2016 IEEE 57th Annual Symposium on Foundations of Computer Science

Extension Complexity of Independent Set Polytopes

[Extended abstract]*

Mika Göös
SEAS
Harvard University

Rahul Jain
Centre for Quantum Technologies and
Department of Computer Science
National University of Singapore and
MajuLab, UMI 3654, Singapore

Thomas Watson
Department of Computer Science
University of Memphis

***Abstract*—We exhibit an n-node graph whose independent set polytope requires extended formulations of size exponential in $\Omega(n/\log n)$. Previously, no explicit examples of n-dimensional $0/1$-polytopes were known with extension complexity larger than exponential in $\Theta(\sqrt{n})$. Our construction is inspired by a relatively little-known connection between extended formulations and (monotone) circuit depth.**

***Keywords*-extended formulations; independent set polytopes; communication complexity; monotone circuit depth;**

I. Introduction

A polytope $P \subseteq \mathbb{R}^n$ with many facets can sometimes admit a concise description as the projection of a higher dimensional polytope $E \subseteq \mathbb{R}^e$ with few facets. This phenomenon is studied in the theory of "extended formulations". The *extension complexity* $\mathrm{xc}(P)$ of a polytope P is defined as the minimum number of facets in any E (called an *extended formulation* for P) such that

$$P = \{x \in \mathbb{R}^n : (x, y) \in E \text{ for some } y\}.$$

Extended formulations are useful for solving combinatorial optimization problems: instead of optimizing a linear function over P, we can optimize it over E—this may be more efficient since the runtime of LP solvers often depends on the number of facets.

Fiorini et al. [2] were the first to show (using methods from communication complexity [3, 4]) exponential extension complexity lower bounds for many explicit polytopes of relevance to combinatorial optimization, thereby solving an old challenge set by Yannakakis [5]. For example, their results include a $2^{\Omega(m)}$ lower bound for the $\binom{m}{2}$-dimensional *correlation/cut polytope*. In another breakthrough, Rothvoß [6] proved a much-conjectured $2^{\Omega(m)}$ lower bound for the $\binom{m}{2}$-dimensional *matching polytope*. By now, many accessible introductions to extended formulations are available; e.g., Roughgarden [7, §5], Kaibel [8], Conforty et al. [9] or their textbook [10, §4.10].

* Most proofs appear in the full version of this work [1].

$\sqrt{n}$-frontier: Both of the results quoted above—while optimal for their respective polytopes—seem to get "stuck" at being exponential in the square root of their dimension. In fact, no explicit n-dimensional 0/1-polytope (convex hull of a subset of $\{0,1\}^n$) was known with extension complexity asymptotically larger than $2^{\Theta(\sqrt{n})}$. In comparison, Rothvoß [11] showed via a counting argument that most n-dimensional 0/1-polytopes have extension complexity $2^{\Omega(n)}$.

A. Our result

Our main result is to construct an explicit 0/1-polytope of near-maximal extension complexity $2^{\Omega(n/\log n)}$. Moreover, the polytope can be taken to be the *independent set polytope* P_G of an n-node graph G, i.e., the convex hull of (the indicator vectors of) the independent sets of G. Previously, a lower bound of $2^{\Omega(\sqrt{n})}$ was known for independent set polytopes [2].

Theorem 1. *There is an (explicit) family of n-node graphs G with $\mathrm{xc}(P_G) \geq 2^{\Omega(n/\log n)}$.*

In fact, our graph family has bounded degree. Hence, using known reductions, we get as a corollary quantitative improvements—from $2^{\Omega(\sqrt{n})}$ to $2^{\Omega(n/\log n)}$—for the extension complexity of, for instance, *3SAT* and *knapsack polytopes*; see [12, 13] for details.

We strongly conjecture that our graph family actually satisfies $\mathrm{xc}(P_G) \geq 2^{\Omega(n)}$, i.e., that the $\log n$ factor in the exponent is an artifact of our proof technique. We give concrete evidence for this by proving an optimal bound for a certain *query complexity* analogue of Theorem 1. In particular, the conjectured bound $\mathrm{xc}(P_G) \geq 2^{\Omega(n)}$ would follow from quantitative improvements to the known query-to-communication simulation theorems ([14] in particular). Incidentally, this also answers a question of Lovász, Naor, Newman, and Wigderson [15]: we obtain a maximal $\Omega(n)$ lower bound on the randomized query complexity of a search problem with constant certificate complexity.

0272-5428/16 $31.00 © 2016 IEEE
DOI 10.1109/FOCS.2016.67

B. Our approach

Curiously enough, an analogous $\sqrt{n}$-frontier existed in the seemingly unrelated field of *monotone circuits*: Raz and Wigderson [16] proved an $\Omega(m)$ lower bound for the depth of any monotone circuit computing the *matching function* on $\binom{m}{2}$ input bits. This remained the largest monotone depth bound for an explicit function until the recent work of Göös and Pitassi [17], who exhibited a function with monotone depth $\Omega(n/\log n)$. In short, our idea is to prove an extension complexity analogue of this latter result.

The conceptual inspiration for our construction is a relatively little-known connection between Karchmer–Wigderson games [18] (which characterize circuit depth) and extended formulations. This "KW/EF connection" (see Section II for details) was pointed out by Hrubeš [19] as a nonnegative analogue of a classic rank-based method of Razborov [20]. In this work, we focus only on the monotone setting. For any monotone $f\colon \{0,1\}^n \to \{0,1\}$ we can study the convex hull of its 1-inputs, namely, the polytope

$$F \;\coloneqq\; \operatorname{conv} f^{-1}(1).$$

The upshot of the KW/EF connection is that extension complexity lower bounds for F follow from a certain type of *strengthening* of monotone depth lower bounds for f. For example, using this connection, it turns out that Rothvoß's result [6] implies the result of Raz and Wigderson [16] in a simple black-box fashion (Section II-C).

Our main technical result is to strengthen the existing monotone depth lower bound from [17] into a lower bound for the associated polytope (though we employ substantially different techniques than were used in that paper). The key communication search problem studied in [17] is a communication version of the well-known *Tseitin* problem (see Section III for definitions), which has especially deep roots in proof complexity (e.g., [4, §18.7]) and has also been studied in query complexity [15]. We use information complexity techniques to prove the required $\Omega(n/\log n)$ communication lower bound for the relevant variant of the Tseitin problem; information theoretic tools have been used in extension complexity several times [21]–[23]. One relevant work is Huynh and Nordström [24] (predecessor to [17]), whose information complexity arguments we extend in this work.

