that $Q(A_1, \dots, A_r) = 0$, as was to be proved. \square

The proof for the general case follows the same structure as the proof of the finite-dimensional case, using Theorem 2 in place of Theorem 1. Other than taking care of *null sets of exceptions*, there are no further differences in the two proofs. At a later stage we will find an application of the SST whose proof for the infinite-dimensional case does require some new ingredients. For now, let us fill in the details of the null-set-of-exceptions argument as a warm-up.

Proof of Lemma 3; general case. Assume the hypotheses of the lemma and let A_1, \dots, A_r be bounded self-adjoint linear operators on \mathcal{H} . Suppose that A_1, \dots, A_r make a fully commuting operator assignment for X_1, \dots, X_r such that the equations $Q_1 = \dots = Q_m = 0$ are satisfied. The operators A_1, \dots, A_r pairwise commute, and since they are self-adjoint they are also normal, so the Strong Spectral Theorem (i.e. Theorem 2) applies to them. Thus, there exist a measure space $(\Omega, \mathcal{M}, \mu)$, a unitary map $U : \mathcal{H} \rightarrow L^2(\Omega, \mu)$ and functions $a_1, \dots, a_r \in L^\infty(\Omega, \mu)$ such that, for the multiplication operators $E_i = T_{a_i}$ of $L^2(\Omega, \mu)$, the relations $A_i = U^{-1}E_iU$ hold for every $i \in [r]$. Equivalently, $UA_iU^{-1} = E_i$. From $A_i^2 = I$ we conclude $E_i^2 = I$. Hence, $a_i(\omega)^2 = 1$ for almost all $\omega \in \Omega$; i.e. formally, $\mu(\{\omega \in \Omega : a_i(\omega)^2 \neq 1\}) = 0$. Thus $a_i(\omega) \in \{\pm 1\}$ for almost all $\omega \in \Omega$. The conditions of Lemma 1 apply, thus $Q_k(A_1, \dots, A_r)$ and $Q_k(E_1, \dots, E_r)$ are similar linear operators for each $k \in [m]$. Since $Q_k(A_1, \dots, A_r) = 0$ and the unique linear operator that is similar to the null operator is the null operator itself, we conclude that $Q_k(E_1, \dots, E_r) = 0$. Now, E_i is the multiplication operator given by the function a_i , so $Q_k(a_1(\omega), \dots, a_r(\omega)) = 0$ for almost all $\omega \in \Omega$. Since for almost all $\omega \in \Omega$ the component $a_i(\omega)$ is in $\{\pm 1\}$ for each $i \in [r]$, the hypothesis of the lemma says that also $Q(a_1(\omega), \dots, a_r(\omega)) = 0$ for almost all $\omega \in \Omega$. Writing $Q(a_1, \dots, a_r)$ for the function that maps ω to $Q(a_1(\omega), \dots, a_r(\omega))$ we have $Q(T_{a_1}, \dots, T_{a_r}) = T_{Q(a_1, \dots, a_r)}$. Thus $Q(E_1, \dots, E_r) = 0$, and another application of Lemma 1 shows that $Q(A_1, \dots, A_r) = 0$, as was to be proved. \square

4. Reductions through primitive positive formulas

Let A be a Boolean constraint language, let r be a positive integer, and let x_1, \dots, x_r be variables ranging over the Boolean domain $\{\pm 1\}$. A primitive positive formula, or pp-formula for short, is a formula of the form

$$\phi(x_1, \dots, x_r) = \exists y_1 \dots \exists y_s (R_1(z_1) \wedge \dots \wedge R_m(z_m)) \quad (9)$$

where each R_i is a relation in A and each z_i is an r_i -tuple of variables or constants from $\{x_1, \dots, x_r\} \cup \{y_1, \dots, y_s\} \cup \{\pm 1\}$, where r_i is the arity of R_i . A relation $R \subseteq \{\pm 1\}^r$ is pp-definable from A if there exists a pp-formula $\phi(x_1, \dots, x_r)$ such that

$$R = \{(a_1, \dots, a_r) \in \{\pm 1\}^r : \phi(x_1/a_1, \dots, x_r/a_r) \text{ is true in } A\}, \quad (10)$$

where the notation x/a refers to the syntactic substitution of each (free) occurrence of the variable x by a . A Boolean constraint language A is pp-definable from another Boolean constraint language B if every relation in A is pp-definable from B . Whenever the constants $+1$ and -1 do not appear in the pp-formulas, we speak of pp-formulas and pp-definability *without constants* or, also, *without parameters*.

In the following we show that if A is pp-definable from B , then every instance \mathcal{I} over A can be translated into an instance \mathcal{J} over B in such a way that the satisfying operator assignments for \mathcal{I} lift to satisfying operator assignments for \mathcal{J} . We make this precise.

4.1. The basic construction

Let A and B be two Boolean constraint languages and assume that every relation in A is pp-definable from B . For R in A , let

$$\phi_R(x_1, \dots, x_r) = \exists y_1 \dots \exists y_t (S_1(w_1) \wedge \dots \wedge S_m(w_m)) \quad (11)$$

be the pp-formula that defines R from B , where S_1, \dots, S_m are relations from B , and w_1, \dots, w_m are tuples of variables or constants in $\{x_1, \dots, x_r\} \cup \{y_1, \dots, y_t\} \cup \{\pm 1\}$ of appropriate lengths. For every instance \mathcal{I} of A we construct an instance \mathcal{J} of B as follows.

Consider a constraint (Z, R) in \mathcal{I} , where $Z = (Z_1, \dots, Z_r)$ is a tuple of variables of \mathcal{I} or constants in $\{\pm 1\}$. In addition to the variables in Z , in \mathcal{J} we add new fresh variables Y_1, \dots, Y_t for the quantified variables y_1, \dots, y_t in ϕ_R . We also add one constraint (W_j, S_j) for each $j \in [m]$, where W_j is the tuple of variables and constants obtained from w_j by replacing the variables in x_1, \dots, x_r by the corresponding components Z_1, \dots, Z_r of Z , replacing any y_i -variable by the corresponding Y_i , and leaving all constants untouched. We do this for each constraint in \mathcal{I} one by one. The collection of variables $Z_1, \dots, Z_r, Y_1, \dots, Y_t$ that are introduced by the constraint (Z, R) of \mathcal{I} is referred to as the *block* of (Z, R) in \mathcal{J} . Note that two blocks of different constraints may intersect, but only on the variables of \mathcal{I} .

This construction is referred to as a *gadget reduction* in the literature. Its main property for satisfiability in the Boolean domain is the following straightforward fact:

Lemma 4. \mathcal{I} is satisfiable in the Boolean domain if and only if \mathcal{J} is.

We omit its very easy proof. Our goal in the rest of this section is to show that one direction of this basic property of gadget reductions is also true for satisfiability via operators, for both finite- and infinite-dimensional Hilbert spaces, and that the other direction is *almost true* in a sense we will make precise in due time.

4.2. Correctness: operator solutions lift

The following lemma shows that the left-to-right direction in Lemma 4 also holds for satisfiability via operators: satisfying operator assignments for \mathcal{I} can be lifted to satisfying operator assignments for \mathcal{J} , over the same Hilbert space.

Lemma 5. *Let \mathcal{I} and \mathcal{J} be as above and let \mathcal{H} be a Hilbert space. For every f that is a satisfying operator assignment for \mathcal{I} over \mathcal{H} , there exists g that extends f and is a satisfying operator assignment for \mathcal{J} over \mathcal{H} . Moreover, g is pairwise commuting on each block of \mathcal{J} .*

As in the proof of Lemma 3 we split into cases.

Proof of Lemma 5, finite-dimensional case. As in the proof of the finite-dimensional case of Lemma 3, we may assume that $\mathcal{H} = \mathbb{C}^d$ for some positive integer d , and that A_1, \dots, A_r are Hermitian $d \times d$ matrices that make a satisfying operator assignment f for \mathcal{I} . We need to define Hermitian matrices for the new variables of \mathcal{J} that were introduced by its construction. We define these matrices simultaneously for all variables Y_1, \dots, Y_t that come from the same constraint (Z, R) of \mathcal{I} .

By renaming the entries in Z if necessary, let us assume without loss of generality that the variables in Z are X_1, \dots, X_r , where r is the arity of R . By the commutativity condition of satisfying operator assignments, the matrices A_1, \dots, A_r pairwise commute. As each A_i is Hermitian, the Strong Spectral Theorem applies to them. Thus, there exist a unitary matrix U and diagonal $d \times d$ matrices E_1, \dots, E_r such that the relations $A_i = U^{-1}E_iU$ hold for each $i \in [r]$. Equivalently, $UA_iU^{-1} = E_i$. From $A_i^2 = I$ we conclude $E_i^2 = I$. Hence, if $a_i(j)$ denotes the j -th diagonal entry of E_i , then $a_i(j)^2 = 1$ for all $j \in [d]$. Thus $a_i(j) \in \{\pm 1\}$ for all $j \in [d]$. The conditions of Lemma 1 apply, thus $P_R(A_1, \dots, A_r)$ and $P_R(E_1, \dots, E_r)$ are similar matrices. Since $P_R(A_1, \dots, A_r) = -I$ and the unique matrix that is similar to $-I$ is $-I$ itself, we conclude that $P_R(E_1, \dots, E_r) = -I$. Now, E_i is the diagonal matrix that has the vector $(a_i(1), \dots, a_i(d))$ in the diagonal, so $P_R(a_1(j), \dots, a_r(j)) = -1$ for all $j \in [d]$. Thus the tuple $a(j) = (a_1(j), \dots, a_r(j))$ belongs to the relation R for all $j \in [d]$. Now we are ready to define the matrices for the variables Y_1, \dots, Y_t .

For each $j \in [d]$, let $b(j) = (b_1(j), \dots, b_t(j)) \in \{\pm 1\}^t$ be a tuple of witnesses to the existentially quantified variables in $\phi_R(x_1/a_1(j), \dots, x_r/a_r(j))$; such a vector of witnesses must exist since the tuple $a(j)$ belongs to R and ϕ_R defines R . Let F_k be the diagonal matrix that has the vector $(b_k(1), \dots, b_k(d))$ in the diagonal, and let Y_k be assigned the matrix $B_k = U^{-1}F_kU$. Since U is unitary, each such matrix is Hermitian and squares to the identity since $b_k(j) \in \{\pm 1\}$ for all $j \in [d]$. Moreover, $E_1, \dots, E_r, F_1, \dots, F_t$ pairwise commute since they are diagonal matrices; thus $A_1, \dots, A_r, B_1, \dots, B_t$ also pairwise commute since they are simultaneously similar via U . Moreover, as each atomic formula in the quantifier-free part of ϕ_R is satisfied by the mapping that sends $x_i \mapsto a_i(j)$ and $y_i \mapsto b_i(j)$ for all $j \in [d]$, another application of Lemma 1 shows that the matrices that are assigned to the variables of this atomic formula make the corresponding indicator polynomial evaluate to -1 . This means that the assignment to the X and Y -variables makes a satisfying operator assignment for the constraints of \mathcal{J} that come from the constraint (Z, R) of \mathcal{I} . As different constraints from \mathcal{I} produce their own sets of Y -variables, these definitions of assignments do not conflict with one another, and the proof of the lemma is complete. \square

The proof for the general case requires some new ingredients. Besides the need to take care of null sets of exceptions as in the proof of Lemma 3, a new complication arises from the need to build the operators for the new variables that are introduced by the reduction. Concretely we need to make sure that the *functions* of witnesses, in contraposition to the *finite tuples* of witnesses in the finite-dimensional case, are bounded and measurable. We go carefully through the argument.

Proof of Lemma 5, general case. Assume that A_1, \dots, A_r are bounded self-adjoint linear operators on \mathcal{H} for the variables of \mathcal{I} . Suppose that the operators A_1, \dots, A_r make a valid satisfying operator assignment for \mathcal{I} . We need to define bounded self-adjoint linear operators for the new variables of \mathcal{J} that were introduced by the construction. We define these operators simultaneously for all variables Y_1, \dots, Y_t that come from the same constraint (Z, R) of \mathcal{I} .

By renaming the components of Z if necessary, assume without loss of generality that the variables in Z are X_1, \dots, X_r , where r is the arity of R . By the commutativity condition of satisfying operator assignments, the operators A_1, \dots, A_r pairwise commute. As each A_i is self-adjoint, it is also normal, and the Strong Spectral Theorem (cf. Theorem 2) applies. Thus, there exist a measure space $(\Omega, \mathcal{M}, \mu)$, a unitary map $U : \mathcal{H} \rightarrow L^2(\Omega, \mu)$ and functions $a_1, \dots, a_r \in L^\infty(\Omega, \mu)$ such that, for the multiplication operators $E_i = T_{a_i}$ of $L^2(\Omega, \mu)$, the relations $A_i = U^{-1}E_iU$ hold for each $i \in [r]$. Equivalently, $UA_iU^{-1} = E_i$. From $A_i^2 = I$ we conclude $E_i^2 = I$. Hence, $a_i(\omega)^2 = 1$ for almost all $\omega \in \Omega$; i.e., formally $\mu(\{\omega \in \Omega : a_i(\omega)^2 \neq 1\}) = 0$. Thus, $a_i(\omega) \in \{\pm 1\}$ for almost all $\omega \in \Omega$. The conditions of Lemma 1 apply, thus $P_R(A_1, \dots, A_r)$ and $P_R(E_1, \dots, E_r)$ are similar linear operators. Since $P_R(A_1, \dots, A_r) = -I$ and the unique linear operator that is similar to $-I$ is $-I$ itself, we conclude that $P_R(E_1, \dots, E_r) = -I$. Now, E_i is the multiplication operator given by a_i , and $a_i(\omega) \in \{\pm 1\}$ for almost all

$\omega \in \Omega$, so $P_R(a_1(\omega), \dots, a_r(\omega)) = -1$ for almost all $\omega \in \Omega$. Thus the tuple $a(\omega) = (a_1(\omega), \dots, a_r(\omega))$ belongs to the relation R for almost all $\omega \in \Omega$. Now we are ready to define the operators for the variables Y_1, \dots, Y_t .

For each $\omega \in \Omega$ for which the tuple $a(\omega)$ belongs to R , let $b(\omega) = (b_1(\omega), \dots, b_t(\omega)) \in \{\pm 1\}^t$ be the *lexicographically smallest* tuple of witnesses to the existentially quantified variables in $\phi_R(x_1/a_1(\omega), \dots, x_r/a_r(\omega))$; such a vector of witnesses must exist since ϕ_R defines R , and the lexicographically smallest exists because R is finite. For every other $\omega \in \Omega$, define $b(\omega) = (b_1(\omega), \dots, b_t(\omega)) = (0, \dots, 0)$.

Note that each function $b_k: \Omega \rightarrow \mathbb{C}$ is bounded since its range is in $\{-1, 0, 1\}$. We claim that such functions of witnesses b_k are also measurable functions of $(\Omega, \mathcal{M}, \mu)$. This will follow from the fact that a_1, \dots, a_r are measurable functions themselves, the fact that R is a finite relation, and the choice of a definite tuple of witnesses of each $\omega \in \Omega$: the lexicographically smallest if $a(\omega)$ is in R , or the all-zero tuple otherwise. We discuss the details.

