# Adventures in Monotone Complexity and TFNP 

Mika Göös ${ }^{1}$<br>Institute for Advanced Study, Princeton, NJ, USA<br>mika@ias.edu

Pritish Kamath ${ }^{2}$<br>Massachusetts Institute of Technology, Cambridge, MA, USA<br>pritish@mit.edu

Robert Robere ${ }^{3}$
Simons Institute, Berkeley, CA, USA
robere@cs.toronto.edu
Dmitry Sokolov ${ }^{4}$
KTH Royal Institute of Technology, Stockholm, Sweden
sokolov.dmt@gmail.com


#### Abstract

Separations: We introduce a monotone variant of Xor-Sat and show it has exponential monotone circuit complexity. Since Xor-Sat is in $N C^{2}$, this improves qualitatively on the monotone vs. non-monotone separation of Tardos (1988). We also show that monotone span programs over $\mathbb{R}$ can be exponentially more powerful than over finite fields. These results can be interpreted as separating subclasses of TFNP in communication complexity.


Characterizations: We show that the communication (resp. query) analogue of PPA (subclass of TFNP) captures span programs over $\mathbb{F}_{2}$ (resp. Nullstellensatz degree over $\mathbb{F}_{2}$ ). Previously, it was known that communication FP captures formulas (Karchmer-Wigderson, 1988) and that communication PLS captures circuits (Razborov, 1995).

2012 ACM Subject Classification Theory of computation $\rightarrow$ Communication complexity, Theory of computation $\rightarrow$ Circuit complexity, Theory of computation $\rightarrow$ Proof complexity

Keywords and phrases TFNP, Monotone Complexity, Communication Complexity, Proof Complexity

Digital Object Identifier 10.4230/LIPIcs.ITCS.2019.38

Related Version A full version of the paper is available at [24], https://eccc.weizmann.ac. il/report/2018/163/.

Acknowledgements We thank Ankit Garg (who declined a co-authorship) for extensive discussions about monotone circuits. We also thank Thomas Watson and anonymous ITCS reviewers for comments.

[^0]
## 1 Our Results

We study the complexity of monotone boolean functions $f:\{0,1\}^{n} \rightarrow\{0,1\}$, that is, functions satisfying $f(x) \leq f(y)$ for every pair $x \leq y$ (coordinate-wise). (An excellent introduction to monotone complexity is the textbook [36].) Our main results are new separations of monotone models of computation and characterizations of those models in the language of query/communication complexity. At the core of these results are two conceptual innovations.

1. We introduce a natural monotone encoding of the usual CSP satisfiability problem (Subsection 1.1). This definition unifies many other monotone functions considered in the literature.
2. We extend and make more explicit an intriguing connection between circuit complexity and total NP search problems (TFNP) via communication complexity. Several prior characterizations [37,55] can be viewed in this light. This suggests a potentially useful organizational principle for circuit complexity measures; see Section 2 for our survey.

### 1.1 Monotone $\mathcal{C}$-Sat

The basic conceptual insight in this work is a new simple definition: a monotone encoding of the usual constraint satisfaction problem (CSP). For any finite set of constraints $\mathcal{C}$, we introduce a monotone function $\mathcal{C}$-Sat. A general definition is given in Section 3, but for now, consider as an example the set $\mathcal{C}=3$ Xor of all ternary parity constraints

$$
3 \mathrm{XoR}:=\left\{\left(v_{1} \oplus v_{2} \oplus v_{3}=0\right),\left(v_{1} \oplus v_{2} \oplus v_{3}=1\right)\right\}
$$

We define 3Xor-SAt ${ }_{n}:\{0,1\}^{N} \rightarrow\{0,1\}$ over $N:=|\mathcal{C}| n^{3}=2 n^{3}$ input bits as follows. An input $x \in\{0,1\}^{N}$ is interpreted as (the indicator vector of) a set of 3XOR constraints over $n$ boolean variables $v_{1}, \ldots, v_{n}$ (there are $N$ possible constraints). We define $3 \operatorname{Xor}-\operatorname{Sat}_{n}(x):=1$ iff the set $x$ is unsatisfiable, that is, no boolean assignment to the $v_{i}$ exists that satisfies all constraints in $x$. This is indeed a monotone function: if we flip any bit of $x$ from 0 to 1 , this means we are adding a new constraint to the instance, thereby making it even harder to satisfy.

Prior work. Our $\mathcal{C}$-Sat encoding generalizes several previously studied monotone functions. (NL) Karchmer and Wigderson [37] (also [29, 49, 56] and textbooks [39, 36]) studied the NLcomplete st-connectivity problem. This is equivalent to a $\mathcal{C}$-Sat problem with $\mathcal{C}$ consisting of a binary implication $\left(v_{1} \rightarrow v_{2}\right)$ and unit clauses $\left(v_{1}\right)$ and $\left(\neg v_{1}\right)$.
(P) Raz and McKenzie [52] (also [14, 15, 25, 18, 56, 48]) studied a certain P-complete generation problem. In hindsight, this is simply Horn-Sat, that is, $\mathcal{C}$ consists of Horn clauses: clauses with at most one positive literal, such as $\left(\neg v_{1} \vee \neg v_{2} \vee v_{3}\right)$.
(NP) Göös and Pitassi [25] and Oliveira [45, §3] (also [47, 48]) studied the NP-complete (dual of) CnF-Sat problem, where $\mathcal{C}$ consists of bounded-width clauses.
These prior works do not exhaust all interesting classes of $\mathcal{C}$, as is predicted by various classification theorems for CSPs [59, 20, 11, 63]. In this work, we focus on linear constraints over finite fields $\mathbb{F}_{p}$ (for example, 3Xor-SAT corresponding to $\mathbb{F}_{2}$ ) and over the reals $\mathbb{R}$.

### 1.2 Separations

First, we show that $3 \mathrm{XoR}-\mathrm{SAT}_{n}$ cannot be computed efficiently with monotone circuits.

- Theorem 1. 3Xor- $\mathrm{SAT}_{n}$ requires monotone circuits of size $2^{n^{\Omega(1)}}$.

This theorem stands in contrast to the fact that there exist fast parallel (non-monotone) algorithms for linear algebra [44]. In particular, 3Xor-SAT is in NC ${ }^{2}$. Consequently, our result improves qualitatively on the monotone vs. non-monotone separation of Tardos [61] who exhibited a monotone function in $P$ (computed by solving a semidefinite program) with exponential monotone circuit complexity. For further comparison, another famous candidate problem to witness a monotone vs. non-monotone separation is the perfect matching function: it is in $\mathrm{RNC}^{2}$ [40] while it is widely conjectured to have exponential monotone circuit complexity (a quasipolynomial lower bound was proved by Razborov [53]).

Span programs. The computational easiness of 3 Xor- SAT $_{n}$ can be stated differently: it can be computed by a linear-size monotone $\mathbb{F}_{2}$-span program. Span programs are a model of computation introduced by Karchmer and Wigderson [38] (see also [36, §8] for exposition) with an extremely simple definition. An $\mathbb{F}$-span program, where $\mathbb{F}$ is a field, is a matrix $M \in \mathbb{F}^{m \times m^{\prime}}$ each row of which is labeled by a literal, $x_{i}$ or $\neg x_{i}$. We say that the program accepts an input $x \in\{0,1\}^{n}$ iff the rows of $M$ whose labels are consistent with $x$ (literals evaluating to true on $x$ ) span the all-1 row vector. The size of a span program is its number of rows $m$. A span program is monotone if all its literals are positive; in this case the program computes a monotone function.

A corollary of Theorem 1 is that monotone $\mathbb{F}_{2}$-span programs cannot be simulated by monotone circuits without exponential blow-up in size. This improves on a separation of Babai, Gál, and Wigderson [3] who showed that monotone circuit complexity can be quasipolynomially larger than monotone $\mathbb{F}_{2}$-span program size.

