Electronic Colloquium on Computational Complexity, Report No. 35 (2020)

No-Signaling MIPs and Faster-Than-Light Communication, Revisited

Justin Holmgren*

Abstract

We revisit one original motivation for the study of no-signaling multi-prover interactive proofs (NS-MIPs): specifically, that two spatially separated provers are guaranteed to be no-signaling by the physical principle that information cannot travel from one point to another faster than light.

We observe that with more than two provers, the physical connection between no-signaling and faster-than-light communication is more nuanced, depending on the relative positioning of the provers. In particular, we observe that provers are guaranteed to be no-signaling if and only if their positions are convexly independent. Other prover positionings provide weaker guaranteees.

We consider a new issue that thus arises only in the many-prover setting: *how precisely must provers be positioned in order to guarantee the soundness of a (NS-)MIP?* We prove that substantially more precision is required to guarantee full no-signaling than to guarantee soundness of a specific NS-MIP for PSPACE implicit in the work of Kalai and Raz (CRYPTO 2009).

*Simons Institute. Email: holmgren@alum.mit.edu.

ISSN 1433-8092

1 Introduction

The notion of no-signaling (but potentially nonlocal) behaviour by multiple spatially separated parties has proven to be a valuable concept in the study of quantum information, as well as in cryptography. If k players (probabilistically) map inputs $q_1, \ldots, q_k$ to outputs $a_1, \ldots, a_k$, they are said to be no-signaling if the *distribution* of $(a_i)_{i \in S}$ for any set $S \subseteq [k]$ depends only on $(q_i)_{i \in S}$. It turns out that this is a more general notion than local behaviour, in which functions $(f_1, \ldots, f_k)$ are jointly sampled independently of $q_1, \ldots, q_k$, and then each a_i is computed as $a_i = f_i(q_i)$. One example of a no-signaling but nonlocal behavior, due to Clauser et al. [CHSH69], maps $q_1, q_2 \in \{0,1\}$ to $a_1, a_2 \in \{0,1\}$ sampled uniformly at random such that $a_1 \oplus a_2 = q_1 \wedge q_2$.

We are primarily interested in multi-prover interactive proofs (MIPs) [BGKW88] that retain soundness even against malicious provers that may use an arbitrary no-signaling strategy. Kalai, Raz, and Rothblum [KRR14] showed that with $k > 2$ provers — in fact $k = n^c$ provers for some constant $c > 0$ — it is possible to construct a k-prover one-round MIP for EXP with soundness against all no-signaling strategies. They were motivated by the fact that under standard cryptographic assumptions, such an MIP can be compiled into a single prover one-round argument for EXP, resolving a long-standing open problem.

Using more than two provers was essential in the result of [KRR14], as Ito [Ito10] has shown that two-prover MIPs with soundness against no-signaling strategies are limited to PSPACE. Subsequently, Holden and Kalai [HK20] have proved that in fact at least $\omega(\sqrt{\log n})$ provers are necessary for any such MIP beyond PSPACE (with negligible soundness error).

One additional folklore motivation (mentioned explicitly by e.g. [KRR14, HK20]) for studying no-signaling MIPs is that

> "*...the principle that information cannot travel faster than light is a central principle in physics, and is likely to remain valid in any future ultimate theory of nature, since its violation means that information could be sent from future to past. Therefore, soundness against no-signaling strategies is likely to ensure soundness against provers that obey a future ultimate theory of physics, and not only the current physical theories that we have, that are known to be incomplete.*" — Kalai, Raz, and Rothblum [KRR14].

In this work, we observe that $\geq$ 3-prover no-signaling is *not* necessarily implied by the principle that information cannot travel between two points faster than light. We then investigate what *is* implied by said principle.

For the rest of this introduction, we abstract out the speed of light by defining a metric on the set of possible prover positions in which the distance between any two points x and y is the minimum time required to transmit information from x to y, which we shall denote by $t(x \rightsquigarrow y)$. If X and Y are sets, we write $t(X \rightsquigarrow Y)$ to denote $\max_{x \in X, y \in Y} t(x \rightsquigarrow y)$. We will exclusively consider Euclidean space $\mathbb{R}^3$, but we allow for the possibility that other metrics may arise.

1.1 Our Results and Techniques

We first recall (Physical Theorem 3.6) a physical "reduction" that uses signaling provers to construct a faster-than-light communicator, and point out that it inherently relies on a condition about the positions of the signaling provers. In particular, suppose that provers $\{P_i\}_{i \in T}$, located at positions $\{x_i\}_{i \in T}$, produce outputs that signal about the inputs to provers $\{P_i\}_{i \in S}$, located at positions $\{x_i\}_{i \in S}$ (we call this $(S \to T)$-signaling). If there are positions a, b such that

$$t(a \rightsquigarrow \{x_i\}_{i \in S}) + t(\{x_i\}_{i \in T} \rightsquigarrow b) < t(a \rightsquigarrow b), \tag{1}$$

then the reduction uses such provers to construct a device that communicates from a to b in time less than $t(a \rightsquigarrow b)$. If Eq. (1) is satisfied, we say that $\{x_i\}_{i \in S}$ and $\{x_i\}_{i \in T}$ are separated.

Physical Theorem 3.6 naturally leads to a generalization of the notion of no-signaling that we call metric no-signaling: For any positions $(x_1, \ldots, x_k)$, we say that P is $(x_1, \ldots, x_k)$-metric no-signaling if for all $S, T \subseteq [k]$

with $\{x_i\}_{i\in S}$ separated from $\{x_i\}_{i\in T}$, P is $(S \to T)$-no-signaling. Our physical reduction implies (under the assumption that point-to-point faster-than-light communication is impossible) that if a collection P of provers is located at positions $x_1, \ldots, x_k$, then P must be $(x_1, \ldots, x_k)$-metric no-signaling.

We show that if $x_1, \ldots, x_k$ are convexly independent, then any $(x_1, \ldots, x_k)$-metric no-signaling strategy is in fact no-signaling. Indeed, if P were $(S \to T)$-signaling for some (disjoint) S, T, then a simple hybrid argument implies that P is $(\{i\} \to T)$-signaling for some $i \in S$. On the other hand, the convex independence of $x_1, \ldots, x_k$ implies that x_i is separated from x_T, which contradicts that P is $(x_1, \ldots, x_k)$-metric no-signaling.