(Instead of using information complexity, an alternative seemingly promising approach would be to "lift" a strong enough query complexity lower bound for Tseitin into communication complexity. Unfortunately, this approach runs into problems due to limitations in existing query-to-communication simulation theorems; we discuss this in Section V.)

Theorem 1 follows by reductions from the result for Tseitin (Section IV). Indeed, it was known that the Tseitin problem reduces to the monotone KW game associated with an $f\colon \{0,1\}^{O(n)} \to \{0,1\}$ that encodes (in a monotone fashion) a certain CSP satisfiability problem. This gives us an extension complexity lower bound for the (explicit) polytope $F \coloneqq \operatorname{conv} f^{-1}(1)$. As a final step, we give a reduction from F to an independent set polytope.

C. Background

Let M be a nonnegative matrix. The *nonnegative rank* of M, denoted $\operatorname{rk}^+(M)$, is the minimum r such that M can be decomposed as a sum $\sum_{i\in[r]} R_i$ where each R_i is a rank-1 nonnegative matrix.

Randomized protocols. Faenza et al. [25] observed that a nonnegative rank decomposition can be naturally interpreted as a type of randomized protocol that computes the matrix M "in expectation". We phrase this connection precisely as follows: $\log \operatorname{rk}^+(M) + \Theta(1)$ is the minimum communication cost of a private-coin protocol Π whose acceptance probability on each input (x, y) satisfies $\mathbb{P}[\Pi(x, y) \text{ accepts}] = \alpha \cdot M_{x,y}$ where $\alpha > 0$ is an absolute constant of proportionality (depending on Π but not on x, y). All communication protocols in this paper are private-coin.

Slack matrices. The extension complexity of a polytope $P = \{x \in \mathbb{R}^n : Ax \geq b\}$ can be characterized in terms of the nonnegative rank of the *slack matrix* $M = M(P)$ associated with P. The entries of M are indexed by (v, i) where $v \in P$ is a vertex of P and i refers to the i-th facet-defining inequality $A_i x \geq b_i$ for P. We define $M_{v,i} \coloneqq A_i v - b_i \geq 0$ as the distance (*slack*) of the i-th inequality from being tight for v. Yannakakis [5] showed that $\operatorname{xc}(P) = \operatorname{rk}^+(M(P))$.

A convenient fact for proving lower bounds on $\operatorname{rk}^+(M)$ is that the nonnegative rank is unaffected by the addition of columns to M that each record the slack between vertices of P and some valid (but not necessarily facet-defining) inequality for P. For notation, let $P \subseteq Q$ be two nested polytopes (in fact, Q can be an unbounded polyhedron). We define $M(P; Q)$ as the slack matrix whose rows correspond to vertices of P and columns correspond to the facets of Q (hence $M(P; P) = M(P)$). We have $\operatorname{rk}^+(M(P)) \geq \operatorname{rk}^+(M(P) \cup M(P; Q)) - 1 \geq \operatorname{rk}^+(M(P; Q)) - 1$ where "$\cup$" denotes concatenation of columns.[1] We summarize all the above in the following.

Fact 2. *For all polytopes $P \subseteq Q$, we have* $\operatorname{xc}(P) = \operatorname{rk}^+(M(P)) \geq \operatorname{rk}^+(M(P; Q)) - 1$.

[1] Specifically, Farkas's Lemma implies that the slack of any valid inequality for P can be written as a nonnegative linear combination of the slacks of the facet-defining inequalities for P, plus a nonnegative constant [26, Proposition 1.9]. Thus if we take $M(P) \cup M(P; Q)$ and subtract off (possibly different) nonnegative constants from each of the "new" columns $M(P; Q)$, we get a matrix each of whose columns is a nonnegative linear combination of the "original" columns $M(P)$ and hence has the same nonnegative rank as $M(P)$. Since we subtracted off a nonnegative rank-1 matrix, we find that $\operatorname{rk}^+(M(P) \cup M(P; Q)) \leq \operatorname{rk}^+(M(P)) + 1$.

II. KW/EF Connection

We now describe the connection showing that EF lower bounds follow from a certain type of strengthening of lower bounds for monotone KW games (and similarly, lower bounds for monotone KW games follow from certain strong enough EF lower bounds). This is not directly used in the proof of Theorem 1, but it serves as inspiration by suggesting the approach we use in the proof.

A. Definitions

Let $f\colon \{0,1\}^n \to \{0,1\}$ be a monotone function. We define $\mathrm{KW}^+(f)$ as the deterministic communication complexity of the following *monotone KW game* associated with f.

——— KW^+-game: ———

Input: Alice gets $x \in f^{-1}(1)$, Bob gets $y \in f^{-1}(0)$.
Output: An index $i \in [n]$ such that $x_i = 1$ and $y_i = 0$.

We often think of x and y as subsets of $[n]$. In this language, a feasible solution for the KW^+-game is an $i \in x \cap \bar{y}$ where $\bar{y} \coloneqq [n] \setminus y$. Given a monotone f, we denote by $F \coloneqq \operatorname{conv} f^{-1}(1)$ the associated polytope. We can express the fact that any pair $(x,y) \in f^{-1}(1) \times f^{-1}(0)$ admits at least one witness $i \in x \cap \bar{y}$ via the following linear inequality:

$$\sum_{i\,:\,y_i=0} x_i \;\geq\; 1. \tag{1}$$

Since (1) is valid for all the vertices $x \in F$, it is valid for the whole polytope F. Define $F_{\mathrm{KW}} \supseteq F$ as the polyhedron whose facets are determined by the inequalities (1), as indexed by 0-inputs y. The (x,y)-th entry in the slack matrix $M(F; F_{\mathrm{KW}})$ is then $\sum_{i\,:\,y_i=0} x_i - 1$. In words, this quantity counts the number of witnesses in the KW^+-game on input (x,y) minus one.

More generally, let $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{Q}$ be *any* communication search problem (not necessarily a KW^+-game, even though any S can be reformulated as such [27, Lemma 2.3]). Here $\mathcal{Q}$ is some set of solutions/witnesses, and letting $S(x,y) \coloneqq \{q \in \mathcal{Q} : (x,y,q) \in S\}$ denote the set of feasible solutions for input (x,y), we assume that $S(x,y) \neq \emptyset$ for all (x,y). We associate with S the following natural *"number of witnesses minus one"* communication game.