Since R is finite, the event $Q = \{\omega \in \Omega : b_k(\omega) = \sigma\}$, for fixed $\sigma \in \{+1, 0, -1\}$, can be expressed as a finite Boolean combination of events of the form $Q_{i,\tau} = \{\omega \in \Omega : a_i(\omega) = \tau\}$, where $i \in [r]$ and $\tau \in \{\pm 1\}$. Here is how: If $\sigma \neq 0$, then

$$Q = \bigcup_{\substack{a \in R: \\ b(a)_k = \sigma}} \left(\bigcap_{i \in [r]} Q_{i,a_i} \right), \tag{12}$$

where $b(a)$ denotes the lexicographically smallest tuple of witnesses in $\{\pm 1\}^t$ for the quantified variables in $\phi_R(x_1/a_1, \dots, x_r/a_r)$. If $\sigma = 0$, then Q is the complement of this set. Each $Q_{i,\tau}$ is a measurable set in the measure space $(\Omega, \mathcal{M}, \mu)$ since a_i is a measurable function and $Q_{i,\tau} = a_i^{-1}(B_{1/4}(\tau))$, where $B_{1/4}(\tau)$ denotes the complex open ball of radius $1/4$ centered at τ , which is a Borel set in the standard topology of \mathbb{C} . Since the range of b_k is in the finite set $\{-1, 0, 1\}$, the preimage $b_k^{-1}(S)$ of each Borel subset S of \mathbb{C} is expressed as a finite Boolean combination of measurable sets, and is thus measurable in $(\Omega, \mathcal{M}, \mu)$.

We just proved that each b_k is bounded and measurable, so its equivalence class under almost everywhere equality is represented in $L^\infty(\Omega, \mu)$. We may assume without loss of generality that b_k is its own representative; else modify it on a set of measure zero in order to achieve so. Let $F_k = T_{b_k}$ be the multiplication operator given by b_k and let Y_k be assigned the linear operator $B_k = U^{-1}F_kU$, which is bounded because b_k is bounded and U is unitary. Also because U is unitary, each such operator is self-adjoint and squares to the identity since $b_k(\omega) \in \{\pm 1\}$ for almost all $\omega \in \Omega$. Moreover, $E_1, \dots, E_r, F_1, \dots, F_t$ pairwise commute since they are multiplication operators; thus $A_1, \dots, A_r, B_1, \dots, B_t$ pairwise commute since they are simultaneously similar via U . Moreover, as each atomic formula in the quantifier-free part of ϕ_R is satisfied by the mapping that sends $x_i \mapsto a_i(\omega)$ and $y_i \mapsto b_i(\omega)$ for almost all $\omega \in \Omega$, another application of Lemma 1 shows that the operators that are assigned to the variables of this atomic formula make the corresponding indicator polynomial evaluate to $-I$. This means that the assignment to the X and Y -variables makes a satisfying operator assignment for the constraints of \mathcal{J} that come from the constraint (Z, R) in \mathcal{I} . As different constraints from \mathcal{I} produce their own sets of Y -variables, these definitions of assignments are not in conflict with each other, and the proof of the lemma is complete. \square

4.3. The extended construction

We proved so far that satisfying operator assignments for \mathcal{I} lift to satisfying operator assignments for \mathcal{J} . We do not know if the converse is true. One could try to just take the restriction of the satisfying assignment for \mathcal{J} to the variables of \mathcal{I} , but there is little chance that this will work because there is no guarantee that the operators that are assigned to any two variables that appear together in a constraint of \mathcal{I} will commute. Instead of trying to modify the assignment, we modify the instance \mathcal{J} . Let us discuss a slightly modified version of \mathcal{J} , over a very minor extension of the constraint language B , that still allows lifting of solutions, and for which the naive projection works for the backward direction. Let us stress now that we plan to use this modified construction over a minor extension of the constraint language merely as a technical device to get other results. At the end of this section we also discuss a possible avenue for avoiding the technical device.

In the following, let T denote the full binary Boolean relation; i.e., $T = \{\pm 1\}^2$. Observe that the indicator polynomial $P_T(X_1, X_2)$ of the relation T is just the constant -1 ; the letter T stands for *true*.

Let A and B be the constraint languages such that A is pp-definable from B . Let \mathcal{I} and \mathcal{J} be the instances over A and B as defined above. The modified version of \mathcal{J} will be an instance over the expanded constraint language $B \cup \{T\}$. We denote it $\hat{\mathcal{J}}$ and it is defined as follows: the variables and the constraints of $\hat{\mathcal{J}}$ are defined as in \mathcal{J} , but we also add all the binary constraints of the form $((X_i, X_j), T)$, $((X_i, Y_k), T)$ or $((Y_k, Y_\ell), T)$, for every four different variables X_i, X_j, Y_k and Y_ℓ that come from the same block in \mathcal{J} .

4.4. Correctness: operator solutions lift and also project

We argue that in this new construction, satisfying assignments not only lift from \mathcal{I} to $\hat{\mathcal{J}}$, but also project from $\hat{\mathcal{J}}$ to \mathcal{I} .

Lemma 6. *Let \mathcal{I} and $\hat{\mathcal{J}}$ be as above and let \mathcal{H} be a Hilbert space. Then the following assertions are true.*

1. *For every f that is a satisfying operator assignment for \mathcal{I} over \mathcal{H} , there exists g that extends f and is a satisfying operator assignment for $\hat{\mathcal{J}}$ over \mathcal{H} ,*

2. For every g that is a satisfying operator assignment for $\hat{\mathcal{J}}$ over \mathcal{H} , the restriction f of g to the variables of \mathcal{I} is a satisfying operator assignment for \mathcal{I} over \mathcal{H} .

Proof. Statement 1 follows from Lemma 5: Fix f that is a satisfying operator assignment for \mathcal{I} and let g be given by Lemma 5. This is also an assignment for the variables of $\hat{\mathcal{J}}$. The constraints of $\hat{\mathcal{J}}$ that are already in \mathcal{J} are of course satisfied by g . Next consider an additional constraint of the form $((X_i, X_j), T)$, $((X_i, Y_k), T)$ or $((Y_k, Y_\ell), T)$, for variables X_i , X_j , Y_k and Y_ℓ coming from the same block in \mathcal{J} . By the “moreover” clause in Lemma 5, the operators A_i , A_j , B_k and B_ℓ associated to X_i , X_j , Y_k and Y_ℓ by g pairwise commute. Moreover, the associated polynomial constraints $P_T(A_i, A_j) = -I$, $P_T(A_i, B_k) = -I$ and $P_T(B_k, B_\ell) = -I$ are trivial (i.e., void) since the indicator polynomial $P_T(X_1, X_2)$ of T is just the constant -1 .

For statement 2, fix g that is a satisfying operator assignment for $\hat{\mathcal{J}}$ over \mathcal{H} , and let f be the restriction of g to the variables of \mathcal{I} . Since g satisfies $\hat{\mathcal{J}}$, for every two variables X_i and X_j that appear together in a constraint (Z, R) of \mathcal{I} , the associated operators $g(X_i)$ and $g(X_j)$ commute since X_i and X_j appear in the same block of \mathcal{J} . Hence $f(X_i)$ and $f(X_j)$ commute. We still need to show that the polynomial constraint $P_R(f(Z)) = -I$ is satisfied for every constraint (Z, R) of \mathcal{I} . To do so, we use Lemma 3 on an appropriately defined system of polynomial equations.

Let r be the arity of R and let ϕ_R be the pp-formula as in (11) that defines R from B . The polynomials we define have variables $X_1, \dots, X_r, Y_1, \dots, Y_t, C_{-1}, C_{+1}$ that correspond to the variables and constants in (11). For every $k \in [m]$, let Q_k be the polynomial $P_{S_k}(W_k) + 1$, so that the equation $Q_k = 0$ ensures $P_{S_k}(W_k) = -1$, where P_{S_k} is the characteristic polynomial of S_k , and W_k is the tuple of components from $X_1, \dots, X_r, Y_1, \dots, Y_t, C_{-1}, C_{+1}$ that appear in the atom $S_k(w_k)$ of (11). Here we use X_i and Y_j in place of x_i and y_j , respectively, and C_{-1} and C_{+1} in place of the constants -1 and $+1$, respectively. Let also Q_{m+1} and Q_{m+2} be the polynomials $C_{-1} + 1$ and $C_{+1} - 1$, so that the equations $Q_{m+1} = Q_{m+2} = 0$ ensure that $C_{-1} = -1$ and $C_{+1} = +1$. Finally, let Q be the polynomial $P_R(X_1, \dots, X_r) + 1$, where P_R is the characteristic polynomial of R . It follows from the definitions that every Boolean assignment that satisfies all equations $Q_1 = \dots = Q_{m+2} = 0$ also satisfies $Q = 0$. Thus Lemma 3 applies, and since g extended to $g(C_{-1}) = -I$ and $g(C_{+1}) = I$ satisfies all equations $Q_1 = \dots = Q_{m+2} = 0$, it also satisfies $Q = 0$. It follows that $P_R(f(Z)) = P_R(g(Z)) = -I$, as was to be proved. \square

4.5. Discussion: the full binary relation and commutativity gadgets

The sole role of the full binary Boolean relation $T = \{\pm 1\}^2$ in the construction of $\hat{\mathcal{J}}$ was to express the commutativity relations that are implied by $((X_i, X_j), T)$, $((X_i, Y_k), T)$ and $((Y_k, Y_\ell), T)$. These are needed for the proof of the second statement in Lemma 6. In some contexts, however, the infusion of the relation T in the target constraint language B can be avoided, as we discuss next.

Consider, for example, the constraint language 3LIN as defined in Section 2, and let R be one of the two relations in 3LIN, say, R is the relation $XYZ = I$. We claim that, whenever the target constraint language is 3LIN, the proof of Lemma 6 goes through with the commutativity relations that involve T in the definition of $\hat{\mathcal{J}}$ replaced by the constraints $((X_i, X_j, U_{ij}), R)$, $((X_i, Y_k, V_{ik}), R)$ and $((Y_k, Y_\ell, W_{k\ell}), R)$, respectively, where U_{ij} , V_{ik} and $W_{k\ell}$ are fresh new variables. This follows by examining the proof of Lemma 6 and by using the following two facts: (i) if A, B, C is an operator assignment for X, Y, Z that satisfies the constraint $((X, Y, Z), R)$, then $AB = BA$; and (ii) if A, B is an operator assignment for X, Y that satisfies $AB = BA$, then there exists C such that A, B, C is an operator assignment for X, Y, Z that satisfies the constraint $((X, Y, Z), R)$. The first of these two facts is an immediate consequence of the definitions. For the second of fact, it suffices to take $C = BA$ and observe that, on one hand $ABC = ABBA = AA = I$ since $A^2 = B^2 = I$, and on the other $AC = ABA = BAA = CA$ and $BC = BBA = BAB = CB$ since $AB = BA$.

More generally, the use of the full binary Boolean relation T can be avoided whenever the target constraint language admits a so-called *commutativity gadget*, as defined next. Let \mathcal{K} be an instance over the target constraint language that contains two distinguished variables X and Y , for a total of $t+2$ variables X, Y, Z_1, \dots, Z_t . We say that \mathcal{K} is a commutativity gadget for X and Y if the following two conditions hold: (i) if A, B, C_1, \dots, C_t is an operator assignment for X, Y, Z_1, \dots, Z_t that satisfies \mathcal{K} , then $AB = BA$; and (ii) if A, B is an operator assignment for X, Y that satisfies $AB = BA$, then there exist C_1, \dots, C_t such that A, B, C_1, \dots, C_t is an operator assignment for X, Y, Z_1, \dots, Z_t that satisfies \mathcal{K} . The claim is that, whenever a commutativity gadget \mathcal{K} exists, the commutativity relations that involve T in the definition of $\hat{\mathcal{J}}$ can be replaced by copies of \mathcal{K} with the appropriate X and Y variables, and fresh new variables Z_1, \dots, Z_t for each copy, and then the proof of Lemma 6 goes through. Again, this follows by examining the proof of Lemma 6 and observing that the only properties of the full Boolean binary relation T used are precisely the ones abstracted in the definition of a commutativity gadget.

The concept of commutativity gadget was introduced by Ji [15] in his proof that the operator version of 3SAT is NP-hard. Obviously, the single-constraint instance $((X, Y), T)$ is always a commutativity gadget. We already argued that, for the constraint language 3LIN, the single-constraint instance $((X, Y, Z), XYZ = I)$ is a commutativity gadget for X and Y . Similarly, it is not hard to see that the constraint language of 3SAT has a commutativity gadget that is also made of the single-constraint instance $((X, Y, Z), X \vee Y \vee Z)$. However, it is not always so easy to show that a constraint language admits a commutativity gadget. For example, consider the constraint language 1-IN-3-SAT introduced in Section 2. In this case, the single-constraint instance $((X, Y, Z), R_{1/3})$, where $R_{1/3} = \{(-1, +1, +1), (+1, -1, +1), (+1, +1, -1)\}$, fails to satisfy the second property of the definition of commutativity gadget. Indeed, the partial assignment $(X, Y) \mapsto (-1, -1)$ satisfies $XY = YX$ but it cannot be extended to a tuple in $R_{1/3}$. It turns out, however, that if we let \mathcal{K} be the more complex

instance $((X, Z_1, Z_4), R_{1/3}), ((Y, Z_2, Z_4), R_{1/3}), ((Z_1, Z_2, Z_3), R_{1/3})$, then both conditions are met. In this case, proving that the second condition in the definition of commutativity gadget holds is easy, but proving that the first condition holds requires the use of some non-trivial identities on commutator polynomials that were discovered by Ji [15] through computer search. This construction will be discussed at length in Section 6, where it will play an important role.

5. Satisfiability gaps via operator assignments

Let A be a Boolean constraint language and let \mathcal{I} be an instance over A . It is easy to see that the following inequalities hold:

$$v(\mathcal{I}) \leq v^*(\mathcal{I}) \leq v^{**}(\mathcal{I}). \quad (13)$$

Indeed, the first inequality holds because if we interpret the field of complex numbers \mathbb{C} as a 1-dimensional Hilbert space, then the only solutions to the equation $X^2 = 1$ are $X = -1$ and $X = +1$. The second inequality is a direct consequence of the definitions. For the same reason, if \mathcal{I} is satisfiable in the Boolean domain, then it is satisfiable via fd-operators, and if it is satisfiable via fd-operators, then it is satisfiable via operators. The converses are, in general, not true; however, finding counterexamples is a non-trivial task. For the Boolean constraint language LIN of affine relations, counterexamples are given by Mermin's magic square [16,17] for the first case, and by Slofstra's recent construction [24] for the second case. These will be discussed at some length in due time. In the rest of this section, we characterize the Boolean constraint languages that exhibit such gaps.