1.1.1 Robust Configurations

We would like MIPs that are not only sound against metric no-signaling strategies, but are "robustly" so. Formally, for a metric space (X, d), we say that a configuration $\vec{x}' = (x'_1, \ldots, x'_k) \in X^k$ is an ϵ-perturbation of $\vec{x} = (x_1, \ldots, x_k)$ if $x'_i = \rho(x_i)$ for some $\rho : X \to X$ such that for all $a, b \in X$, it holds that $e^{-\epsilon} \cdot d(a, b) \leq d\big(\rho(a), \rho(b)\big) \leq e^{\epsilon} \cdot d(a, b)$. When an MIP is sound against $\vec{x}'$-metric no-signaling strategies for all ϵ-perturbations $\vec{x}'$ of some fixed $\vec{x}$, we say that that the MIP is ϵ-robustly metric no-signaling sound.

One might hope that for some positions $(x_1, \ldots, x_k)$, it in fact holds for all ϵ-perturbations $(x'_1, \ldots, x'_k)$ that the notion of $(x'_1, \ldots, x'_k)$-metric no-signaling is equivalent to (full) no-signaling. Indeed, this would imply that any no-signaling MIP is in fact ϵ-robustly metric no-signaling sound. Unfortunately, we show that such positions cannot exist unless $\epsilon \leq \frac{1}{\Omega(k^{1/3})}$. Loosely speaking, Descarte's theorem on total angular defect (a 3-dimensional analogue of the fact that the exterior angles of a convex polygon sum to 2π) states that some vertex of a convex polyhedron has small "curvature". Unlike in two dimensions, a vertex v having small curvature does *not* mean that v has neighboring vertices a and b such that the angle $\angle avb$ is large. However, if this is not the case we show that it is possible to flatten the polyhedron at v with only a mild deformation.

Finally, we show that a no-signaling MIP for PSPACE (with a super-constant number of provers) implicit in a work of Kalai and Raz [KR09] is $\Omega(1)$-robustly metric no-signaling sound. Intuitively, in this MIP each prover corresponds to a round of a public-coin interactive proof, and soundness only relies on the provers corresponding to rounds $1, \ldots, i$ not signaling about the inputs to prover j for any $j > i$. To guarantee this, it suffices to arrange the provers sequentially along a line.

So as not to oversell this result, we mention that existing 2-prover NS-MIPs for PSPACE [IKM08] are automatically ϵ-robustly metric no-signaling sound *for all* ϵ. It remains an interesting open question whether there is any robustly metric no-signaling sound MIP beyond PSPACE.

2 Preliminaries

Definition 2.1. A metric space is a pair (X, d), where X is a set, and $d : X \times X \to \mathbb{R}_{\geq 0}$ is a function (called a metric on X) such that for all $x, y, z \in X$:

1. $d(x, y) = 0$ if and only if $x = y$.
2. $d(x, y) = d(y, x)$
3. $d(x, z) \leq d(x, y) + d(y, z)$.

Definition 2.2. A strategy $P : \mathcal{Q}^k \xrightarrow{\$} \mathcal{A}^k$ is said to be no-signaling if for all $q, q' \in \mathcal{Q}^k$ and all $S \subseteq [k]$ with $q_S = q'_S$, it holds that $P(q)_S$ and $P(q')_S$ are identically distributed.

2.1 Geometry

We use some facts and notions from 3D Euclidean and spherical geometry. Rather than reproduce them here, we refer the reader to any textbook, and simply state the notation we will use here.

If a, b, and c, are distinct points, we write $\overrightarrow{ab}$ to denote the ray with endpoint a that passes through b, we write $\overleftrightarrow{ab}$ to denote the line passing through a and b, and we write $\overline{ab}$ to denote the line segment with endpoints a and b. We write $\angle abc$ to refer to the angle with vertex b enclosed by the rays $\overrightarrow{ba}$ and $\overrightarrow{bc}$. We write $m\angle abc$ to denote the measure (in radians) of $\angle abc$. We write Δabc to denote the triangle with sides $\overline{ab}$, $\overline{bc}$, and $\overline{ca}$. We use the same notation for spherical lines, line segments, angles, and triangles.

2.2 Games and Proofs

The rest of Section 2 is taken verbatim (with author permission) from [Hol19].

It will be convenient for us to consider separately from interactive proofs (which are associated with a language $\mathcal{L}$, involve an input x, and have completeness / soundness properties depending on whether $x \in \mathcal{L}$) a notion of an interactive game, which has no input.

We think of an interactive game as something that is played by a single player in r rounds. At the beginning of the i^{th} round, the player must specify a message $\alpha_i \in \{0,1\}^*$. Then, a message β_i is sampled uniformly from $\{0,1\}^{\ell_i}$ for some ℓ_i that is pre-specified independently of any of the player's choices. At the end of the r^{th} round, a predicate W is applied to $(\alpha_1, \beta_1, \ldots, \alpha_r, \beta_r)$ to determine whether the player wins.

More formally:

Definition 2.3 (Interactive Game)**.** An (r-round) public-coin interactive game is a tuple $(\ell_1, \ldots, \ell_r, W)$, where each $\ell_i \in \mathbb{Z}^+$ and $W \subseteq \{0,1\}^*$ is an "acceptance" set. A strategy is a function $s : \{0,1\}^* \to \{0,1\}^*$.

If $\mathcal{G} = (\ell_1, \ldots, \ell_r, W)$ is a public-coin interactive game and s is a strategy, then the value of $\mathcal{G}$ with respect to s (alternatively the probability with which s wins $\mathcal{G}$) is

$$v[s](\mathcal{G}) \stackrel{\text{def}}{=} \Pr_{\substack{\beta_1 \leftarrow \{0,1\}^{\ell_1} \\ \cdots \\ \beta_r \leftarrow \{0,1\}^{\ell_r}}} \left[(\alpha_1, \beta_1, \ldots, \alpha_r, \beta_r) \in W \right],$$

where each α_i is defined to be $s(\beta_1, \ldots, \beta_{i-1})$. The value of $\mathcal{G}$, denoted $v(\mathcal{G})$, is $\sup_s v[s](\mathcal{G})$.

Definition 2.4 (Interactive Proof)**.** An ($r(\cdot)$-round) public-coin interactive proof for a language $\mathcal{L}$ with soundness error $\epsilon(\cdot)$ is a pair (P, V), where V is a polynomial-time algorithm mapping any string $x \in \{0,1\}^*$ to an $r(|x|)$-round single-player game with the following properties:

- (Completeness) If $x \in \mathcal{L}$, then $P(x)$ is a strategy that wins $V(x)$ with probability 1.
- (Soundness) If $x \notin \mathcal{L}$, then *all* strategies P^* win $V(x)$ with probability at most $\epsilon(|x|)$.

The interactive proof is said to be public-coin if each $V(x)$ is public-coin.