——— $(\#\exists - 1)$-game: ———

Input: Alice gets $x \in \mathcal{X}$, Bob gets $y \in \mathcal{Y}$.
Output: Accept with probability proportional to $|S(x,y)| - 1$

The communication complexity of this game is simply $\log \mathrm{rk}^+(M^S) + \Theta(1)$ where $M^S_{x,y} \coloneqq |S(x,y)| - 1$.

B. The connection

What Hrubeš [19, Proposition 4] observed was that an efficient protocol for a search problem S implies an efficient protocol for the associated $(\#\exists - 1)$-game. In particular, for KW^+-games,

$$\log \mathrm{rk}^+(M(F; F_{\mathrm{KW}})) \;\leq\; O(\mathrm{KW}^+(f)). \tag{KW/EF}$$

The private-coin protocol for $M(F; F_{\mathrm{KW}})$ computes as follows. On input $(x,y) \in f^{-1}(1) \times f^{-1}(0)$ we first run the optimal deterministic protocol for the KW^+-game for f to find a particular $i \in [n]$ witnessing $x_i = 1$ and $y_i = 0$. Then, Alice uses her private coins to sample a $j \in [n] \setminus \{i\}$ uniformly at random, and sends this j to Bob. Finally, the two players check whether $x_j = 1$ and $y_j = 0$ accepting iff this is the case. The acceptance probability of this protocol is proportional to the number of witnesses minus one, and the protocol has cost $\mathrm{KW}^+(f) + \log n + O(1) \leq O(\mathrm{KW}^+(f))$ (where we assume w.l.o.g. that f depends on all of its input bits so that $\mathrm{KW}^+(f) \geq \log n$).

C. Example: Matchings

Rothvoß vs. Raz–Wigderson. Consider the monotone function $f\colon \{0,1\}^{\binom{m}{2}} \to \{0,1\}$ that outputs 1 iff the input, interpreted as a graph on m nodes (m even), contains a perfect matching. Then $F \coloneqq \operatorname{conv} f^{-1}(1)$ is the perfect matching polytope. The inequalities (1) for f happen to include the so-called "odd set" inequalities, which were exploited by Rothvoß [6] in showing that $\log \mathrm{rk}^+(M(F; F_{\mathrm{KW}})) \geq \Omega(m)$. Applying the (KW/EF) connection to Rothvoß's lower bound implies in a black-box fashion that $\mathrm{KW}^+(f) \geq \Omega(m)$, which is the result of Raz and Wigderson [16].

Converse to (KW/EF)*?* It is interesting to compare the above with the case of *bipartite* perfect matchings. Consider a monotone $f\colon \{0,1\}^{m \times m} \to \{0,1\}$ that takes a bipartite graph as input and outputs 1 iff the graph contains a perfect matching. It is well-known that $F \coloneqq \operatorname{conv} f^{-1}(1)$ admits a polynomial-size extended formulation [28, Theorem 18.1]. By contrast, the lower bound $\mathrm{KW}^+(f) \geq \Omega(m)$ from [16] continues to hold even in the bipartite case. This example shows that the converse inequality to (KW/EF) does not hold in general. Hence, a lower bound for the $(\#\exists - 1)$-game can be a strictly stronger result than a similar lower bound for the KW^+-game.

D. Minterms and maxterms

A *minterm* $x \in f^{-1}(1)$ is a minimal 1-input in the sense that flipping any 1-entry of x into a 0 will result in a 0-input. Analogously, a *maxterm* $y \in f^{-1}(0)$ is a maximal 0-input. It is a basic fact that solving the KW^+-game for minterms/maxterms is enough to solve the search problem on any input: Say that Alice's input x is not a minterm. Then Alice can replace x with any minterm $x' \subseteq x$ and run the protocol on x'. A witness $i \in [n]$ for (x', y) works also for (x, y). A similar fact holds for the $(\#\exists - 1)$-game: we claim that the nonnegative rank does not change by much when restricted to minterms/maxterms. Say that Alice's input

x is not a minterm. Then Alice can write $x = x' \cup x''$ (disjoint union) where x' is a minterm. Then $|x \cap \bar{y}| - 1 = (|x' \cap \bar{y}| - 1) + |x'' \cap \bar{y}|$ where the first term is the $(\#\exists{-}1)$-game for (x', y) and the second term has nonnegative rank at most n. (A similar argument works if Bob does not have a maxterm.)

III. Tseitin Problem

A. Query version

Fix a connected node-labeled graph $G = (V, E, \ell)$ where $\ell \in \mathbb{Z}_2^V$ has *odd weight*, i.e., $\sum_{v \in V} \ell(v) = 1$ where the addition is modulo 2. For any edge-labeling $z \in \mathbb{Z}_2^E$ and a node $v \in V$ we write concisely $z(v) \coloneqq \sum_{e \ni v} z(e)$ for the mod-2 sum of the edge-labels adjacent to v.

Tseitin problem: TSE_G

Input: Labeling $z \in \mathbb{Z}_2^E$ of the edges.
Output: A node $v \in V$ containing a *parity violation* $z(v) \neq \ell(v)$.

As a sanity check, we note that on each input z there must exist at least one node with a parity violation. This follows from the fact that, since each edge has two endpoints, the sum $\sum_v z(v)$ is even, whereas we assumed that the sum $\sum_v \ell(v)$ is odd.

Basic properties: The above argument implies more generally that the set of violations $\text{viol}(z) \coloneqq \{v \in V : z(v) \neq \ell(v)\}$ is always of odd size. Conversely, for any odd-size set $S \subseteq V$ we can design an input z such that $\text{viol}(z) = S$. To see this, it is useful to understand what happens when we *flip a path* in an input z. Formally, suppose $p \in \mathbb{Z}_2^E$ is (an indicator vector of) a path. Define z^p as z with bits on the path p flipped (note that $z^p = z + p \in \mathbb{Z}_2^E$; however, the notation z^p will be more convenient later). Flipping p has the effect of flipping whether each endpoint of p is a violation. More precisely, the violated nodes in z^p are related to those in z as follows: (i) if both endpoints of p are violated in z then the flip causes that pair of violations to disappear; (ii) if neither endpoint of p is violated in z, then the flip introduces a pair of new violations; (iii) if precisely one endpoint of p was violated in z, then the flip moves a violation from one endpoint of p to the other. By applying (i)–(iii) repeatedly in a connected graph G, we can design an input z where $\text{viol}(z)$ equals any prescribed odd-size set S.