We distinguish three types of gaps. Specifically, we say that an instance \mathcal{I} witnesses

1. a *satisfiability gap of the first kind* if $v(\mathcal{I}) < 1$ and $v^*(\mathcal{I}) = 1$;
2. a *satisfiability gap of the second kind* if $v(\mathcal{I}) < 1$ and $v^{**}(\mathcal{I}) = 1$;
3. a *satisfiability gap of the third kind* if $v^*(\mathcal{I}) < 1$ and $v^{**}(\mathcal{I}) = 1$.

As a mnemonic rule, count the number of stars $*$ that appear in the defining inequalities in 1, 2 or 3 to recall what kind the gap is.

We say that a Boolean constraint language A has a *satisfiability gap of the i -th kind*, $i = 1, 2, 3$, if there is at least one instance \mathcal{I} over A that witnesses such a gap. Clearly, a gap of the first kind or a gap of the third kind implies a gap of the second kind. In other words, if A has no gap of the second kind, then A has no gap of the first kind and no gap of the third kind. A priori no other relationships seem to hold. We show that, in a precise sense, either A has no gaps of any kind or A has a gap of every kind. Recall from Section 4 that T denotes the full binary Boolean relation; i.e. $T = \{\pm 1\}^2$. We are now ready to state and prove the main result of this section.

Theorem 3. *Let A be a Boolean constraint language. Then the following statements are equivalent.*

1. A does not have a satisfiability gap of the first kind.
2. A does not have a satisfiability gap of the second kind.
3. $A \cup \{T\}$ does not have a satisfiability gap of the third kind.
4. A is 0-valid, or A is 1-valid, or A is bijunctive, or A is Horn, or A is dual Horn.

The proof of Theorem 3 has two main parts. In the first part, we show that if A satisfies at least one of the conditions in the fourth statement, then A has no satisfiability gaps of the first kind or the second kind, and $A \cup \{T\}$ has no satisfiability gaps of the third kind. In the second part, we show that, in all other cases, A has satisfiability gaps of the first kind and the second kind, and $A \cup \{T\}$ has satisfiability gaps of the third kind. The ingredients in the proof of the second part are the existence of gaps of all three kinds for LIN, results about Post's lattice [21], and *gap-preserving* reductions that use the results about pp-definability established in Section 4.

5.1. No gaps of any kind

Assume that A satisfies at least one of the conditions in the fourth statement in Theorem 3. First, we observe that the full relation T is 0-valid, 1-valid, bijunctive, Horn, and dual Horn. Indeed, T is obviously 0-valid and 1-valid. Moreover, it is bijunctive, Horn, and dual Horn because it is equal to the set of satisfying assignments of the Boolean formula $(x \vee \neg x) \wedge (y \vee \neg y)$, which is bijunctive, Horn, and dual Horn. Therefore, to prove that the fourth statement in Theorem 3 implies the other three statements, it suffices to prove that if A satisfies at least one of the conditions in the fourth statement, then A has no gaps of any kind. Towards this goal, we argue by cases.

We start with the trivial cases in which A is 0-valid or 1-valid. If an instance \mathcal{I} of A contains a constraint of the form (Z, F) , where F is an empty relation (of some arity), then \mathcal{I} is not satisfiable by any operator assignment. Otherwise, \mathcal{I} is satisfiable in the Boolean domain, hence it is satisfiable by assigning the identity operator I to every variable, if A is 0-valid, or by assigning the operator $-I$ to every variable, if A is 1-valid.

Next, we have to show that if A is bijunctive or Horn or dual Horn, then A has no gaps of any kind. As discussed earlier, it suffices to show that A does not have a gap of the second kind (since a gap of the first kind or a gap of the third kind implies a gap of the second kind).

Ji [15] proved that if \mathcal{I} is a 2SAT instance or a HORN SAT instance that is satisfiable via fd-operators, then \mathcal{I} is also satisfiable in the Boolean domain. In other words, Ji showed that 2SAT and HORN SAT have no gaps of the first kind. This is quite close to what we have to prove, but there are two differences. The first difference is that we cannot apply Ji's argument directly to bijunctive or Horn constraint languages. The reason is that it is by no means obvious that if an operator assignment satisfies a constraint expressed by a bijunctive relation (or a Horn relation) R , then this operator assignment also satisfies every clause of the 2CNF-formula (or of the Horn formula) that defines R . While the preceding property is true, proving it requires an application of the Strong Spectral Theorem, which is somewhat delicate in the infinite dimensional case. For this, however, the full generality of the statement of Lemma 3 will come in handy. The second and main difference is that, at first glance, Ji's proof for 2SAT and HORN SAT does not seem to extend to operator assignments of arbitrary (finite or infinite) dimension. The reason for this is that Ji's argument relies on the existence of eigenvalues and associated orthogonal eigenspaces for the linear operators, which are not guaranteed to exist in the infinite-dimensional case, even for self-adjoint bounded linear operators. Note, however, that in our case we have the additional requirement that the operators satisfy $A^2 = I$, and in such a case their eigenvalues and associated eigenspaces can be reinstated. This observation could perhaps be used to give a proof along the lines of Ji's that 2SAT and HORN SAT have no gaps of the second kind. Nonetheless, we prefer to give an alternative and more direct proof that does not rely at all on the existence of eigenvalues. Our proof is based on the manipulation of non-commutative polynomial identities, a method that has been called *the substitution method* (see, e.g., [7]).

Lemma 7. *Let \mathcal{I} be a 2SAT instance or a HORN SAT instance or a DUAL HORN SAT instance. Then the following statements are equivalent.*

1. \mathcal{I} is satisfiable in the Boolean domain;
2. \mathcal{I} is satisfiable via fd-operators;
3. \mathcal{I} is satisfiable via operators.

We split the proof into two: one for 2SAT and another one for HORN SAT; the proof for DUAL HORN SAT is analogous to the proof for HORN SAT, and it is omitted.

Proof of Lemma 7 for 2SAT. Let \mathcal{I} be a 2CNF-formula. The implications $1 \implies 2$ and $2 \implies 3$ follow from the definitions. To prove the implication $3 \implies 1$, assume that f is a satisfying operator assignment for \mathcal{I} over a (finite-dimensional or infinite-dimensional) Hilbert space \mathcal{H} , and, towards a contradiction, assume that \mathcal{I} is unsatisfiable in the Boolean domain. We will make use of the well-known characterization of unsatisfiable in the Boolean domain 2SAT instances in terms of a reachability property of their associated *implication graph*. For \mathcal{I} , the implication graph is the directed graph G that has one vertex for each literal x or $\neg x$ of every variable x in \mathcal{I} , and two directed edges for each clause $(\ell_1 \vee \ell_2)$ of \mathcal{I} , one edge from $\overline{\ell_1}$ to ℓ_2 , and another one from $\overline{\ell_2}$ to ℓ_1 . The well-known characterization states that \mathcal{I} is unsatisfiable in the Boolean domain if and only if there exists a variable x and two directed paths in G , one from the variable x to the literal $\neg x$, and another one from the literal $\neg x$ to the variable x (see, e.g., [19]). Accordingly, let ℓ_1, \dots, ℓ_r and m_1, \dots, m_s be literals such that $x, \ell_1, \dots, \ell_r, \neg x$ and $\neg x, m_1, \dots, m_s, x$ are the vertices in the paths from x to $\neg x$ and from $\neg x$ to x , respectively, in the order they are traversed.

The existence of the path $x, \ell_1, \dots, \ell_r, \neg x$ from the variable x to the literal $\neg x$ in the implication graph G means that the clauses

$$(\neg x \vee \ell_1), (\overline{\ell_1} \vee \ell_2), \dots, (\overline{\ell_{r-1}} \vee \ell_r), (\overline{\ell_r} \vee \neg x) \quad (14)$$

are clauses of the instance \mathcal{I} . Symmetrically, the existence of the path $\neg x, m_1, \dots, m_s, x$ from the literal $\neg x$ to the variable x in the implication graph G means that the clauses

$$(x \vee m_1), (\overline{m_1} \vee m_2), \dots, (\overline{m_{s-1}} \vee m_s), (\overline{m_s} \vee x) \quad (15)$$

are clauses of the instance \mathcal{I} .

In the case of satisfiability in the Boolean domain, one reasons that the instance \mathcal{I} is unsatisfiable, because if it were satisfiable by some truth assignment, then the path of implications from x to $\neg x$ forces x to be set to *false*, while the path of implications from $\neg x$ to x forces x to be set to *true*. In what follows, we will show that, with some care, essentially the same reasoning can be carried out for operator assignments that satisfy the instance \mathcal{I} .

Extend the operator assignment f to all literals by setting $f(\ell) = \text{sg}(\ell)f(x)$, where x the variable underlying ℓ . Since f is a quantum satisfying assignment for \mathcal{I} , Lemma 2 implies that

$$(I - f(x))(I + f(\ell_1)) = 0 \quad (16)$$

$$(I - f(\ell_i))(I + f(\ell_{i+1})) = 0, \quad 1 \leq i \leq r - 1. \quad (17)$$

$$(I - f(\ell_r))(I - f(x)) = 0 \quad (18)$$

We now claim that

$$(I - f(x))(I + f(\ell_i)) = 0, \quad 1 \leq i \leq r. \quad (19)$$

We prove the claim by induction on i . For $i = 1$, what we need is just equation (16). By induction, assume now that

$$(I - f(x))(I + f(\ell_{i-1})) = 0 \quad (20)$$

holds for some i with $2 \leq i \leq r - 1$. By (17), we have that

$$(I - f(\ell_{i-1}))(I + f(\ell_i)) = 0 \quad (21)$$

holds. First, by multiplying equation (20) from the right by $(I + f(\ell_i))$, we get

$$(I - f(x))(I + f(\ell_{i-1}))(I + f(\ell_i)) = 0 \quad (22)$$

Second, by multiplying equation (21) from the left by $(I - f(x))$, we get

$$(I - f(x))(I - f(\ell_{i-1}))(I + f(\ell_i)) = 0 \quad (23)$$

By adding equations (22) and (23), we obtain

$$(I - f(x))(I + f(\ell_i)) = 0, \quad (24)$$

as desired. In particular, by considering the case $i = r$, we get

$$(I - f(x))(I + f(\ell_r)) = 0, \quad (25)$$

which, after multiplying out the left-hand side, becomes

$$I + f(\ell_r) - f(x) - f(x)f(\ell_r) = 0. \quad (26)$$

Furthermore, by multiplying out the left-hand side of equation (18), we get

$$I - f(x) - f(\ell_r) + f(\ell_r)f(x) = 0. \quad (27)$$

Since the variable x and the literal ℓ_r appear in the same clause of the instance \mathcal{I} , namely, the clause $(\overline{\ell_r} \vee \neg x)$, we have that $f(x)f(\ell_r) = f(\ell_r)f(x)$. Therefore, by adding equations (26) and (27), we get that $2I - 2f(x) = 0$, which implies that $f(x) = I$.

An entirely symmetric argument using the path from $\neg x$ to x , instead of the path from x to $\neg x$, gives $f(x) = -I$, which contradicts the previous finding that $f(x) = I$. \square

Proof of Lemma 7 for HORN SAT. Let \mathcal{I} be a Horn formula. As with the proof for 2SAT, the only non-trivial direction is $3 \implies 1$. To prove the implication $3 \implies 1$, assume that f is a satisfying operator assignment for \mathcal{I} over a (finite-dimensional or infinite-dimensional) Hilbert space \mathcal{H} , and, towards a contradiction, assume that \mathcal{I} is unsatisfiable in the Boolean domain. As in the proof for 2SAT, let f be extended to all literals by $f(\ell) = \text{sg}(\ell)f(x)$, where x is the variable underlying ℓ . We will make use of the characterization of unsatisfiability in the Boolean domain Horn instances in terms of unit resolution. For this, we need to first introduce some terminology and notation. If C and C' are two clauses such that C contains a literal ℓ and C' contains the complementary literal $\overline{\ell}$ of ℓ , then the *resolution rule* produces in one step the *resolvent* clause D that is the disjunction of all literals in the *premises* C and C' other than ℓ and $\overline{\ell}$. The *unit resolution rule* is the special case of the resolution rule in which (at least) one of the clauses C and C' is a single literal. It is well known (see, e.g., [23]) that a Horn formula \mathcal{I} is unsatisfiable if and only if there is a *unit resolution derivation* of the empty clause from the clauses of \mathcal{I} , i.e., there is a sequence C_1, \dots, C_m of clauses such that, for each $i \in \{1, \dots, m\}$, we have that C_i is one of the clauses of \mathcal{I} or C_i is obtained from earlier clauses C_j and C_k in the sequence via the unit resolution rule. Clearly, in a unit resolution derivation of the empty clause, the last application of the unit resolution rule involves two clauses each of which is the complementary literal of the other.

In what follows, we will show that a unit resolution derivation can be “simulated” by a sequence of equations involving operator assignments. We begin by formulating and proving the following claim.

Claim 1: Let $(\ell_1 \vee \dots \vee \ell_r)$ be clause and let $\overline{\ell_j}$ be the complementary literal of some literal ℓ_j in that clause. If f satisfies both the clause $(\ell_1 \vee \dots \vee \ell_r)$ and the literal $\overline{\ell_j}$, then f also satisfies the resolvent $(\ell_1 \vee \dots \vee \ell_{j-1} \vee \ell_{j+1} \vee \dots \vee \ell_r)$ of $(\ell_1 \vee \dots \vee \ell_r)$ and $\overline{\ell_j}$; equivalently, the operators $\{f(\ell_i) : i \neq j\}$ pairwise commute and

$$\prod_{i=1}^{j-1} (I + f(\ell_i)) \prod_{i=j+1}^r (I + f(\ell_i)) = 0. \quad (28)$$

Observe that for the unit resolution rule, as is the case here, the resolvent is always a subclause of one of the premises. In particular, since f satisfies both premises, all the operators involved in the premises commute, and so do the ones involved in the resolvent clause. To complete the proof of the claim observe that, since f satisfies both the clause $(\ell_1 \vee \dots \vee \ell_r)$ and the literal $\bar{\ell}_j$, the corresponding operators commute, and the identity of polynomials in commuting variables of Lemma 2 implies that

$$\prod_{i=1}^r (I + f(\ell_i)) = 0 \quad (29)$$

$$(I - f(\ell_j)) = 0. \quad (30)$$

By multiplying equation (30) by $\prod_{i=1}^{j-1} (I + f(\ell_i))$ from the left, and by $\prod_{i=j+1}^r (I + f(\ell_i))$ from the right, we get

$$\left(\prod_{i=1}^{j-1} (I + f(\ell_i)) \right) (I - f(\ell_j)) \left(\prod_{i=j+1}^r (I + f(\ell_i)) \right) = 0. \quad (31)$$

By adding equations (29) and (31), we get (28), which completes the proof of Claim 1.

Consider now a unit resolution derivation C_1, \dots, C_m of the empty clause from the clauses of \mathcal{I} . Since the operator assignment f satisfies every clause of \mathcal{I} , we can apply Claim 1 repeatedly and, by induction, show that f satisfies each clause in this derivation. Since C_m is the empty clause, it must have been derived via the unit resolution rule from two earlier clauses each of which is the complementary literal of the other, say, ℓ and $\bar{\ell}$. So, we must have $f(\ell) = -I$ and $f(\bar{\ell}) = -I$, which is a contradiction since $f(\ell) = -f(\bar{\ell})$. \square

In what follows, we will use Lemma 7 to show that if A is bijunctive or Horn or dual Horn, then A has no gaps of any kind.