Definition 2.5 (Game Transcript)**.** If $\mathcal{G} = (\ell_1, \ldots, \ell_r, W)$ is a public-coin interactive game, then a (complete) transcript for $\mathcal{G}$ is $\alpha_1|\beta_1|\cdots|\alpha_r|\beta_r$ with each $\beta_i \in \{0,1\}^{\ell_i}$ and $\alpha_i \in \{0,1\}^*$. An accepting transcript is one that is contained in W. A transcript prefix is any $\alpha_1|\beta_1|\cdots|\alpha_i|\beta_i$ for $i \in \{0, \ldots, r\}$.

Definition 2.6 (Game Suffix)**.** If $\mathcal{G} = (\ell_1, \ldots, \ell_r, W)$ is an r-round public-coin interactive game and $\alpha_1|\beta_1|\cdots|\alpha_i|\beta_i$ is a transcript prefix for $\mathcal{G}$, we denote by $\mathcal{G}|_\tau$ the game $(\ell_{i+1}, \ldots, \ell_r, W|_\tau)$, where $W|_\tau$ is the set of strings of the form $\alpha_{i+1}|\beta_{i+1}|\cdots|\alpha_r|\beta_r$ for which $\alpha_1|\beta_1|\cdots|\alpha_r|\beta_r \in W$.

We refer to $\mathcal{G}|_\tau$ as the suffix of $\mathcal{G}$ following τ.

3 Physical Axioms and Metric No-Signaling

In this section, we aim to relate the purely mathematical notion of no-signaling to physical theories of reality. In so doing, we necessarily sacrifice some definitional rigor. Still, wherever possible we strive as much as possible to work with well-defined purely mathematical notions.

Physics Axiom 3.1. *Space consists of a set of positions X, endowed with a metric d, such that for all $a, b \in X$, it is possible to transmit any message from a to b in time $d(a, b)$, and it is impossible to transmit any information from a to b faster than this.*

Example 3.2. In special relativity, the space of all possible positions is $\mathbb{R}^3$ with the usual Euclidean metric (up to some constant that depends on the chosen units and the speed of light).

Definition 3.3. Say that a strategy $P : \mathcal{Q}^k \xrightarrow{\$} \mathcal{A}^k$ is $(S \to T)$-no-signaling if for all $q, q' \in \mathcal{Q}^k$ satisfying $q_i = q'_i$ for all $i \notin S$, it holds that $P(q)_T$ and $P(q')_T$ are identically distributed.

We observe that the standard notion of no-signaling is equivalent to requiring $(S \to T)$-no-signaling for all disjoint S, T. We show that Physics Axiom 3.1 implies $(S \to T)$-no-signaling for a different set of (S, T) that depends on the positions of the provers.

Definition 3.4. For any metric space (X, d), we say that non-empty subsets $S, T \subseteq X$ are separated if there exist $s, t \in X$ such that

$$\sup_{x \in S} d(s, x) + \sup_{x \in T} d(x, t) < d(s, t).$$

One reason that Definition 3.4 is natural is because of Proposition 4.4, which states that in $\mathbb{R}^n$ our notion of separation is equivalent to hyperplane separation. A basic fact about separated sets is that they are always disjoint.

Definition 3.5. Say that a strategy $P : \mathcal{Q}^k \xrightarrow{\$} \mathcal{A}^k$ is $(x_1, \ldots, x_k)$-metric no-signaling if P is $(S \to T)$-no-signaling for all $S, T \subseteq [k]$ for which $\{x_i\}_{i \in S}$ and $\{x_i\}_{i \in T}$ are separated.

Physical Theorem 3.6. *Let D be a device that at a pre-specified time (say $t = 0$) takes inputs $q_1, \ldots, q_k \in \mathcal{Q}$ at respective positions $x_1, \ldots, x_k$, and at the same positions produces outputs $a_1, \ldots, a_k \in \mathcal{A}$. Then either D implements a $(x_1, \ldots, x_k)$-metric no-signaling strategy, or it can be used to build a faster-than-light communicator.*

Proof. Suppose that D does not implement a $(x_1, \ldots, x_k)$-metric no-signaling strategy. Then there exist sets $S, T \subseteq [k]$ such that:

- $\{x_i\}_{i \in S}$ and $\{x_i\}_{i \in T}$ are separated; that is, there exist $s, t \in X$ such that

$$\max_{i \in S} d(s, x_i) + \max_{i \in T} d\big(x_i, t\big) < d(s, t).$$

- There exist $q^{(0)}, q^{(1)} \in \mathcal{Q}^k$ such that:
  - $q_i^{(0)} = q_i^{(1)}$ for all $i \notin S$, but
  - $D(q^{(0)})_T$ and $D(q^{(1)})_T$ have statistical distance $\epsilon > 0$.

Imagine placing a device at each position x_i that acts as follows. If $i \in S$, it listens for a bit b at time 0, and immediately provides $q_i^{(b)}$ as the i^{th} input to D. If $i \notin S$, then our device simply provides the fixed value $q_i^{(0)} = q_i^{(1)}$ as the i^{th} input to D. If $i \in T$, then upon receiving an output a_i at x_i from D, our device transmits a_i to position t.

Suppose that Alice receives a random bit $b \in \{0, 1\}$ at time $-\max_{i \in S} d(s, x_i)$ and then, for each $i \in S$, Alice transmits b at time $-d(s, x_i)$ to position x_i. For each $i \in T$, Bob receives a message a_i from position x_i at time $d(x_i, t)$. By time $\max_{i \in T} d(x_i, t)$, Bob has received such messages $(a_i)_{i \in T}$. If $(a_i)_{i \in T}$ has higher probability in $D(q^{(0)})_T$ than in $D(q^{(1)})_T$, then Bob guesses that $b = 0$; otherwise, Bob guesses that $b = 1$. Clearly Bob guesses correctly with probability $\frac{1+\epsilon}{2}$. On the other hand, the elapsed time between when Alice receives b (at position s) and when Bob guesses b (at position t) is $\max_{i \in S} d(s, x_i) + \max_{i \in T} d(x_i, t)$, which is less than $d(s, t)$, which is a contradiction. □

4 No-Signaling vs. Metric No-Signaling

In this section, we show that it is possible to arrange provers in Euclidean space so that metric no-signaling implies the standard notion of no-signaling. Conversely, we show that with some (non-degenerate) arrangements, metric no-signaling is a strictly weaker requirement than no-signaling.