If z and z' have the same set of violations, $\text{viol}(z) = \text{viol}(z')$, then their difference $q \coloneqq z - z' \in \mathbb{Z}_2^E$ satisfies $q(v) = 0$ for all $v \in V$. That is, q is an *eulerian* subgraph of G. On the other hand, for any eulerian graph q, the inputs z and z^q have the same violations. Consequently, to generate a random input with the same set of violations as some fixed z, we need only pick a random eulerian graph q and output z^q. (Eulerian graphs form a subspace of $\mathbb{Z}_2^E$, sometimes called the *cycle space* of G.)

B. Communication version

The communication version of the Tseitin problem is obtained by composing (or *lifting*) TSE_G with a constant-size two-party gadget $g \colon \mathcal{X} \times \mathcal{Y} \to \{0, 1\}$. In the lifted problem $\text{TSE}_G \circ g^n$, where $n \coloneqq |E|$, Alice gets $x \in \mathcal{X}^n$ as input, Bob gets $y \in \mathcal{Y}^n$ as input, and their goal is to find a node $v \in V$ that is violated for

$$z \coloneqq g^n(x, y) = (g(x_1, y_1), \ldots, g(x_n, y_n)).$$

We define our gadget precisely in the full version [1]. For this extended abstract—in particular, for the reductions presented in the next section—the only important property of our gadget is that $|\mathcal{X}|, |\mathcal{Y}| \leq O(1)$.

C. Statement of result

We prove that there is a family of bounded-degree graphs G such that the $(\#\exists - 1)$-game associated with $\text{TSE}_G \circ g^n$ requires $\Omega(n/\log n)$ bits of communication. We prove our lower bound assuming only that $G = (V, E)$ is well-connected enough as captured by the following definition (also used in [17]). A graph G is *k-routable* iff there is a set of $2k + 1$ nodes $T \subseteq V$ called *terminals* such that for any *pairing* $\mathcal{P} \coloneqq \{\{s_i, t_i\} : i \in [\kappa]\}$ (set of pairwise disjoint pairs) of 2κ terminals ($\kappa \leq k$), there exist κ edge-disjoint paths (called *canonical* paths for $\mathcal{P}$) such that the i-th path connects s_i to t_i. Furthermore, we tacitly equip G with an arbitrary odd-weight node-labeling.

The following is proved in the full version [1].

Theorem 3. *There is a constant-size g such that for every k-routable graph G with n edges, the $(\#\exists - 1)$-game for $\text{TSE}_G \circ g^n$ requires $\Omega(k)$ bits of communication.*

If we choose G to be a sufficiently strong expander graph, we may take $k = \Theta(n/\log n)$ as shown by Frieze et al. [29, 30]. Alternative constructions with $k = \Theta(n/\log n)$ exist based on bounded-degree "butterfly" graphs; see [31, §5] for an exposition.

Corollary 4. *There is a constant-size g and an explicit bounded-degree graph G with n edges such that the $(\#\exists{-}1)$-game for $\text{TSE}_G \circ g^n$ requires $\Omega(n/\log n)$ bits of communication.*

As a bonus, we also prove that the *query complexity* of the $(\#\exists - 1)$-game for TSE_G is $\Omega(n)$ on any expander G (see Section V).

IV. Reductions

The goal of this section is to show, via reductions, that a lower bound on the $(\#\exists - 1)$-game for $\text{TSE}_G \circ g^n$ (where $G = (V, E)$ is of bounded degree and $n \coloneqq |E|$) translates directly into a lower bound on the extension complexity of P_K for an $O(n)$-node bounded-degree graph K.

A. Definition: Monotone CSP-SAT

We start by describing a way of representing constraint satisfaction problems (CSP) as a monotone function; this was introduced in [17] and further studied by Oliveira [32, Chapter 3]. The function is defined relative to some finite alphabet Σ and a fixed constraint topology determined by a bipartite graph $H \coloneqq (L \cup R, E)$. The left nodes L are thought of as *variables* (taking values in Σ) and the right nodes R correspond to *constraints*. For a constraint $c \in R$, let $\mathrm{var}(c) \subseteq L$ denote the variables involved in c. Let d denote the maximum degree of a node in R. The function $\textsc{Sat} = \textsc{Sat}_{\Sigma,H} \colon \{0,1\}^m \to \{0,1\}$, where $m \leq |R| \cdot |\Sigma|^d$, is now defined as follows. An input $x \in \{0,1\}^m$ defines a CSP instance by specifying, for each $c \in R$, a truth table $\Sigma^{\mathrm{var}(c)} \to \{0,1\}$ that records which assignments to the variables $\mathrm{var}(c)$ satisfy c. Then $\textsc{Sat}(x) \coloneqq 1$ iff there is some global assignment $L \to \Sigma$ that satisfies all the constraints as specified by x. This is monotone: if we flip any 0 into a 1 in the truth table of a constraint, we are only making the constraint easier to satisfy.

B. From Tseitin to CSP-SAT

For completeness, we present the reduction (due to [17, §5.1]) from the search problem $\textsc{Tse}_G \circ g^n$ to the KW^+-game for $\textsc{Sat} = \textsc{Sat}_{\mathcal{X},H} \colon \{0,1\}^m \to \{0,1\}$. Here the alphabet is $\mathcal{X}$ and the bipartite graph H is defined on $E(G) \cup V(G)$ such that there is an edge $(e,v) \in E(H)$ iff $v \in e$. Note that $m \leq O(n)$ provided that $|\mathcal{X}| \leq O(1)$ and that G is of bounded degree.

On input (x,y) to $\textsc{Tse}_G \circ g^n$ the two players proceed as follows:

- Alice maps her $x \in \mathcal{X}^{E(G)}$ into a CSP whose sole satisfying assignment is x. Namely, for each constraint $v \in V(G)$, the truth table $\mathcal{X}^{\mathrm{var}(v)} \to \{0,1\}$ is all-0 except for a unique 1 in position $x|_{\mathrm{var}(v)}$ (restriction of x to coordinates in $\mathrm{var}(v)$).
- Bob maps his $y \in \mathcal{Y}^{E(G)}$ into an unsatisfiable CSP. Namely, for each constraint $v \in V(G)$, the truth table $t_v \colon \mathcal{X}^{\mathrm{var}(v)} \to \{0,1\}$ is given by $t_v(\hat{x}) \coloneqq 1$ iff $(g(\hat{x}_e, y_e))_{e \in \mathrm{var}(v)} \in \{0,1\}^{\mathrm{var}(v)}$ is a partial edge-labeling of G that does *not* create a parity violation on v.