Furthermore, Theorem 1 holds more generally over any field $\mathbb{F}$ : an appropriately defined function $3 \operatorname{Lin}(\mathbb{F})-\mathrm{SAT}_{n}$ (ternary $\mathbb{F}$-linear constraints; see Section 3) is easy for monotone $\mathbb{F}$-span programs, but exponentially hard for monotone circuits. No such separation, even superpolynomial, was previously known for fields of characteristic other than 2.

This brings us to our second theorem.

- Theorem 2. $3 \operatorname{Lin}(\mathbb{R})-$ SAT $_{n}$ requires monotone $\mathbb{F}_{p}$-span programs of size $2^{n^{\Omega(1)}}$ for any prime $p$.

In other words: monotone $\mathbb{R}$-span programs can be exponentially more powerful than monotone span programs over finite fields. This separation completes the picture for the relative powers of monotone span programs over distinct fields, since the remaining cases were exponentially separated by Pitassi and Robere [48].

Finally, our two results above yield a bonus result in proof complexity as a byproduct: the Nullstellensatz proof system (over any field) can be exponentially more powerful than the Cutting Planes proof system (see Subsection 4.2).

Techniques. The new lower bounds are applications of the lifting theorems for monotone circuits [23] and monotone span programs [48]. We show that, generically, if some unsatisfiable formula composed of $\mathcal{C}$ constraints is hard to refute for the Resolution (resp. Nullstellensatz) proof system, then the $\mathcal{C}$-SAT problem is hard for monotone circuits (resp. span programs). Hence we can invoke (small modifications of) known Resolution and Nullstellensatz lower bounds $[8,10,1]$. The key conceptual innovation here is a reduction from unsatisfiable $\mathcal{C}$-CSPs (or their lifted versions) to the monotone Karchmer-Wigderson game for $\mathcal{C}$-Sat. This reduction is extremely slick, which we attribute to having finally found the "right" definition of $\mathcal{C}$-SAT.

### 1.3 Characterizations

There are two famous "top-down" characterizations of circuit models (both monotone and nonmonotone variants) using the language of communication complexity; these characterizations are naturally related to communication analogues of subclasses of TFNP.
(FP) Karchmer and Wigderson [37] showed that the logarithm of the (monotone) formula complexity of a (monotone) function $f:\{0,1\}^{n} \rightarrow\{0,1\}$ is equal, up to constant factors, to the communication complexity of the (monotone) Karchmer-Wigderson game:

$$
\begin{array}{r}
\text { Search problem KW }(\boldsymbol{f}) \\
{\left[\operatorname{resp} . \mathbf{K W}^{+}(\boldsymbol{f})\right]}
\end{array}=\begin{array}{r}
\text { input: a pair }(x, y) \in f^{-1}(1) \times f^{-1}(0) \\
\text { output: } \quad \text { an } i \in[n] \text { with } x_{i} \neq y_{i}\left[\text { resp. } x_{i}>y_{i}\right]
\end{array}
$$

We summarize this by saying that the communication analogue of FP captures formulas. Here FP $\subseteq$ TFNP is the classical (Turing machine) class of total NP search problems efficiently solved by deterministic algorithms [42].
(PLS) Razborov [55] (see also [50, 60]) showed that the logarithm of the (monotone) circuit complexity of a function $f:\{0,1\}^{n} \rightarrow\{0,1\}$ is equal, up to constant factors, to the least cost of a PLS-protocol solving the $\mathrm{KW}(f)\left(\right.$ or $\left.\mathrm{KW}^{+}(f)\right)$ search problem. Here a PLS-protocol (Definition 14 in Appendix A) is a natural communication analogue of PLS $\subseteq$ TFNP [35]. We summarize this by saying that the communication analogue of PLS captures circuits.

We contribute a third characterization of this type: the communication analogue of PPA captures $\mathbb{F}_{2}$-span programs. The class PPA [46] is a well-known subclass of TFNP embodying the combinatorial principle "every graph with an odd degree vertex has another". Informally, a search problem is in PPA if for every $n$-bit input $x$ we may describe implicitly an undirected graph $G_{x}=(V, E)$ (typically of size exponential in $n$; the edge relation is computed by a polynomial-size circuit) such that $G$ has degree at most 2 , there is a distinguished degree-1 vertex $v^{*} \in V$, and every other degree- 1 vertex $v \in V$ is associated with a feasible solution to the instance $x$ (that is, the solution can be efficiently computed from $v$ ).



Communication PPA. The communication analogue of PPA is defined canonically by letting the edge relation be computed by a (deterministic) communication protocol. Specifically, first fix a two-party search problem $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{O}$, that is, Alice gets $x \in \mathcal{X}$, Bob gets $y \in \mathcal{Y}$, and their goal is to find a feasible solution in $S(x, y):=\{o \in \mathcal{O}:(x, y, o) \in S\}$. A PPA-protocol $\Pi$ solving $S$ consists of a vertex set $V$, a distinguished vertex $v^{*} \in V$, and for each vertex $v \in V$ there is an associated solution $o_{v} \in \mathcal{O}$ and a protocol $\Pi_{v}$ (taking inputs from $\mathcal{X} \times \mathcal{Y})$. Given an input $(x, y)$, the protocols $\Pi_{v}$ implicitly describe a graph $G=G_{x, y}$
on the vertex set $V$ as follows. The output of protocol $\Pi_{v}$ on input $(x, y)$ is interpreted as a subset $\Pi_{v}(x, y) \subseteq V$ of size at most 2 . We define $\{u, v\} \in E(G)$ iff $u \in \Pi_{v}(x, y)$ and $v \in \Pi_{u}(x, y)$. The correctness requirements are:
(C1) if $\operatorname{deg}\left(v^{*}\right) \neq 1$, then $o_{v^{*}} \in S(x, y)$.
(C2) if $\operatorname{deg}(v) \neq 2$ for $v \neq v^{*}$, then $o_{v} \in S(x, y)$.
The cost of $\Pi$ is defined as $\log |V|+\max _{v}\left|\Pi_{v}\right|$ where $\left|\Pi_{v}\right|$ is the communication cost of $\Pi_{v}$. Finally, define $\mathrm{PPA}^{\mathrm{cc}}(S)$ as the least cost of a PPA-protocol that solves $S$.

For a (monotone) function $f$, define $\mathrm{SP}_{\mathbb{F}}(f)$ (resp. $\mathrm{mSP}_{\mathbb{F}}(f)$ ) as the least size of a (monotone) $\mathbb{F}$-span program computing $f$. Our characterization is in terms of $S:=\mathrm{KW}(f)$.

- Theorem 3. For any boolean function $f$, we have $\log \operatorname{SP}_{\mathbb{F}_{2}}(f)=\Theta\left(\operatorname{PPA}^{c c}(\mathrm{KW}(f))\right)$. Furthermore, if $f$ is monotone, we have $\log \operatorname{mSP}_{\mathbb{F}_{2}}(f)=\Theta\left(\operatorname{PPA}^{\mathrm{cc}}\left(\mathrm{KW}^{+}(f)\right)\right)$.

Query PPA. Our second characterization concerns the Nullstellensatz proof system; see Section 3 for the standard definition. Span programs and Nullstellensatz are known to be connected via interpolation [51] and lifting [48]. Given our first characterization (Theorem 3), it is no surprise that a companion result should hold in query complexity: the query complexity analogue of PPA captures the degree of Nullstellensatz refutations over $\mathbb{F}_{2}$.