Theorem 4.1. *If $x_1, \ldots, x_k \in \mathbb{R}^n$ are convexly independent, then any $(x_1, \ldots, x_k)$-metric no-signaling strategy $P : \mathcal{Q}^k \xrightarrow{\$} \mathcal{A}^k$ is no-signaling.*

Proof. This follows as a simple corollary of Propositions 4.3 and 4.4 below. Specifically, by the definition of convex independence, no x_i is contained in the convex hull of $\{x_j\}_{j \neq i}$. Thus there is a hyperplane separating x_i from $\{x_j\}_{j \neq i}$. By Proposition 4.4 and the definition of metric no-signaling, P must be $(\{i\} \to [k] \setminus \{i\})$-no-signaling. Since this holds for all i, P must be no-signaling by Proposition 4.3. □

Corollary 4.2. *For any k, there exist $x_1, \ldots, x_k \in \mathbb{R}^n$ such that $(x_1, \ldots, x_k)$-metric no-signaling is equivalent to no-signaling.*

Proof. This follows from the fact that there exist polytopes with arbitrarily many vertices. □

Proposition 4.3. *$P : \mathcal{Q}^k \xrightarrow{\$} \mathcal{A}^k$ is no-signaling if and only if it is $(\{i\} \to [k] \setminus \{i\})$-no-signaling for al $i \in [k]$.*

Proof. By definition, P is no-signaling if and only if it is $(S \to T)$-no-signaling for all disjoint $S, T \subseteq [k]$. By a simple hybrid argument, this is equivalent to P being $(\{i\} \to T)$-no-signaling for all i, T satisfying $i \notin T$. In turn, this is clearly equivalent to P being $(\{i\} \to [k] \setminus \{i\})$-no-signaling for all $i \in [k]$. □

Proposition 4.4. *In $\mathbb{R}^n$, finite sets S and T are separated if and only if there is a hyperplane H separating S from T.*

Proof. Suppose that S and T are separated. That is, there exist points s and t such that $\max_{p \in S} d(s,p) + \max_{p \in T} d(p,t) < d(s,t)$. Then S is contained in a ball of radius $\max_{p \in S} d(s,p)$ centered at s, and T is contained in a ball of radius $\max_{p \in T} d(p,t)$ centered at t. These two balls are convex and disjoint, and thus the "only if" direction follows from the well-known hyperplane separation theorem.

For the "if" direction, suppose that S and T are separated by a hyperplane $H = \{p : u \cdot p = 0\}$ for a unit vector u (by changing coordinates we may assume without loss of generality that H passes through the origin). Let $\epsilon > 0$ be such that $u \cdot p \leq -\epsilon$ for all $p \in S$ and $u \cdot p \geq \epsilon$ for all $p \in T$.

By the Pythagorean theorem, it holds for all $p \in S$ that

$$\lim_{\lambda \to \infty} \lambda - d(-\lambda u, p) = d(p, H) \geq \epsilon.$$

Similarly for $p \in T$,

$$\lim_{\lambda \to \infty} \lambda - d(\lambda u, p) = d(p, H) \geq \epsilon.$$

Thus for sufficiently large λ, it holds for all $p \in S$ that $d(-\lambda u, p) \leq \lambda - \epsilon/2$, and for all $p \in T$ that $d(\lambda u, p) \leq \lambda - \epsilon/2$. Thus with $s = -\lambda u$ and $t = \lambda u$,

$$\max_{p \in S} d(s,p) + \max_{p \in T} d(p,t) \leq 2\lambda - \epsilon < 2\lambda = d(s,t).$$ □

Theorem 4.5. *For any $x_1, \ldots, x_k \in \mathbb{R}^n$ that are not convexly independent, there exists a $(x_1, \ldots, x_k)$-metric no-signaling strategy that is not no-signaling.*

Proof. There must exist some i such that x_i is in the convex hull of $\{x_j\}_{j \neq i}$. Without loss of generality, suppose that $i = k$. Let $P : \{0,1\}^k \xrightarrow{\$} \{0,1\}^k$ denote a strategy that on input $(q_1, \ldots, q_k)$ outputs random $(a_1, \ldots, a_k)$ such that $a_1 \oplus \cdots \oplus a_{k-1} = q_k$.

We first show that P is $(x_1, \ldots, x_k)$-metric no-signaling. Because of how we constructed the output distribution of P, it can only possibly fail to be $(S \to T)$-no-signaling when T contains $\{1, \ldots, k-1\}$. But we only require $(S \to T)$-no-signaling when S and T are separated (and in particular disjoint and no-empty). The only candidate is $S = \{k\}$ and $T = \{1, \ldots, k\}$, but these are not separated (x_k is contained in the convex hull of $\{x_1, \ldots, x_k\}$).

On the other hand, P is not no-signaling because the distribution of $(a_1, \ldots, a_{k-1})$ clearly depends on q_k. □

5 Robustness and Boundedness in Metric No-Signaling

In real life, it is impossible to ensure that provers are positioned *exactly* as specified. We consider perturbations of prover positions that may modify the distance between any two points by at most a small multiplicative factor.

Definition 5.1. An ϵ-isometry of a metric space (X, d) is a function $\rho : X \to X$ such that for all $x, y \in X$, it holds that $e^{-\epsilon} \cdot d(x, y) \leq d\big(\rho(x), \rho(y)\big) \leq e^{\epsilon} \cdot d(x, y)$.

Definition 5.2. A class of k-prover strategies $\mathcal{S}$ is said to be ϵ-robustly instantiable via metric no-signaling (in a metric space (X, d)) if there exist positions $x_1, \ldots, x_k \in X$ such that for all ϵ-isometries ρ, it holds that any $\big(\rho(x_1), \ldots, \rho(x_k)\big)$-metric no-signaling strategy is contained in $\mathcal{S}$.

Theorem 5.3. *There is a constant $c > 0$ such that the class of k-prover no-signaling strategies is at most $O(k^{-1/3})$-robustly instantiable via metric no-signaling (in $\mathbb{R}^3$).*

Proof. By Theorem 4.5, it suffices to show that for all points $\vec{v}_1, \ldots, \vec{v}_k \in \mathbb{R}^3$, there exists an ϵ-isometry $\rho : \mathbb{R}^3 \to \mathbb{R}^3$ such that $\rho(\vec{v}_1), \ldots, \rho(\vec{v}_k)$ are not convexly independent. If $\vec{v}_1, \ldots, \vec{v}_k$ are not already vertices of a convex polyhedron K, we are done.

By Descartes' theorem on total angular defect, there exists a vertex $\vec{v}^*$ of this polyhedron with angular defect at most $\frac{4\pi}{k}$. Without loss of generality say that $\vec{v}^*$ is at the origin, and let $\vec{v}_1, \ldots, \vec{v}_d$ be its neighboring vertices.

For $i \in [d]$, define the plane $P_i = \mathsf{span}(\vec{v}_i, \vec{v}_j)$ for $j \in [d]$ such that $j \equiv i + 1 \pmod{d}$. Let $\hat{n}^{(i)}$ denote the unit normal vector to P_i, such that K is contained in the half-space $H_i \stackrel{\mathsf{def}}{=} \{\vec{x} : \vec{x} \cdot \hat{n}^{(i)} \leq 0\}$.