Assume that A is bijunctive. Note that we cannot apply Lemma 7 directly to conclude that A has no gaps of any kind, because the relations in the constraint-language A are defined by conjunctions of 2-clauses, but need not be defined by individual 2-clauses. In order to be able to apply Lemma 7, first we need to verify the following claim. Assume that (Z, R) is a constraint in which R is a relation in A defined by a conjunction $C_1 \wedge \dots \wedge C_m$, where each C_i is a 2-clause on the variables in Z . Then a satisfying operator assignment for the instance consisting of the single constraint (Z, R) will also satisfy each of the 2-clause constraints $(W_1, C_1), \dots, (W_r, C_r)$ individually, where $W_i = (Z_{c_i}, Z_{d_i})$ is the tuple of components of $Z = (Z_1, \dots, Z_r)$ that appear in C_i . To prove this claim, first note that the commutativity condition on the operators assigned to the variables in W_i is guaranteed by the commutativity condition on the variables in Z . Thus, we just need to check that the characteristic polynomial of C_i evaluates to $-I$, and to do so we use Lemma 3 for an appropriately defined system of polynomial equations. In the remaining, fix $i \in [m]$.

Our polynomials have variables X_1, \dots, X_r . Let Q_1 be the polynomial $P_R(X_1, \dots, X_r) + 1$, so that the equation $Q_1 = 0$ ensures $P_R(X_1, \dots, X_r) = -1$, where P_R is the characteristic polynomial of R . Let Q be the polynomial $P_{C_i}(X_{c_i}, X_{d_i}) + 1$, so that the equation $Q = 0$ ensures $P_{C_i}(X_{c_i}, X_{d_i}) = -1$, where P_{C_i} is the characteristic polynomial of C_i , and c_i and d_i are the indices of the components of Z in W_i . Then, every Boolean assignment that satisfies the equation $Q_1 = 0$ belongs to R , from which it follows that the Boolean assignment satisfies the conjunct C_i in the bijunctive definition of R , and hence it also satisfies the equation $Q = 0$. Thus, Lemma 3 applies and every operator assignment that satisfies $P_R(Z) = -I$ also satisfies $P_{C_i}(W) = -I$, as was to be proved.

We are now ready to complete the proof that if A is bijunctive, then A has no gaps. Let \mathcal{I} be an instance over A that is satisfiable via operators. The preceding paragraph shows that the 2SAT instance that results from replacing each constraint in the instance \mathcal{I} by its defining conjunction of 2-clauses is also satisfiable via operators. By Lemma 7, this 2SAT instance is also satisfiable in the Boolean domain. But then \mathcal{I} itself is satisfiable in the Boolean domain, as was to be shown.

If A is Horn or dual Horn, then the proof is entirely analogous.

5.2. Background on Post's lattice

Before we start with the second part in the proof of Theorem 3, we need to introduce some basic terminology and basic results from universal algebra; we devote this section to that.

Let $R \subseteq \{\pm 1\}^r$ be a Boolean relation of arity r and let $f : \{\pm 1\}^m \rightarrow \{\pm 1\}$ be a Boolean operation of arity m . The relation R is *invariant* under f if, for all sequences of m many r -tuples $(a_{1,1}, \dots, a_{1,r}), \dots, (a_{m,1}, \dots, a_{m,r})$ in $\{\pm 1\}^r$, the following holds:

$$\begin{aligned} &\text{if } (a_{1,1}, \dots, a_{1,r}), \dots, (a_{m,1}, \dots, a_{m,r}) \text{ are tuples in } R, \\ &\text{then } (f(a_{1,1}, \dots, a_{m,1}), \dots, f(a_{1,r}, \dots, a_{m,r})) \text{ is also a tuple in } R. \end{aligned} \quad (32)$$

Note that the tuple in the second line is obtained by applying the m -ary operation f to the m many tuples in the first line *componentwise*. If A is a Boolean constraint language, we say that A is invariant under f if every relation in A is invariant under f . Whenever A is invariant under f we also say that f is a *closure operation* of A .

The importance of the closure operations of a constraint language stems from the fact that they completely determine the relations that are pp-definable from it. This semantic characterization of the syntactic notion of pp-definability was discovered by Geiger [13] and, independently, Bodnarchuk et al. [3], for all constraint languages of arbitrary but finite domain. Here we state the special case of this characterization for the Boolean domain, since only this special case is needed in our applications.

Theorem 4 ([13,3]). *Let A be a Boolean constraint language and let R be a Boolean relation. The following statements are equivalent:*

1. R is pp-definable from A by a pp-formula without constants,
2. R is invariant under all Boolean closure operations of A .

In the following we refer to Theorem 4 as Geiger's Theorem.

Recall from Section 4 that a pp-formula without constants is one in which the constants $+1$ and -1 do not appear in its quantifier-free part of the formula. Although it will not be used until a later section, it is worth pointing out here that a similar characterization of pp-definability *with* constants exists. Indeed, it is easy to see that Geiger's Theorem implies that a Boolean relation R is pp-definable from the Boolean constraint language A by a pp-formula with constants if and only if R is invariant under all *idempotent* Boolean closure operations of A , or equivalently, invariant under all Boolean closure operations of the Boolean constraint language A^+ that is obtained from A by adding the two unary *singleton* relations $\{+1\}$ and $\{-1\}$; i.e., $A^+ = A \cup \{\{+1\}, \{-1\}\}$. We return to the issue of definability with constants in Section 7.

For every set F of Boolean operations, let $\text{Inv}(F)$ denote the set of all Boolean relations that are invariant under all operations in F . Conversely, for every set of Boolean relations A , let $\text{Pol}(A)$ denote the set of all Boolean operations under which all relations in A are invariant. Geiger's Theorem implies that the mappings $A \mapsto \text{Pol}(A)$ and $F \mapsto \text{Inv}(F)$ are the lower and upper adjoints of a Galois connection [8] between the partially ordered set of sets of Boolean relations ordered by inclusion, and the partially ordered set of sets of Boolean operations, also ordered by inclusion.

Note that for every constraint language A , the set $\text{Pol}(A)$ contains all *projection operations*: all operations $f : \{\pm 1\}^r \rightarrow \{\pm 1\}$ for which there exists an index $i \in [r]$ such that $f(x_1, \dots, x_r) = x_i$ for all $(x_1, \dots, x_r) \in \{\pm 1\}^r$. Also, $\text{Pol}(A)$ is closed under *compositions*: if $f : \{\pm 1\}^s \rightarrow \{\pm 1\}$ and $g_1, \dots, g_s : \{\pm 1\}^r \rightarrow \{\pm 1\}$ are operations from $\text{Pol}(A)$, then the operation $h = f \circ (g_1, \dots, g_s)$ defined by $h(x_1, \dots, x_r) = f(g_1(x_1, \dots, x_r), \dots, g_s(x_1, \dots, x_r))$ for all $(x_1, \dots, x_r) \in \{\pm 1\}^r$ is also in $\text{Pol}(A)$. Any set of operations that contains all projection operations and that is closed under compositions is called a *clone*.

Post [21] analyzed the collection of all clones of Boolean operations and completely determined the inclusions between them. In particular, he showed that this collection forms a lattice under inclusion, which is known as Post's lattice. In denoting clones in Post's lattice, we will follow the notation and terminology used by Böhler et al. [4]. The lattice is represented by the diagram in Fig. 1, which is also borrowed from [4] (we thank Steffen Reith for allowing us to reproduce the diagram here).

Each circle in the diagram of Fig. 1 represents a clone of Boolean operations, and a line between two circles denotes inclusion of the clone of the lower circle into the clone of the upper circle. Post showed that every clone of Boolean operations is represented in the diagram. Post also identified a *finite basis* of operations for each clone, which means that the clone is the smallest class of operations that contains the operations in the basis and all the projections operations, and that is closed under composition. For our application, we need only the bases for the eight clones called I_2 , I_0 , I_1 , D_2 , E_2 , V_2 , L_2 and N_2 . These are listed in the table in Fig. 2.

The final ingredient we need from Post's lattice is a characterization of the tractable Boolean constraint languages from Schaefer's Theorem in terms of their closure operations.

Theorem 5 (see Section 1.1 in [5]). *Let A be a Boolean constraint language. The following statements hold.*

1. A is 0-valid if and only if A is invariant under the constant false operation.
2. A is 1-valid if and only if A is invariant under the constant true operation.
3. A is *bijunctive* if and only if A is invariant under $(x \wedge y) \vee (x \wedge z) \vee (y \wedge z)$.
4. A is *Horn* if and only if A is invariant under $x \wedge y$.
5. A is *dual Horn* if and only if A is invariant under $x \vee y$.
6. A is *affine* if and only if A is invariant under $x \oplus y \oplus z$.

For the connection with Post's lattice, note that, by Fig. 2, the six conditions listed on the right of the entries 1 through 6 in Theorem 5 are equivalent to $\text{Pol}(A)$ containing the clones I_0 , I_1 , D_2 , E_2 , V_2 and E_2 , respectively.

As a stepping stone, we will use the known fact that LIN has gaps of every kind. We now discuss the proof of this fact and give the appropriate references to the literature.

Recall that LIN is the class of all affine relations, i.e., Boolean relations that are the set of solutions of a system of linear equations over the two-element field. In the ± 1 -representation, every such equation is a parity equation of the form $\prod_{i=1}^r x_i = y$, where $y \in \{\pm 1\}$.

Mermin [16,17] considered the following system \mathcal{M} of parity equations:

$$\begin{aligned} X_{11}X_{12}X_{13} &= 1 & X_{11}X_{21}X_{31} &= 1 \\ X_{21}X_{22}X_{23} &= 1 & X_{12}X_{22}X_{32} &= 1 \\ X_{31}X_{32}X_{33} &= 1 & X_{13}X_{23}X_{33} &= -1. \end{aligned} \tag{33}$$

Graphically, this system of equations can be represented by a square, where each equation on the left of (33) comes from a row, and each equation on the right of (33) comes from a column.

X_{11}	X_{12}	X_{13}	+1
X_{21}	X_{22}	X_{23}	+1
X_{31}	X_{32}	X_{33}	+1
+1	+1	-1	

It is easy to see that this system of equations has no solutions in the Boolean domain. Indeed, by multiplying the left-hand sides of all equations, we get 1 because every variable X_{ij} occurs twice in the system and $X_{ij}^2 = 1$. At the same time, by multiplying the right-hand sides of all equations, we get -1 , hence the system has no solutions in the Boolean domain. Observe, however, that this argument used the assumption that variables commute pairwise, even if they do not appear in the same equation. Thus, this argument does not go through if one assumes only that variables occurring in the same equation commute pairwise. Mermin [16,17] showed that the system \mathcal{M} has a solution consisting of linear operators on a Hilbert space of dimension four. Thus, in our terminology, Mermin established the following result.

Theorem 6 ([16,17]). \mathcal{M} witnesses a satisfiability gap of the first kind for LIN.

Cleve and Mittal [7, Theorem 1] have shown that a system of parity equations has a solution consisting of linear operators on a finite-dimensional Hilbert space if and only if there is a perfect strategy in a certain non-local game in the tensor-product model. Cleve, Liu, and Slofstra [6, Theorem 4] have shown that a system of parity equations has a solution consisting of linear operators on a (finite-dimensional or infinite-dimensional) Hilbert space if and only if there is a perfect strategy in a certain non-local game in the commuting-operator model. Slofstra [24] obtained a breakthrough result that has numerous consequences about these models. In particular, Corollary 3.2 in Slofstra’s paper [24] asserts that there is a system \mathcal{S} of parity equations whose associated non-local game has a perfect strategy in the commuting-operator model, but not in the tensor-product model. Thus, by combining Theorem 1 in [7], Theorem 4 in [6], and Corollary 3.2 in [24], we obtain the following result.

Theorem 7 ([6,7,24]). \mathcal{S} witnesses a satisfiability gap of the third kind for LIN.

LIN has a rather special place among all classes of Boolean relations that are not 0-valid, are not 1-valid, are not bi-junctive, are not Horn, and are not dual Horn. This special role is captured by the next lemma, which follows from Post’s analysis of the lattice of clones of Boolean functions from Section 5.2.

Lemma 8. Let A be a Boolean constraint language. If A is not 0-valid, not 1-valid, not bi-junctive, not Horn, and not dual Horn, then LIN is pp-definable from A .

Proof. Assume that A is a Boolean constraint language satisfying the hypothesis of Lemma 8. We consider the clone $\text{Pol}(A)$ and distinguish several cases using Post’s lattice.

If $\text{Pol}(A)$ is the smallest clone I_2 in Post’s lattice, then $\text{Pol}(A)$ contains only the projection functions; hence, every Boolean relation is closed under every function in $\text{Pol}(A)$. Geiger’s Theorem implies that every Boolean relation and, in particular, every relation in LIN, is pp-definable from A (and, in fact, it is pp-definable without using constants).

If $\text{Pol}(A)$ is not the smallest clone I_2 in Post’s lattice, then it must contain one of the seven minimal clones $I_0, I_1, D_2, E_2, V_2, L_2, N_2$ that contain I_2 . Recall that these clones have bases of operations as described in Fig. 2. Since A is not i -valid,

where $i = 0, 1$, and since the clone I_i is generated by the constant function $c_i(x) = i$, it must be the case that $\text{Pol}(A)$ does not contain the clone I_0 or the clone I_1 . Since A is not bijunctive, there is a relation in A that is not closed under the majority function $\text{maj}(x, y, z) = (x \wedge y) \vee (x \wedge z) \vee (y \wedge z)$. Since the clone D_2 is generated by the function $\text{maj}(x, y, z)$, it must be the case that $\text{Pol}(A)$ does not contain the clone D_2 . Since A is not Horn, there is a relation in A that is not closed under the function $\text{and}(x, y) = x \wedge y$. Since the clone E_2 is generated by the function $\text{and}(x, y)$, it must be the case that $\text{Pol}(A)$ does not contain the clone E_2 . Since A is not dual Horn, there is a relation in A that is not closed under the function $\text{or}(x, y) = x \vee y$. Since the clone V_2 is generated by the function $\text{or}(x, y)$, it must be the case that $\text{Pol}(A)$ does not contain the clone V_2 .