The query analogue of PPA is defined in the same way as the communication analogue, except we replace protocols by (deterministic) decision trees. In fact, query PPA was already studied by Beame et al. [6] who separated query analogues of different subclasses of TFNP. To define it, first fix a search problem $S \subseteq\{0,1\}^{n} \times \mathcal{O}$, that is, on input $x \in\{0,1\}^{n}$ the goal is to find a feasible solution in $S(x):=\{o \in \mathcal{O}:(x, o) \in S\}$. A PPA-decision tree $\mathcal{T}$ solving $S$ consists of a vertex set $V$, a distinguished vertex $v^{*} \in V$, and for each vertex $v \in V$ there is an associated solution $o_{v} \in \mathcal{O}$ and a decision tree $\mathcal{T}_{v}$ (querying bits of an $n$-bit input). Given an input $x \in\{0,1\}^{n}$, the decision trees $\mathcal{T}_{v}$ implicitly describe a graph $G=G_{x}$ on the vertex set $V$ as follows. The output of $\mathcal{T}_{v}$ on input $x$ is interpreted as a subset $\mathcal{T}_{v}(x) \subseteq V$ of size at most 2. We then define $\{u, v\} \in E(G)$ iff $u \in \mathcal{T}_{v}(x)$ and $v \in \mathcal{T}_{u}(x)$. The correctness requirements are the same as before, 1 and 2 . The cost of $\mathcal{T}$ is defined as the maximum over all $v \in V$ and all inputs $x$ of the number of queries made by $\mathcal{T}_{v}$ on input $x$. Finally, define $\mathrm{PPA}^{\mathrm{dt}}(S)$ as the least cost of a PPA-decision tree that solves $S$.

With any unsatisfiable $n$-variate boolean CSP $F$ one can associate a canonical search problem:

```
CSP search problem S(F)= input: an n-variate truth assignment }x\in{0,1\mp@subsup{}}{}{n
    output: constraint C of F falsified by x (i.e., }C(x)=0\mathrm{ )
```

- Theorem 4. The $\mathbb{F}_{2}$-Nullstellensatz degree of an $k$-CNF formula $F$ equals $\Theta\left(\operatorname{PPA}^{\mathrm{dt}}(S(F))\right)$.

The easy direction of this characterization is that Nullstellensatz degree lower bounds PPA ${ }^{\mathrm{dt}}$. This fact was already observed and exploited by Beame et al. [6] to prove lower bounds for PPA ${ }^{\mathrm{dt}}$. Our contribution is to show the other (less trivial) direction.

Let us finally mention a related result in Turing machine complexity due to Belovs et al. [9]: a circuit-encoded version of Nullstellensatz is PPA-complete. Their proof is highly nontrivial whereas our characterizations admit relatively short proofs, owing partly to us working with simple nonuniform models of computation.


Figure 1 The landscape of communication search problem classes (uncluttered by the usual 'cc' superscripts). A solid arrow $C_{1} \rightarrow C_{2}$ denotes $C_{1} \subseteq C_{2}$, and a dashed arrow $C_{1} \rightarrow C_{2}$ denotes $C_{1} \nsubseteq C_{2}$ (in fact, an exponential separation). Some classes can characterize other models of computation (printed in blue). See Appendix A for definitions.

## 2 Survey: Communication TFNP

Given the results in Section 1, it is natural to examine other communication analogues of subclasses of TFNP. The goal in this section is to explain the current state of knowledge as summarized in Figure 1. The formal definitions of the communication classes appear in Appendix A.

TFNP. As is customary in structural communication complexity [2, 30, 27] we formally define TFNP ${ }^{c c}$ (resp. PLS ${ }^{c c}$, PPA $^{c c}$, etc.) as the class of all total two-party $n$-bit search problems that admit a nondeterministic protocol ${ }^{5}$ (resp. PLS-protocol, PPA-protocol, etc.) of communication cost polylog$(n)$. For example, Karchmer-Wigderson games KW $(f)$ and $\mathrm{KW}^{+}(f)$, for an $n$-bit boolean function $f$, have efficient nondeterministic protocols: guess a $\log n$-bit coordinate $i \in[n]$ and check that $x_{i} \neq y_{i}$ or $x_{i}>y_{i}$. Hence these problems are in TFNP ${ }^{\text {cc }}$. In fact, a converse holds: any total two-party search problem with nondeterministic complexity $c$ can be reduced to $\mathrm{KW}^{+}(f)$ for some $2^{c}$-bit partial monotone function $f$, see $[21$, Lemma 2.3]. In summary, the study of total NP search problems in communication complexity is equivalent to the study of monotone Karchmer-Wigderson games for partial monotone functions.

[^1]Sometimes a partial function $f$ can be canonically extended into a total one $f^{\prime}$ without increasing the complexity of $\mathrm{KW}(f)$ (or $\mathrm{KW}^{+}(f)$ ). This is possible whenever $\mathrm{KW}(f)$ lies in a communication class that captures some associated model of computation. For example, if KW $(f)$ is solved by a deterministic protocol (resp. PLS-protocol, PPA-protocol) then the Karchmer-Wigderson connection can build us a corresponding formula (resp. circuit, $\mathbb{F}_{2}$-span program) that computes some total extension $f^{\prime}$ of $f$. Consequently, separating two communication classes that capture two monotone models is equivalent to separating the monotone models themselves.

FP. Raz and McKenzie [52] showed an exponential separation between monotone formula size and monotone circuit size. This can be rephrased as PLS ${ }^{c c} \nsubseteq \mathrm{FP}^{c c}$. Their technique is much more general: they develop a query-to-communication lifting theorem for deterministic protocols (see also [26] for exposition). By plugging in known query complexity lower bounds against the class EoML (combinatorial subclass of CLS [17] introduced by [32, 19]), one can obtain a stronger separation EoML ${ }^{\mathrm{cc}} \nsubseteq \mathrm{FP}^{\mathrm{cc}}$.

A related question is whether randomization helps in solving TFNP ${ }^{c c}$ problems. Lower bounds against randomized protocols have applications in proof complexity [34, 7, 33, 25] and algorithmic game theory [58, 5, 28, 57, 22, 4]. In particular, some of these works (for finding Nash equilibria) have introduced a communication analogue of the PPAD-complete End-of-Line problem, which we will continue to study in Subsection 4.2.

PLS. Razborov's [54] famous monotone circuit lower bound for the clique/coloring problem (which is in $\mathrm{PPP}^{\mathrm{cc}}$ ) can be interpreted as an exponential separation $\mathrm{PPP}^{c c} \nsubseteq \mathrm{PLS}^{\mathrm{cc}}$. We show a stronger separation $\mathrm{PPAD}^{\mathrm{cc}} \nsubseteq \mathrm{PLS}^{c \mathrm{cc}}$ using the End-OF-Line problem in Subsection 4.2. Note that this is even slightly stronger than Theorem 1, which only implies PPA ${ }^{c c} \nsubseteq$ PLS $^{c c}$.

PPApD. In light of our characterization of PPA $^{\mathrm{cc}}$, we may interpret the inability of monotone $\mathbb{F}_{2}$-span program to efficiently simulate monotone circuits [48] as a separation PLS ${ }^{c c} \nsubseteq \mathrm{PPA}^{c c}$. We show an incomparable separation PPADS ${ }^{c c} \nsubseteq$ PPA $^{c c}$ in Subsection 4.3.

In the other direction, prior work implies $\mathrm{PPA}^{\mathrm{cc}} \nsubseteq \mathrm{PPAD}^{\mathrm{cc}}$ as follows. Pitassi and Robere [48] exhibit a monotone $f$ (in hindsight, one can take $f:=3 \mathrm{Xor}^{-\mathrm{SAT}_{n}}$ ) computable with a small monotone $\mathbb{F}_{2}$-span program (hence $\mathrm{KW}^{+}(f) \in \mathrm{PPA}^{\mathrm{cc}}$ ) and such that $\mathrm{KW}^{+}(f)$ has an exponentially large $\mathbb{R}$-partition number (see Section 3 for a definition); however, we observe that all problems in PPAD ${ }^{c c}$ have a small $\mathbb{R}$-partition number (see Remark 9).