Let us explain why Bob really produces a 0-input of SAT. Suppose for contradiction that there is an $\hat{x} \in \mathcal{X}^{E(G)}$ that satisfies all of Bob's constraints: $t_v(\hat{x}|_{\mathrm{var}(v)}) = 1$ for all v. By definition, this means that $z \coloneqq g^n(\hat{x}, y)$ is an input to $\textsc{Tse}_G$ without any violated nodes—a contradiction.

This reduction is parsimonious: it maps witnesses to witnesses in 1-to-1 fashion. Indeed, a node v is violated for $\textsc{Tse}_G \circ g^n$ if and only if Alice's truth table for v has its unique 1 in a coordinate where Bob has a 0. In conclusion, the $(\#\exists{-}1)$-game associated with (the KW^+-game for) SAT is at least as hard as the $(\#\exists{-}1)$-game for $\textsc{Tse}_G \circ g^n$.

C. From CSP-SAT to independent sets

As a final step, we start with $\textsc{Sat} = \textsc{Sat}_{\Sigma,H} \colon \{0,1\}^m \to \{0,1\}$ and construct an m-node graph K such that a slack matrix of the independent set polytope P_K embeds the $(\#\exists{-}1)$-game for SAT (restricted to minterms). Let $H \coloneqq (L \cup R, E)$ (as above) and define $n \coloneqq |R|$ (above we had $n = |L|$, but in our case $|L| = \Theta(|R|)$ anyway).

The m-node graph K is defined as follows (this is reminiscent of a reduction from [33]).

- The nodes of K are in 1-to-1 correspondence with the input bits of SAT. That is, for each constraint $c \in R$ we have $|\Sigma^{\mathrm{var}(c)}|$ many nodes in K labeled with assignments $\mathrm{var}(c) \to \Sigma$.
- There is an edge between any two nodes whose assignments are *inconsistent* with one another. (Here $\phi_i \colon \mathrm{var}(c_i) \to \Sigma$, $i \in \{1,2\}$, are inconsistent iff there is some $e \in \mathrm{var}(c_1) \cap \mathrm{var}(c_2)$ such that $\phi_1(e) \neq \phi_2(e)$.) In particular, the truth table of each constraint becomes a clique.

(It can be seen that K has bounded degree if H has bounded left- and right-degree, which it does after our reduction from Tseitin for a bounded-degree G.)

The key property of this construction is the following:

The minterms of SAT *are precisely the (indicator vectors of) maximal independent sets of K.*

Indeed, the minterms $x \in \textsc{Sat}^{-1}(1)$ correspond to CSPs with a unique satisfying assignment $\phi \colon L \to \Sigma$; there is a single 1-entry in each of the n truth tables (so that $|x| = n$) consistent with ϕ. Such an x, interpreted as a subset of nodes, is independent in K as it only contains nodes whose labels are consistent with ϕ. Conversely, because every independent set $x \subseteq V(K)$ can only contain pairwise consistently labeled nodes, x naturally defines a partial assignment $L' \to \Sigma$ for some $L' \subseteq L$. A maximal independent set x corresponds to picking a node from each of the n constraint cliques consistent with some total assignment $\phi \colon L \to \Sigma$. Hence x is a 1-input to SAT with unique satisfying assignment ϕ.

Our goal is now to exhibit a set of valid inequalities for the independent set polytope P_K whose associated slack matrix embeds the $(\#\exists{-}1)$-game for SAT. Let $x \subseteq V(K)$ be an independent set and $y \in \textsc{Sat}^{-1}(0)$. We claim that the following inequalities (indexed by y) are valid:

$$|x \cap y| \;=\; \sum_{i \,:\, y_i = 1} x_i \;\leq\; n - 1. \tag{2}$$

Clearly (2) holds whenever $|x| \leq n-1$. Since it is impossible to have $|x| \geq n+1$, assume that x is maximal: $|x| = n$. As argued above, x is a minterm of SAT. Hence (x,y) is a valid pair of inputs to the KW^+-game, and so they admit a witness: $|x \cap \bar{y}| \geq 1$. Therefore $|x \cap y| = n - |x \cap \bar{y}| \leq n-1$. This shows that (2) is valid. The slack matrix associated with

inequalities (2) has entries

$$n - 1 - |x \cap y| \;=\; |x \cap \bar{y}| - 1,$$

for any minterm x and any $y \in \textsc{Sat}^{-1}(0)$. But this is just the $(\#\exists{-}1)$-game for SAT with Alice's input restricted to minterms.

D. Proof of Theorem 1

Here we simply string the above reductions together. By Corollary 4 there is a constant-size g and a bounded-degree G with n edges such that the $(\#\exists{-}1)$-game for $\textsc{Tse}_G \circ g^n$ requires $\Omega(n/\log n)$ bits of communication. By the reduction of Section IV-B this implies an $\Omega(n/\log n)$ lower bound for the $(\#\exists - 1)$-game associated with (the KW$^+$-game for) a monotone function $\textsc{Sat}\colon \{0,1\}^{O(n)} \to \{0,1\}$. As discussed in Section II-D, the complexity of the $(\#\exists{-}1)$-game for SAT is affected only by $\pm \log n$ when restricted to minterms. Thus the minterm-restricted $(\#\exists - 1)$-game for SAT still has complexity $\Omega(n/\log n)$. (Alternatively, one can note that the reduction from Tseitin to CSP-SAT produced only minterms.) Hence the nonnegative rank of the matrix for that game is $2^{\Omega(n/\log n)}$. By the reduction of Section IV-C there is a bounded-degree $O(n)$-node graph K and a system of valid inequalities (2) for the independent set polytope P_K such that the slack matrix $M(P_K; Q)$, where Q is the polyhedron with facets determined by (2), embeds the matrix for the minterm-restricted $(\#\exists - 1)$-game for SAT. Thus $\log \mathrm{rk}^+(M(P_K; Q)) \geq \Omega(n/\log n)$. By Fact 2 we have $\log \mathrm{xc}(P_K) = \log \mathrm{rk}^+(M(P_K)) \geq \log\big(\mathrm{rk}^+(M(P_K; Q)) - 1\big) \geq \Omega(n/\log n)$.

