


A Quantum Pigeonhole Principle and Two Semidefinite Relaxations of Communication Complexity

Pavel Dvořák 

Charles University, Prague, Czech Republic

Bruno Loff 

LASIGE, Faculty of Sciences of the University of Lisbon, Portugal

Suhail Sherif 

LASIGE, Faculty of Sciences of the University of Lisbon, Portugal

Abstract

We are interested in what happens when we take a Π_1 combinatorial statement, write its negation as a homogeneous quadratic feasibility problem (HQFP), and relax the problem into a positive semidefinite feasibility problem. This question is particularly interesting owing to the fact that any statement written as a PSD feasibility problem can be proven or disproven using a short proof. We investigate this for one very simple and one very complicated statement.

The simple statement we look at is the pigeonhole principle. We prove that the relaxed negation of the PHP remains unsatisfiable and we thus obtain a new “quantum” pigeonhole principle (QPHP) which is a stronger statement than the vanilla PHP. It states that if we take n copies of the same state, and measure each copy using a measurement with only $n - 1$ outcomes (the measurement can be different for different copies), then there will be an outcome j and two copies i_1, i_2 where the resulting states, obtained when the outcome is j for both copies, are not orthogonal.

We then look at the statement “the deterministic communication complexity of f is $\leq k$ ”, where f could be either a function or a relation. We write this statement in two equivalent ways, using two different HQFPs. By relaxing to PSD feasibility, we increase the set of available protocols, and thus we always get a communication model which is stronger than deterministic communication complexity. An argument from proof complexity shows that any model obtained in this way will solve all Karchmer–Wigderson games efficiently. However, the argument is very indirect and does not give us an explicit protocol that solves the Karchmer–Wigderson games. We then work to find such protocols in the two communication models obtained by relaxing our two formulations.

When relaxing the first of the two formulations we obtain a *structured* variant of the γ_2 norm. This communication model is to subunit γ_2 norm matrices like deterministic protocols are to rectangles, and so we call the protocols in this model γ_2 protocols. We show that log-inverse-discrepancy is a lower-bound for this model. We then show how to compute equality (deterministically) using $O(1)$ bits of γ_2 -communication, which implies that KW games are easy in the model.

When relaxing the second of the two formulations we obtain what we call *quantum lab protocols*. This model happens to have a functional description, wherein Alice and Bob communicate solely via the outcomes of binary measurements of a shared quantum state (whose initial state is independent of the inputs). They are required to give the correct output with zero error probability. We use our QPHP to prove a lower-bound of n against two-round quantum lab protocols for equality. However we also show that *any* Boolean function f can be computed in three rounds and four measurements.

2012 ACM Subject Classification Theory of computation \rightarrow Proof complexity; Mathematics of computing \rightarrow Semidefinite programming; Theory of computation \rightarrow Communication complexity; Theory of computation \rightarrow Quantum communication complexity

Keywords and phrases Proofs, Semidefinite Programs, Quantum Pigeonhole Principle, Communication Complexity

Digital Object Identifier 10.4230/LIPIcs.STACS.2026.35



© Pavel Dvořák, Bruno Loff, and Suhail Sherif;
licensed under Creative Commons License CC-BY 4.0

43rd International Symposium on Theoretical Aspects of Computer Science (STACS 2026).

Editors: Meena Mahajan, Florin Manea, Annabelle McIver, and Nguyễn Kim Thăng

Article No. 35; pp. 35:1–35:20



Leibniz International Proceedings in Informatics

LIPIcs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Related Version This paper omits some proofs because of length constraints.

Full Version: <https://arxiv.org/abs/2409.04592> [11]

Funding This work was funded by the European Union (ERC, HOFGA, 101041696). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

It was also supported by FCT through the LASIGE Research Unit, ref. UIDB/00408/2020 and ref. UIDP/00408/2020, and by CMAFcIO, FCT Project UIDB/04561/2020, <https://doi.org/10.54499/UIDB/04561/2020>.

P. Dvořák was supported by Czech Science Foundation GAČR grant #22-14872O.

Acknowledgements The authors would like to thank Carlos Florentino for fun conversations around this topic.

1 Introduction

The good thing about existentially-quantified (Σ_1) statements is that proving them amounts to finding a witness, after which the proof is a routine verification. But – if we assume that $\text{NP} \neq \text{coNP}$ – there will necessarily exist universally-quantified (Π_1) statements which cannot be proven in this way. Simultaneously, there exists a small number of situations when a particular class of Σ_1 statements is known to be closed under negation, meaning, every statement in this class can be either proven or disproven by finding an explicit, easy-to-verify witness. Of course, this includes all “easy” statements (decidable in P), but beyond that the exhaustive list is quite short. To our knowledge, the following list includes all problems that are known to be in $\text{NP} \cap \text{coNP}$,¹ but not known to be in P: arithmetical statements around factoring and discrete log, conic feasibility [24], which includes semidefinite feasibility, (approximate) lattice problems [1], stochastic games (see, e.g. [28]), the unknotting problem [17], and the arrival problem [10]. In this paper, we focus on semidefinite feasibility problems (SDFPs), which are a particular kind of conic feasibility.

Our original interest in looking at SDFPs was: **to study semidefinite relaxations of communication complexity**. Why should one care to do this? Well, the existence of a protocol for solving a given function is our Σ_1 statement, and by duality of SDFPs, every lower-bound for any such model is explicitly witnessed by a small object, which is desirable if one wishes to prove lower-bounds. A previous result by Karchmer, Kushilevitz and Nisan [15] attempted to do this with linear programming – and showed that the natural way of doing it gives models that are too powerful, they can solve all Karchmer–Wigderson games in $O(\log n)$ communication. To attempt to do this using semidefinite programming is a *natural, obvious* thing to try.² We show, however, that the conclusion is essentially the same.

¹ More precisely, conic feasibility is known to be in $\text{NP}(\mathbb{R}) \cap \text{coNP}(\mathbb{R})$, as there are issues with the bitlength of solutions, which appear unavoidable. For example, one can construct a semidefinite feasibility problem (A, b) , with polynomially-many bits of precision, which is satisfiable, but any solution x must be specified with exponentially-many bits of precision. This is a classical example by Khachiyan, see [23].

² It may seem like the linear relaxations of communication protocols being too powerful would imply that the semidefinite relaxations are too. However, the semidefinite constraints used to define what a valid relaxation is are also stronger than the linear constraints, and hence the linear relaxations can’t be turned into valid semidefinite relaxations.

This paper consists of the following.

- We begin by studying a natural semidefinite relaxation of the pigeonhole principle, which we call the “Quantum Pigeonhole Principle” (QPHP), a statement which is strictly stronger than the classical PHP, and which we consider to be an interesting linear algebraic theorem in its own right. We show several proofs of the QPHP, from a loose, easy-to-prove bound, to the best possible quantitative statement.
- We then study two different computational models, both obtained as semidefinite relaxations of communication complexity:
 - One such model is a structured generalization of the γ_2 norm. It relates with γ_2 norm in an analogous way to how communication protocols are related with rectangles. We prove a simple lower-bound for such a model, based on the γ_2 norm, and we prove that, in this model, equality can be solved in $O(\log n)$ bits of communication.
 - We call the other model “Quantum lab protocols”. We prove that such protocols can compute every function in 3 rounds with 4 bits of communication. As an application of the QPHP we show that, in 2 rounds, such protocols cannot compute equality with less than $n + 1$ bits of communication.
- Finally, we show that these results are not an artifact of the particular semidefinite program which we used to model communication complexity, that any program which attempts to do this, and obeys some mild assumptions, must necessarily fail. We call this a “no-go theorem”. This theorem also applies to linear-programming relaxations, since they are special case of SDPs, and even in this restricted setting our no-go theorem is very different to the “complexity is not convex” results of [15] and the follow-up paper [12] (see [14], Section 6.9).

We will now overview our reasoning in greater detail. Supplementary definitions on various topics touched on in the introduction can be found in Section 6.

Homogeneous Quadratic and Semidefinite Feasibility Problems (HQFP and SDFP)

In a linear feasibility problem, we are given a linear map $\mathcal{A} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and a vector $b \in \mathbb{R}^m$, and we wish to know if there exists $x \in \mathbb{R}_{\geq 0}^n$ such that $\mathcal{A}(x) = b$. As it turns out, many (but not all) of the properties of linear programming generalize to the case where the non-negative orthant $\mathbb{R}_{\geq 0}^n$ is replaced by a closed, convex cone \mathcal{K} , namely, a subset of \mathbb{R}^n closed under limits, sums and multiplication by non-negative scalars.

In a semidefinite feasibility problem, we are given a linear map $\mathcal{A} : \mathbb{R}^{\frac{n(n+1)}{2}} \rightarrow \mathbb{R}^m$ from the set of all symmetric matrices to \mathbb{R}^m , and a vector $b \in \mathbb{R}^m$, and we wish to know if there exists a positive semidefinite matrix M such that $\mathcal{A}(M) = b$. I.e., we replace the non-negative orthant $\mathbb{R}_{\geq 0}^n$ with the cone of *positive semidefinite* $n \times n$ matrices $\text{PSD}_n \subseteq \mathbb{R}^{\frac{n(n+1)}{2}}$ (such matrices are symmetric). This set can be alternatively characterized as the set of symmetric matrices with non-negative eigenvalues, or as the set of *Gram* matrices, i.e., matrices equal to $A^t A$ for some $n \times m$ matrix A , or in other words, matrices M of inner products, given by a family of vectors a_1, \dots, a_n (the columns of A), so that $M_{ij} = \langle a_i | a_j \rangle$.