The preceding analysis shows that there are just two possibilities: $\text{Pol}(A)$ contains the clone L_2 or $\text{Pol}(A)$ contains the clone N_2 . Assume first that $\text{Pol}(A)$ contains the clone L_2 . Since L_2 is generated by the *exclusive or* function $\oplus(x, y, z) = x \oplus y \oplus z$ and since a relation is affine if and only if it is closed under the function \oplus , Geiger's Theorem implies that a relation is pp-definable without constants from A if and only if it is an affine relation. Thus, LIN is pp-definable from A (and, in fact, it is pp-definable without constants). Finally, assume that $\text{Pol}(A)$ contains the clone N_2 . Since $\text{Pol}(A)$ is generated by the function $\text{not}(x) = \neg x$, Geiger's Theorem implies that a relation is pp-definable without constants from A if and only if it is closed under the function $\text{not}(x)$. In particular, for every $n \geq 1$ and for $i = 0, 1$, the affine relation that is the set of solutions of the equation $x_1 + \dots + x_{2n} = i \pmod 2$ is pp-definable without constants from A . By using the constant 0 in these equations, we have that for every $n \geq 1$ and for every $i = 0, 1$, the affine relation that is the set of solutions of the equation $x_1 + \dots + x_{2n-1} = i \pmod 2$ is pp-definable from A (recall that pp-definitions allow constants). It follows that LIN is pp-definable from A . \square

The final lemma in this section asserts that reductions based on pp-definitions preserve satisfiability gaps upwards.

Lemma 9. *Let B and C be Boolean constraint languages such that B is pp-definable from C .*

1. *If B has a satisfiability gap of the first kind, then so does C .*
2. *If B has a satisfiability gap of the third kind, then so does $C \cup \{\text{T}\}$.*

Proof. For the first part, assume that B is pp-definable from C and that \mathcal{I} is an instance that witnesses a satisfiability gap of the first kind for B . Thus, \mathcal{I} is satisfiable via fd-operators, but is not satisfiable in the Boolean domain. Let \mathcal{J} be the instance over C as defined in Section 4.1. On the one hand, by Lemma 5, the instance \mathcal{J} is also satisfiable via fd-operator. On the other hand, by Lemma 4, the instance \mathcal{J} is also not satisfiable in the Boolean domain. Thus, \mathcal{J} witnesses a satisfiability gap of the first kind for C .

For the second part, assume that B is pp-definable from C and that \mathcal{I} is an instance that witnesses a satisfiability gap of the third kind for B . Thus, \mathcal{I} is satisfiable via operators, but it is not satisfiable via fd-operators. Let $\hat{\mathcal{J}}$ be the instance over $C \cup \{\text{T}\}$ as defined in Section 4.3. By Lemma 6, the instance $\hat{\mathcal{J}}$ is satisfiable via operators, but it is not satisfiable via fd-operators. Thus, $\hat{\mathcal{J}}$ witnesses a satisfiability gap of the third kind for $C \cup \{\text{T}\}$. \square

We now have all the machinery needed to put everything together.

Let A be a Boolean constraint language that is not 0-valid, not 1-valid, not bijunctive, not Horn, and not dual Horn. By Lemma 8, we have that LIN is pp-definable from A . Since, by Theorem 6, LIN has a satisfiability gap of the first kind, the first part of Lemma 9 implies that A has a satisfiability gap of the first kind. Since, by Theorem 7, LIN has a satisfiability gap of the third kind, the second part of Lemma 9 implies that A has a satisfiability gap of the third kind. The proof of Theorem 3 is now complete.

6. Further applications

In this section we discuss two applications of the results from Sections 4 and 5. The first application is about classification theorems in the style of Schaefer. The second application builds on Slofstra's results to answer some open questions from [1] on the quantum realizability of contextuality scenarios. While these open questions were solved earlier by Fritz also using Slofstra's results (see [12]), our alternative perspective may still add some value since, as we will see, we obtain improved, and indeed optimal, parameters.

6.1. Dichotomy theorems

For a Boolean constraint language A , let $\text{SAT}(A)$ denote the following decision problem:

Given an instance \mathcal{I} over A , is \mathcal{I} satisfiable in the Boolean domain?

Similarly, let $\text{SAT}^*(A)$ and $\text{SAT}^{**}(A)$ be the versions of the problem in which the questions are whether \mathcal{I} is satisfiable via an operator assignment on a finite-dimensional Hilbert space, or on an arbitrary Hilbert space, respectively. We say that a

problem poly- m -reduces to another if there is a polynomial-time computable function that transforms instances of the first problem into instances of the second in such a way that the answer is preserved.

Recall that T denotes the full binary Boolean relation $\{\pm 1\}^2$. The construction in Section 4.3 and Lemma 6 give the following:

Lemma 10. *Let A and B be Boolean constraint languages and let $A' = A \cup \{T\}$ and $B' = B \cup \{T\}$. If A is pp-definable from B , then*

1. $\text{SAT}(A')$ poly- m -reduces to $\text{SAT}(B')$, $\text{SAT}^*(B')$, and $\text{SAT}^{**}(B')$,
2. $\text{SAT}^*(A')$ poly- m -reduces to $\text{SAT}^*(B')$.
3. $\text{SAT}^{**}(A')$ poly- m -reduces to $\text{SAT}^{**}(B')$.

Slofstra's Corollary 3.3 in [24] in combination with Theorem 4 in [6] gives the undecidability of $\text{SAT}^{**}(\text{LIN})$ which, from now on we denote by LIN SAT^{**} .

Theorem 8 ([24], [6]). *LIN SAT^{**} is undecidable.*

In combination with Lemmas 7, 10, and 8, we get the following dichotomy theorem:

Theorem 9. *Let A be a Boolean constraint language and let $A' = A \cup \{T\}$. Then, exactly one of the following holds:*

1. $\text{SAT}^{**}(A')$ is decidable in polynomial time,
2. $\text{SAT}^{**}(A')$ is undecidable.

Moreover, the first case holds if and only if A is 1-valid, or A is 0-valid, or A is bijunctive, or A is Horn, or A is dual Horn.

Proof. If A is 1-valid, 0-valid, bijunctive, Horn, or dual Horn, then A' is also of the same type; indeed T is both 1-valid and 0-valid, and it is also bijunctive, Horn and dual Horn since it is defined by the empty conjunction of any kind of clauses. Thus $\text{SAT}^{**}(A')$ is the same problem as $\text{SAT}(A')$ by Lemma 7, which is solvable in polynomial time.

If on the contrary A is neither 1-valid, nor 0-valid, nor bijunctive, nor Horn, nor dual Horn, then Lemma 8 applies and LIN has a pp-definition from A . In such a case Lemma 10 applies and $\text{SAT}^{**}(\text{LIN}')$ reduces to $\text{SAT}^{**}(A')$, where LIN' denotes $\text{LIN} \cup \{T\}$. Since every instance of LIN SAT^{**} is also an instance of $\text{SAT}^{**}(\text{LIN}')$, the undecidability of $\text{SAT}^{**}(A')$ follows from Theorem 8. \square

Note that, in case 2, Theorem 9 states that $\text{SAT}^{**}(A')$ is undecidable but it says nothing about $\text{SAT}^{**}(A)$. Luckily, in most cases it is possible to infer the undecidability of $\text{SAT}^{**}(A)$ from the undecidability of $\text{SAT}^{**}(A')$. This is the case, for example, for both

$$3\text{SAT} = \{\{\pm 1\}^3 \setminus \{(a_1, a_2, a_3) : a_1, a_2, a_3 \in \{\pm 1\}\},$$

$$3\text{LIN} = \{(a_1, a_2, a_3) \in \{\pm 1\}^3 : a_1 a_2 a_3 = b\} : b \in \{\pm 1\}\}.$$

In the following we write 3LIN SAT^* and 3LIN SAT^{**} to denote the problems $\text{SAT}^*(A)$ and $\text{SAT}^{**}(A)$ for $A = 3\text{LIN}$. Similarly, we use 3SAT^* and 3SAT^{**} to denote $\text{SAT}^*(A)$ and $\text{SAT}^{**}(A)$ for $A = 3\text{SAT}$.

Theorem 10. *3LIN SAT^{**} and 3SAT^{**} are both undecidable.*

Proof. Let A be the Boolean constraint language of 3LIN or 3SAT . It follows from Theorem 9 that $\text{SAT}^{**}(A')$ is undecidable. Now we reduce this problem to $\text{SAT}^{**}(A)$. Take any instance \mathcal{I} over A' and replace each constraint of the type $((Z_1, Z_2), T)$ by an equation $Z_1 Z_2 Y = -1$ in the case of 3LIN , and a clause $Z_1 \vee Z_2 \vee Y$ in the case of 3SAT , where Y is a fresh variable not used anywhere else in the instance. Let \mathcal{J} be the resulting instance. If f is a satisfying operator assignment for \mathcal{I} , then we claim that an appropriate extension g of f is a satisfying operator assignment for \mathcal{J} . For 3LIN , set $g(Y) = -f(Z_2)f(Z_1)$. For 3SAT , set $g(Y) = -I$. To see that this works, first note that $g(Z_1) = f(Z_1)$ and $g(Z_2) = f(Z_2)$ commute since they appear together in a constraint of \mathcal{I} . Thus, in both cases $g(Z_1)$, $g(Z_2)$ and $g(Y)$ pairwise commute. Moreover, in the 3LIN case the assignment $g(Y) = -f(Z_2)f(Z_1)$ is chosen so that the equation $g(Z_1)g(Z_2)g(Y) = -I$ is satisfied; to check this, multiply $g(Y) = -f(Z_2)f(Z_1)$ by $g(Z_1)g(Z_2)$ from the right and use $g(Z_2)f(Z_2) = f(Z_2)^2 = I$ and $g(Z_1)f(Z_1) = f(Z_1)^2 = I$. Also, in the 3SAT case the assignment $g(Y) = -I$ annihilates the product in the expression of the characteristic polynomial of the clause $Z_1 \vee Z_2 \vee Y$ in see Lemma 2, which makes the characteristic polynomial evaluate to $-I$ regardless of what $g(Z_1)$ and $g(Z_2)$ are. Thus, the new constraints are satisfied by g and the claim is proved.

Conversely, if g is a satisfying operator assignment for \mathcal{J} , then the restriction of g to the variables of \mathcal{I} is a satisfying operator assignment for \mathcal{I} , just because the commutativity of $f(Z_1)$ and $f(Z_2)$ is enforced by the fact that they

appear together in the constraint $Z_1 Z_2 Y = -I$ or $Z_1 \vee Z_2 \vee Y$ of \mathcal{J} , and because the characteristic polynomial of T is the constant -1 . \square

The same construction and argument that we used in Theorem 10 starting at a gap instance over the constraint language $3SAT \cup \{T\}$ gives a gap instance over $3SAT$ that will be useful later on.

Corollary 1. *There is an instance over the Boolean constraint language $3SAT$ that witnesses a satisfiability gap of the third kind; it is satisfiable via operator assignments over some Hilbert space but not over a finite-dimensional Hilbert space.*

6.2. Quantum realizability of contextuality scenarios

We follow the terminology in the paper by Acín, Fritz, Leverrier and Sainz [1]. A *contextuality scenario* is a hypergraph H with set $V(H)$ of vertices and set $E(H) \subseteq 2^{V(H)}$ of edges such that $\bigcup_{e \in E(H)} e = V(H)$. Given a contextuality scenario H , a quantum model for it is, informally, an assignment of probabilities to the vertices of H that are reproduced as the observation probabilities of a collection of projective measurements associated to the edges of H , when the measurements are applied to some quantum state. When a contextuality scenario has at least one quantum model, one says that H allows quantum models. As argued in [1], this can be equivalently stated formally, without any reference to measurements or quantum states, as follows.

We say that a contextuality scenario H *allows a quantum model*, or is *quantum realizable*, if there exists a Hilbert space \mathcal{H} and an assignment of bounded linear operators P_v on \mathcal{H} to each vertex v in $V(H)$ in such a way that:

1. P_v is self-adjoint,
2. $P_v^2 = P_v$ for each $v \in V(H)$,
3. $\sum_{v \in e} P_v = I$ for each $e \in E(H)$.

Note that 1 and 2 together say that each P_v is an orthogonal projection operator,³ and 3 says that the projection operators associated to the vertices of each edge resolve the identity. In [1] the question was raised whether there exist contextuality scenarios that are quantum realizable but only over infinite-dimensional Hilbert spaces (see Problem 8.3.2 in [1]). A related computational question was also raised: Is it decidable whether a contextuality scenario given as input allows a quantum state? (see Conjecture 8.3.3 in [1]). Following the notation in [1], this problem is called *ALLOWS-QUANTUM*. The restriction of the problem in which the input hypergraph has edges of cardinality at most k we call *k-ALLOWS-QUANTUM*. See [1] for a discussion on why these problems are important, and their relationship to Connes Embedding Conjecture in functional analysis.

Soon after Slofstra published his results, both questions raised in [1] were answered by Fritz by reduction from Slofstra's Theorems 7 and 8 (see [12]). In particular, Fritz proved that *ALLOWS-QUANTUM* is undecidable. In the following we illustrate the methods developed in the previous sections to give alternative proofs of these results. As a bonus, our proof also gives optimal parameters; we get hypergraphs with edges of size at most 3 that separate infinite-dimensional realizability from finite-dimensional realizability, and we show that already *3-ALLOWS-QUANTUM* is undecidable. In contrast, Fritz' reduction incurs an exponential loss in the size of the edges of the hypergraphs with respect to the arity of the constraints in Slofstra's result, which is a priori not bounded, and the best it can achieve is size 4 anyway. Moreover, as we will see, our 3 in the maximum size of the edges is optimal since it turns out that *2-ALLOWS-QUANTUM* is decidable (and even solvable in polynomial time).

Next we show how our methods can be used to answer these questions. First, notice that there is a clear similarity between the requirements 1, 2 and 3 in the definition of quantum realization of H and the requirements that an operator assignment for a collection of variables $\{X_v : v \in V(H)\}$ associated to the vertices of H must satisfy. For one thing, if we define $A_v = I - 2P_v$ for every $v \in V(H)$, then each A_v is a bounded self-adjoint linear operator such that $A_v^2 = I$. Moreover, the fact that the projections associated to an edge of H resolve the identity implies that they pairwise commute. Thus, the operators A_v associated to the vertices of e also pairwise commute for every edge e of H . This means that the assignment $X_v \mapsto A_v$ thus defined is a valid operator assignment to any instance with constraint scopes given by the hyperedges of H .

However, the condition $\sum_{v \in e} \frac{1}{2}(A_v - I) = I$ implied by condition 3 through the inverse transformation $P_v = \frac{1}{2}(A_v - I)$ does not correspond directly to a constraint of the form $P_R(A_v : v \in e) = -I$ for any Boolean relation R . This means that we cannot interpret the quantum realizability problem directly as an instance of a satisfiability problem via operator assignments over a Boolean constraint language. However, as it turns out, the problem that we called *3-ALLOWS-QUANTUM* is literally the same as the arbitrary Hilbert space version *1-IN-3 SAT** of the problem called *1-IN-3 SAT** by Ji.⁴ Ji proved

³ Acín et al. refer to orthogonal projection operators as projections. In the following we will also refer to them as projections in order to avoid confusion with the fact that two orthogonal projection operators P and Q could fail to satisfy $PQ = 0$. It may also be worth pointing out that linear-algebraic projection operators of this section are unrelated to the universal-algebraic projection operations from Section 5.2.