V. Query Lower Bound

An alternative approach for proving a lower bound for the $(\#\exists{-}1)$-game for $\textsc{Tse}_G \circ g^n$ is:

1) Prove an appropriate *query complexity* lower bound for $\textsc{Tse}_G$.
2) Use a query-to-communication simulation theorem like [14, 34, 35].

In this section, we carry out the first step by proving an optimal $\Omega(n)$ lower bound (which in particular answers a question from [15])—this proof is a lot simpler than our proof for the $\Omega(n/\log n)$ communication lower bound in the full version [1]. Unfortunately, as we discuss below, it is not known how to perform the second step for constant-size gadgets g.

The result of this section can be interpreted as evidence that the right bound in Theorem 1 is $2^{\Omega(n)}$ and the right bound in Corollary 4 is $\Omega(n)$, and also as motivation for further work to improve parameters for simulation theorems.

A. Query-to-communication

The query complexity analogue of nonnegative rank decompositions (nonnegative combinations of nonnegative rank-1 matrices) are *conical juntas*: nonnegative combinations of conjunctions of literals (input bits or their negations). We write a conical junta as $h = \sum_C w_C C$ where $w_C \geq 0$ and C ranges over all conjunctions $C\colon \{0,1\}^n \to \{0,1\}$. The *degree* of h is the maximum number of literals in a conjunction C with $w_C > 0$. Each conical junta naturally computes a nonnegative function $h\colon \{0,1\}^n \to \mathbb{R}_{\geq 0}$. Hence we may study $(\#\exists - 1)$-games in query complexity. In particular, the query complexity of the $(\#\exists{-}1)$-game for $\textsc{Tse}_G$ is the least degree of a conical junta h that on input z outputs $h(z) = |\mathrm{viol}(z)| - 1$.

The main result of [14] is a simulation of randomized protocols (or nonnegative rank decompositions) by conical juntas: a cost-d protocol for a lifted problem $F \circ g^n$ can be simulated by a degree-$O(d)$ conical junta (approximately) computing F. While F here is arbitrary, the result unfortunately assumes that $g := \mathrm{IP}_b$ is a logarithmic-size, $b := \Theta(\log n)$, inner-product function $\mathrm{IP}_b\colon \{0,1\}^b \times \{0,1\}^b \to \{0,1\}$ given by $\mathrm{IP}_b(x, y) := \langle x, y\rangle \bmod 2$.

Plugging b-bit gadgets into the reductions of Section IV would blow up the number of input bits of CSP-SAT exponentially in b. This is not only an artifact of our particular reduction! Consider more generally any reduction from a communication search problem $S \circ g^n$ to a KW$^+$-game for a monotone $f\colon \{0,1\}^m \to \{0,1\}$. Since the KW$^+$-game has *nondeterministic* communication complexity $\log m$ (number of bits the players must nondeterministically guess to find a witness), the reduction would imply $c \leq \log m$ where c is the nondeterministic communication complexity of $S \circ g^n$. If merely computing g requires b bits of nondeterministic communication, then clearly $c \geq b$ so that $m \geq 2^b$.

B. A linear lower bound

Theorem 5. *There is a family of n-node bounded-degree graphs G such that the $(\#\exists{-}1)$-game for $\textsc{Tse}_G$ requires query complexity $\Omega(n)$.*

Relation to [15]: An analogue of the (KW/EF) connection holds for query complexity: if there is a deterministic decision tree of height d that solves the search problem $\textsc{Tse}_G$, we can convert this into a degree-$(d + O(1))$ conical junta for the associated $(\#\exists{-}1)$-game. Moreover, if we only have a *randomized* ϵ-error decision tree for the search problem, then the connection gives us a conical junta h that *approximately* solves the $(\#\exists{-}1)$-game: $h(z) \in (|\mathrm{viol}(z)| - 1) \cdot (1 \pm \epsilon)$ for all z.

Our proof below is robust enough that the $\Omega(n)$ bound holds even for conical juntas that merely approximately solve the $(\#\exists{-}1)$-game. Hence we get a randomized $\Omega(n)$ lower bound for $\textsc{Tse}_G$, which was conjectured by [15, p. 125]; note however that the earlier work [17] already got a near-optimal $\Omega(n/\log n)$ bound. In any case, to our knowledge, this is the first $O(1)$-vs-$\Omega(n)$ separation between certificate complexity and randomized query complexity for search problems.

The proof: Fix an n-node bounded-degree expander $G = (V, E)$. That is, for any subset $U \subseteq V$ of size $|U| \leq n/2$, the number of edges leaving U is $\Theta(|U|)$. We tacitly equip G with an arbitrary odd-weight node-labeling. Assume for the sake of contradiction that there is a conical junta $h = \sum w_C C$ of degree $o(n)$ for the $(\#\exists{-}1)$-game for $\textsc{Tse}_G$. Let C be a conjunction with $w_C > 0$. Denote by $S \subseteq E$ the set of edges that C reads; hence $|S| \leq o(n)$. Below, we write $G \setminus S$ for the graph induced on the edges $E \setminus S$ (deleting nodes that become isolated).

Claim 6. *We may assume w.l.o.g. that $G \setminus S$ is connected.*

Proof: If $G \setminus S$ is not connected, we may replace C with a conjunction (actually, a sum of them) that reads more input variables; namely, we let C read a larger set of edges $S' \supseteq S$ including all edges from connected components of $G \setminus S$ of "small" size $\leq n/2$. When adding some small component $K \subseteq E$ to S' we note that, because G is expanding, the size of K is big-O of the size of the edge boundary of K (which is contained in S). On the other hand, every edge in S lies on the boundary of at most two components. It follows that $|S'| = O(|S|)$, i.e., we increased the degree of h only by a constant factor. Now in $G \setminus S'$ we have only components of size $> n/2$, but there can only be one such component. ■

Claim 7. *We may assume w.l.o.g. that C witnesses at least two fixed nodes with a parity violation (i.e., C reads all the edge labels incident to the two nodes).*