PPP. There are no lower bounds against PPP $^{c c}$ for an explicit problem in TFNP ${ }^{c c}$. However, we can show non-constructively the existence of $\mathrm{KW}(f) \in \mathrm{TFNP}^{c c}$ such that $\mathrm{KW}(f) \notin \mathrm{PPP}^{c c}$, which implies $\mathrm{PPP}^{c c} \neq$ TFNP ${ }^{c c}$. Indeed, we argue in Remark 8 that every $S$ reduces to $\mathrm{KW}^{+}\left(3 \mathrm{CnF}^{-\mathrm{SAT}_{N}}\right)$ over $N:=\exp \left(O\left(\operatorname{PPP}^{c c}(S)\right)\right)$ variables. Applying this to $S:=\mathrm{KW}(f)$ for an $n$-bit $f$, we conclude that $f$ is a (non-monotone) projection of $3 \mathrm{CNF}^{2} \mathrm{SAT}_{N}$ for $N:=\exp \left(O\left(\operatorname{PPP}^{c c}(\mathrm{KW}(f))\right)\right)$. In particular, if $\mathrm{KW}(f) \in \operatorname{PPP}^{c c}$ (i.e., $\operatorname{PPP}^{c c}(\mathrm{KW}(f)) \leq$ $\operatorname{polylog}(n))$, then $f$ is in non-uniform quasipoly-size NP. Therefore $\mathrm{KW}(f) \notin \mathrm{PPP}^{c c}$ for a random $f$.

EoML, SoML, and comparator circuits. One prominent circuit model that currently lacks a characterization via a TFNP ${ }^{c c}$ subclass is comparator circuits $[41,16]$. These circuits are composed only of comparator gates (taking two input bits and outputting them in sorted order) and input literals (positive literals in the monotone case).

We can show an upper bound better than PLS ${ }^{c c}$ for comparator circuits. Indeed, we introduce a new class SoML generalizing EoML [32, 19] as follows. Recall that EoML is the class of problems reducible to End-of-Metered-Line: we are given a directed graph of in/out-degree at most 1 with a distinguished source vertex $v^{*}$ (in-degree 0), and moreover, each vertex is labeled with an integer "meter" that is strictly decreasing along directed paths; a solution is any sink or source distinct from $v^{*}$. The complete problem defining SoML is Sink-of-Metered-Line, which is the same as End-of-Metered-Line except only sinks count as solutions. It is not hard (left as an exercise) to adapt the characterization of circuits via $\mathrm{PLS}^{c c}[55,50,60]$ to show that $\mathrm{KW}(f)$ is in $\mathrm{SoML}^{c c}$ if $f$ is computed by a small comparator circuit. However, we suspect that the converse (SoML-protocol for KW $(f)$ implies a comparator circuit) is false.

### 2.1 Open problems

In query complexity, the relative complexities of TFNP subclasses are almost completely understood [ $6,12,43]$. In communication complexity, by contrast, there are huge gaps in our understanding as can be gleaned from Figure 1. For example:
(1) There are no lower bounds against classes PPADS ${ }^{c c}$ and PPP $^{c c}$ for an explicit problem in TFNP ${ }^{c c}$. For starters, show PLS $^{c c} \nsubseteq$ PPADS $^{c c}$ or PPA $^{c c} \nsubseteq$ PPADS $^{c c}$.
(2) Find computational models captured by EoML ${ }^{c c}$, SoML ${ }^{c c}, \mathrm{PPAD}^{c c}$, PPADS $^{c c}$, PPP $^{c c}$.
(3) Query-to-communication lifting theorems are known for FP [52], PLS [23], PPA [48]. Prove more. (This is one way to attack Question 1 if proved for PPADS.)
(4) Prove more separations. For example, can our result PPADS ${ }^{c c} \nsubseteq$ PPA $^{c c}$ be strengthened to SoML ${ }^{\mathrm{cc}} \nsubseteq \mathrm{PPA}^{\mathrm{cc}}$ ? This is closely related to whether monotone comparator circuits can be more powerful than monotone $\mathbb{F}_{2}$-span programs (no separation is currently known).

## 3 Preliminaries

$\mathcal{C}$-Sat. Fix an alphabet $\Sigma$ (potentially infinite, e.g., $\Sigma=\mathbb{R}$ ). Let $\mathcal{C}$ be a finite set of $k$-ary predicates over $\Sigma$, that is, each $C \in \mathcal{C}$ is a function $C: \Sigma^{k} \rightarrow\{0,1\}$. We define a monotone function $\mathcal{C}-$ SAT $_{n}:\{0,1\}^{N} \rightarrow\{0,1\}$ over $N=|\mathcal{C}| n^{k}$ input bits as follows. An input $x \in\{0,1\}^{N}$ is interpreted as a $\mathcal{C}$-CSP instance, that is, $x$ is (the indicator vector of) a set of $\mathcal{C}$-constraints, each applied to a $k$-tuple of variables from $v_{1}, \ldots, v_{n}$. We define $\mathcal{C}-\operatorname{SAT}_{n}(x):=1$ iff the $\mathcal{C}$-CSP $x$ is unsatisfiable: no assignment $v \in \Sigma^{n}$ exists such that $C(v)=1$ for all $C \in x$.

For a field $\mathbb{F}$, we define $k \operatorname{Lin}(\mathbb{F})$ as the set of all $\mathbb{F}$-linear equations of the form

$$
\sum_{i \in[k]} a_{i} v_{i}=a_{0}, \quad \text { where } a_{i} \in\{0, \pm 1\}
$$

In particular, we recover $3 \mathrm{Xor}-\mathrm{Sat}_{n}$ defined in Section 1 essentially as $3 \operatorname{Lin}\left(\mathbb{F}_{2}\right)$ - $\mathrm{SAT}_{n}$. We could have allowed the $a_{i}$ to range over $\mathbb{F}$ when $\mathbb{F}$ is finite, but we stick with the above convention as it ensures that the set $k \operatorname{Lin}(\mathbb{R})$ is always finite.

Boolean alphabets. We assume henceforth that all alphabets $\Sigma$ contain distinguished elements 0 and 1 . We define $\mathcal{C}_{\text {bool }}$ to be the constraint set obtained from $\mathcal{C}$ by restricting each $C \in \mathcal{C}$ to the boolean domain $\{0,1\}^{k} \subseteq \Sigma^{k}$. Moreover, if $F$ is a $\mathcal{C}$-CSP, we write $F_{\text {bool }}$ for the $\mathcal{C}_{\text {bool- }}$ CSP obtained by restricting the constraints of $F$ to boolean domains. Consequently, any $S\left(F_{\text {bool }}\right)$ associated with a $\mathcal{C}$-CSP $F$ is a boolean search problem.

Algebraic partitions. We say that a subset $A \subseteq \mathcal{X} \times \mathcal{Y}$ is monochromatic for a two-party search problem $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{O}$ if there is some $o \in \mathcal{O}$ such that $o \in S(x, y)$ for all $(x, y) \in A$. Moreover, if $M \in \mathbb{F}^{\mathcal{X}} \times \mathcal{Y}$ is a matrix, we say $M$ is monochromatic if the support of $M$ is monochromatic. For any field $\mathbb{F}$, an $\mathbb{F}$-partition of a search problem $S$ is a set $\mathcal{M}$ of rank-1 matrices $M \in \mathbb{F}^{\mathcal{X} \times \mathcal{Y}}$ such that $\sum_{M \in \mathcal{M}} M=\mathbb{1}$ and each $M \in \mathcal{M}$ is monochromatic for $S$. The size of the partition is $|\mathcal{M}|$. The $\mathbb{F}$-partition number $\chi_{\mathbb{F}}(S)$ is the least size of an $\mathbb{F}$-partition of $S$. In the following characterization, recall that we use $\mathrm{SP}_{\mathbb{F}}$ and $\mathrm{mSP}_{\mathbb{F}}$ to denote (monotone) span program complexity.

- Theorem 5 ([21]). For any boolean function $f$ and any field $\mathbb{F}, \mathrm{SP}_{\mathbb{F}}(f)=\chi_{\mathbb{F}}(\mathrm{KW}(f))$. Furthermore, if $f$ is monotone then $\operatorname{mSP}_{\mathbb{F}}(f)=\chi_{\mathbb{F}}\left(\mathrm{KW}^{+}(f)\right)$.