⁴ There is an unfortunate clash in notation in that the problem *1-IN-3 SAT** studied by Ji [15] is not the same as the problem that we would call *SAT*(1-IN-3 SAT)*, where *1-IN-3 SAT* is the Boolean relation $\{((-1, +1, +1), (+1, -1, +1), (+1, +1, -1))\}$. Note that $P_{1\text{-IN-3 SAT}}(X_1, X_2, X_3) = \frac{3}{4}X_1X_2X_3 +$

that 3SAT* reduces to 1-IN-3 SAT*, and in view of Theorem 10, the question arises whether 3SAT** also reduces to 1-IN-3 SAT**, or to 3-ALLOWS-QUANTUM, which is the same. We show that it does.

Before we can do it, though, we need the following lemma that Ji proved for finite-dimensional Hilbert spaces (see Lemma 5 in [15]), and that we prove for all Hilbert spaces:

Lemma 11. *Let \mathcal{H} a Hilbert space. For every two projection operators P_1 and P_2 of \mathcal{H} that commute, there exist projection operators Q_1, Q_2, Q_3 and Q_4 of \mathcal{H} such that*

$$\begin{aligned} P_1 + Q_1 + Q_4 &= I \\ P_2 + Q_2 + Q_4 &= I \\ Q_1 + Q_2 + Q_3 &= I. \end{aligned}$$

Conversely, if $P_1, P_2, Q_1, Q_2, Q_3, Q_4$ are projection operators of \mathcal{H} that satisfy these equations, then P_1 and P_2 commute.

Proof. To prove the first claim, consider the pp-formula

$$\phi(Z_1, Z_2) = \exists U_1 \exists U_2 \exists U_3 \exists U_4 (R_{1/3}(Z_1, U_1, U_4) \wedge R_{1/3}(Z_2, U_2, U_4) \wedge R_{1/3}(U_1, U_2, U_3)), \tag{34}$$

where $R_{1/3} = \{(-1, +1, +1), (+1, -1, +1), (+1, +1, -1)\}$. It is straightforward to check that this formula defines the full binary Boolean relation $T = \{\pm 1\}^2$. Now, let \mathcal{I} be the instance $((Z_1, Z_2), T)$ and let \mathcal{J} be the instance obtained from \mathcal{I} as in Section 4.1. Let f be defined by $f(Z_1) = 1 - 2P_1$ and $f(Z_2) = 1 - 2P_2$. Since P_1 and P_2 commute and the characteristic polynomial of T is the constant -1 , the assignment f is a satisfying operator assignment for the instance $((Z_1, Z_2), T)$. By Lemma 5, there exists g that extends f and is a satisfying operator assignment for \mathcal{J} over \mathcal{H} . Moreover, g is pairwise commuting on each block of \mathcal{J} . Take $Q_i = (1 - g(U_i))/2$ for $i = 1, 2, 3, 4$. Then Q_1, \dots, Q_4 are projection operators, and $P_1, P_2, Q_1, \dots, Q_4$ pairwise commute. We claim that they satisfy the equations in the lemma. To see this we apply Lemma 3. Since the equation $P_{R_{1/3}}(Z_1, U_1, U_4) = -1$ entails the equation $(1 - Z_1)/2 + (1 - U_1)/2 + (1 - U_4)/2 = 1$ over the Boolean domain $\{\pm 1\}$, and at the same time P_1, Q_1, Q_4 pairwise commute, the equation $P_{R_{1/3}}(g(Z_1), g(U_1), g(U_4)) = -1$ implies $P_1 + Q_1 + Q_4 = I$ by Lemma 3. For the other two equations, the argument is the same.

For the converse, we use the following easy to verify identities discovered via a computer search by Ji (see the proof of Lemma 5 in [15]):

$$\begin{aligned} [P_1 + Q_1 + Q_4 - I, -P_1 + Q_1 + Q_3] &= [P_1, Q_3] + [Q_4, Q_3] \\ [P_2 + Q_2 + Q_4 - I, -P_1] &= [P_1, P_2] + [P_1, Q_2] \\ [Q_1 + Q_2 + Q_3 - I, P_1 + Q_4] &= [Q_2, P_1] + [Q_3, P_1] + [Q_3, Q_4], \end{aligned}$$

where $[X, Y]$ denotes the commutator polynomial $XY - YX$. Note that the equations in the lemma imply that the left-hand sides are all 0. On the other hand, using the identity $[X, Y] + [Y, X] = 0$, the sum of the right-hand sides is $[P_1, P_2]$. This gives $[P_1, P_2] = 0$ and thus P_1 and P_2 commute. \square

Lemma 12. *3SAT** poly- m -reduces to 3-ALLOWS-QUANTUM.*

Proof. Schaefer proved that 3SAT is pp-definable from the constraint language given by the single relation $R_{1/3} = \{(-1, +1, +1), (+1, -1, +1), (+1, +1, -1)\}$. If in addition to $R_{1/3}$ we allow also the relations $R_{1/2} = \{(-1, +1), (+1, -1)\}$ and $R_{1/1} = \{-1\}$, then the pp-definition can be assumed to have the property that each atom involves different variables and no constants. For example, an atom of the form $R_{1/3}(X, X, Z)$ can be replaced by $R_{1/3}(X, Y, Z) \wedge R_{1/2}(X, Y') \wedge R_{1/2}(Y', Y)$, where Y and Y' are fresh quantified variables that do not appear anywhere else in the formula.

We use this for the construction in Section 4.3. Let \mathcal{I} be a 3SAT instance and let $\hat{\mathcal{J}}$ be the instance over the Boolean constraint language $A = \{R_{1/3}, R_{1/2}, R_{1/1}, T\}$ given by the construction in Section 4.3, using the pp-definition of 3SAT from A. Starting at $\hat{\mathcal{J}}$ we produce an instance of 3-ALLOWS-QUANTUM as follows: Each variable in $\hat{\mathcal{J}}$ becomes a vertex in the hypergraph. Each constraint of the type $((Z_1, Z_2, Z_3), R_{1/3})$ becomes a hyperedge $\{Z_1, Z_2, Z_3\}$, each constraint of the type $((Z_1, Z_2), R_{1/2})$ becomes a hyperedge $\{Z_1, Z_2\}$, each constraint of the type $(Z, R_{1/1})$ becomes a singleton hyperedge $\{Z\}$, and each constraint of the type $((Z_1, Z_2), T)$ introduces four fresh vertices U_1, U_2, U_3, U_4 and three hyperedges $\{Z_1, U_1, U_4\}$, $\{Z_2, U_2, U_4\}$ and $\{U_1, U_2, U_3\}$ in correspondence with the equations of Lemma 11 with Z_1, Z_2 playing the role of P_1, P_2 , and U_1, U_2, U_3, U_4 playing the role of Q_1, Q_2, Q_3, Q_4 . Let H be the hypergraph that results from this construction. We claim

$\frac{1}{4}X_1X_2 + \frac{1}{4}X_2X_3 + \frac{1}{4}X_1X_3 - \frac{1}{4}X_1 - \frac{1}{4}X_2 - \frac{1}{4}X_3 + \frac{1}{4}$, so the difference is that, even though the characteristic polynomial equation $P_{1\text{-IN-}3\text{ SAT}}(X_1, X_2, X_3) = -1$ is satisfied by an operator assignment if and only if the resolution of the identity equation $\frac{1}{2}(1 - X_1) + \frac{1}{2}(1 - X_2) + \frac{1}{2}(1 - X_3) = -1$ is satisfied by the same operator assignment, the two polynomials $P_{1\text{-IN-}3\text{ SAT}}(X_1, X_2, X_3)$ and $\frac{1}{2}(1 - X_1) + \frac{1}{2}(1 - X_2) + \frac{1}{2}(1 - X_3)$ are by no means the same.

that for every Hilbert space \mathcal{H} , the instance \mathcal{I} is satisfiable via operator assignments over \mathcal{H} if and only if the hypergraph H is quantum realizable over \mathcal{H} .

In the forward direction, let f be a satisfying operator assignment for \mathcal{I} over \mathcal{H} . By Lemma 6, there is a g that extends f and is a satisfying operator assignment for $\hat{\mathcal{J}}$ over \mathcal{H} . Recall now that each vertex of H is indeed a variable of $\hat{\mathcal{J}}$, or an additional vertex of the type U_1, U_2, U_3, U_4 introduced by a constraint of the form $((Z_1, Z_2), T)$. For each v of the first type, let P_v be the projection operator given by $(1 - g(v))/2$. For each v of the second type, let P_v be the projection given by Lemma 11 for the projection assignment $P_1 = P_{Z_1}$ and $P_2 = P_{Z_2}$ with U_1, U_2, U_3, U_4 corresponding to Q_1, Q_2, Q_3, Q_4 . Note that P_1 and P_2 commute, since Z_1 and Z_2 appear together in $((Z_1, Z_2), T)$ and hence $g(Z_1)$ and $g(Z_2)$ commute, so the lemma applies. We claim that this assignment of operators does the job.

We just need to check that the projection operators resolve the identity on every edge of H . For edges of the type $\{Z_1, Z_2, Z_3\}$ introduced by a constraint $((Z_1, Z_2, Z_3), R_{1/3})$ we show this with an application of Lemma 3: the equation $P_{R_{1/3}}(Z_1, Z_2, Z_3) = -1$ entails the equation $(1 - Z_1)/2 + (1 - Z_2)/2 + (1 - Z_3)/2 = 1$ over the Boolean domain $\{\pm 1\}$, and therefore, since $g(Z_1), g(Z_2), g(Z_3)$ pairwise commute, the equation $P_{R_{1/3}}(g(Z_1), g(Z_2), g(Z_3)) = -1$ implies $P_{Z_1} + P_{Z_2} + P_{Z_3} = I$ by Lemma 3. For edges of the types $\{Z_1, Z_2\}$ or $\{Z\}$ introduced by constraints of the types $((Z_1, Z_2), R_{1/2})$ or $(Z, R_{1/2})$, respectively, the argument is the same. Finally, for the three edges that come from a constraint of the form $((Z_1, Z_2), T)$, the claim follows from Lemma 11. This completes one direction of the reduction.

For the other direction, let $v \mapsto P_v$ be an assignment of projection operators of \mathcal{H} that witnesses that H is quantum realizable. Recall again that each vertex v of H is a variable of $\hat{\mathcal{J}}$, or an additional vertex of the type U_1, U_2, U_3, U_4 coming from a T-constraint. For each v of the first type, let $A_v = I - 2P_v$. Each A_v is a self-adjoint bounded linear operator that squares to the identity. Moreover, any two variables of $\hat{\mathcal{J}}$ that appear together in a constraint that is not a T-constraint appear together as vertices in some edge of H . Therefore the corresponding operators belong to the resolution of the identity of that edge, and a set of projection operators that resolve the identity are pairwise orthogonal and hence commute. Also, for any two variables of $\hat{\mathcal{J}}$ that appear together in a constraint of the form $((Z_1, Z_2), T)$, the corresponding operators commute thanks to the “conversely” clause in Lemma 11. Thus, the only thing left to do is checking that each constraint of $\hat{\mathcal{J}}$ is satisfied.

For constraints of the type $((Z_1, Z_2, Z_3), R_{1/3})$ this follows also from an application of Lemma 3: the equation $(1 - Z_1)/2 + (1 - Z_2)/2 + (1 - Z_3)/2 = 1$ entails the equation $P_{R_{1/3}}(Z_1, Z_2, Z_3) = -1$ over the Boolean domain $\{\pm 1\}$, and since $A_{Z_1}, A_{Z_2}, A_{Z_3}$ pairwise commute, the equation $P_{Z_1} + P_{Z_2} + P_{Z_3} = I$ implies $P_{R_{1/3}}(A_{Z_1}, A_{Z_2}, A_{Z_3}) = -1$ by Lemma 3. For constraints of the type $((Z_1, Z_2), R_{1/2})$ and $(Z, R_{1/2})$ the argument is the same. \square

In combination with Theorem 10 we get the following.

Corollary 2. 3-ALLOWS-QUANTUM and ALLOWS-QUANTUM are undecidable.

The same construction as in Lemma 12 starting from Corollary 1 gives the next.

Corollary 3. There exists a hypergraph with edges of size at most three that is quantum realizable on some Hilbert space, but not on a finite-dimensional Hilbert space.

It was mentioned earlier that 2-ALLOWS-QUANTUM is decidable in polynomial time. One way to see this is by arguing that a hypergraph with edges of size two (i.e., a graph) is quantum realizable if and only if it is bipartite. Another is by reduction to 2SAT**, which is decidable in polynomial time by Theorem 9. A close look reveals that, indeed, both proofs are the same.

Theorem 11. 2-ALLOWS-QUANTUM is decidable in polynomial time.

Proof. We reduce to 2SAT**. Given a hypergraph H , build the 2SAT instance that has one variable X_v for each vertex in $V(H)$, two clauses $X_u \vee X_v$ and $\neg X_u \vee \neg X_v$ for every edge $\{u, v\} \in E(H)$, and one unit clause X_u for each singleton edge $\{u\}$ in $E(H)$. It is straightforward to check that this reduction works through the usual conversion from projection operators to involutions $P_v \mapsto 1 - 2P_v$, and the usual conversion from involutions to projection operators $A_v \mapsto (1 - A_v)/2$. \square

7. Closure operations

In this section we develop a generalization of the concept of closure operation from Section 5.2 for sets of operator assignments. For every Boolean r -ary relation R , let R^* denote the set of fully commuting r -variable operator assignments over finite-dimensional Hilbert spaces that satisfy the equation $P_R(X_1, \dots, X_r) = -I$. We show that every closure operation for R gives a suitable closure operation for R^* . As an application, we show that the set of Boolean relations that are pp-definable from a Boolean constraint language is not enlarged when we allow the existential quantifiers to range over operator assignments.

7.1. Closure operations and pp*-definitions

Let A be a Boolean constraint language and let R be a Boolean relation of arity r . Let $\psi = R_1(z_1) \wedge \dots \wedge R_m(z_m)$ be a conjunction of atoms with relations from A ; i.e. each R_i is a relation from A , and each z_i denotes a tuple of the appropriate arity made of first-order variables or constants in $\{\pm 1\}$. Each such formula can be thought of as an instance over A . Concretely, it can be thought of as the instance $\mathcal{I} = ((Z_1, R_1), \dots, (Z_m, R_m))$, where each Z_i is obtained from the corresponding z_i by replacing each first-order variable x by a corresponding variable X , and leaving all constants untouched.