Say that a plane P through the origin, defined by orthogonality to a unit normal vector $\hat{n}$, is an ϵ-good projection plane (for some $\epsilon > 0$ to be determined later) if $\hat{n} \cdot \hat{n}^{(i)} \geq 1 - \epsilon$ for each $i \in [d]$, and if $\vec{v}^*$ lies in the convex hull of the orthogonal projections of $\vec{v}_1, \ldots, \vec{v}_d$ onto P.

Case 1: There exists a $(10\pi k)^{-2/3}$-good projection plane Let $P = \{\vec{x} : \vec{x} \cdot \hat{n} = 0\}$ be an ϵ-good projection plane for $\epsilon = (10\pi k)^{-2/3}$.

We now work in a coordinate system in which $\hat{n}$ is the azimuthal (z-) axis, $\vec{v}^*$ is the origin, and the x- and y- axes are arbitrary. For $i \in [d]$, let (r_i, θ_i, z_i) denote the cylindrical coordinates of $\vec{v}_i$. We define a mapping $\rho : \mathbb{R}^3 \to \mathbb{R}^3$ in these cylindrical coordinates as follows. On input (r, θ, z), let $i, j \in [d]$ and $\lambda \in [0, 1]$ be such that $\theta \equiv \lambda\theta_i + (1 - \lambda)\theta_j \pmod{2\pi}$ and $j \equiv i + 1 \pmod{d}$ (this uniquely determines i, j, and λ). Then we define

$$\rho(r, \theta, z) = \Big(r, \theta, z - r \cdot \big(\lambda \frac{z_i}{r_i} + (1 - \lambda)\frac{z_j}{r_j}\big)\Big).$$

By construction, ρ acts as orthogonal projection onto P when restricted to the points $\vec{v}_1, \ldots, \vec{v}_d$. It remains to show that ρ is an approximate isometry. We first establish the following claim about the cylindrical coordinates of $\{\vec{v}_i\}$.

Claim 5.4. *For all $i, j \in [d]$,*

$$\left|\frac{z_i}{r_i}\right| \leq O(k^{-1/3}) \tag{2}$$

and

$$\left|\frac{z_i}{r_i} - \frac{z_j}{r_j}\right| \le |\theta_i - \theta_j| \cdot O(k^{-1/3}). \tag{3}$$

Proof. The point $\vec{v}_i$ lies on a plane P_i defined by orthogonality to a unit vector $\hat{n}^{(i)}$ that is close to $\hat{n}$. In Cartesian coordinates, $\hat{n}^{(i)} = (n_x^{(i)}, n_y^{(i)}, n_z^{(i)})^\top$ for $n_z^{(i)} \ge 1 - \epsilon$ (and thus $\|(n_x^{(i)}, n_y^{(i)})^\top\|_2 \le \sqrt{2\epsilon}$). In cylindrical coordinates, P_i can be described as

$$P_i = \left\{(r, \theta, z) : r \cdot \left(n_x^{(i)} \cos\theta + n_y^{(i)} \sin\theta\right) + n_z^{(i)} \cdot z = 0\right\}. \tag{4}$$

This implies that for all points $(r, \theta, z) \in P_i$ with $r \neq 0$, we have $\frac{z}{r} = -\frac{n_x^{(i)} \cos\theta + n_y^{(i)} \sin\theta}{n_z^{(i)}}$, so

$$\left|\frac{z}{r}\right| \le \frac{\|(n_x^{(i)}, n_y^{(i)})^\top\|_2}{n_z^{(i)}} \le \frac{\sqrt{2\epsilon}}{1-\epsilon} \le O(\sqrt{\epsilon}) = O(k^{-1/3}).$$

Eq. (3) follows from Eq. (4) and the fact that $n_x^{(i)} \cos\theta + n_y^{(i)} \sin\theta$ is an $O(k^{-1/3})$-Lipschitz function of θ. □

To see that ρ is an $O(k^{-1/3})$-isometry, it suffices to show that $\rho(\vec{x}) - \vec{x}$ is $O(k^{-1/3})$-Lipschitz. In cylindrical coordinates, ρ always applies a displacement in the z-direction that depends only on (r, θ), so it suffices to show that the displacement is $O(k^{-1/3})$-Lipschitz. By a hybrid argument, it suffices to show that $\rho(\cdot, \theta, z)$ and $\rho(r, \cdot, z)$ are $O(k^{-1/3})$-Lipschitz for fixed (r, θ, z). This follows from Claim 5.4.

Case 2: There is no $(10\pi k)^{-2/3}$-good projection plane In this case, the polyhedron K is contained in the intersection of two half-spaces $H = \{\vec{x} : \vec{x} \cdot \hat{n} \le 0\}$ and $H' = \{\vec{x} : \vec{x} \cdot \hat{n}' \le 0\}$ whose bounding planes P and P' form a dihedral angle of α with $\cos(\alpha) \ge -1 + \epsilon$ for $\epsilon = (10\pi k)^{-2/3}$. Equivalently, $\hat{n} \cdot \hat{n}' \le 1 - \epsilon$.

Claim 5.5. *There exist $i, j \in [d]$ such that $m\angle\vec{v}_i\vec{v}^*\vec{v}_j \ge \pi - O(k^{-1/3})$.*

Proof. Let $\hat{a}$ and $-\hat{a}$ be two opposite unit vectors in $P \cap P'$. We will show the existence of $i, j \in [d]$ such that

$$m\angle\hat{a}\vec{v}^*\vec{v}_i \le \varphi \tag{5}$$

and

$$m\angle(-\hat{a})\vec{v}^*\vec{v}_j \le \varphi, \tag{6}$$

where $\varphi = \frac{10\pi}{\epsilon k} = O(k^{-1/3})$. Suppose for contradiction that Eqs. (5) and (6) cannot both be satisfied. Without loss of generality, suppose that Eq. (6) that cannot be satisfied.

We consider the radial projection of K onto a unit sphere centered at $\vec{v}^*$. The fact that $\vec{v}^*$ has small angular defect means that the projection is a (convex) spherical polygon with perimeter at least $2\pi - \frac{4\pi}{k}$, whose corners are the projections of $\vec{v}_1, \ldots, \vec{v}_d$.

We will obtain a contradiction by bounding the projection of K within a shape of perimeter less than $2\pi - \frac{4\pi}{k}$. The non-existence of a vertex $\vec{v}_i$ with $m\angle(-\hat{a})\vec{v}^*\vec{v}_i \le \varphi$ means that the projection of K avoids a spherical cap of polar angle φ centered at $-\hat{a}$. Let P'' be a plane passing through $\vec{v}^*$ such that the lines $P'' \cap P$ and $P'' \cap P'$ each intersect $P \cap P'$ at an angle of φ. Let $\hat{b}$ and $\hat{c}$ be (the unique) unit vectors in $P'' \cap P \cap H'$ and $P'' \cap P' \cap H$, respectively. The projection of K then is contained in the spherical triangle $\Delta\hat{a}\hat{b}\hat{c}$.