It follows that a SDFP is asking whether there exist vectors a_1, \dots, a_n obeying a given system of linear equations on their inner products $\langle a_i | a_j \rangle$. (With linear programming being the special case where the linear equations only depend on the diagonal entries of M .) One can easily show that the dimension m can be made to be $\leq n$. Hence, the matrix A serves as a short, easy-to-verify witness that a given SDFP is feasible.

Now, suppose we further restrict the solution M to have rank 1, i.e., the vectors a_i and a_j are now scalars. We then obtain a system of linear equations on degree-2 products $a_i \cdot a_j$, and we wish to know if some choice of scalars satisfies these equations. This is a different kind of problem, called a Homogeneous Quadratic Feasibility Problem (HQFP), and it is easily shown to be NP-hard.

Being Relaxed about the Truth Helps in Finding Short Proofs

It then follows that it is possible to take any Σ_1 combinatorial statement Ψ , write it down as a HQFP Q , and then relax it by dropping the rank-1 restriction, to obtain a SDFP P .

A radical transformation always happens in this process. The statement “ Q is feasible” is equivalent to Ψ , and by relaxation it always implies “ P is feasible”. However, there is a fundamental result of Ramana [24] saying that given any SDFP P we can efficiently construct a different “dual” SDFP P' , such that “ P is *not* feasible” if and only if “ P' is feasible”. Hence, if P is not feasible, we can always prove that P is not feasible by presenting a short, easy witness – the witness that P' is feasible. So if Ψ is true, “ P is feasible” remains true, and if Ψ is false, then either “ P is feasible” becomes true (we relaxed too much), or “ P is feasible” is also false. In the latter case, there exists a short, easy witness that proves “ P is not feasible”, and hence also proves that Ψ is false. In other words, the relaxation map sends instances of an NP-complete problem to instances of a problem in $\text{NP} \cap \text{coNP}$. Understandably, then, not all false Σ_1 statements Ψ will remain false after relaxation, but when they do, we are guaranteed to have short proofs of falsity.

Now suppose there exists a particular Π_1 statement Ψ we wish to prove. Maybe it is a tautological combinatorial principle, or even a complexity lower-bound. We then write $\neg\Psi$ as a HQFP Q and relax it into the SDFP P and try to prove that P is false by constructing a solution for P' . If we succeed, it then follows that $\neg\Psi$ is false, i.e., Ψ is true, and this is witnessed by a short, easy-to-verify object. Or maybe, encouraged by the guaranteed existence of a short proof of P' , we may try to prove that P is false in another way, without necessarily aiming for a “canonical” proof.

In this paper, we report on what happens when we carry out the above approach, for two different Π_1 statements: the pigeonhole principle, and communication complexity lower-bounds. The whole approach can be seen as trying to express Π_1 statements in a very simple proof system, and we will have more to say below on the connection with proof complexity.

2 A Quantum Pigeonhole Principle

In proof complexity, more specifically in a proof system called Polynomial Calculus, the negation of the pigeonhole principle is sometimes formalized as the following quadratic feasibility problem:

$$\begin{array}{ll}
 \text{There exist } \lambda \in \mathbb{R} & \\
 v_{i,j} \in \mathbb{R} & \forall i \in [p], j \in [h] \\
 \text{such that} & \\
 \lambda^2 = 1 & \\
 \sum_{j=1}^h v_{i,j} = \lambda & \forall i \in [p] \\
 v_{i,j} \cdot v_{i,j'} = 0 & \forall i \forall j \neq j' \\
 v_{i,j} \cdot v_{i',j} = 0 & \forall j \forall i \neq i'
 \end{array}$$

This system is not homogeneous, so it is not immediate how to express it as a SDFP. Nonetheless, we can attempt to naively relax this program to higher dimensions, by replacing real numbers with vectors, and products with inner products. This gives us exactly the negation of the QPHP (Theorem 1):

There exists a vector space V

and vectors $\lambda \in V$

$$v_{i,j} \in V \quad \forall i \in [p], j \in [h]$$

such that

$$\|\lambda\|^2 = 1$$

$$\sum_{j=1}^h v_{i,j} = \lambda \quad \forall i \in [p] \quad (1)$$

$$\langle v_{i,j} \mid v_{i,j'} \rangle = 0 \quad \forall i \forall j \neq j' \quad (2)$$

$$\langle v_{i,j} \mid v_{i',j} \rangle = 0 \quad \forall j \forall i \neq i' \quad (3)$$

Again it is not immediate that this is a SDFP, since (1) is not directly an equation about inner-products. However, we can replace (1) with:

$$\sum_{j=1}^h \|v_{i,j}\|^2 = \|\lambda\|^2 \quad \forall i \in [p] \quad (1a)$$

$$\sum_{j=1}^h \langle v_{i,j} \mid \lambda \rangle = \|\lambda\|^2 \quad \forall i \in [p] \quad (1b)$$

To see the equivalence, notice that for each fixed $i \in [p]$, (2) states that the $v_{i,j}$ are orthogonal. Under such orthogonality, it is obvious that 1 implies 1a and 1b, by Pythagoras' Theorem. Conversely, let $\lambda'_i = \sum_j v_{i,j}$. Then 1b states that $\langle \lambda'_i, \lambda \rangle = \|\lambda\|^2$ and, under orthogonality, Pythagoras' Theorem says $\|\lambda'_i\|^2 = \sum_j \|v_{i,j}\|^2$, and so (3) is saying that $\|\lambda'_i\|^2 = \|\lambda\|^2$. These two together imply, by the equality case of Cauchy-Schwarz, that $\lambda'_i = \lambda$.

It is now clear that we have a semidefinite feasibility problem. It can also be seen that taking constraints (1)-(4), and further restricting $\lambda, v_{i,j}$ to have dimension 1, gives us a HQFP, which is equivalent to the negation of the PHP.

It is also possible to state the QPHP using only quantum language, as follows. Suppose that we have p quantum registers $1, \dots, p$, which are all initialized in the same state: $|\psi_1\rangle = \dots = |\psi_p\rangle$. We then apply an h -outcome measurement to each of the registers. The specific measurement which we make may be different for different registers. Regardless, the measurements cause the registers to collapse to possibly-different states $|\phi_1\rangle, \dots, |\phi_p\rangle$. The QPHP states that, if $h < p$, there will always exist an outcome j and two registers $i \neq i'$, such that there is a non-zero probability of obtaining the same outcome j after measuring both registers i and i' , and when this happens the resulting states $|\phi_i\rangle$ and $|\phi_{i'}\rangle$ are not orthogonal.

We show that the Quantum Pigeonhole Principle is true.

► **Theorem 1 (QPHP).** *Let $\{\lambda\} \cup \{v_{i,j} \mid i \in [p], j \in [h]\} \subseteq \mathcal{H}$ be a family of vectors in a Hilbert space \mathcal{H} , such that*

$$\begin{aligned} \|\lambda\|^2 &= 1 \\ \sum_{j=1}^h v_{i,j} &= \lambda & \forall i \in [p] \\ \langle v_{i,j}, v_{i,j'} \rangle &= 0 & \forall i \forall j \neq j' \end{aligned}$$

I.e., each family $V_i = \{v_{i,j} \mid j \in [h]\}$ decomposes the same unit vector λ as a sum of h -many orthogonal vectors (We have p copies of λ – the “pigeons” – and divide each pigeon among h “holes”). Suppose that $h < p$. Then, there exists $j \in [h]$ and $i \neq i'$ in $[p]$, such that

$$\langle v_{i,j}, v_{i',j} \rangle \neq 0$$

(one of the holes must have more than one pigeon).

In the full version of our paper [11] we show proofs of three versions of the QPHP:

- Using the AM-GM inequality we give a short and simple proof of a weaker version of the QPHP, only showing non-orthogonality if the number of holes h is significantly smaller than the number p of pigeons, namely $h < \frac{1}{4}\sqrt{p}$.
- We also prove a quantitatively stronger version of the QPHP, which allows for the initial states λ_i to be different for different pigeons, and gives a tight lower-bound on the maximal overlap $\langle v_{i,j}, v_{i',j} \rangle$, as a function of the average initial overlap $\beta := \frac{1}{p(p-1)} \sum_{i \neq i'} \langle \lambda_i \mid \lambda_{i'} \rangle$. Namely we get that there exists $j \in [h]$ and $i \neq i' \in [p]$ such that

$$\langle v_{i,j}, v_{i',j} \rangle \geq \frac{1}{h^2} \left(\beta - \frac{h-1}{p-1} \right)$$

- Finally we also prove the QPHP via one of the short “canonical” proofs which are guaranteed to exist via duality. Namely, we derive a feasibility problem dual to the relaxed negation of the QPHP, and give an explicit solution for it.