Nullstellensatz. Let $P:=\left\{p_{1}=0, p_{2}=0, \ldots, p_{m}=0\right\}$ be an unsatisfiable system of polynomial equations in $\mathbb{F}\left[z_{1}, z_{2}, \ldots, z_{n}\right]$ for a field $\mathbb{F}$. An $\mathbb{F}$-Nullstellensatz refutation of $P$ is a sequence of polynomials $q_{1}, q_{2}, \ldots, q_{m} \in \mathbb{F}\left[z_{1}, z_{2}, \ldots, z_{n}\right]$ such that $\sum_{i=1}^{m} q_{i} p_{i}=1$ where the equality is syntactic. The degree of the refutation is $\max _{i} \operatorname{deg}\left(q_{i} p_{i}\right)$. The $\mathbb{F}$-Nullstellensatz degree of $P$, denoted $\mathrm{NS}_{\mathbb{F}}(P)$, is the least degree of an $\mathbb{F}$-Nullstellensatz refutation of $P$.

Moreover, if $F$ is a $k$-CNF formula (or a boolean $k$-CSP), we often tacitly think of it as a polynomial system $P_{F}$ by using the standard encoding (e.g., $\left.\left(z_{1} \vee \neg z_{2}\right) \rightsquigarrow\left(1-z_{1}\right) z_{2}=0\right)$ and also including the boolean axioms $z_{i}^{2}-z_{i}=0$ in $P_{F}$ if we are working over $\mathbb{F} \neq \mathbb{F}_{2}$.

Lifting theorems. Let $S \subseteq\{0,1\}^{n} \times \mathcal{O}$ be a boolean search problem and $g: \mathcal{X} \times \mathcal{Y} \rightarrow\{0,1\}$ a two-party function, usually called a gadget. The composed search problem $S \circ g^{n} \subseteq$ $\mathcal{X}^{n} \times \mathcal{Y}^{n} \times \mathcal{O}$ is defined as follows: Alice holds $x \in \mathcal{X}^{n}$, Bob holds $y \in \mathcal{Y}^{n}$, and their goal is to find an $o \in S(z)$ where $z:=g^{n}(x, y)=\left(g\left(x_{1}, y_{1}\right), \ldots, g\left(x_{n}, y_{n}\right)\right)$. We focus on the usual index gadget $\mathrm{Ind}_{m}:[m] \times\{0,1\}^{m} \rightarrow\{0,1\}$ given by $\operatorname{InD}_{m}(x, y):=y_{x}$.

The main results of $[23,48]$ can be summarized as follows (we define more terms below).

- Theorem 6. Let $k \geq 1$ be a constant and let $m=m(n):=n^{C}$ for a large enough constant $C \geq 1$. Then for any an unsatisfiable boolean n-variate $k$-CSP $F$,

$$
\begin{array}{ll}
\text { [23]: } & \operatorname{PLS}^{c c}\left(S(F) \circ \mathrm{IND}_{m}^{n}\right)=\operatorname{PLS}^{\mathrm{dt}}(S(F)) \cdot \Theta(\log n), \\
\text { [48]: } & \operatorname{PPA}^{c \mathrm{c}}\left(S(F) \circ \mathrm{IND}_{m}^{n}\right)=\operatorname{PPA}^{\mathrm{dt}}(S(F)) \cdot \Theta(\log n), \\
\text { [48]: } & \log \chi_{\mathbb{F}}\left(S(F) \circ \mathrm{IND}_{m}^{n}\right)=\mathrm{NS}_{\mathbb{F}}(F) \cdot \Theta(\log n), \quad \forall \mathbb{F} \in\left\{\mathbb{F}_{p}, \mathbb{R}\right\} .
\end{array}
$$

For aesthetic reasons, we have used $\operatorname{PLS}^{\mathrm{dt}}(S(F))$ here to denote the Resolution width of $F$ (introduced in [10]), which is how the result of [23] was originally stated. (But one can check that the query analogue of PLS, obtained by replacing protocols with decision trees in Definition 14, is indeed equivalent to Resolution width.) We also could not resist incorporating our new characterizations of $\mathrm{PPA}^{c c}$ and $P$ PA $^{\text {dt }}$ to interpret the result of [48] specialized to $\mathbb{F}_{2}$.

## 4 Proofs of Separations

In this section, we show lower bounds for $\mathcal{C}$-SAT against monotone circuits (Theorem 1) and monotone span programs (Theorem 2), plus some bonus results (PPAD ${ }^{c c} \nsubseteq$ PLS $^{c c}$, PPADS ${ }^{c c} \nsubseteq$ PPA $^{\text {cc }}$, Nullstellensatz degree vs. Cutting Planes).

### 4.1 Reduction

The key to our lower bounds is a new reduction. We show that a lifted version of $S\left(F_{\text {bool }}\right)$, where $F$ is an unsatisfiable $\mathcal{C}$-CSP, reduces to the monotone Karchmer-Wigderson game for $\mathcal{C}$-Sat. Note that we require $F$ to be unsatisfiable over its original alphabet $\Sigma$, but the reduction is from the booleanized (and hence easier-to-refute) version of $F$.

- Lemma 7. Let $F$ be an unsatisfiable $\mathcal{C}$-CSP. Then $S\left(F_{\mathrm{bool}}\right) \circ \mathrm{IND}_{m}^{n}$ reduces to $\mathrm{KW}^{+}\left(\mathcal{C}-\mathrm{SAT}_{n m}\right)$.
Proof. Suppose the $\mathcal{C}$-CSP $F$ consists of $k$-ary constraints $C_{1}, \ldots, C_{t}$ applied to variables $z_{1}, \ldots, z_{n}$. We reduce $S\left(F_{\text {bool }}\right) \circ \operatorname{IND}_{m}^{n} \subseteq[m]^{n} \times\left(\{0,1\}^{m}\right)^{n} \times[t]$ to the problem $\mathrm{KW}^{+}(f) \subseteq$ $f^{-1}(1) \times f^{-1}(0) \times[N]$ where $f:=\mathcal{C}$-SAT $_{m n}$ over $N:=|\mathcal{C}|(m n)^{k}$ input bits. The two parties compute locally as follows.
Alice: Given $\left(x_{1}, \ldots, x_{n}\right) \in[m]^{n}$, Alice constructs a $\mathcal{C}$-CSP over variables $\left\{v_{i, j}:(i, j) \in\right.$ $[n] \times[m]\}$ that is obtained from $F$ by renaming its variables $z_{1}, \ldots, z_{n}$ to $v_{1, x_{1}}, \ldots, v_{n, x_{n}}$ (in this order). Since $F$ was unsatisfiable, so is Alice's variable-renamed version of it. Thus, when interpreted as an indicator vector of constraints, Alice has constructed a 1-input of $\mathcal{C}$-SAT $\mathrm{Sm}_{m}$.
Bob: Given $y \in\left(\{0,1\}^{m}\right)^{n}$, Bob constructs a $\mathcal{C}$-CSP over variables $\left\{v_{i, j}:(i, j) \in[n] \times[m]\right\}$ as follows. We view $y$ naturally as a boolean assignment to the variables $v_{i, j}$. Bob includes in his $\mathcal{C}$-CSP instance all possible $\mathcal{C}$-constraints $C$ applied to the $v_{i, j}$ such that $C$ is satisfied under the assignment $y$ (i.e., $C(y)=1$ ). This is clearly a satisfiable $\mathcal{C}$-CSP instance, as the assignment $y$ satisfies all Bob's constraints. Thus, when interpreted as an indicator vector of constraints, Bob has constructed a 0 -input of $\mathcal{C}$ - Sat $_{m n}$.