Let \mathcal{H} be a finite-dimensional Hilbert space. We say that R is pp*-definable from A over \mathcal{H} if there is a pp-formula $\phi(x_1, \dots, x_r) = \exists y_1 \dots \exists y_s (\psi(x_1, \dots, x_r, y_1, \dots, y_s))$ over A , where ψ is a conjunction as above, such that, for every $a_1, \dots, a_r \in \{\pm 1\}$, the tuple (a_1, \dots, a_r) is in R if and only if the instance

$$\psi(x_1/a_1, \dots, x_r/a_r, y_1/Y_1, \dots, y_s/Y_s) \tag{35}$$

is satisfiable via operator assignments over \mathcal{H} . We say that R is pp*-definable from A if it is pp*-definable from A over a finite-dimensional Hilbert space. One of the goals of this section is to prove the following conservativity theorem:

Theorem 12. *Let A be a Boolean constraint language and let R be a Boolean relation. If R is pp*-definable from A , then R is pp-definable from A .*

In order to prove this we need to develop the concept of closure operation for sets of operator assignments. Let r be a positive integer and let X_1, \dots, X_r be variables. Recall that an operator assignment for X_1, \dots, X_r is a tuple A_1, \dots, A_r of self-adjoint bounded linear operators that square to the identity, all taken from a common Hilbert space. The operator assignment is fully commuting if $A_i A_j = A_j A_i$ holds for each $i, j \in [r]$. A relation of operator assignments of arity r is a set of fully commuting operator assignments for a fixed set of r variables. Note that we do not require that all operator assignments come from the same Hilbert space. The relation is called *Boolean* if all assignments in it come from a Hilbert space of dimension 1; i.e., from \mathbb{C} . If \mathcal{H} is a Hilbert space and $R \subseteq \{\pm 1\}^r$ is a Boolean relation of arity r , we write $R^{\mathcal{H}}$ for the set of fully commuting operator assignments for X_1, \dots, X_r over \mathcal{H} that satisfy the polynomial equation $P_R(X_1, \dots, X_r) = -I$, where P_R is the characteristic polynomial of R . We write R^* for the union of $R^{\mathcal{H}}$ over all finite-dimensional Hilbert spaces. If A is a set of Boolean relations, define $A^* = \{R^* : R \in A\}$.

Let $\mathcal{H}_1, \dots, \mathcal{H}_m$ and \mathcal{H} be Hilbert spaces, and let f be a function that takes as inputs m many linear operators, one on each \mathcal{H}_i , and produces as output a linear operator on \mathcal{H} . We say that f is an *operation* if the following conditions are satisfied.

1. If A_i is a 1-variable operator assignment over \mathcal{H}_i for $i = 1, \dots, m$, then $f(A_1, \dots, A_m)$ is a one-variable operator assignment over \mathcal{H} .
2. If $A_{i,1}, A_{i,2}$ is a fully commuting 2-variable operator assignment over \mathcal{H}_i for $i = 1, \dots, m$, then $f(A_{1,1}, \dots, A_{m,1})$, $f(A_{1,2}, \dots, A_{m,2})$ is a fully commuting 2-variable operator assignment over \mathcal{H} .

Let R be a relation of operator assignments of arity r and let F be a collection of operations as above. We say that R is *invariant* under F if for each $f \in F$ the following additional condition is also satisfied.

3. If $A_{i,1}, \dots, A_{i,r}$ is a fully commuting r -variable operator assignment over \mathcal{H}_i that belongs to R for $i = 1, \dots, m$, then $f(A_{1,1}, \dots, A_{m,1}), \dots, f(A_{1,r}, \dots, A_{m,r})$ is a fully commuting r -variable operator assignment over \mathcal{H} that belongs to R .

If A is a set of relations of operator assignments, we say that A is *invariant* under F if every relation in A is invariant under F . We also say that F is a *closure operation* of A . A *Boolean closure operation* of A is one in which the dimensions of all Hilbert spaces involved are 1; i.e., they are \mathbb{C} . Before we prove the main technical result of this section, we work out a motivating example.

7.2. Example: LIN

In this section we study whether R^* for $R = \text{LIN}$ has some closure operation. In the 0-1-representation of Boolean values, the function $(X_1, X_2, X_3) \mapsto X_1 \oplus X_2 \oplus X_3$ is a Boolean closure operation of LIN. In the ± 1 -representation of Boolean values, this is $(X_1, X_2, X_3) \mapsto X_1 X_2 X_3$. It is tempting to think that the map $(X_1, X_2, X_3) \mapsto X_1 X_2 X_3$ applied to linear operators on a Hilbert space could already be a closure operation for LIN*. However, the solution to the Mermin-Peres magic square equations (33) is a counterexample: each row equation is a parity equation with even right-hand side that is satisfied, but the composition of columns by the operation $X_1 X_2 X_3$ gives an operator assignment that satisfies a parity equation with odd right-hand side.

It turns out that the correct way of generalizing the Boolean closure operation is not by taking ordinary products, but Kronecker products. Let F be the function that takes any three linear operators X_1, X_2, X_3 over the same finite-dimensional Hilbert space and is defined by

$$F(X_1, X_2, X_3) = X_1 \otimes X_2 \otimes X_3. \tag{36}$$

Now let (A_1, \dots, A_r) , (B_1, \dots, B_r) and (C_1, \dots, C_r) be three fully commuting r -variable operator assignments over a finite-dimensional Hilbert space, say \mathbb{C}^d . We think of all operators as matrices. Take $D_i = F(A_i, B_i, C_i)$ for $i = 1, \dots, r$. These are Hermitian matrices since the operations of conjugate transposition and Kronecker product commute. Also

$$D_i D_j = (A_i A_j) \otimes (B_i B_j) \otimes (C_i C_j) = (A_j A_i) \otimes (B_j B_i) \otimes (C_j C_i) = D_j D_i \tag{37}$$

so D_1, D_2, D_3 pairwise commute. Equation (37) also gives $D_i^2 = (A_i^2) \otimes (B_i^2) \otimes (C_i^2) = I \otimes I \otimes I = I$, so (D_1, \dots, D_r) is a fully commuting r -variable operator assignment. Next we consider a relation in LIN, say $R = \{(a_1, \dots, a_r) \in \{\pm 1\}^r : a_1 \cdots a_r = b\}$, with $b \in \{\pm 1\}$. Note that its characteristic polynomial is $P_R(X_1, \dots, X_r) = -b \cdot X_1 \cdots X_r$. We show that if $P_R(A_1, \dots, A_r) = P_R(B_1, \dots, B_r) = P_R(C_1, \dots, C_r) = -I$, then also $P_R(D_1, \dots, D_r) = -I$. We have

$$\prod_{i=1}^r D_i = \left(\prod_{i=1}^r A_i \right) \otimes \left(\prod_{i=1}^r B_i \right) \otimes \left(\prod_{i=1}^r C_i \right) = (bI) \otimes (bI) \otimes (bI) = b^3 I = bI. \tag{38}$$

Hence $P_R(D_1, \dots, D_r) = -b^2 I = -I$. This shows that F is a closure operation of LIN*.

One consequence of the existence of F as a closure operation of LIN* is that the binary OR relation $\text{OR}_2 = \{\pm 1\}^2 \setminus \{(+1, +1)\}$ is not pp*-definable from LIN.

Theorem 13. OR_2 is not pp*-definable from LIN.

Note that this follows from the more general statement in Theorem 12 since it is known that the Boolean relation OR_2 is not pp-definable from LIN. Indeed, OR_2 is not closed under the (idempotent) Boolean closure operation $(X_1, X_2, X_3) \mapsto X_1 X_2 X_3$ of LIN, since $(-1, -1)$, $(+1, -1)$ and $(-1, +1)$ are all three in the relation OR_2 , but $(+1, +1)$ is not in OR_2 . The undefinability of OR_2 from LIN by a pp-formula (with or without constants) follows from the easy direction in Geiger’s Theorem 4. Since we prove Theorem 12 below, we omit a proof of Theorem 13 at this point.

7.3. Generalization

We show that every Boolean closure operation gives a closure operation for relations of operator assignments over finite-dimensional Hilbert spaces. In the following, if X_i is a linear operator on a Hilbert space, X_i^0 and X_i^1 are to be interpreted as the identity operator and X_i itself, respectively. If S is a set, we write $S(i)$ for the 0-1-indicator of the fact that i is in S ; i.e. $S(i) = 1$ if i is in S , and $S(i) = 0$ if i is not in S .

Theorem 14. Let A be a Boolean constraint language and let $f : \{\pm 1\}^m \rightarrow \{\pm 1\}$ be a Boolean closure operation of A . Then the function on linear operators on finite-dimensional Hilbert spaces defined by

$$F(X_1, \dots, X_m) = \sum_{S \subseteq [m]} \hat{f}(S) \bigotimes_{i \in [m]} X_i^{S(i)} \tag{39}$$

is a closure operation of A^* . Moreover, $F(a_1 I, \dots, a_m I) = f(a_1, \dots, a_m) I$ holds for every $(a_1, \dots, a_m) \in \{\pm 1\}^m$.

Proof. First we show that F is an operation; i.e., it satisfies conditions 1 and 2 in the definition of operation. Let X_1, \dots, X_m be 1-variable operator assignments over $\mathcal{H}_1, \dots, \mathcal{H}_m$. In particular, X_1, \dots, X_m are all self-adjoint linear operators. Thus, for $S \subseteq [m]$ we have

$$\left(\bigotimes_{i \in [m]} (X_i)^{S(i)} \right)^* = \bigotimes_{i \in [m]} (X_i^*)^{S(i)} = \bigotimes_{i \in [m]} (X_i)^{S(i)}. \tag{40}$$

From this it follows that $F(X_1, \dots, X_m)$ is self-adjoint since each $\hat{f}(S)$ is a real number. Next we want to show that $F(X_1, \dots, X_m)^2 = I$. First note that for $S, T \subseteq [m]$, their symmetric difference $U = S \Delta T$ and their intersection $V = S \cap T$, we have

$$\left(\bigotimes_{i \in [m]} X_i^{S(i)} \right) \left(\bigotimes_{i \in [m]} X_i^{T(i)} \right) = \left(\bigotimes_{i \in [m]} X_i^{U(i)} \right) \left(\bigotimes_{i \in [m]} (X_i^2)^{V(i)} \right) = \left(\bigotimes_{i \in [m]} X_i^{U(i)} \right), \tag{41}$$

where the last equality follows from the fact that $X_i^2 = I$ for all $i \in [m]$. Now we can expand $F(X_1, \dots, X_m)^2$ as follows

$$F(X_1, \dots, X_m)^2 = \sum_{S \subseteq [m]} \sum_{T \subseteq [m]} \hat{f}(S) \hat{f}(T) \left(\bigotimes_{i \in [m]} X_i^{S(i)} \right) \left(\bigotimes_{i \in [m]} X_i^{T(i)} \right) \tag{42}$$

$$= \sum_{S \subseteq [m]} \sum_{U \subseteq [m]} \hat{f}(S) \hat{f}(S \Delta U) \left(\bigotimes_{i \in [m]} X_i^{U(i)} \right) \tag{43}$$

$$= \sum_{U \subseteq [m]} \sum_{S \subseteq [m]} \hat{f}(S) \hat{f}(S \Delta U) \left(\bigotimes_{i \in [m]} X_i^{U(i)} \right) \tag{44}$$

$$= \sum_{U \subseteq [m]} \left(\bigotimes_{i \in [m]} X_i^{U(i)} \right) \sum_{S \subseteq [m]} \hat{f}(S) \hat{f}(S \Delta U). \tag{45}$$

By the Convolution Formula (3) we have

$$\sum_{S \subseteq [m]} \hat{f}(S) \hat{f}(S \Delta U) = \widehat{f^2}(U). \tag{46}$$

Since the range of f is $\{\pm 1\}$, the function f^2 is identically 1, from which it follows that

$$\widehat{f^2}(U) = \begin{cases} 1 & \text{if } U = \emptyset \\ 0 & \text{if } U \neq \emptyset \end{cases} \tag{47}$$

by the uniqueness of the Fourier transform. Back into (45), this gives

$$F(X_1, \dots, X_m)^2 = \left(\bigotimes_{i \in [m]} X_i^{\emptyset(i)} \right) = I \tag{48}$$

as was to be proved. Finally, if $S \subseteq [m]$ and $(X_1, Y_1), \dots, (X_m, Y_m)$ are such that X_i and Y_i commute for every $i \in [m]$, then

$$\left(\bigotimes_{i \in [m]} X_i^{S(i)} \right) \left(\bigotimes_{i \in [m]} Y_i^{S(i)} \right) = \bigotimes_{i \in [m]} (X_i Y_i)^{S(i)} = \tag{49}$$

$$= \bigotimes_{i \in [m]} (Y_i X_i)^{S(i)} = \left(\bigotimes_{i \in [m]} Y_i^{S(i)} \right) \left(\bigotimes_{i \in [m]} X_i^{S(i)} \right). \tag{50}$$

It follows that $F(X_1, \dots, X_m)$ and $F(Y_1, \dots, Y_m)$ commute. This completes the proof that F is an operation.

Next we show that for every relation R in A , the operator assignment relation R^* is invariant under F . Let r be the arity of R and let $P_R(X_1, \dots, X_r)$ be the characteristic polynomial of R . Let $A_{i,1}, \dots, A_{i,r}$ be an r -variable operator assignment over a finite-dimensional Hilbert space \mathcal{H}_i for $i = 1, \dots, m$. We may assume that $\mathcal{H}_i = \mathbb{C}^{d_i}$ where d_i is the dimension of \mathcal{H}_i . From now on we switch to the language of matrices.

Assume that all the assignments $(A_{1,1}, \dots, A_{1,r}), \dots, (A_{m,1}, \dots, A_{m,r})$ are in R^* . In particular, each sequence $A_{i,1}, \dots, A_{i,r}$ is a fully commuting assignment of Hermitian matrices and $P_R(A_{i,1}, \dots, A_{i,r}) = -I$. The Strong Spectral Theorem (i.e. Theorem 1) applies, so $A_{i,1}, \dots, A_{i,r}$ simultaneously diagonalize. Let U_i be a unitary matrix of \mathcal{H}_i that achieves that, and let $D_{i,j} = U_i A_{i,j} U_i^*$ for $j \in [r]$ be the resulting diagonal matrices. From $A_{i,j}^2 = I$ and $U_i^* U_i = U_i U_i^* = I$ we conclude that $D_{i,j}^2 = I$ and hence each entry in the diagonal of $D_{i,j}$ is $+1$ or -1 . For $c \in [d_i]$, let $D_{i,j}(c)$ denote the entry in position c of the diagonal of $D_{i,j}$. The hypotheses of Lemma 1 apply to the pairs $(A_{i,1}, D_{i,1}), \dots, (A_{i,r}, D_{i,r})$, so $P_R(A_{i,1}, \dots, A_{i,r})$ and $P_R(D_{i,1}, \dots, D_{i,r})$ are similar matrices. As $P_R(A_{i,1}, \dots, A_{i,r}) = -I$, and the only matrix that is similar to $-I$ is $-I$ itself, we get $P_R(D_{i,1}, \dots, D_{i,r}) = -I$. In particular

$$P_R(D_{i,1}(c), \dots, D_{i,r}(c)) = -1 \tag{51}$$

for every $c \in [d_i]$. This will be of use later.