Related work by Aharonov, Colombo, Popescu, Sabadini, Struppa and Tollaksen. While reading our paper, a colleague pointed us to the paper “The quantum pigeonhole principle and the nature of quantum correlations”, by Aharonov et al. [2]. In this paper, a quantum experiment is described, and an argument is made with the idea that this experiment should be interpreted as saying that quantum mechanics somehow avoids the pigeonhole principle. In the interest of staying on focus, we will not describe here the experiment, or the argument which purports to show that this experiment is evidence of some “quantum violation of the PHP”. Other researchers have pointed out that the argument is flawed in several ways³, see [29, 9].

3 Connection with Natural Proofs and Proof Complexity

Sections 4 and 5 of the paper apply the above approach to statements of the form “the communication complexity of f is $> k$ ”. This is a Π_1 statement when the two-player function (or relation) f is given as a communication matrix. Indeed, the statement “the communication complexity of f is $\leq k$ ” is easily seen to be Σ_1 , by taking an existential quantifier over all protocols of cost at most k .

When starting this project over two years ago, our naive hope was that maybe we could use semidefinite programming to prove some new lower-bounds against Karchmer–Wigderson games. This would follow a long, successful tradition of using convex optimization to prove

³ If you do read the paper by Aharonov et al., note the following. Their main argument is that the state $(|LL\rangle + |RR\rangle)|+i\rangle$, where the first two particles are in the same box, is orthogonal to $|+i\rangle|+i\rangle|+i\rangle$, and by symmetry this holds for any two of the three particles. This they interpret as somehow meaning that the state $|+i\rangle|+i\rangle|+i\rangle$ does not have any two particles on the same box. This proposed interpretation completely breaks down by noticing (as in [29]) that the state $|LLL\rangle + |RRR\rangle$, which following their thinking would represent the state where all three particles are in the same box, is *not* orthogonal to $|+i\rangle|+i\rangle|+i\rangle$.

lower-bounds: approximate and threshold degree [8], the quantum adversary bound [19], and the γ_2 norm [20] are all examples of complexity measures which relax classical measures in one way or another, and which have been used to prove lower-bounds on classical and quantum query complexity, communication complexity, proof complexity, data structures, *etc.*

But also, such attempts have systematically failed against more powerful computational models, such as Boolean circuits and formulas. A famous result by Karchmer, Kushilevitz and Nisan [15] (CCC'92) shows that the smooth partition bound is small for every Karchmer–Wigderson relation.⁴ A smooth partition is a linear-programming relaxation of an integer program defining the partition number, which is the smallest number of monochromatic rectangles needed to partition a communication matrix, itself a relaxation of the number of leaves in a communication protocol. KKN were hoping [15, page 2] that such a linear relaxation would help them prove lower bounds on the communication complexity of Karchmer–Wigderson relations, and hence lower bounds on the depth of Boolean formulas. Sadly, they could only report on a failed attempt. A few years later, Razborov and Rudich presented their natural-proofs barrier [25] (STOC'95), which strongly suggests that no linear programming relaxation, or any other efficiently computable quantity, will be able to approximate the computational complexity of any model which is powerful enough to contain pseudorandom function generators.

One might think that the natural proofs barrier applies here, but one would be subtly mistaken. Indeed, semidefinite feasibility is not known to be in P , and there is significant evidence that it is actually a hard problem [30].⁵ However, semidefinite feasibility *is* in $NP(\mathbb{R}) \cap coNP(\mathbb{R})$, and one can formulate a sufficiently strong cryptographic conjecture, which would imply the existence of a natural proofs barrier that would apply here.⁶ One could argue whether such a strong cryptographic assumption is believable, but such a discussion will soon become irrelevant to our purpose.

Because shortly after we started working on this, Austrin and Risse [5] showed that the sum of squares proof system (SOS) needs degree roughly S to prove, for any given function f , that f needs circuits of size S . Carefully checking their proof, and doing the necessary adaptations, it also follows from their results that SOS needs degree roughly 2^d to prove a depth- d lower-bound on Boolean formulas. And it is possible to formalize the Karchmer–Wigderson theorem in the SOS proof system, and hence it will follow that SOS needs degree roughly 2^d to prove a lower-bound of d on the communication complexity of a

⁴ This result was generalized by Hrubeš et al. [12], to show that any “convex rectangle measure” assigns small complexity to KW relations.

⁵ We are referring to a result by Tarasov and Vyalyi, showing that any algorithm for solving semidefinite feasibility could be used to compare numbers represented by arithmetic circuits. Note that here we do not have a bound on the degree of the circuits, which could then be exponential in the size of the circuit, and efficiently comparing the (possibly doubly-exponentially large) numbers output by such arithmetic circuits is an old, longstanding problem, which includes the infamous sum-of-square-roots problem as a special case, and which may well not be polynomial-time solvable.

⁶ In a follow-up to his and Razborov’s natural-proofs result [26], Rudich extended the natural proofs barrier as follows. Clearly no pseudorandom generator can fool NP , since in order to distinguish a random from a pseudorandom string, one can always guess the preimage. In his work, Rudich considers the possibility that there exist pseudorandom generators that fool $coNP$. In other words, he conjectures that there exist pseudorandom generators such that no family of short, efficiently recognizable $\{0, 1\}^*$ -valued objects serve to witness that a given string is *not* pseudorandom (not even for a non-negligible fraction of all strings). One could extend Rudich’s conjecture from $coNP$ distinguishers to $coNP(\mathbb{R})$ distinguishers: that no family of low-dimensional, efficiently recognizable real-valued objects could serve to witness that a given string is not pseudorandom. Under this generalization of Rudich’s conjecture, it necessarily follows that all attempts at approximating complexity using semidefinite feasibility are doomed to fail, since the real-valued dual witnesses could ultimately be used to witness that a given string is not pseudorandom.

Karchmer–Wigderson relation. However, a satisfying instance of a semidefinite feasibility problem can be verified in the SOS proof system using a degree-2 proof! It must then follow that, if we define a communication model using our approach, i.e., we generalize communication complexity by formalizing the existence of a deterministic protocol using a HQFP, and relaxing it to a SDFP, then either (1) the proof that our communication model is stronger than the usual deterministic protocols cannot be shown by low-degree SOS proofs (“our formalization of communication complexity is weird”), or (2) our generalized communication model can actually solve every single Karchmer–Wigderson game. This follows because our generalized model is such that we always have short, low-degree proofs of any true lower-bound.

The above considerations lead to a *no-go theorem*, which (informally stated) says that, unless a weird “high-degree” ingredient is introduced somewhere in the formalization (of communication complexity as a HQFP), the model obtained by semidefinite relaxation will be too strong, and will solve all Karchmer–Wigderson relations. We found it remarkable that statements in proof complexity about lengths of proofs imply the existence of algorithms for Karchmer–Wigderson relations, in a large class of computational models!

► **Theorem 2** (No-go Theorem, informal). *Let $f : \{0, 1\}^n \rightarrow \{0, 1\}$, and take $d \geq \log^c n$ for a large enough constant c . Let $Q_{f,d}$ be a HQFP that formalizes “the communication complexity of KW_f is at most d ”, and whose proofs reduce, via low degree polynomials, to proofs of $Q_{f,d}^{\text{ckt}}$ (a specific formulation of “an alternating circuit of depth d can compute f ”). Then, if the SDFP relaxation of $Q_{f,d}$ obeys the Berman–Ben-Israel criterion, it must have a solution (i.e. a cost d protocol for KW_f in the relaxed model).*

This no-go theorem should be seen as a natural, expected consequence of the results of Austrin and Risse. But, perhaps owing to our inexperience with proof complexity, it was not easy for us to verify that the formal connection is really there, and so in the full version of our paper [11] we provide a formalization and proof of this no-go theorem.

In light of such a result, one should ask: is it still worthwhile to pursue the stated aim, of formalizing communication complexity using a HQFP, relaxing to a SDFP, and studying the resulting communication model? As it turned out, we went through this formalize-and-relax process twice, and in both times there was something interesting to be found on the other side. In one case we ended up with a communication model which is a kind of structured version of the well-known and well-studied γ_2 norm. In the other case, we ended up with a communication model that has a natural, physical description, and understanding this model revealed to us something non-obvious about the nature of quantum measurements.

And although it is now expected that both models can solve all Karchmer–Wigderson relations, the above no-go theorem is not constructive, and gives us no explicit description of the algorithms in the model that actually do this. So it is still worthwhile to give a constructive proof of this, i.e., to find algorithms in the model for solving Karchmer–Wigderson relations.

4 γ_2 Communication

Our first attempt to express a communication protocol as an HQFP proceeds as follows. We view a two-party communication protocol computing a function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ as defined by two parts: the structure \mathcal{T} of the protocol tree (i.e. the nodes in the protocol tree and which player speaks at which node) and the specification of the state of the protocol at each node in the protocol tree. The second part can be specified by a combinatorial rectangle of $\mathcal{X} \times \mathcal{Y}$ at every node.

Now, for a given tree structure \mathcal{T} we will design an HQFP Q_{protocol} such that solutions to Q_{protocol} are in 1-1 correspondence with protocols of structure \mathcal{T} for computing f , i.e., associations of rectangles to the nodes of \mathcal{T} that form a valid protocol for computing f . Then, there will exist a protocol with structure \mathcal{T} computing the function f if and only if there is a solution to Q_{protocol} .