Proof: Suppose for contradiction that C witnesses at most one violation. Then we may fool C into accepting an input (and hence h into outputting a positive value on that input) where the number of violations is 1, which is a contradiction to the definition of the $(\#\exists{-}1)$-game. Indeed, let z be some input accepted by C. Then we may modify z freely on the connected graph $G \setminus S$ (by Claim 6) without affecting C's acceptance: we may eliminate pairs of violations from z by flipping paths (as in Section III) until only one remains. (This is possible since by definition, all the non-witnessed violations of z remain in $G \setminus S$.) ■

Let μ_i (i odd) denote the distribution on inputs that have i violations at a random set of i nodes, and are otherwise random with this property. We may generate an input from μ_i as follows:

1) Choose an i-set $T_i \subseteq V$ of nodes at random.
2) Let $z \in \mathbb{Z}_2^E$ be any fixed input with $\mathrm{viol}(z) = T_i$.
3) Let $q \in \mathbb{Z}_2^E$ be a random eulerian graph.
4) Output $z + q$.

Theorem 5 follows from the following lemma. Here we identify C with the set (subcube) of inputs it accepts.

Lemma 8. $\mu_5(C) \geq (10/3 - o(1)) \cdot \mu_3(C)$.

Indeed, consider the expected output value $\mathbb{E}_{z_i \sim \mu_i}[h(z_i)]$. This should be 2 for $i = 3$, and 4 for $i = 5$, i.e., a factor 2 increase. However, the above lemma implies that the output value gets multiplied by more than a factor 3, which is the final contradiction.

Proof of Lemma 8: By Claim 7 let $\{v_1, v_2\}$ be a pair of nodes where C witnesses two violations. For $i = 3, 5$, let $z_i \sim \mu_i$ and denote by T_i the i-set of its violations. Then

$$\begin{aligned}\mu_3(C) &= \mathbb{P}[C(z_3) = 1] \\ &= \mathbb{P}[C(z_3) = 1 \text{ and } T_3 \supseteq \{v_1, v_2\}] \\ &= \tbinom{n-2}{1} / \tbinom{n}{3} \cdot \mathbb{P}[C(y_3) = 1], \\ &\qquad \text{(for } y_3 \coloneqq (z_3 \,|\, T_3 \supseteq \{v_1, v_2\})\text{)}\end{aligned}$$

$$\begin{aligned}\mu_5(C) &= \mathbb{P}[C(z_5) = 1] \\ &= \mathbb{P}[C(z_5) = 1 \text{ and } T_5 \supseteq \{v_1, v_2\}] \\ &= \tbinom{n-2}{3} / \tbinom{n}{5} \cdot \mathbb{P}[C(y_5) = 1]. \\ &\qquad \text{(for } y_5 \coloneqq (z_5 \,|\, T_5 \supseteq \{v_1, v_2\})\text{)}\end{aligned}$$

So their ratio is

$$\frac{\mu_5(C)}{\mu_3(C)} = \frac{10}{3} \cdot \frac{\mathbb{P}[C(y_5) = 1]}{\mathbb{P}[C(y_3) = 1]}.$$

Hence the following claim concludes the proof of Lemma 8. ■

Claim 9. $\mathbb{P}[C(y_5) = 1] / \mathbb{P}[C(y_3) = 1] \geq 1 - o(1)$.

Proof: We can generate y_3 and y_5 jointly as follows:

y_3: Choose $v_3 \in V \setminus \{v_1, v_2\}$ uniformly random and let x_3 be some input with $\mathrm{viol}(x_3) = \{v_1, v_2, v_3\}$. Output $y_3 \coloneqq x_3 + q$ where q is a random eulerian graph.

y_5: Continuing from the above, choose $\{v_4, v_5\} \subseteq V \setminus \{v_1, v_2, v_3\}$ at random. If possible, let p be a path in $G \setminus S$ joining $\{v_4, v_5\}$ (a "good" event), otherwise let p be any path joining $\{v_4, v_5\}$. Output $y_5 \coloneqq x_3 + p + q$.

It suffices to prove the claim conditioned on any particular v_3 (and hence also on x_3). By Claim 6 we have $\mathbb{P}[\text{"good"} \,|\, v_3] = \mathbb{P}\big[v_4, v_5 \in G \setminus S \,\big|\, v_3\big] \geq 1 - o(1)$ since $|S| \leq o(n)$. If the "good" event occurs, then C cannot distinguish between $y_3 = x_3 + q$ and $y_5 = x_3 + p + q$ so that $\mathbb{P}[C(y_3) = 1 \,|\, v_3] = \mathbb{P}\big[C(y_5) = 1 \,\big|\, \text{"good"}, v_3\big]$. The claim follows as

$$\begin{aligned}\mathbb{P}[C(y_5) = 1 \,|\, v_3] &\geq \mathbb{P}\big[C(y_5) = 1 \text{ and "good"} \,\big|\, v_3\big] \\ &= \mathbb{P}\big[C(y_5) = 1 \,\big|\, \text{"good"}, v_3\big] \\ &\quad \cdot \mathbb{P}[\text{"good"} \,|\, v_3] \\ &= \mathbb{P}[C(y_3) = 1 \,|\, v_3] \cdot \mathbb{P}[\text{"good"} \,|\, v_3] \\ &\geq \mathbb{P}[C(y_3) = 1 \,|\, v_3] \cdot (1 - o(1)).\end{aligned}$$

■

Acknowledgements

Thanks to Denis Pankratov, Toniann Pitassi, and Robert Robere for discussions and to anonymous referees for comments. We also thank Samuel Fiorini and Raghu Meka for e-mail correspondence. M.G. admits to having a wonderful time at IBM while learning about extended formulations with T.S. Jayram and Jan Vondrak.

Part of this research was done while M.G. and R.J. were attending the *Semidefinite and Matrix Methods for Optimization and Communication* program at the Institute for Mathematical Sciences, National University of Singapore in 2016. This research was supported in part by NSERC, and in part by the Singapore Ministry of Education and the National Research Foundation, also through the Tier 3 Grant *Random numbers from quantum processes* MOE2012-T3-1-009. M.G. is partially supported by the Simons Award for Graduate Students in Theoretical Computer Science (#360891).