It remains to argue that any solution to $\mathrm{KW}^{+}\left(\mathcal{C}-\mathrm{SAT}_{m n}\right)$ gives rise to a solution to $S\left(F_{\text {bool }}\right) \circ \operatorname{IND}_{m}^{n}$. Indeed, a solution to $\mathrm{KW}^{+}\left(\mathcal{C}-\mathrm{SAT}_{m n}\right)$ corresponds to a $\mathcal{C}$-constraint $C$ that is present in Alice's $\mathcal{C}$-CSP but not in Bob's. By Bob's construction, such a $C$ must be violated by the assignment $y$ (i.e., $C(y)=0$ ). Since all Alice's constraints involve only variables $v_{1, x_{1}}, \ldots, v_{n, x_{n}}$, the constraint $C$ must in fact be violated by the partial assignment to the said variables, which is $z=\operatorname{InD}_{m}^{n}(x, y)$. Thus the constraint of $F$ from which $C$ was obtained via renaming is a solution to $S\left(F_{\text {bool }}\right) \circ \operatorname{IND}_{m}^{n}$.

- Remark 8 (Generic reduction to Cnf-SAT). We claim that any problem $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{O}$ that lies in one of the known subclasses of TFNP ${ }^{\text {cc }}$ (as listed in Section 2) reduces efficiently to $\mathrm{KW}^{+}\left(k \mathrm{CNF}^{2} \mathrm{SAT}_{n}\right)$ for constant $k$ (one can even take $k=3$ by standard reductions). Let us sketch the argument for $S \in \mathrm{PPP}^{c c}$; after all, better reductions are known for $\mathrm{PLS}^{c c}$ and PPA ${ }^{\text {cc }}$, namely to Horn-Sat and 3Xor-Sat.

Proof Sketch. Let $\Pi:=\left(V, v^{*}, o_{v}, \Pi_{v}\right)$ be a PPP-protocol solving $S$ of cost $c:=\operatorname{PPP}^{c c}(S)$. We may assume wlog that all the $\Pi_{v}$ have constant communication cost $k \leq O(1)$ by embedding the protocol trees of the $\Pi_{v}$ as part of the implicitly described bipartite graph. In particular, we view each $\Pi_{v}$ as a function $\mathcal{X} \times \mathcal{Y} \rightarrow\{0,1\}^{k}$ where the output is interpreted according to some fixed map $\{0,1\}^{k} \rightarrow V$. Consider a set of $n:=k|V|\left(|V| \leq 2^{c}\right)$ boolean variables $\left\{z_{v, i}:(v, i) \in V \times[k]\right\}$ with the intuitive interpretation that $z_{v, i}$ is the $i$-th output bit of $\Pi_{v}$. We may encode the correctness conditions for $\Pi$ as an unsatisfiable $2 k$-CNF formula $F$ over the $z_{v, i}$ that has, for each $\{v, u\} \in\binom{V}{2}$, clauses requiring that the outputs of $\Pi_{v}$ and $\Pi_{u}$ (as encoded by the $z_{v, i}$ ) should point to distinct vertices. Finally, we note that computing the $i$-th output bit $\left(\Pi_{v}\right)_{i}: \mathcal{X} \times \mathcal{Y} \rightarrow\{0,1\}$ reduces to a large enough constant-size index gadget $\mathrm{IND}_{O_{(1)}}$ (which embeds any two-party function of communication complexity $k \leq O(1))$. Therefore $S$ naturally reduces to $S(F) \circ \operatorname{IND}_{O(1)}^{n}$, which by Lemma 7 reduces to $\mathrm{KW}^{+}\left(2 k{\left.\mathrm{CnF}-\mathrm{SAT}_{O(n)}\right)}\right.$, as desired.

### 4.2 Monotone circuit lower bounds

Xor-Sat. The easiest result to prove is Theorem 1: an exponential monotone circuit lower bound for 3Xor-Sat $n_{n}$. By the characterization of [55] it suffices to show

$$
\begin{equation*}
\operatorname{PLS}^{\mathrm{cc}}\left(\mathrm{KW}^{+}\left(3 \mathrm{XoR}^{-\mathrm{SAT}_{n}}\right)\right) \geq n^{\Omega(1)} \tag{1}
\end{equation*}
$$

Urquhart [62] exhibited unsatisfiable $n$-variate 3Xor-CSPs $F$ (aka Tseitin formulas) requiring linear Resolution width, that is, $\mathrm{PLS}^{\mathrm{dt}}(S(F)) \geq \Omega(n)$ in our notation. Hence Theorem 6 implies that $\mathrm{PLS}^{c \mathrm{cc}}\left(S(F) \circ \operatorname{InD}_{m}^{n}\right) \geq \Omega(n)$ for some $m=n^{O(1)}$. By the reduction in Lemma 7, we get that $\operatorname{PLS}^{c c}\left(\mathrm{KW}^{+}\left(3 \mathrm{Xor}-\mathrm{SAT}_{n m}\right)\right) \geq \Omega(n)$. (Note that 3Xor has a boolean alphabet, so $F=F_{\text {bool }}$.) This yields the claim (1) by reparameterizing the number of variables.
$\operatorname{Lin}(\mathbf{F})$-Sat. More generally, we can prove a similar lower bound over any field $\mathbb{F} \in\left\{\mathbb{F}_{p}, \mathbb{R}\right\}$ :

$$
\begin{equation*}
\operatorname{PLS}^{c c}\left(\mathrm{KW}^{+}\left(3 \operatorname{LIN}(\mathbb{F})-\operatorname{SAT}_{n}\right)\right) \geq n^{\Omega(1)} \tag{2}
\end{equation*}
$$

Fix such an $\mathbb{F}$ henceforth. This time we start with a $k \operatorname{LiN}(\mathbb{F})$-CSP introduced in [13] for $\mathbb{F}=\mathbb{F}_{p}$ (aka mod-p Tseitin formulas), but the definition generalizes to any field. The CSP is constructed based on a given directed graph $G=(V, E)$ that is regular: in-deg $(v)=$ out-deg $(v)=k / 2$ for all $v \in V$. Fix also a distinguished vertex $v^{*} \in V$. Then $F=F_{G, \mathbb{F}}$ is defined as the following $k \operatorname{LIN}(\mathbb{F})$-CSP over variables $\left\{z_{e}: e \in E\right\}$ :

$$
\begin{equation*}
\forall v \in V: \quad \sum_{(v, u) \in E} z_{(v, u)}-\sum_{(u, v) \in E} z_{(u, v)}=\mathbb{1}_{v^{*}}(v), \tag{G,F}
\end{equation*}
$$

where $\mathbb{1}_{v^{*}}\left(v^{*}\right)=1$ and $\mathbb{1}_{v^{*}}(v)=0$ for $v \neq v^{*}$. This system is unsatisfiable because the sum over $v \in V$ of the RHS equals 1 whereas the sum of the LHS equals 0 (each variable appears once with a positive sign, once with a negative sign).

We claim that the booleanized $k$-CSP $F_{\text {bool }}$ (more precisely, its natural $k$-CNF encoding) has linear Resolution width, that is, $\mathrm{PLS}^{\mathrm{dt}}\left(S\left(F_{\mathrm{bool}}\right)\right) \geq \Omega(n)$ in our notation. Indeed, the constraints of $F_{\text {bool }}$ are $k / 2$-robust in the sense that if a partial assignment $\rho \in\{0,1, *\}^{k}$ fixes the value of a constraint of $F_{\text {bool }}$, then $\rho$ must set more than $k / 2$ variables. Alekhnovich et al. [1, Theorem 3.1] show that if $k$ is a large enough constant, there exist regular expander graphs $G$ such that $F_{\text {bool }}$ (or any $k$-CSP with $\Omega(k)$-robust constraints) has Resolution width $\Omega(n)$, as desired.