The central feature of Q_{protocol} is that we have one variable $A_t(x)$ for each node t and each input x of Alice, and one variable $B_t(y)$ for each node t and each input y of Bob, so that the product $A_t(x) \cdot B_t(y)$ is to be interpreted as an indicator of whether the input (x, y) belongs to the rectangle associated with node t . Given this particular choice of variables, the constraints are the most obvious possible. We give a detailed description of the HQFP in the full version of our paper [11] and over here we directly provide the SDFP relaxation P_{protocol} . It will follow, then, that one can view a solution of P_{protocol} as a generalization of a protocol computing f and we refer to the solutions of P_{protocol} as “ γ_2 protocols” due to their relationship with the γ_2 norm that we will describe shortly.

Recall that a SDFP asks whether there exists vectors obeying some linear constraints on the inner products.

A (binary, two-player) γ_2 deterministic protocol is a tuple $\pi = (\mathcal{X} \times \mathcal{Y}, \mathcal{T}, d, \alpha, \beta)$, where $\mathcal{X} \times \mathcal{Y}$ is a finite product set of inputs, \mathcal{T} is a protocol structure, and α and β are collections of maps $\alpha_t : \mathcal{X} \rightarrow \mathbb{R}^d$ and $\beta_t : \mathcal{Y} \rightarrow \mathbb{R}^d$, for each node $t \in \mathcal{T}$, satisfying a number of constraints below – that arise from relaxation of the standard protocol constraints described above, where we replace the multiplication by the standard inner product $\langle \cdot, \cdot \rangle$ in \mathbb{R}^d .

Root constraints. For the root λ of \mathcal{T} we have the following constraints.

$$\begin{aligned} \langle \alpha_\lambda(x), \alpha_\lambda(x') \rangle &= 1 & \forall x, x' \in \mathcal{X} \\ \langle \beta_\lambda(y), \beta_\lambda(y') \rangle &= 1 & \forall y, y' \in \mathcal{Y} \\ \langle \alpha_\lambda(x), \beta_\lambda(y) \rangle &= 1 & \forall (x, y) \in \mathcal{X} \times \mathcal{Y} \end{aligned}$$

This implies that every $\alpha_\lambda(x)$ and $\beta_\lambda(x)$ is the same unit-length vector (in ℓ_2 norm).

Alice’s nodes constraints. Let $t \in \mathcal{T}$ be an Alice’s node with children t_0, t_1 . We impose the following constraints which are equivalent to saying (using the Cauchy-Schwarz inequality and the Pythagorean theorem) that for any $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, we have that $\alpha_t(x) = \alpha_{t_0}(x) + \alpha_{t_1}(x)$, $\alpha_{t_0}(x)$ and $\alpha_{t_1}(x)$ are orthogonal, and $\beta_t(y) = \beta_{t_0}(y) = \beta_{t_1}(y)$. (Note that if we restrict to 1-dimensional vectors, this is exactly how a rectangle gets partitioned into subrectangles at an Alice node.)

$$\begin{aligned} \|\alpha_{t_0}(x)\|^2 + \|\alpha_{t_1}(x)\|^2 &= \|\alpha_t(x)\|^2 & \forall x \in \mathcal{X} \\ \langle \alpha_{t_0}(x), \alpha_t(x) \rangle + \langle \alpha_{t_1}(x), \alpha_t(x) \rangle &= \|\alpha_t(x)\|^2 & \forall x \in \mathcal{X} \\ \langle \alpha_{t_0}(x), \alpha_{t_1}(x) \rangle &= 0 & \forall x \in \mathcal{X} \\ \|\beta_{t_0}(y)\|^2 &= \|\beta_t(y)\|^2 & \forall y \in \mathcal{Y} \\ \|\beta_{t_1}(y)\|^2 &= \|\beta_t(y)\|^2 & \forall y \in \mathcal{Y} \\ \langle \beta_{t_0}(y), \beta_t(y) \rangle &= \|\beta_t(y)\|^2 & \forall y \in \mathcal{Y} \\ \langle \beta_{t_1}(y), \beta_t(y) \rangle &= \|\beta_t(y)\|^2 & \forall y \in \mathcal{Y} \end{aligned}$$

Bob’s nodes constraints. The constraints for Bob’s nodes are analogous to Alice’s node constraints.

Let $f \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ be a relation with output set $\mathcal{Z} \subseteq \{0, 1\}^k$ and let $\pi = (\mathcal{X} \times \mathcal{Y}, \mathcal{T}, d, \alpha, \beta)$ be a γ_2 protocol. We say that π computes f if the depth of every leaf $\ell \in \mathcal{T}$ is at least k , and the collections α and β satisfy the following constraints.

Computational constraints. For every leaf $\ell \in \mathcal{T}$ of the form $\ell = tz$ for some $z \in \{0, 1\}^k$ we have the following constraints:

$$\langle \alpha_\ell(x), \beta_\ell(y) \rangle = 0 \quad \forall (x, y) \in \mathcal{X} \times \mathcal{Y} \text{ s.t. } (x, y, z) \notin f$$

The *deterministic γ_2 communication complexity* of f , $\Gamma_2 D^{\text{cc}}(f)$, is the smallest depth of a protocol structure \mathcal{T} such that there exists a γ_2 deterministic protocol $\pi = (\mathcal{X} \times \mathcal{Y}, \mathcal{T}, d, \alpha, \beta)$ that computes f .

4.1 Relation to the γ_2 norm

The γ_2 norm was introduced to the TCS community by Linial et al. [20] to study sign matrices.

► **Definition 3.** Let $A \in \mathbb{R}^{m \times n}$ be a matrix. Then,

$$\gamma_2(A) = \min\{r(X)r(Y) \mid A = XY^t\},$$

where $r(M)$ is the largest ℓ_2 norm of a row of the matrix M .

One can see a matrix A with $\gamma_2(A) \leq 1$ as a generalization of a rectangle. Let \mathcal{X} and \mathcal{Y} be sets and $R = A \times B$ be a rectangle, where $A \subseteq \mathcal{X}$ and $B \subseteq \mathcal{Y}$. Let $M_R = \{0, 1\}^{\mathcal{X} \times \mathcal{Y}}$ be a matrix representing the rectangle R , i.e., $M_R[x, y] = 1$ if and only if $(x, y) \in R$. We can decompose the matrix M_R as $M_R = uv^t$, where $u \in \{0, 1\}^{\mathcal{X}}$ and $v \in \{0, 1\}^{\mathcal{Y}}$ are the characteristic vectors of the sets A and B , respectively. Clearly, $r(u) = r(v) = 1$, if we take the vectors u and v as matrices with one column. Thus, $\gamma_2(M_R) \leq 1$. From the Cauchy-Schwarz inequality, it follows that $\gamma_2(M_R) = 1$. Hence, one can think of matrices with $\gamma_2(M) \leq 1$ as a generalization of the notion of a combinatorial rectangle. (In fact, it is possible to write down a HQFP whose solutions are precisely indicator matrices of combinatorial rectangles, and whose semidefinite relaxations are precisely matrices of subunit γ_2 norm.) This line of thought bore many fruits in the study of communication complexity, such as lower bounds, lifting theorems, the ability to approximate PP-communication-complexity using semidefinite programming, etc, see [18] for a survey.

Our definition of γ_2 protocols are akin to deterministic protocols in that it decomposes the space $\mathcal{X} \times \mathcal{Y}$ in a “tree-structured” manner into “monochromatic” subunit γ_2 norm matrices. It is this structure that is not captured by the linear relaxations mentioned in the introduction.

4.2 The power of γ_2 protocols

In the full version of the paper [11], we show two results about γ_2 protocols.

Via the relation to the γ_2 norm we are able to show limits on the power of γ_2 protocols. For a relation $f \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$, let γ_2 leaf complexity $\Gamma_2 L^{\text{cc}}(f)$ denote the smallest number of leaves of the protocol structure of a γ_2 protocol that computes f . It clearly holds that

$$\Gamma_2 D^{\text{cc}}(f) \geq \log \Gamma_2 L^{\text{cc}}(f).$$

We show the following lower bound analogous to the rank lower bound in communication complexity.

► **Theorem 4.** For any Boolean function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$, it holds that

$$\Gamma_2 L^{\text{cc}}(f) \geq \gamma_2(f).$$

So, for example, the inner-product mod-2 function cannot be computed by γ_2 protocols of depth $o(n)$.

On the other hand, we also design a two-round γ_2 protocol for the equality function.

Let \mathcal{T}_ℓ be the following protocol structure:

1. The root λ of \mathcal{T}_ℓ is an Alice's node and has degree ℓ .
2. Each child of the root is a Bob's node and has exactly 2 children.

► **Theorem 5.** *Let $\ell \geq 11$ and $d > 0$ be integers. Then, there is γ_2 protocol with the structure \mathcal{T}_ℓ computing the equality function EQ_d .*

By the usual binary-search reduction of Karchmer–Wigderson relations to equality, it follows that every Karchmer–Wigderson relation can be solved in γ_2 communication $O(\log n)$.