Our next goal is to show that $P_R(F(A_{1,1}, \dots, A_{m,1}), \dots, F(A_{1,r}, \dots, A_{m,r})) = -I$ and we do so by showing that

$$\sum_{T \subseteq [r]} \hat{R}(T) \prod_{j \in T} F(A_{1,j}, \dots, A_{m,j}) = -I. \tag{52}$$

For fixed $T \subseteq [r]$, let $A_T = \prod_{j \in T} F(A_{1,j}, \dots, A_{m,j})$ be the matrix product appearing in the left-hand side of (52). Let $t = |T|$. By first expanding on the definition of F and then distributing the product over the sum we get

$$A_T = \prod_{j \in T} \sum_{S \subseteq [m]} \hat{f}(S) \bigotimes_{i \in [m]} (A_{i,j})^{S(i)} = \sum_{S: T \rightarrow 2^{[m]}} \prod_{j \in T} \hat{f}(S(j)) \prod_{i \in [m]} \bigotimes_{j \in T} (A_{i,j})^{S(i)}. \tag{53}$$

For fixed $T \subseteq [r]$ and $S : T \rightarrow 2^{[m]}$, let $B_{T,S} = \prod_j \bigotimes_i (A_{i,j})^{S(i)}$ be the matrix product appearing in the right-hand side of (53). By distributing \prod over \bigotimes and applying $A_{i,j} = U_i^* D_{i,j} U_i$ in (53) we get

$$B_{T,S} = \bigotimes_{i \in [m]} \left(\prod_{j \in T} (U_i^* D_{i,j} U_i) \right)^{S(i)} = \bigotimes_{i \in [m]} \left(U_i^* \left(\prod_{j \in T} D_{i,j} \right)^{S(i)} U_i \right). \tag{54}$$

Hence

$$A_T = U^* \left(\sum_{S: T \rightarrow 2^{[m]}} \prod_{j \in T} \hat{f}(S(j)) \bigotimes_{i \in [m]} \left(\prod_{j \in T} D_{i,j} \right)^{S(i)} \right) U, \tag{55}$$

for $U = \bigotimes_{i \in [m]} U_i$. Let M denote the matrix sitting within U^* and U in line (55). As each $D_{i,j}$ is a $d_i \times d_i$ diagonal matrix, M is a $d \times d$ diagonal matrix with $d = \prod_{i \in [m]} d_i$. We think of the entries in the diagonal of M as indexed by tuples $c = (c_1, \dots, c_m)$ from $[d_1] \times \dots \times [d_m]$. Let $M(c)$ denote the entry in position c of the diagonal of M . Then

$$M(c) = \sum_{S: T \rightarrow 2^{[m]}} \prod_{j \in T} \hat{f}(S(j)) \prod_{i \in [m]} \prod_{j \in T} (D_{i,j}(c_i))^{S(i)}. \tag{56}$$

Factoring the product over $j \in T$ back, the right-hand side in (56) reads

$$\prod_{j \in T} \sum_{S \subseteq [m]} \hat{f}(S) \prod_{i \in [m]} (D_{i,j}(c_i))^{S(i)} = \prod_{j \in T} f(D_{1,j}(c_1), \dots, D_{m,j}(c_m)). \tag{57}$$

For fixed $j \in [r]$ and $c \in [d_1] \times \dots \times [d_m]$, let $X_{j,c} = f(D_{1,j}(c_1), \dots, D_{m,j}(c_m))$ so that equations (56) and (57) give $M(c) = \prod_{j \in T} X_{j,c}$. From (51) and the fact that f is a Boolean closure operator of R , the tuple $(X_{1,c}, \dots, X_{r,c})$ belongs to the relation R . Thus

$$\sum_{T \subseteq [r]} \hat{R}(T) M(c) = \sum_{T \subseteq [r]} \hat{R}(T) \prod_{j \in T} X_{j,c} = P_R(X_{1,c}, \dots, X_{r,c}) = -1. \tag{58}$$

Since this holds for every diagonal entry of M , we get $\sum_{T \subseteq [r]} \hat{R}(T) M = -I$. Putting it all together, the left-hand side of our goal (52) evaluates to

$$\sum_{T \subseteq [r]} \hat{R}(T) U^* M U = U^* \left(\sum_{T \subseteq [r]} \hat{R}(T) M \right) U = U^* (-I) U = -I. \tag{59}$$

This gives (52) as desired.

In order to prove the ‘moreover’ clause of the theorem, observe that if $S \subseteq [m]$ and $(a_1, \dots, a_m) \in \{\pm 1\}^m$, then

$$\bigotimes_{i \in [m]} (a_i I)^{S(i)} = \left(\prod_{i \in S} a_i \right) I, \tag{60}$$

where in the left hand side the identity matrices have dimensions d_1, \dots, d_m , respectively, and in the right-hand side the identity matrix has dimension $d \times d$ for $d = \prod_{i \in [m]} d_i$. It follows that

$$F(a_1 I, \dots, a_m I) = \left(\sum_{S \subseteq [m]} \hat{f}(S) \prod_{i \in S} a_i \right) I = f(a_1, \dots, a_m) I. \tag{61}$$

This completes the proof of the theorem. \square

7.4. Finale

Before we prove Theorem 12, we need the following straightforward fact about the role of constants in pp-definitions.

Lemma 13. *Let A be a Boolean constraint language, let R be Boolean a relation, and let $A^+ = A \cup \{+1, -1\}$. The following two statements hold.*

1. R is pp-definable from A if and only if it is pp-definable without constants from A^+ .

2. R is pp^* -definable from A if and only if it is pp^* -definable without constants from A^+ .

Proof. In both cases, for the ‘only if’ part it suffices to replace each occurrence of a constant in the quantifier-free part of the pp -formula by a new existentially quantified variable Z , and force it to belong to the corresponding new unary relation in A^+ by an additional conjunct: if Z replaces the constant -1 , we force Z to belong $\{-1\}$ by a new conjunct, and if Z replaces the constant $+1$, we force it to belong $\{+1\}$ by a new conjunct. In both cases too, the ‘if’ part follows from the reverse construction: replace each occurrence of a variable that appears within the scope of one of the new unary relations in A^+ by the corresponding constant, and remove the conjuncts that involve the new unary relations. That these transformations are correct follows directly from the definitions and the fact that both I and $-I$ commute with any operator. \square

We are ready to prove Theorem 12.

Proof of Theorem 12. Assume R is pp^* -definable from A . By Lemma 13, the relation R is also pp^* -definable without constants from $A^+ = A \cup \{\{+1\}, \{-1\}\}$. Let r be the arity of R and let $\phi(x_1, \dots, x_r)$ be the pp -formula without constants that pp^* -defines R from A^+ . By Geiger’s Theorem 4 and Lemma 13 it suffices to show that R is invariant under all Boolean closure operations of A^+ .

Let $f : \{\pm 1\}^m \rightarrow \{\pm 1\}$ be a Boolean closure operation of A^+ . By Theorem 14, the function F is a closure operation of A^{+*} . Let $(a_{1,1}, \dots, a_{1,r}), \dots, (a_{m,1}, \dots, a_{m,r})$ be tuples in R and let $a_j = f(a_{1,j}, \dots, a_{m,j})$ for every $j \in [m]$. We need to show that (a_1, \dots, a_r) is also in R . Let $\psi(x_1, \dots, x_r, y_1, \dots, y_s)$ be the quantifier-free part of ϕ and consider the instance over A^+ that is given by

$$\psi(x_1/a_{i,1}, \dots, x_r/a_{i,r}, y_1/Y_1, \dots, y_s/Y_s) \quad (62)$$

as described in the beginning of this section. Since the tuple $(a_{i,1}, \dots, a_{i,r})$ is in R and ϕ pp^* -defines R , the instance in (62) is satisfiable via operator assignments over a finite-dimensional Hilbert space for every $i \in [m]$. Let $B_{i,1}, \dots, B_{i,s}$ be such a satisfying operator assignment for every $i \in [m]$. Since I and $-I$ commute with any operator, this means that $a_{i,1}I, \dots, a_{i,r}I, B_{i,1}, \dots, B_{i,s}$ is a satisfying operator assignment of

$$\psi(x_1/X_1, \dots, x_r/X_r, y_1/Y_1, \dots, y_s/Y_s) \quad (63)$$

for every $i \in [m]$. Let $A_j = F(a_{1,j}I, \dots, a_{m,j}I)$ and $B_j = F(B_{1,j}, \dots, B_{m,j})$. As F is a closure operation of A^{+*} , the tuple $A_1, \dots, A_r, B_1, \dots, B_s$ is a satisfying operator assignment for (63). Moreover, from the ‘moreover’ clause in Theorem 14 we know that $A_j = f(a_{1,j}, \dots, a_{m,j})I = a_jI$ for every $j \in [m]$. Thus, the instance

$$\psi(x_1/a_1, \dots, x_r/a_r, y_1/Y_1, \dots, y_s/Y_s) \quad (64)$$

is satisfiable via operator assignments over a finite-dimensional Hilbert space; the finite-dimensional operator assignment B_1, \dots, B_s satisfies it. As ϕ pp^* -defines R , it follows that (a_1, \dots, a_r) is in R , as was to be shown. \square

8. Concluding remarks

Motivated by considerations in the foundations of quantum mechanics, we investigated a relaxation of Boolean satisfiability in which the variables are assigned linear operators on some Hilbert space as values. Among other results, we characterized completely the classes of Boolean constraint languages in Schaefer’s framework for which there are instances that are unsatisfiable in the Boolean domain, but become satisfiable under this relaxation. Moreover, we explored the difference between satisfiability via linear operators on some Hilbert space of finite dimension and satisfiability via linear operators on some Hilbert space of arbitrary dimension. The main tool used was a careful adaptation of the notion of pp -definability, which gave rise to reductions that preserve gaps between satisfiability in the Boolean domain and satisfiability via linear operators on some Hilbert space.

Several technical questions remain unanswered, including the following two.

In Section 4.5, we discussed the role of the full binary relation T and how its use can be avoided when the Boolean constraint language at hand admits a commutativity gadget. If it were the case that every Boolean constraint language admits a commutativity gadget, then all the results of this paper that require the inclusion of T to the constraint language would hold also without this requirement. Is it true that every Boolean constraint language admits a commutativity gadget?

In Section 7 we proved Theorem 12 which asserts that, as regards Boolean constraint languages and Boolean relations, there is no difference between pp -definability and pp^* -definability, where in pp^* -definability the existential quantifiers range over operators on a finite-dimensional Hilbert space. Is there a difference between pp -definability and pp^{**} -definability, where in pp^{**} -definability the existential quantifiers range over operators on a Hilbert space of arbitrary (finite or infinite) dimension?

More broadly, the work reported here paves the way for the investigation of general constraint satisfaction problems under this relaxation of satisfiability. In particular, do the concepts and the results (for example, Theorem 3) about Boolean constraint satisfaction problems obtained here extend to q -valued constraint satisfaction problems, for $q \geq 3$?

Declaration of Competing Interest

There is no conflict of interest.

Acknowledgments

We are grateful to Heribert Vollmer for sharing with us Steffen Reith's diagram of Post's lattice (Fig. 1). This work was initiated and part of the research was carried out while all three authors were in residence at the Simons Institute for the Theory of Computing during the fall of 2016, where they participated in the program on Logical Structures in Computation. The research of Albert Aterias was partially funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme, grant agreement ERC-2014-CoG 648276 (AUTAR), and by MINECO through TIN2013-48031-C4-1-P (TASSAT2); the research of Simone Severini was partially funded by The Royal Society, Engineering and Physical Sciences Research Council (EPSRC), and the National Natural Science Foundation of China (NSFC).

References

- [1] Antonio Acín, Tobias Fritz, Anthony Leverrier, Ana Belén Sainz, A combinatorial approach to nonlocality and contextuality, *Commun. Math. Phys.* 2 (334) (2015) 533–628.
- [2] John S. Bell, On the problem of hidden variables in quantum mechanics, *Rev. Mod. Phys.* 38 (3) (1966) 447.
- [3] V.G. Bodnarchuk, L.A. Kaluzhnin, V.N. Kotov, B.A. Romov, Galois theory for Post algebras, I, *Cybernetics* 5 (3) (1969) 243–252.
- [4] Elmar Böhler, Nadia Creignou, Steffen Reith, Heribert Vollmer, Playing with Boolean blocks, part I: post's lattice with applications to complexity theory, in: *ACM SIGACT-Newsletter*, 2003.
- [5] Elmar Böhler, Nadia Creignou, Steffen Reith, Heribert Vollmer, Playing with Boolean blocks, part II: constraint satisfaction problems, in: *ACM SIGACT-Newsletter*, 2004.
- [6] Richard Cleve, Li Liu, William Slofstra, Perfect commuting-operator strategies for linear system games, preprint, arXiv:1606.02278, 2016.
- [7] Richard Cleve, Rajat Mittal, Characterization of binary constraint system games, in: *Automata, Languages, and Programming - 41st International Colloquium, Proceedings, Part I, ICALP 2014, Copenhagen, Denmark, July 8–11, 2014, 2014*, pp. 320–331.
- [8] Brian A. Davey, Hilary A. Priestley, *Introduction to Lattices and Order*, Cambridge University Press, 2002.
- [9] Tomás Feder, Moshe Y. Vardi, The computational structure of monotone monadic SNP and constraint satisfaction: a study through Datalog and group theory, *SIAM J. Comput.* 28 (1) (1998) 57–104.
- [10] Gerald B. Folland, *A Course in Abstract Harmonic Analysis, Studies in Advanced Mathematics*, Taylor & Francis, 1994.
- [11] Gerald B. Folland, *Real Analysis: Modern Techniques and Their Applications, Pure and Applied Mathematics: A Wiley Series of Texts, Monographs and Tracts.*, Wiley, 2013.
- [12] Tobias Fritz, Quantum logic is undecidable, preprint, arXiv:1607.05870, 2016.
- [13] David Geiger, Closed systems of functions and predicates, *Pac. J. Math.* 27 (1) (1968) 95–100.
- [14] Paul R. Halmos, *Introduction to Hilbert Space and the Theory of Spectral Multiplicity, Benediction Classics*, 2016.
- [15] Zhengfeng Ji, Binary constraint system games and locally commutative reductions, preprint, arXiv:1310.3794, 2013.
- [16] N. David Mermin, Simple unified form for the major no-hidden-variables theorems, *Phys. Rev. Lett.* 65 (27) (1990) 3373.
- [17] N. David Mermin, Hidden variables and the two theorems of John Bell, *Rev. Mod. Phys.* 65 (3) (1993) 803.
- [18] Ryan O'Donnell, *Analysis of Boolean Functions*, Cambridge University Press, 2014.
- [19] Christos H. Papadimitriou, *Computational Complexity*, Addison Wesley, 1994.
- [20] Asher Peres, Incompatible results of quantum measurements, *Phys. Lett. A* 151 (3–4) (1990) 107–108.
- [21] Emil L. Post, *The Two-Valued Iterative Systems of Mathematical Logic, Annals of Mathematical Studies*, vol. 5, Princeton University Press, 1941.
- [22] Thomas J. Schaefer, The complexity of satisfiability problems, in: *Proceedings of the 10th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, May 1–3, 1978, 1978*, pp. 216–226.
- [23] Uwe Schöning, *Logic for Computer Scientists*, Springer Science & Business Media, 2008.
- [24] William Slofstra, Tsirelson's problem and an embedding theorem for groups arising from non-local games, arXiv:1606.03140, 2016.