$\Delta\hat{a}\hat{b}\hat{c}$ is obtained from the perimeter-2π spherical quadrilateral with vertices $\hat{a}$, $\hat{b}$, $-\hat{a}$, and $\hat{c}$ by "short-cutting" from $\hat{b}$ to $\hat{c}$. Thus $\Delta\hat{a}\hat{b}\hat{c}$ has perimeter

$$2\pi - m\angle\hat{b}\vec{v}^*(-\hat{a}) - m\angle(-\hat{a})\vec{v}^*\hat{c} + m\angle\hat{c}\vec{v}^*\hat{b} = 2\pi - 2\varphi + m\angle\hat{c}\vec{v}^*\hat{b}.$$

Letting γ denote $m\angle\hat{c}\vec{v}^*\hat{b}$, the spherical law of cosines gives

$$\begin{aligned}\cos\gamma &= \cos^2\varphi + \sin^2\varphi\cdot\cos\alpha \\ &= \frac{1+\cos(2\varphi)}{2} + \cos\alpha\cdot\frac{1-\cos(2\varphi)}{2} \\ &= \frac{1+\cos\alpha}{2} + \cos(2\varphi)\cdot\frac{1-\cos\alpha}{2} \\ &= \cos^2\left(\frac{\alpha}{2}\right) + \cos(2\varphi)\cdot\sin^2\left(\frac{\alpha}{2}\right) \\ &\geq \frac{\epsilon}{2} + (1-\frac{\epsilon}{2})\cdot\cos(2\varphi).\end{aligned}$$

Using the approximation $1-\frac{\theta^2}{2} \leq \cos(\theta) \leq 1-\frac{\theta^2}{2}+\frac{\theta^2}{24}$, we obtain

$$\begin{aligned}\frac{\gamma^4}{24} - \frac{\gamma^2}{2} &\geq -(1-\epsilon/2)\cdot 2\varphi^2. \\ \left(\frac{\gamma^2}{\sqrt{24}} - \frac{\sqrt{24}}{4}\right)^2 &\geq \frac{3}{2} - (1-\epsilon/2)\cdot 2\varphi^2 \\ \left|\frac{\gamma^2}{\sqrt{24}} - \frac{\sqrt{24}}{4}\right| &\geq \sqrt{\frac{3}{2} - (1-\epsilon/2)\cdot 2\varphi^2} \\ \frac{\gamma^2}{\sqrt{24}} &\leq \frac{\sqrt{24}}{4} - \sqrt{\frac{3}{2} - (1-\epsilon/2)\cdot 2\varphi^2},\end{aligned}$$

where in the last step we use the fact that γ is small. Thus,

$$\begin{aligned}\gamma^2 &\leq 6 - 6\cdot\sqrt{1-(1-\epsilon/2)\cdot\frac{4\varphi^2}{3}} \\ &\leq (1-\epsilon/2)\cdot 4\varphi^2\end{aligned}$$

and finally

$$\gamma \leq \sqrt{1-\epsilon/2}\cdot 2\varphi \leq 2\varphi - \frac{\varphi\epsilon}{2}.$$

The above calculation shows that the perimeter of $\Delta\hat{a}\hat{b}\hat{c}$ is at most $2\pi - \frac{\varphi\epsilon}{2}$. But $\frac{\varphi\epsilon}{2} > \frac{4\pi}{k}$, which is a contradiction.

An identical argument shows that there must exist a vertex $\vec{v}_j$ with $m\angle(-\hat{a})\vec{v}^*\vec{v}_j \leq \varphi$. □

Let i and j be as guaranteed to exist by Claim 5.5. Consider a 3-dimensional Cartesian coordinate system in which $\vec{v}^*$ is the origin, $\vec{v}_i$ lies on the negative x-axis and $\vec{v}_j$ lies in the upper half of the x-y plane. With θ denoting $\pi - m\angle\vec{v}_i\vec{v}^*\vec{v}_j$, we define a mapping ρ so that $\rho(x,y,z) = (x,y',z)$ where

$$y' = \begin{cases} y & \text{if } x \leq 0 \\ y - x\tan\theta & \text{otherwise.}\end{cases}$$

This is a $\tan(\theta)(= O(\epsilon'))$-isometry that maps $\angle\vec{v}_i\vec{v}^*\vec{v}_j$ into a straight line. □

In our feasibility results, we also consider the *boundedness* of a prescribed prover arrangement $(\vec{v}_1,\ldots,\vec{v}_k)$, normalized so that the minimum distance between any $\vec{v}_i$, $\vec{v}_j$ for $i \neq j$ is 1.

Definition 5.6. We say that a configuration $(\vec{v}_1,\ldots,\vec{v}_k)$ is B-bounded if

$$\frac{\max_{i,j} d(\vec{v}_i,\vec{v}_j)}{\min_{i\neq j} d(\vec{v}_i,\vec{v}_j)} \leq B.$$

5.1 Robustly Instantiating the Kalai-Raz NS-MIP

We now show that there do exist useful relaxations of no-signaling that are $\Omega(\frac{1}{\log n})$-robustly instantiable via metric no-signaling.

Recall the following multi-prover interactive proof $(P_{\mathsf{MIP}}, V_{\mathsf{MIP}})$ implicit in the work of Kalai and Raz [KR09]. Let $\Pi = (P_{\mathsf{IP}}, V_{\mathsf{IP}})$ be an r-round public-coin interactive proof with negligible soundness error $\epsilon = \epsilon(n)$, where each verifier message is of length $\ell = \ell(n)$. The MIP, which we will denote by $\mathsf{KR}[\Pi]$, works as follows: There are r provers. On input x, V_{MIP} samples $B_1, \ldots, B_r \leftarrow \{0,1\}^\ell$, and sends $B_{[i]} \stackrel{\mathsf{def}}{=} (B_1, \ldots, B_{i-1})$ to the i^{th} prover. The verifier accepts if the i^{th} prover returns A_i such that $(A_1, B_1, \ldots, A_r, B_r)$ is an accepting transcript for x in Π.