5 Quantum Lab Protocols

Let us begin by contrasting what we will do in Section 5 with what we have done in Section 4. As before, we will formulate the existence of a two-party deterministic protocol computing f as a HQFP. In the previous section, for every node t in the protocol tree we had variables $\{A_t(x)\}_{x \in \mathcal{X}} \cup \{B_t(y)\}_{y \in \mathcal{Y}}$. The different starting point here is that our HQFP will instead have $\{C_t(x, y)\}_{x \in \mathcal{X}, y \in \mathcal{Y}}$. Before, we interpreted $A_t(x) \cdot B_t(y) \in \{0, 1\}$ as indicating whether (x, y) is in the rectangle associated with t . Now, instead, we let $C_t(x, y) \in \{0, 1\}$ indicate the same thing. The constraints of the new program are again designed in the most obvious way possible, so as to ensure that the HQFP is feasible if and only if f can be computed by a deterministic communication protocol with the given structure. We will then relax the quadratic program to a semidefinite program and see what we get.

Notice the difference in approach. In the previous section we had a rationale to obtain the semidefinite program which we obtained: we wanted to add structure to a known rectangle-like notion, the γ_2 norm, in a similar way to how protocols are obtained from rectangles. The previous model can thus be justified on technical grounds, as, *what happens when we add structure to the γ_2 norm?* In contrast, the work in this section began by simply trying to make a different set of constraints where the variables are organized differently. It was surprising to us, then, to discover that the resulting computational model has a natural, functional definition, which is why we call it the “quantum lab” model. We provide the semidefinite program below and then elaborate on the quantum lab interpretation.

A *deterministic quantum-lab protocol* is a tuple $\pi = (\mathcal{X} \times \mathcal{Y}, \mathcal{T}, d, \psi)$, where $\mathcal{X} \times \mathcal{Y}$ is a finite product set of *inputs*, \mathcal{T} is a protocol structure, and ψ is a collections of maps $\psi_t : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}^d$, for each node $t \in \mathcal{T}$, satisfying the following constraints.

Root constraints. For the root λ of \mathcal{T} we have:

$$\psi_\lambda(x, y) \cdot \psi_\lambda(x', y') = 1 \quad \forall x, x' \in \mathcal{X}, y, y' \in \mathcal{Y}$$

This implies that every $\psi_\lambda(x, y)$ is the same unit-length vector.

Alice's nodes constraints. For $t \in \mathcal{T}$ an Alice node with children t_0, t_1 :

$$\begin{aligned} \|\psi_{t_0}(x, y)\|^2 + \|\psi_{t_1}(x, y)\|^2 &= \|\psi_t(x, y)\|^2 & \forall x \in \mathcal{X} \\ \langle \psi_{t_0}(x, y), \psi_t(x, y) \rangle + \langle \psi_{t_1}(x, y), \psi_t(x, y) \rangle &= \|\psi_t(x, y)\|^2 & \forall x \in \mathcal{X} \\ \langle \psi_{t_0}(x, y), \psi_{t_1}(x, y') \rangle &= 0 & \forall x \in \mathcal{X}, y, y' \in \mathcal{Y} \end{aligned}$$

While the first two constraints above are just an orthogonal decomposition, the third constraint is more subtle. It is a relaxation of a constraint in the HQFP that maintains that when Alice makes a decision on what to communicate on input x , the same decision is made regardless of what the value of y is.

Bob's nodes constraints. The constraints for Bob's nodes are analogous to Alice's node constraints.

We are only missing the constraints that define when a protocol computes a relation. So let $f \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Z}$ be a relation with output set $\mathcal{Z} \subseteq \{0, 1\}^k$ and let $\pi = (\mathcal{X} \times \mathcal{Y}, \mathcal{T}, d, \psi)$ be a quantum lab protocol. We say that π *computes* f if the depth of every leaf $\ell \in \mathcal{T}$ is at least k , and ψ satisfies:

Computational constraints. For every leaf $\ell \in \mathcal{T}$ of the form $\ell = tz$ for some $z \in \{0, 1\}^k$ we have the following constraints:

$$\|\psi_\ell(x, y)\|^2 = 0 \quad \forall (x, y) \in \mathcal{X} \times \mathcal{Y} \text{ s.t. } (x, y, z) \notin f$$

5.1 The Quantum Lab

In the above program, let us think of each $\psi_t(x, y)$ as an (unnormalized) quantum state. Then the root constraints say that the initial state, at the root λ , is the same for all (x, y) . The constraints at an Alice node say that $\psi_{t0}(x, y)$ and $\psi_{t1}(x, y)$ are an orthogonal decomposition of $\psi_t(x, y)$, but furthermore every quantum state $\psi_{t0}(x, y)$ is orthogonal to every $\psi_{t1}(x, y')$. This implies that there exists a pair of orthogonal projections $\Pi_{t,x,0}, \Pi_{t,x,1}$ such that $\psi_{ti}(x, y) = \Pi_{t,x,i}\psi_t(x, y)$ (e.g. $\Pi_{t,x,0}$ projects onto the span of every $(\psi_{t0}(x, y))_{y \in \mathcal{Y}}$, and $\Pi_{t,x,1}$ projects to its orthogonal complement). In other words, to each t and each x corresponds a measurement, and $\psi_{ti}(x, y)$ is the (unnormalized) state obtained by measuring $\psi_t(x, y)$. Likewise, the constraints at Bob's nodes are equivalent to the existence of such a measurement $(\Pi_{y,0}, \Pi_{y,1})$ depending only on t and y . This leads us to the following physically realizable interpretation of this model.

Alice and Bob work in a idealized quantum laboratory. In this quantum lab, they can prepare any quantum state that they wish, and they can manipulate it without any error using the available equipment. With this lab at their disposal, they play the following “communication” game. Before they receive their respective inputs, Alice and Bob are allowed to go to the lab together, and prepare a quantum system in some initial state $|\psi_0\rangle$, known to both. Then they are separated, Alice receives an input $x \in X$, and Bob receives an input $y \in Y$. Their goal is now to compute $f(x, y)$. For this purpose, Alice and Bob take separate turns going to the lab. When one of them is in the lab, she or he is allowed to perform a binary measurement on the quantum system, and write the outcome, 0 or 1, in the lab's whiteboard. The measurement that is performed by each player can depend on the input known to her or him, and on the *transcript* of all previous measurement outcomes, which are written in the whiteboard. The question is then: how many times (in the worst case) must Alice and Bob make a measurement in the lab, in order to discover $f(x, y)$? Note that, unusually for a quantum model, here we require that Alice and Bob learn $f(x, y)$ without any error. To this minimum number we could call the (*deterministic*) *quantum-lab complexity* of f .

The first observation is that Alice and Bob can simulate a deterministic protocol. Indeed, if they prepare the two qubit state $|01\rangle$, then Alice can “communicate” a 0 to Bob by measuring the first qubit, which will always be 0, and she can communicate a 1 by measuring the second qubit. So this shows, for example, that the two-round quantum-lab complexity of

any Boolean function is at most $n + 1$, since Alice can communicate their entire input to Bob, and Bob replies with $f(x, y)$. The question is now: can Alice and Bob do better if the lab is quantum? ⁷

On our part, after discovering this functional description of the model, we were possessed of the following strong intuition: *the measurement that a player is allowed to make depends on her/his input and on the current state $|\psi\rangle$, but if it is a binary measurement, then it cannot reveal more than 1 bit of information about her/his input, and hence there should exist some kind of information-theoretic lower-bound on the quantum-lab complexity.* We were hoping to prove, at least, that the quantum information complexity [31] would serve as a lower-bound for quantum-lab complexity.

This intuition, however, turned out to be spectacularly wrong. We were first encouraged by a proof that equality requires $\Omega(n)$ bits to be computed by a two-round quantum-lab protocol (in a two-round protocol Alice does several measurements, then Bob, after which the answer must be known). There is in fact a very simple proof of this using the quantum pigeonhole principle. This early result was encouraging but highly misleading. After a lot of effort trying to prove a lower-bound for 3 rounds, we eventually discovered that equality has a 3-round quantum lab protocol with $O(1)$ complexity. Perhaps this is not surprising, since the information complexity of equality is $O(1)$, and the no-go theorem implies that KW-games will all be easy in the model.

However, a small adjustment to the same protocol revealed that *every Boolean function* can be solved in three rounds with $O(1)$ measurements! This, we did find very surprising, as did everyone to whom we explained the result. On the nature of quantum measurements, we can conclude that although each measurement in the quantum lab can only reveal one bit of information (about x to Bob, and about y to Alice), measurements alone can manipulate the state so that *any* joint bit $f(x, y)$ is revealed.

Perhaps here the reader is tempted to try and solve the puzzle themselves, for which we give the structure of the protocol as a clue: Alice goes to the lab, makes a 1-bit measurement depending on x , then Bob goes and makes a two-bit measurement depending on y and on the outcome of Alice’s measurement, and then Alice returns to the lab, and does one final 1-bit measurement (depending on x and the previous outcomes) whose answer will be exactly $f(x, y)$. This same protocol structure works for computing any Boolean function f , it is only the chosen measurements that vary. Our solution appears in Section 5.3.

5.2 A 2-round Lower Bound for Equality

Here we show that the equality function on n bits needs $n + 1$ bits to be computed by a two-round quantum lab protocol, i.e., a quantum lab protocol where Alice speaks, and then Bob speaks, with his last measurement giving the answer.