REFERENCES

[1] M. Göös, R. Jain, and T. Watson, "Extension complexity of independent set polytopes," Electronic Colloquium on Computational Complexity (ECCC), Tech. Rep. TR16-070, 2016, Full version. [Online]. Available: http://eccc.hpi-web.de/report/2016/070/

[2] S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf, "Exponential lower bounds for polytopes in combinatorial optimization," *Journal of the ACM*, vol. 62, no. 2, pp. 17:1–17:23, 2015.

[3] E. Kushilevitz and N. Nisan, *Communication Complexity*. Cambridge University Press, 1997.

[4] S. Jukna, *Boolean Function Complexity: Advances and Frontiers*, ser. Algorithms and Combinatorics. Springer, 2012, vol. 27.

[5] M. Yannakakis, "Expressing combinatorial optimization problems by linear programs," *Journal of Computer and System Sciences*, vol. 43, no. 3, pp. 441–466, 1991.

[6] T. Rothvoß, "The matching polytope has exponential extension complexity," in *Proceedings of the 46th Symposium on Theory of Computing (STOC)*. ACM, 2014, pp. 263–272.

[7] T. Roughgarden, "Communication complexity (for algorithm designers)," arXiv, Tech. Rep., 2015.

[8] V. Kaibel, "Extended formulations in combinatorial optimization," arXiv, Tech. Rep., 2011.

[9] M. Conforti, G. Cornuéjols, and G. Zambelli, "Extended formulations in combinatorial optimization," *4OR*, vol. 8, no. 1, pp. 1–48, 2010.

[10] M. Conforti, G. Cornuéjols, and G. Zambelli, *Integer Programming*. Springer, 2014.

[11] T. Rothvoß, "Some 0/1 polytopes need exponential size extended formulations," *Mathematical Programming*, vol. 142, no. 1, pp. 255–268, 2012.

[12] D. Avis and H. R. Tiwary, "On the extension complexity of combinatorial polytopes," *Mathematical Programming*, vol. 153, no. 1, pp. 95–115, 2014.

[13] S. Pokutta and M. Van Vyve, "A note on the extension complexity of the knapsack polytope," *Operations Research Letters*, vol. 41, no. 4, pp. 347–350, 2013.

[14] M. Göös, S. Lovett, R. Meka, T. Watson, and D. Zuckerman, "Rectangles are nonnegative juntas," in *Proceedings of the 47th Symposium on Theory of Computing (STOC)*. ACM, 2015, pp. 257–266, (Full version: http://eccc.hpi-web.de/report/2014/147/).

[15] L. Lovász, M. Naor, I. Newman, and A. Wigderson, "Search problems in the decision tree model," *SIAM Journal on Discrete Mathematics*, vol. 8, no. 1, pp. 119–132, 1995.

[16] R. Raz and A. Wigderson, "Monotone circuits for matching require linear depth," *Journal of the ACM*, vol. 39, no. 3, pp. 736–744, 1992.

[17] M. Göös and T. Pitassi, "Communication lower bounds via critical block sensitivity," in *Proceedings of the 46th Symposium on Theory of Computing (STOC)*. ACM, 2014, pp. 847–856.

[18] M. Karchmer and A. Wigderson, "Monotone circuits for connectivity require super-logarithmic depth," in *Proceedings of the 20th Symposium on Theory of Computing (STOC)*. ACM, 1988, pp. 539–550.

[19] P. Hrubeš, "On the nonnegative rank of distance matrices," *Information Processing Letters*, vol. 112, no. 11, pp. 457–461, 2012.

[20] A. Razborov, "Applications of matrix methods to the theory of lower bounds in computational complexity," *Combinatorica*, vol. 10, no. 1, pp. 81–93, 1990.

[21] M. Braverman and A. Moitra, "An information complexity approach to extended formulations," in *Proceedings of the 45th Symposium on Theory of Computing (STOC)*. ACM, 2013, pp. 161–170.

[22] G. Braun and S. Pokutta, "Common information and unique disjointness," in *Proceedings of the 54th Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2013, pp. 688–697.

[23] ——, "The matching polytope does not admit fully-polynomial size relaxation schemes," in *Proceedings of the 26th Symposium on Discrete Algorithms (SODA)*. ACM–SIAM, 2015, pp. 837–846.

[24] T. Huynh and J. Nordström, "On the virtue of succinct proofs: Amplifying communication complexity hardness to time–space trade-offs in proof complexity," in *Proceedings of the 44th Symposium on Theory of Computing (STOC)*. ACM, 2012, pp. 233–248.

[25] Y. Faenza, S. Fiorini, R. Grappe, and H. R. Tiwary, "Extended formulations, nonnegative factorizations, and randomized communication protocols," *Mathematical Programming*, vol. 153, no. 1, pp. 75–94, 2014.

[26] G. Ziegler, *Lectures on Polytopes*, ser. Graduate Texts in Mathematics. Springer, 1995, vol. 152.

[27] A. Gál, "A characterization of span program size and improved lower bounds for monotone span programs," *Computational Complexity*, vol. 10, no. 4, pp. 277–296, 2001.

[28] A. Schrijver, *Combinatorial Optimization: Polyhedra and Efficiency*, ser. Algorithms and Combinatorics. Springer, 2003, vol. 24.

[29] A. Frieze and L. Zhao, "Optimal construction of edge-disjoint paths in random regular graphs," *Combinatorics, Probability and Computing*, vol. 9, no. 3, pp. 241–263, 2000.

[30] A. Frieze, "Edge-disjoint paths in expander graphs," *SIAM Journal on Computing*, vol. 30, no. 6, pp. 1790–1801, 2001.

[31] J. Nordström, "New wine into old wineskins: A survey of some pebbling classics with supplemental results," KTH Royal Institute of Technology, Tech. Rep., 2015.

[32] I. Oliveira, "Unconditional lower bounds in complexity theory," Ph.D. dissertation, Columbia University, 2015.

[33] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy, "Interactive proofs and the hardness of approximating cliques," *Journal of the ACM*, vol. 43, no. 2, pp. 268–292, 1996.

[34] S. O. Chan, J. Lee, P. Raghavendra, and D. Steurer, "Approximate constraint satisfaction requires large LP relaxations," in *Proceedings of the 54th Symposium on Foundations of Computer Science (FOCS)*. IEEE, 2013, pp. 350–359.

[35] J. Lee, P. Raghavendra, and D. Steurer, "Lower bounds on the size of semidefinite programming relaxations," in *Proceedings of the 47th Symposium on Theory of Computing (STOC)*. ACM, 2015, pp. 567–576.