Theorem 5.7. *For any ϵ-sound r-round public-coin interactive proof Π, the MIP $\mathsf{KR}[\Pi]$ is ϵ-sound against all strategies that are $(\{j\} \to \{1, \ldots, i\})$-no-signaling for all $i < j$.*

Proof. Let $\mathsf{KR}[\Pi] = (P_{\mathsf{MIP}}, V_{\mathsf{MIP}})$. For any $x \in \{0,1\}^n \setminus \mathcal{L}$, consider a strategy P^* that wins in the game $V_{\mathsf{MIP}}(x)$ with probability greater than $\epsilon(n)$. That is, in the probability space defined by sampling $B_1, \ldots, B_r \leftarrow \{0,1\}^\ell$ and computing $(A_1, \ldots, A_r) \leftarrow P^*(B_{[0]}, \ldots, B_{[r-1]})$,

$$\Pr\left[(A_1, B_1, \ldots, A_r, B_r) \text{ is accepting}\right] > \epsilon(n).$$

Then in particular there must exist $i \in [r]$ such that

$$\mathbb{E}\left[v\left(V_{\mathsf{IP}}(x)|_{A_1|B_1|\cdots|A_i|B_i}\right)\right] > \mathbb{E}\left[v\left(V_{\mathsf{IP}}(x)|_{A_1|B_1|\cdots|A_{i-1}|B_{i-1}}\right)\right].$$

By definition of $v(V_{\mathsf{IP}}(x)|_{A_1|B_1|\cdots|A_{i-1}|B_{i-1}})$, the random variables B_i and $(A_1, B_1, \ldots, A_{i-1}, B_{i-1}, A_i)$ must not be independent. Thus there are values $b_i^{(0)}, b_i^{(1)}$ such that the distribution of $(A_1, B_1, \ldots, A_{i-1}, B_{i-1}, A_i)$ conditioned on $B_i = b_i^{(0)}$ is different than conditioned on $B_i = b_i^{(1)}$. This means that P^* is not $(\{i+1, \ldots, r\} \to \{1, \ldots, i\})$-no-signaling. Then, by a simple hybrid argument, there must exist some $j > i$ such that P^* is not $(\{j\} \to \{1, \ldots, i\})$-no-signaling. □

Motivated by Theorem 5.7, we make the following definition.

Definition 5.8. We say that a k-prover strategy P is KR-no-signaling if it is $(\{j\} \to \{1, \ldots, j-1\})$-no-signaling for all $j \in [k]$.

Theorem 5.9. *n-prover KR-no-signaling is $\Omega(\frac{1}{\log(n)})$-robustly n-boundedly instantiable via metric no-signaling in $\mathbb{R}^3$.*

Proof. For all $i \in \mathbb{Z}$, let x_i denote the point $(i, 0, 0)$. We will show that the arrangement $(x_1, \ldots, x_n)$ suffices.

Let $\epsilon < \frac{1}{\log(n)}$, let ρ be an arbitrary ϵ-isometry, and let $y_i = \rho(x_i)$ for all $i \in \mathbb{Z}$.

We begin by establishing lower bounds on the measures of the angles $\angle y_{i_1} y_{i_2} y_{i_3}$ and $\angle y_{i_2} y_{i_3} y_{i_4}$. Recall that for an angle $\angle ABC$, we write $m\angle ABC$ to denote the measure of $\angle ABC$ (which is between 0 and π).

Definition 5.10. For any $a, b \in \mathbb{Z}^+$, let $\theta_{a,b}$ denote

$$\sup_{\rho, i} \left[\pi - m\angle y_{i-a} y_i y_{i+b}\right],$$

where the supremum is over all ϵ-isometries ρ and all $i \in [n]$.

Lemma 5.11. *For all $a, b \in \mathbb{Z}^+$, it holds that $\theta_{a,b} \leq O\big(\sqrt{\epsilon} \cdot \log(a+b)\big)$.*

Before proving Lemma 5.11, we demonstrate that it suffices for Theorem 5.9.

Suppose that for some $j \in [n]$, y_j is in the convex hull of $\{y_1, \ldots, y_{j-1}\}$. By Carathéodory's theorem, there must exist $i_1 < i_2 < i_3 < i_4$ such that y_j is in the convex hull of $\{y_{i_1}, y_{i_2}, y_{i_3}, y_{i_4}\}$.

Thus in the tetrahedron with vertices y_{i_1}, y_{i_2}, y_{i_3}, y_{i_4}, it holds that y_{i_2}, y_{i_3}, and y_{i_4} each have angular defect at most 2ϵ. Thus by Descartes' theorem on total angular defect, y_{i_1} and y_{i_4} both have angular defect at least $2\pi - 2\epsilon$.

Replacing y_{i_1} by y_j does not increase the angular defect of y_{i_4}. Hence the angle $\angle y_{i_1} y_{i_4} y_j$ is at most ϵ, which is a contradiction.

Proof of Lemma 5.11. We begin with the special case in which $a = b$, in which we obtain the stronger bound $O(\sqrt{\epsilon})$ (independent of a).

Claim 5.12. *For any $a \in \mathbb{Z}^+$, $\theta_{a,a} \leq O(\sqrt{\epsilon})$.*

Proof. We know that for all i, $\|y_i - y_{i-a}\|$ and $\|y_{i+a} - y_i\|$ are at most $e^{\epsilon} \cdot a$, while $\|y_{i+a} - y_{i-a}\|$ is at least $2e^{-\epsilon} \cdot a$

By the law of cosines, we have

$$\begin{aligned}
\cos(m\angle y_{i-a} y_i y_{i+a}) &= \frac{\|y_i - y_{i-a}\|^2 + \|y_{i+a} - y_i\|^2 - \|y_{i+a} - y_{i-a}\|^2}{2 \cdot \|y_i - y_{i-a}\| \cdot \|y_{i+a} - y_i\|} \\
&\leq \frac{2e^{2\epsilon}a^2 - 4e^{-2\epsilon}a^2}{2e^{22\epsilon}a^2} \\
&= \frac{e^{2\epsilon} - 2e^{-2\epsilon}}{e^{2\epsilon}} \\
&= 1 - 2e^{-4\epsilon} \\
&\leq -1 + O(\epsilon) \qquad \text{for small } \epsilon.
\end{aligned}$$

Thus

$$\cos(\pi - m\angle y_{i-a} y_i y_{i+a}) \geq 1 - O(\epsilon).$$

By the Taylor expansion $\cos(\delta) \approx 1 - \frac{\delta^2}{2}$ for small δ, we have

$$m\angle y_{i-a} y_i y_{i+a} \leq O(\sqrt{\epsilon}).$$ □

We also note that $\theta_{a,b}$ is symmetric.