⁷ As a passing remark, we note that we could have given the very same definition above, but for a *classical laboratory*. In a classical lab, Alice and Bob can prepare any classical state (a distribution over basic states), and measurements correspond to orthogonal projections on a fixed basis, followed by renormalization in the ℓ_1 norm. One can get a sense for the model by imagining a lab made of mechanical contraptions that toss random coins and pull strings and send metal spheres rolling down rails and so on. Every day Alice or Bob go to the lab, and do a “orthogonal measurement in a fixed basis”, meaning they partition the set of possible outcomes into two, and ask in which of the two sets is the state of the lab. (One can imagine that they look through a window to learn one bit about the state.) As it turns out, this model corresponds to the completely positive relaxation of our HQFP, and it can be shown that, if we require the output to be correct with probability at least $\varepsilon \in [0, 1]$, our program gives us exactly the ε -error randomized communication complexity.

Indeed, if Alice has input x and makes k measurements, then the initial state ψ_λ is broken into an orthogonal decomposition, which does not depend on y since Bob did not speak yet:

$$\psi_\lambda = \sum_t \psi_t(x, y) = \sum_t \psi_t(x) \quad \langle \psi_t(x), \psi_{t'}(x) \rangle = 0$$

Now, if $2^k < 2^n$, the QPHP (Theorem 1) states that there must exist some message t , and two inputs x, x' , such that

$$\langle \psi_t(x), \psi_t(x') \rangle \neq 0.$$

Now Bob comes along and does some measurements. Suppose he has input x . Since $\psi_t(x)$ and $\psi_t(x')$ are not orthogonal, then no matter which measurement he does, there must be an outcome i such that $\psi_{ti}(x, x)$ and $\psi_{ti}(x', x)$ are both non-zero. It follows that ti is not monochromatic, i.e., the computational constraints associated with leaf ti are not obeyed.

5.3 Model Collapse – All Functions Are Easy

► **Theorem 6.** *Given any function $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$, there is a 3-round Quantum Lab protocol using 4 bits of communication that computes f .*

Proof. In our protocol given below the root node is a Bob node, The nodes at depths 1 and 2 are Alice nodes, the nodes at depth 3 are Bob nodes and the depth 4 nodes are leaves. We refer to nodes using their partial transcripts (i.e. elements of $\{0, 1\}^{\leq 4}$ with ε being the empty string). We refer to the state in the quantum lab at a node v on inputs x and y as $|\psi_v^{xy}\rangle$.

The state in the quantum lab has 3 registers, which we number $1'$, $2'$ and 3. Register 3 is 2-dimensional with basis states $|0\rangle$ and $|1\rangle$ (i.e. the register consists of one qubit) and registers $1'$ and $2'$ are $|\mathcal{X}| + |\mathcal{Y}| + 1$ -dimensional with their basis states being $|\perp\rangle$, $|x\rangle$ and $|y\rangle$ for each $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. We now provide the (unnormalized) states in the quantum lab at each node for the first three bits of communication.

$$\begin{aligned} \blacksquare \quad & |\psi_\varepsilon^{xy}\rangle = |0\rangle_3 |\perp\rangle_{1'} |\perp\rangle_{2'} \\ \blacksquare \quad & |\psi_0^{xy}\rangle = \frac{1}{2} |0\rangle_3 (|\perp\rangle_{1'} + |y\rangle_{1'}) |\perp\rangle_{2'} \\ & |\psi_1^{xy}\rangle = \frac{1}{2} |0\rangle_3 (|\perp\rangle_{1'} - |y\rangle_{1'}) |\perp\rangle_{2'} \\ \blacksquare \quad & |\psi_{00}^{xy}\rangle = \frac{1}{4} |0\rangle_3 (|\perp\rangle_{1'} + |y\rangle_{1'}) (|\perp\rangle_{2'} + |x\rangle_{2'}) \\ & |\psi_{01}^{xy}\rangle = \frac{1}{4} |0\rangle_3 (|\perp\rangle_{1'} + |y\rangle_{1'}) (|\perp\rangle_{2'} - |x\rangle_{2'}) \\ \blacksquare \quad & |\psi_{000}^{xy}\rangle = \frac{1}{2} \left(\frac{1}{4} |0\rangle_3 (|\perp\rangle_{1'} + |y\rangle_{1'}) (|\perp\rangle_{2'} + |x\rangle_{2'}) + \frac{1}{2\sqrt{2}} |1\rangle_3 (|x\rangle_{1'} + (-1)^{f(x,y)} |y\rangle_{1'}) |\perp\rangle_{2'} \right) \\ & |\psi_{001}^{xy}\rangle = \frac{1}{2} \left(\frac{1}{4} |0\rangle_3 (|\perp\rangle_{1'} + |y\rangle_{1'}) (|\perp\rangle_{2'} + |x\rangle_{2'}) - \frac{1}{2\sqrt{2}} |1\rangle_3 (|x\rangle_{1'} + (-1)^{f(x,y)} |y\rangle_{1'}) |\perp\rangle_{2'} \right) \end{aligned}$$

We will address the last bit of communication after analyzing the above. We have only specified the relevant states along the all-0 transcript, and we will show that these can be realized by a quantum lab protocol. The states that appear along the other transcripts are the same up to some sign changes and so can also be realized similarly. As an example of how the states differ along different transcripts, here is the state at a node of depth 3:

$$\begin{aligned} |\psi_{b_1 b_2 b_3}^{xy}\rangle = & \frac{1}{2} \left(\frac{1}{4} |0\rangle_3 (|\perp\rangle_{1'} + (-1)^{b_1} |y\rangle_{1'}) (|\perp\rangle_{2'} + (-1)^{b_2} |x\rangle_{2'}) \right. \\ & \left. + (-1)^{b_3} \frac{1}{2\sqrt{2}} |1\rangle_3 (|x\rangle_{1'} + (-1)^{b_1} (-1)^{f(x,y)} |y\rangle_{1'}) |\perp\rangle_{2'} \right) \end{aligned}$$

To show that the above quantum states can be realized by a quantum lab protocol, we will verify that the quantum lab protocol constraints are satisfied by these. For each node $v \in \{\varepsilon, 0, 00\}$ it suffices to verify the following.

- $\psi_v^{xy} = \psi_{v0}^{xy} + \psi_{v1}^{xy}$.

This constraint is easy to verify.

- At an Alice node v , $\langle \psi_{v0}^{xy}, \psi_{v1}^{x'y'} \rangle = 0$ for all x, y, y' .

This constraint is easy to verify for $v = 0$. For $v = 00$, this inner product is

$$\frac{1}{4} \left(\frac{1}{16} \cdot 1 \cdot (1 + [y = y']) \cdot 2 - \frac{1}{8} \cdot 1 \cdot (1 + (-1)^{f(x,y)+f(x,y')} [y = y']) \cdot 1 \right)$$

where $[y = y']$ is 1 if $y = y'$ and 0 otherwise. Note that this is 0 both when $y \neq y'$ and when $y = y'$.

- At a Bob node v , $\langle \psi_{v0}^{xy}, \psi_{v1}^{x'y} \rangle = 0$ for all x, x', y

Since the only Bob node in the first three bits is ε , we only need to ensure that $\langle \psi_0^{xy}, \psi_1^{x'y} \rangle = 0$. This is again easy to verify.

The final bit of communication. We now make an additional observation about the state that we have reached after 3 bits of communication. Namely, fix any $y \in \{0, 1\}^n$ and let x, x' be two inputs such that $f(x, y) \neq f(x', y)$. Then

$$\langle \psi_{000}^{xy}, \psi_{000}^{x'y} \rangle = \frac{1}{4} \left(\frac{1}{16} \cdot 1 \cdot 2 \cdot 1 + \frac{1}{8} \cdot 1 \cdot (-1) \cdot 1 \right) = 0.$$

As a consequence $V_0^y := \text{span}(\{\psi_{000}^{xy}\}_{x:f(x,y)=0})$ is orthogonal to $V_1^y := \text{span}(\{\psi_{000}^{xy}\}_{x:f(x,y)=1})$. So now Bob can perform the measurement $\{\Pi_{V_0^y}, I - \Pi_{V_0^y}\}$. The output of the measurement is the value of $f(x, y)$. ◀

5.4 Future directions

We have proposed a specific way of generalizing Π_1 statements. We would like to suggest a few questions for the future.

- What other combinatorial principles can be relaxed by the above approach? An interesting avenue is to investigate the several different combinatorial principles that lie at the basis of TFNP classes, write each of them down by a HQFP, relax to a SDFP, and see what is there. Does this work often? Do we get interesting quantum versions of known principles? In other words, we have an (incomplete) proof system for Σ_1 and Π_1 statements, such that every statement or its negation has short proofs. What other interesting theorems can it prove?
- Could we take a similar approach using lattice duality? E.g. we could try to express Σ_1 statements using the closest vector problem (which is NP-hard), and then relax the approximation factor to \sqrt{n} , which puts the problem in $\text{NP} \cap \text{coNP}$ [1], and see if the statement is still meaningful.
- Could we take a similar approach using stochastic games? Here we have no suggestion for which NP-hard problem could be used, that has stochastic games as a relaxation.
- We have proven that any KW game can be solved by γ_2 protocols of depth $\leq \log(11 \times 2) \cdot \log n \approx 4.45 \log n$, i.e. size $\approx n^{4.45}$. However, the best known lower-bounds on formula size are (roughly) cubic [13]. Although it seems like a long shot, perhaps one can still prove a super-cubic lower-bound on formula size by constructing an explicit dual to the SDFPs defining γ_2 protocol for the Karchmer–Wigderson game of some explicit function?