Combining the above with the lifting theorem in Theorem 6 and the reduction in Lemma 7 yields $\operatorname{PLS}^{c c}\left(k \operatorname{Lin}(\mathbb{F})-\operatorname{SAT}_{n}\right) \geq n^{\Omega(1)}$ for large enough $k$. Finally, we can reduce the arity from $k$ to 3 by a standard trick. For example, given the linear constraint $a_{1} v_{1}+a_{2} v_{2}+a_{3} v_{3}+a_{4} v_{4}=$ $a_{0}$ we can introduce a new auxiliary variable $u$ and two equations $a_{1} v_{1}+a_{2} v_{2}+u=0$ and $-u+a_{3} v_{3}+a_{4} v_{4}=a_{0}$. In general, we replace each equation on $k>3$ variables with a collection of $k-2$ equations by introducing $k-3$ auxiliary variables to create an equisatisfiable instance. This shows that $k \operatorname{LIN}(\mathbb{F})-\mathrm{SAT}_{n}$ reduces to (i.e., is a monotone projection of) $3 \operatorname{Lin}(\mathbb{F})$ - $\mathrm{SAT}_{k n}$, which concludes the proof of (2).
$\mathbf{P P A D}^{\mathbf{c c}} \nsubseteq \mathbf{P L S}^{\mathbf{c c}}$ via End-of-Line. Consider the $\mathbb{R}$-linear system $F=F_{G, \mathbb{R}}$ defined above. We observe that $S\left(F_{\mathrm{bool}}\right)$ is in fact equivalent to (a query version of) the PPAD-complete End-of-Line problem. In the End-of-Line problem, we are given a directed graph of in/out-degree at most 1 and a distinguished source vertex $v^{*}$ (in-degree 0 ); the goal is to find a sink or a source distinct from $v^{*}$ (cf. Definition 15). On the other hand, in $S\left(F_{\text {bool }}\right)$
we are given a boolean assignment $z \in\{0,1\}^{E}$, which can be interpreted as (the indicator vector of) a subset of edges defining a (spanning) subgraph $G_{z}$ of $G$; the goal is to find a vertex $v \in V$ such that either
(1) $v=v^{*}$ and out-deg $(v) \neq \operatorname{in}-\operatorname{deg}(v)+1$ in $G_{z}$; or
(2) $v \neq v^{*}$ and out- $\operatorname{deg}(v) \neq \operatorname{in}-\operatorname{deg}(v)$ in $G_{z}$.

The only essential difference between $S\left(F_{\mathrm{bool}}\right)$ and End-of-Line is that the graph $G_{z}$ can have in/out-degree a large constant $k / 2$ rather than 1 . But there is a standard reduction between the two problems [46]: we may locally interpret a vertex $v \in V\left(G_{z}\right)$ with out-deg $(v)=$ in- $\operatorname{deg}(v)=\ell$ as $\ell$ distinct vertices of in/out-degree 1 . This reduction also shows that the lifted problem $S\left(F_{\text {bool }}\right) \circ \mathrm{InD}_{m}$ for $m=n^{O(1)}$ admits a $O(\log n)$-cost PPAD-protocol, and is thus in $\mathrm{PPAD}^{c c}$. By contrast, we proved above that this problem is not in $\mathrm{PLS}^{c c}$ (for appropriate $G$ ).

- Remark 9 (Algebraic partitions for $\mathrm{PPAD}^{c \mathrm{c}}$ ). We claim that every problem $S \in \mathrm{PPAD}^{\text {cc }}$ admits a small $\mathbb{Z}$-partition, and hence a small $\mathbb{F}$-partition over any field $\mathbb{F}$. More precisely, we argue that $\log \chi_{\mathbb{Z}}(S) \leq O\left(\operatorname{PPAD}^{c c}(S)\right)$. Indeed, let $\Pi:=\left(V, v^{*}, o_{v}, \Pi_{v}\right)$ be an optimal PPADprotocol for $S$. We define a $\mathbb{Z}$-partition $\mathcal{M}$ by describing it as a nondeterministic protocol for $S$ whose accepting computations output weights in $\mathbb{Z}$ (interpreted as values of the entries of an $M \in \mathcal{M}$ ): On input $(x, y)$, guess a vertex $v \in V$; if $v$ is a sink in $G_{x, y}$, accept with weight 1 ; if $v$ is a source distinct from $v^{*}$, accept with weight -1 ; otherwise reject (i.e., weight 0 ). This protocol accepts with overall weight \#(sinks) - \#(non-distinguished sources) $=1$ on every input $(x, y)$, as desired.

A similar argument yields an analogous query complexity bound $\mathrm{NS}_{\mathbb{Z}}(F) \leq O\left(\operatorname{PPAD}^{\mathrm{dt}}(S(F))\right)$ where $\operatorname{PPAD}^{\mathrm{dt}}(S)$ is the least cost of a PPAD-decision tree (Definition 15) solving $S$.

Nullstellensatz vs. Cutting Planes. By the above remark, $F_{\text {bool }}$ for $F=F_{G, \mathbb{F}}$ admits a lowdegree - in fact, constant-degree - Nullstellensatz refutation over any field $\mathbb{F}$. Nullstellensatz degree behaves well with respect to compositions: if we compose $F_{\text {bool }}$ with a gadget $\mathrm{InD}_{m}^{n}$, $m=n^{O(1)}$ (see, e.g., $[23, \S 8]$ how this can be done), the Nullstellensatz degree can only increase by the query complexity of the gadget, which is $O(\log n)$ for $\operatorname{InD}_{m}^{n}$. This gives us an $n^{O(1)}$-variate boolean $k$-CSP $F^{\prime}:=F_{\text {bool }} \circ \operatorname{IND}_{m}^{n}$ (where $k$ is constant $\left.[23, \S 8]\right)$ such that $\mathrm{NS}_{\mathbb{F}}\left(F^{\prime}\right) \leq O(\log n)$. On the other hand, we can invoke the strong version of the main result of [23]: if $F$ has Resolution width $w$, then $F \circ \mathrm{InD}_{m}^{n}$ requires Cutting Planes refutations of length $n^{\Omega(w)}$. In summary, $F^{\prime}$ witnesses that $\mathbb{F}$-Nullstellensatz can be exponentially more powerful than log of Cutting Planes length.

### 4.3 Monotone span program lower bounds

Let us prove Theorem 2: $3 \mathrm{Lin}(\mathbb{R})$ - $\mathrm{SAT}_{n}$ requires exponential-size monotone $\mathbb{F}_{p}$-span programs, that is,

$$
\begin{equation*}
\chi_{\mathbb{F}_{p}}\left(\mathrm{KW}^{+}\left(3 \operatorname{Lin}(\mathbb{R})-\mathrm{SAT}_{n}\right)\right) \geq n^{\Omega(1)} \tag{3}
\end{equation*}
$$

Using Theorem 6 and Lemma 7 similarly as in Subsection 4.2, it suffices to show that $\mathrm{NS}_{\mathbb{F}_{p}}\left(F_{\text {bool }}\right) \geq n^{\Omega(1)}$, for some unsatisfiable $k \operatorname{Lin}(\mathbb{R})$-CSP $F$ where $k$ is a constant. To this end, we consider an $\mathbb{R}$-linear system $F=F_{G, U, \mathbb{R}}$ that generalizes $F_{G, \mathbb{R}}$ defined above:

$$
\begin{equation*}
\forall v \in V: \quad \sum_{(v, u) \in E} z_{(v, u)}-\sum_{(u, v) \in E} z_{(u, v)}=\mathbb{1}_{U}(v) \tag{G,U,R}
\end{equation*}
$$

where $\mathbb{1}_{U}: V \rightarrow\{0,1\}$ is the indicator function for $U \subseteq V$. This is unsatisfiable as long as $U \neq \emptyset$. Combinatorially, the boolean search problem $S\left(F_{\text {bool }}\right)$ can be interpreted as an End-OF- $\ell$-Lines problem for $\ell:=|U|$ : given a graph with distinguished source vertices $U$, find a sink or a source not in $U$. It is important to have many distinguished sources, $|U| \geq n^{\Omega(n)}$, as otherwise $S\left(F_{\text {bool }}\right)$ is in $\operatorname{PPAD}^{\mathrm{dt}}[31]$ and hence $F_{\text {bool }}$ has too low an $\mathbb{F}_{p}$-Nullstellensatz degree (by Remark 9).