Claim 5.13. *For all a, b, it holds that $\theta_{a,b} = \theta_{b,a}$.*

Proof Sketch. Follows by the symmetry of reversing the ordering of $x_1, \ldots, x_n$. □

Claim 5.14. *For all $a, c > 0$ and all $0 < b < c$, it holds that $\theta_{a,c} \leq \theta_{a,b} + \theta_{b,c-b}$.*

Corollary 5.15. *For all $a, b > 1$, it holds that $\theta_{a,b} \leq \theta_{a-1,1} + \theta_{1,1} + \theta_{1,b-1}$.*

Proof. Assuming Claim 5.14, we have $\theta_{a,b} \leq \theta_{a,1} + \theta_{1,b-1} = \theta_{1,a} + \theta_{1,b-1} \leq \theta_{1,1} + \theta_{1,a-1} + \theta_{1,b-1}$. □

Claim 5.16. *For all $a, b \in \mathbb{Z}^+$, it holds that $\theta_{a,b} \leq \theta_{2a,b} + \theta_{a,a}$.*

In our proofs of Claims 5.14 and 5.16, we will repeatedly use the following "spherical triangle inequality", which we state without proof.

Fact 5.17. *For any points $O, A, B, C \in \mathbb{R}^3$, it holds that*

$$m\angle AOC \leq m\angle AOB + m\angle BOC.$$

Proof of Claim 5.14. For any i, we have

$$\begin{aligned}
m\angle y_{i-a} y_i y_{i+c} &\geq m\angle y_{i-a} y_i y_{i+b} - m\angle y_{i+b} y_i y_{i+c} && \text{(Fact 5.17)} \\
&\geq m\angle y_{i-a} y_i y_{i+b} - (\pi - m\angle y_i y_{i+b} y_{i+c})
\end{aligned}$$

So $\pi - m\angle y_{i-a} y_i y_{i+c} \leq \pi - m\angle y_{i-a} y_i y_{i+b} + (\pi - m\angle y_i y_{i+b} y_{i+c}) \leq \theta_{a,b} + \theta_{b,c-b}$, and the claim follows. □

Proof of Claim 5.16. For any i and ρ, we have

$$\begin{aligned} m\angle y_{i-a}y_iy_{i+b} &\geq m\angle y_{i-2a}y_iy_{i+b} - m\angle y_{i-2a}y_iy_{i-a} && \text{(Fact 5.17)} \\ &\geq m\angle y_{i-2a}y_iy_{i+b} - (\pi - m\angle y_{i-2a}y_{i-a}y_i) \end{aligned}$$

So $\pi - m\angle y_{i-a}y_iy_{i+b} \leq \pi - m\angle y_{i-2a}y_iy_{i+b} + (\pi - m\angle y_{i-2a}y_{i-a}y_i) \leq \theta_{2a,b} + \theta_{a,a}$, and the claim follows. □

Using Claims 5.12 to 5.14 and 5.16, we easily see that $\theta_{1,a} \leq O\big(\sqrt{\epsilon} \cdot \log(a)\big)$. For example,

$$\begin{aligned} \theta_{1,11} &\leq \theta_{1,1} + \theta_{1,10} && \text{(Claim 5.14)} \\ &\leq 2 \cdot \theta_{1,1} + \theta_{2,10} && \text{(Claim 5.16)} \\ &\leq 2 \cdot \theta_{1,1} + \theta_{2,2} + \theta_{2,8} && \text{(Claim 5.14)} \\ &\leq 2 \cdot \theta_{1,1} + 2 \cdot \theta_{2,2} + \theta_{4,8} && \text{(Claim 5.16)} \\ &\leq 2 \cdot \theta_{1,1} + 2 \cdot \theta_{2,2} + 2 \cdot \theta_{4,4} && \text{(Claim 5.14)} \\ &\leq 6 \cdot O(\sqrt{\epsilon}) && \text{(Claim 5.12).} \end{aligned}$$

The general bound is obtained in a similar way following the binary decomposition of a.

Given this bound on $\theta_{1,a}$, we apply Corollary 5.15 to obtain $\theta_{a,b} \leq \theta_{a-1,1} + \theta_{1,1} + \theta_{1,b-1} \leq O\big(\sqrt{\epsilon} \cdot \log(a + b)\big)$. □

By the discussion following the statement of Lemma 5.11, we have proved Theorem 5.9. □

It is possible to $\Omega(1)$-robustly instantiate n-prover KR-no-signaling, at the cost of using an arrangement that is only B-bounded for $B \geq e^{\Omega(n)}$.

Theorem 5.18. *For $\epsilon > 0$, it holds that n-prover KR-no-signaling is ϵ-robustly $e^{O(n)}$-boundedly instantiable via metric no-signaling.*

Proof. Let $\alpha > 2\epsilon$, let $x_i = \big(e^{\alpha i}, 0, 0\big)$ for $i \in \mathbb{Z}^+$, and let ρ be an arbitrary ϵ-isometry. Let y_i denote $\rho(x_i)$.

For any i, we have $e^{(2i-1)\epsilon} < e^{\alpha i - \epsilon} \leq \|y_i - \rho(0)\| \leq e^{\alpha i + \epsilon} < e^{(2i+1)\epsilon}$. Thus $\{y_1, \ldots, y_i\}$ are contained in a ball of radius $e^{(2i+1)\epsilon}$ centered at the origin, while the point y_{i+1} is not. □

Acknowledgments

We thank Ofer Grossman for helpful discussion regarding Theorem 5.9.

References

[BGKW88] Michael Ben-Or, Shafi Goldwasser, Joe Kilian, and Avi Wigderson, *Multi-prover interactive proofs: How to remove intractability assumptions*, STOC, ACM, 1988, pp. 113–131.

[CHSH69] John F. Clauser, Michael A. Horne, Abner Shimony, and Richard A. Holt, *Proposed experiment to test local hidden-variable theories*, Phys. Rev. Lett. **23** (1969), 880–884.

[HK20] Dhiraj Holden and Yael Kalai, *No-signaling proofs with $o(\sqrt{\log n})$ provers are in PSPACE*, STOC, ACM, 2020.

[Hol19] Justin Holmgren, *On round-by-round soundness and state restoration attacks*, Cryptology ePrint Archive, Report 2019/1261, 2019, https://eprint.iacr.org/2019/1261.

[IKM08] Tsuyoshi Ito, Hirotada Kobayashi, and Keiji Matsumoto, *Oracularization and two-prover one-round interactive proofs against nonlocal strategies*, CoRR **abs/0810.0693** (2008).

[Ito10] Tsuyoshi Ito, *Polynomial-space approximation of no-signaling provers*, ICALP (1), Lecture Notes in Computer Science, vol. 6198, Springer, 2010, pp. 140–151.

[KR09] Yael Tauman Kalai and Ran Raz, *Probabilistically checkable arguments*, CRYPTO, Lecture Notes in Computer Science, vol. 5677, Springer, 2009, pp. 143–159.

[KRR14] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum, *How to delegate computations: the power of no-signaling proofs*, STOC, ACM, 2014, pp. 485–494.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il