- We chose not include the details in this write-up, but it is possible to relax the HQFPs using the completely positive cone, instead of the semidefinite cone. The semidefinite cone is the cone of matrices of inner products of vectors in the entire space, and the completely positive cone is the cone of matrices inner products of vectors in the non-negative orthant. When doing so, one systematically obtains *randomized* versions of the statements, instead of *quantum* versions. We did not explore this much, because completely positive feasibility is still an NP-complete problem. But it might be interesting to see what one gets by such relaxation: maybe new randomized versions of known combinatorial principles?

6 Supplementary Definitions

We assume that the reader is familiar with Boolean formulas, Boolean circuits, and communication complexity. Recall that the Karchmer–Wigderson theorem states that the minimum depth of a Boolean circuit or formula that computes a given Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$, is equal to the communication complexity of the Karchmer–Wigderson relation KW_f , where Alice is given $x \in f^{-1}(1)$ and Bob is given $y \in f^{-1}(0)$, and they wish to find some i such that $x_i \neq y_i$. A proof can be found in [16, Section 10.2, see also Chapters 5 & 10].

Discrepancy

A well-known lower bound for the communication complexity of several models is the discrepancy of a function f (see, e.g., [16, Section 3.5]). Informally speaking, if a function f has a small discrepancy, then any large rectangle R is almost balanced (the number of 1's and 0's in R is roughly the same).

► **Definition 7.** Let $f : \mathcal{X} \times \mathcal{Y} \rightarrow \{0, 1\}$ be a function, $R \subseteq \mathcal{X} \times \mathcal{Y}$ be a rectangle, and μ be a distribution over $\mathcal{X} \times \mathcal{Y}$. Denote

$$\text{disc}_\mu(R, f) = \left| \Pr_{(x,y) \sim \mu} [f(x, y) = 0, (x, y) \in R] - \Pr_{(x,y) \sim \mu} [f(x, y) = 1, (x, y) \in R] \right|.$$

The discrepancy of f according to μ is

$$\text{disc}_\mu(f) = \max_R \text{disc}_\mu(R, f),$$

where the maximum is over all rectangles $R \subseteq \mathcal{X} \times \mathcal{Y}$. The discrepancy of f is

$$\text{disc}(f) = \min_\mu \text{disc}_\mu(f).$$

The notation Σ_1 , Π_1 , NP, coNP, NP(\mathbb{R}) and coNP(\mathbb{R})

We use Σ_1 and Π_1 to informally refer to existential and universal statements, respectively. When precision is required, we will use NP and coNP for the well-known Boolean complexity classes, and NP(\mathbb{R}) and coNP(\mathbb{R}) for the low-degree Blum-Shub-Smale (BSS) variants. The definition is rather technical, but here it is: The BSS model is a variant of the multitape Turing machine where each tape cell holds a real number, and at each step the machine can read the numbers under some of the tape heads, apply a multilinear polynomial to the numbers (which polynomial depends on the state), and write the result back; it can also branch on comparisons between cells, or between a cell and a fixed constant. The low-degree polytime variant imposes the restriction that the computation is syntactically

polynomial-degree, meaning that the machine runs in polynomial time, but furthermore: at any given time, for each possible branching that happened before time t , the contents of each cell will be a polynomial in the real numbers x_1, \dots, x_n placed in the tape at the start of the computation, and we then require that the degree of this polynomial to also be $\text{poly}(n)$ -bounded (in principle the degree after t steps could be 2^t by repeated squaring). Then $\text{NP}(\mathbb{R})$ is the class of languages $L \subseteq \mathbb{R}^*$ for which there exists a low-degree polytime BSS machine M such that $(x_1, \dots, x_n) \in L \iff \exists(y_1, \dots, y_m) \in \mathbb{R}^{\text{poly}(n)} M(\bar{x}, \bar{y}) = 1$.⁸

Conic feasibility problems

Here we discuss duality for conic feasibility problems.

► **Definition 8.** Let $S, T \subseteq \mathcal{H}$ denote arbitrary, non-empty subsets of a finite-dimensional real Hilbert space \mathcal{H} . I.e., $\mathcal{H} = \mathbb{R}^d$ for some d , but equipped with a possibly non-standard inner-product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$.

- We let $\text{cl}(S)$, the closure of S , be the set of points $x \in \mathcal{H}$ for which there exists a sequence $(x_i)_{i \in \mathbb{N}}$ of points in S such that $\|x_i - x\|_{\mathcal{H}} \rightarrow 0$. We call S closed if $S = \text{cl}(S)$.
- For $\lambda \in \mathbb{R}$, we denote $\lambda S = \{\lambda s \mid s \in S\}$, $S + T = \{s + t \mid s \in S, t \in T\}$.
- A set S is called convex if it contains all the line segments between its points, i.e., $\alpha S + (1 - \alpha)S \subseteq S$ for every $0 \leq \alpha \leq 1$.
- S is called a cone if $\lambda S \subseteq S$ for all $\lambda \geq 0$. A cone S will be convex iff $S + S \subseteq S$. A cone is called pointed if $S \cap -S = \{0\}$.
For example, a subspace is a closed convex cone. The non-negative orthant is a closed, convex, pointed cone.
- The polar of S , denoted S^* , is the set

$$S^* = \{y \in \mathcal{H}^* \mid \forall x \in S \langle x, y \rangle_{\mathcal{H}} \geq 0\}.$$

Examples. The following sets are closed, convex, pointed cones:

- The non-negative orthant $\mathbb{R}_{\geq 0}^n$. It is self-dual, meaning $(\mathbb{R}_{\geq 0}^n)^* = \mathbb{R}_{\geq 0}^n$.
- The set of positive semidefinite $n \times n$ matrices PSD_n , which is a subset of the space $\mathbb{R}^{\frac{n(n+1)}{2}}$ of symmetric matrices, with the inner product $\langle M, N \rangle = \sum_{i,j} M_{i,j} N_{i,j}$.
This set can be alternatively characterized as the set of symmetric matrices with non-negative eigenvalues, or as the set of Gram matrices, i.e., matrices equal to AA^t for some $n \times m$ matrix A , i.e., matrices M of inner products, given by a family of vectors a_1, \dots, a_n (the rows of A), so that $M_{ij} = \langle a_i \mid a_j \rangle$. It is also self-dual.
- The set of completely positive $n \times n$ matrices $\text{CP}_n \subseteq \mathbb{R}^{\frac{n(n+1)}{2}}$ (also symmetric). This set can be alternatively characterized as the set of symmetric matrices with non-negative eigenvalues whose eigenvectors are entrywise non-negative in the standard basis, or the matrices of the form $M = AA^t$ for some $n \times m$ matrix A with non-negative entries, or matrices of inner-products of vectors in the non-negative orthant. Its dual cone is the cone of co-positive matrices, but we will not define it or mention it again.

► **Definition 9.** Let $\mathcal{K} \subseteq \mathbb{R}^n$ be a closed, convex, pointed cone. A conic feasibility problem over \mathcal{K} is defined by a linear map $\mathcal{A} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and a point $b \in \mathbb{R}^m$. The problem asks whether there exists an element $Z \in \mathcal{K}$ such that $\mathcal{A}(Z) = b$. Such a Z is called a solution. If a solution exists, we say that the problem (\mathcal{A}, b) is feasible, or satisfiable, and otherwise we say that the problem (\mathcal{A}, b) is infeasible, or unsatisfiable.

⁸ If the reader is wondering why the low-degree restriction, it is because polytime BSS machines without degree constraints can do things that seem too powerful, such as factoring [27].

Examples. A linear feasibility problem is a conic feasibility over the non-negative orthant. A semidefinite feasibility problem (SDFP) is a conic feasibility problem over the cone of positive semidefinite matrices.

Duality for SDFPs

The feasibility of a conic feasibility problem over \mathcal{K} is an existential statement, in fact it is a Σ_1 statement provided that $Z \in \mathcal{K}$ is itself a Σ_1 statement. A remarkable general fact about conic feasibility is that the *infeasibility* of a conic feasibility problem can *also* be formulated as a Σ_1 statement. This fact is really non-obvious: it was first proven for SDFPs by Ramana [24] (see [22] for a simplified treatment), and for general conic feasibility by [21]. This result is an instance of the general phenomenon of *convex duality*, which is also the source of the $\text{NP} \cap \text{coNP}$ inclusions of approximate lattice problems [1] and stochastic games (e.g. [3, 4], although here convexity is over the tropical semiring).

The precise statement which is equivalent to the infeasibility of a conic optimization problem, the so called *dual problem*, is not easy to describe in general. It is usually a Σ_1 statement with another cone as an oracle, usually the polar cone \mathcal{K}^* over a larger dimension, or another related cone.

However, in some cases, a dual problem exists which *is* easy to describe, whose flavor is similar to Farkas' lemma of linear feasibility, and indeed gives exactly Farkas' lemma when applied to the non-negative orthant. It was proven long ago by Ben-Israel:

► **Theorem 10** (Ben-Israel [6]). *Let $\mathcal{K} \subseteq \mathbb{R}^n$ be a closed convex cone. Let $\mathcal{A} : \mathbb{R}^n \rightarrow \mathbb{R}^m$ be a linear map, and $b \in \mathbb{R}^m$. Suppose that $\ker(\mathcal{A}) + \mathcal{K}$ is a closed set (Ben-Israel's criterion). Then exactly one of the following two things are true:*

- (i) *Either there exists $Z \in \mathcal{K}$ such that $\mathcal{A}(Z) = b$,*
- (ii) *Or there exists $w \in \mathbb{R}^m$ such that $\mathcal{A}^t(w) \in \mathcal{K}^*$ and $\langle w, b \rangle < 0$.*

A sufficient condition for the closure of $\ker(\mathcal{A}) + \mathcal{K}$ is given by the following lemma. It appears in a paper by Berman and Ben-Israel [7], and there the proof is attributed to A. Charnes and A. Lent.

► **Lemma 11** (Berman–Ben-Israel criterion). *If $L \subseteq \mathbb{R}^n$ is a linear subspace, $S \subseteq \mathbb{R}^n$ is a closed convex cone, and $L \cap S$ is a linear subspace, then $L + S$ is closed. Hence, a sufficient condition for Ben-Israel's criterion to hold is that $\ker(\mathcal{A}) \cap \mathcal{K}$ is a linear subspace, for example, $\ker(\mathcal{A}) \cap \mathcal{K} = \{0\}$.*

In all the SDFPs we will consider, we will have the simplest of conditions $\ker(\mathcal{A}) \cap \mathcal{K} = \{0\}$.

HQFPs, and their relaxation

A SDFP asks whether there exists a positive semidefinite (symmetric) $n \times n$ matrix Z such that $\mathcal{A}(Z) = b$, where \mathcal{A} is a linear map in the entries of Z and $b \in \mathbb{R}^m$. In other words, $\mathcal{A}(Z) = (\langle A_1, Z \rangle, \dots, \langle A_m, Z \rangle)$ for some symmetric real matrices A_1, \dots, A_m . Since positive semidefinite matrices are matrices of inner-products, we can rephrase this question as follows: We wish to know whether there exist vectors $a_1, \dots, a_n \in \mathbb{R}^n$ obeying a set of linear equations in their inner-products $\langle a_i, a_j \rangle$.

We can now consider the same problem, with the additional constraint that the vectors a_1, \dots, a_n are scalars (i.e. come from the same 1-dimensional subspace). This is equivalent to requiring that the solution Z has rank 1. With this additional constraint, we have a system of linear equations in the quadratic products $a_i \cdot a_j$, and we wish to know whether there

exists some choice of scalars that satisfy the system. We call such a problem a *Homogeneous Quadratic Feasibility Problem* (HQFP). Naturally, we can take *any* HQFP and relax it to a SDFP by dropping the rank-1 restriction, i.e. by replacing scalars with vectors and products with inner-products.

References

- 1 Dorit Aharonov and Oded Regev. Lattice problems in $np \cap comp$. *Journal of the ACM (JACM)*, 52(5):749–765, 2005.
- 2 Yakir Aharonov, Fabrizio Colombo, Sandu Popescu, Irene Sabadini, Daniele C Struppa, and Jeff Tollaksen. Quantum violation of the pigeonhole principle and the nature of quantum correlations. *Proceedings of the National Academy of Sciences*, 113(3):532–535, 2016.
- 3 M. Akian, S. Gaubert, and A. Guterman. Tropical polyhedra are equivalent to mean payoff games. *International Journal of Algebra and Computation*, 22(1), 2012. doi:10.1142/S0218196711006674.
- 4 X. Allamigeon, S. Gaubert, and M. Skomra. Solving generic nonarchimedean semidefinite programs using stochastic game algorithms. *Journal of Symbolic Computation*, 85:25–54, 2018. doi:10.1016/j.jsc.2017.07.002.
- 5 Per Austrin and Kilian Risse. Sum-of-squares lower bounds for the minimum circuit size problem. In *Proceedings of CCC*, 2023.
- 6 Adi Ben-Israel. Linear equations and inequalities on finite dimensional, real or complex, vector spaces: A unified theory. *Journal of Mathematical Analysis and Applications*, 27(2):367–389, 1969.
- 7 Abraham Berman and Adi Ben-Israel. More on linear inequalities with applications to matrix theory. *Journal of Mathematical Analysis and Applications*, 33(3):482–496, 1971.
- 8 Mark Bun, Justin Thaler, et al. Approximate degree in classical and quantum computing. *Foundations and Trends in Theoretical Computer Science*, 15(3-4):229–423, 2022. doi:10.1561/0400000107.
- 9 Raul Corrêa and Pablo L Saldanha. Apparent quantum paradoxes as simple interference: Quantum violation of the pigeonhole principle and exchange of properties between quantum particles. *Physical Review A*, 104(1):012212, 2021.
- 10 Jérôme Dohrau, Bernd Gärtner, Manuel Kohler, Jiří Matoušek, and Emo Welzl. Arrival: a zero-player graph game in $np \cap comp$. In *A Journey Through Discrete Mathematics: A Tribute to Jiří Matoušek*, pages 367–374. Springer, 2017.
- 11 Pavel Dvořák, Bruno Loff, and Suhail Sherif. A quantum pigeonhole principle and two semidefinite relaxations of communication complexity, 2024. doi:10.48550/arXiv.2409.04592.
- 12 Pavel Hrubeš, Stasys Jukna, Alexander Kulikov, and Pavel Pudlak. On convex complexity measures. *Theoretical Computer Science*, 411(16-18):1842–1854, 2010. doi:10.1016/J.TCS.2010.02.004.
- 13 Johan Håstad. The shrinkage exponent of de morgan formulas is 2. *SIAM Journal on Computing*, 27(1):48–64, 1998. doi:10.1137/S0097539794261556.
- 14 Stasys Jukna. *Boolean function complexity: advances and frontiers*. Springer, 2012. doi:10.1007/978-3-642-24508-4.
- 15 Mauricio Karchmer, Eyal Kushilevitz, and Noam Nisan. Fractional covers and communication complexity. *SIAM Journal on Discrete Mathematics*, 8(1):76–92, 1995. doi:10.1137/S0895480192238482.
- 16 Eyal Kushilevitz and Noam Nisan. *Communication complexity*. Cambridge University Press, 1997.
- 17 Marc Lackenby. The efficient certification of knottedness and thurston norm. *Advances in Mathematics*, 387:107796, 2021.

- 18 Troy Lee, Adi Shraibman, et al. Lower bounds in communication complexity. *Foundations and Trends in Theoretical Computer Science*, 3(4):263–399, 2009. doi:10.1561/0400000040.
- 19 Lily Li and Morgan Shirley. The general adversary bound: A survey. *arXiv preprint arXiv:2104.06380*, 2021.
- 20 Nati Linial, Shahar Mendelson, Gideon Schechtman, and Adi Shraibman. Complexity measures of sign matrices. *Combinatorica*, 27(4):439–463, 2007. doi:10.1007/S00493-007-2160-5.
- 21 Minghui Liu and Gábor Pataki. Exact duals and short certificates of infeasibility and weak infeasibility in conic linear programming. *Mathematical Programming*, 167:435–480, 2018. doi:10.1007/S10107-017-1136-5.
- 22 Bruno F Lourenço and Gábor Pataki. A simplified treatment of ramana’s exact dual for semidefinite programming. *Optimization Letters*, 17(2):219–243, 2023. doi:10.1007/S11590-022-01898-2.
- 23 Gábor Pataki and Aleksandr Touzov. How do exponential size solutions arise in semidefinite programming? *SIAM Journal on Optimization*, 34(1):977–1005, 2024. doi:10.1137/21M1434945.
- 24 Motakuri V. Ramana. An exact duality theory for semidefinite programming and its complexity implications. *Mathematical Programming*, 77:129–162, 1997. doi:10.1007/BF02614433.
- 25 Alexander A Razborov and Steven Rudich. Natural proofs. *Journal of Computer and System Sciences*, 1(55):24–35, 1997. doi:10.1006/JCSS.1997.1494.
- 26 Steven Rudich. Super-bits, demi-bits, and np/qpoly-natural proofs. In *Proceedings of RANDOM/APPROX*, 1997.
- 27 Adi Shamir. Factoring numbers in $o(\log n)$ arithmetic steps. *Information Processing Letters*, 8(1):28–31, 1979. doi:10.1016/0020-0190(79)90087-5.
- 28 Mateusz Skomra. *Spectraèdres tropicaux: application à la programmation semi-définie et aux jeux à paiement moyen*. PhD thesis, Université Paris Saclay (COmUE), 2018.
- 29 Bengt EY Svensson. Quantum weak values and logic: an uneasy couple. *Foundations of Physics*, 47(3):430–452, 2017.
- 30 Sergey P Tarasov and Mikhail N Vyalyi. Semidefinite programming and arithmetic circuit evaluation. *Discrete Applied Mathematics*, 156(11):2070–2078, 2008. doi:10.1016/J.DAM.2007.04.023.
- 31 Dave Touchette. Quantum information complexity. In *Proceedings of STOC*, 2015.