Nullstellensatz lower bound. To show $\mathrm{NS}_{\mathbb{F}_{p}}\left(F_{\mathrm{bool}}\right) \geq n^{\Omega(1)}$ for an appropriate $F=F_{G, U, \mathbb{R}}$, we adapt a result of Beame and Riis [8]. They proved a Nullstellensatz lower bound for a related bijective pigeonhole principle $P_{n}$ whose underlying graph has unbounded degree; we obtain a bounded-degree version of their result by a reduction.

- Lemma 10 ([8,§8]). Fix a prime $p$. The following system of polynomial equations over variables $\left\{x_{i j}:(i, j) \in D \times R\right\}$, where $|D|=n$ and $|R|=n-n^{\Omega(1)}$, requires $\mathbb{F}_{p}$-Nullstellensatz degree $n^{\Omega(1)}$ :
$\forall i \in D: \quad \sum_{j \in R} x_{i j}=1 \quad$ "each pigeon occupies a hole",
$\forall j \in R: \quad \sum_{i \in D} x_{i j}=1 \quad$ "each hole houses a pigeon",
(iv) $\forall j \in R,\left\{i, i^{\prime}\right\} \in\binom{D}{2}: \quad x_{i j} x_{i^{\prime} j}=0 \quad$ "no hole houses two pigeons".

We construct a natural bounded-degree version $G$ of the complete bipartite graph $D \times R$ and show that each constraint of $F_{\text {bool }}$ for $F=F_{G, U, \mathbb{R}}$ is a low-degree $\mathbb{F}_{p}$-Nullstellensatz consequence of $P_{n}$. Hence, if $F_{\text {bool }}$ admits a low-degree $\mathbb{F}_{p}$-Nullstellensatz proof, so does $P_{n}$ (see, e.g., [13, Lemma 1] for composing proofs), which contradicts Lemma 10.

The directed graph $G=(V, E)$ is obtained from the complete bipartite graph $D \times R$ as illustrated in Figure 2 (for $|D|=4$ and $|R|=3$ ). Specifically, each vertex of degree $d$ in $D \times R$ is replaced with a binary tree of height $\log d$. The result is a layered graph with the first and last layers identified with $D$ and $R$, respectively. We also add a "feedback" edge from each vertex in $R$ to a vertex in $D$ according to some arbitrary injection $R \rightarrow D$ (dashed edges in Figure 2). The vertices in $D$ not incident to feedback edges will form the set $U$ (singleton in Figure 2).

This defines a boolean 3-CSP $F_{\mathrm{bool}}$ for $F=F_{G, U, \mathbb{R}}$ over variables $\left\{z_{e}: e \in E\right\}$. In order to reduce $P_{n}$ to $F_{\text {bool }}$, we define an affine map between the variables $x_{i j}$ of $P_{n}$ and $z_{e}$ of $F_{\text {bool }}$. Namely, for a feedback edge $e$ we set $z_{e}:=1$, and for every other $e=(v, u)$ we set

$$
z_{(v, u)}:=\sum_{i \in D_{v}} x_{i j},
$$

where $\quad D_{v}:=\{i \in D: v$ is reachable from $i$ without using feedback edges $\}$,
$R_{u}:=\{j \in R: j$ is reachable from $u$ without using feedback edges $\}$.
Note in particular that this map naturally identifies the edge-variables $z_{e}$ in the middle of $G$ (yellow edges) with the variables $x_{i j}$ of $P_{n}$. The other variables $z_{e}$ are simply affinely dependent on the middle edge-layer. We then show that from the equations of $P_{n}$ we can derive each constraint of $F_{\text {bool }}$. Recall that the constraint for $v \in V$ requires that the out-flow $\sum_{(v, u) \in E} z_{(v, u)}$ equals the in-flow $\sum_{(u, v) \in E} z_{(u, v)}$ (plus 1 iff $\left.v \in U\right)$.
$v \notin \boldsymbol{D} \cup \boldsymbol{R}$ : Suppose $v$ is on the left side of $G$ (right side is handled similarly) so that $z_{(v, u)}=\sum_{j \in R_{u}} x_{i j}$ for some fixed $i \in D$. The out-flow is

$$
\begin{equation*}
\sum_{(v, u) \in E} z_{(v, u)}=\sum_{(v, u) \in E} \sum_{j \in R_{u}} x_{i j}=\sum_{j \in R_{v}} x_{i j} \tag{4}
\end{equation*}
$$



Figure 2 Graph $G=(V, E)$, a bounded-degree version of the biclique $D \times R$.

On the other hand, $v$ has a unique incoming edge $\left(u^{*}, v\right)$ so the in-flow is $\sum_{(u, v) \in E} z_{(u, v)}=$ $z_{\left(u^{*}, v\right)}=\sum_{j \in R_{v}} x_{i j}$, which equals (4).
$v \in D:$ (Case $v \in R$ is handled similarly). The in-flow equals 1 (either $v \in U$ so that we have the +1 term from $\mathbb{1}_{U}(v)$; or $v \notin U$ and the value of a feedback-edge variable gives +1 ). The out-flow equals $\sum_{j \in R_{v}} x_{i j}=\sum_{j \in R} x_{i j}=1$ by (4), $R_{v}=R$, and (ii).
Finally, we can verify the boolean axioms $z_{e}^{2}=z_{e}$. This holds trivially for feedback edges $e$. Let $e=(v, u)$ be an edge in the left side of $G$ (right side is similar) so that $z_{e}=\sum_{j \in R_{u}} x_{i j}$ for some fixed $i \in D$. We have $z_{e}^{2}=\left(\sum_{j \in R_{u}} x_{i j}\right)^{2}=\sum_{j \in R_{u}} x_{i j}^{2}=\sum_{j \in R_{u}} x_{i j}=z_{e}$ by (iii) and the boolean axioms for $P_{n}$.

This concludes the reduction and hence the proof of (3).

PPADS $^{\text {cc }} \nsubseteq$ PPA $^{\text {cc }}$ via End-of- $\boldsymbol{\ell}$-Lines. It is straightforward to check that $F_{\text {bool }}$ for $F=$ $F_{G, U, \mathbb{R}}$ is in the query class PPADS ${ }^{\mathrm{dt}}$ (Definition 16). In particular, in the PPADS-decision tree, we can define the distinguished vertex $v^{*}$ as being associated with any vertex from $U$. Similarly, the lifted problem $S^{\prime}:=S\left(F_{\mathrm{bool}}\right) \circ \mathrm{IND}_{n}^{m}$ for $m=n^{O(1)}$ is in the communication class PPADS ${ }^{c c}$. By contrast, we just proved that $\chi_{\mathbb{F}_{2}}\left(S^{\prime}\right) \geq n^{\Omega(1)}$, which implies that $S^{\prime} \notin \mathrm{PPA}^{\mathrm{cc}}$.

## 5 Proofs of Characterizations

Due to space constraints, the proofs of Theorem 3 and Theorem 4 are omitted from this extended abstract. See the full version [24] for complete proofs.

## References

1 Michael Alekhnovich, Eli Ben-Sasson, Alexander Razborov, and Avi Wigderson. Pseudorandom Generators in Propositional Proof Complexity. SIAM Journal on Computing, 34(1):67-88, 2004. doi:10.1137/S0097539701389944.
2 László Babai, Peter Frankl, and Janos Simon. Complexity Classes in Communication Complexity Theory. In Proceedings of the 27th Symposium on Foundations of Computer Science (FOCS), pages 337-347. IEEE, 1986. doi:10.1109/SFCS.1986.15.
3 László Babai, Anna Gál, and Avi Wigderson. Superpolynomial Lower Bounds for Monotone Span Programs. Combinatorica, 19(3):301-319, 1999. doi:10.1007/s004930050058.

4 Yakov Babichenko, Shahar Dobzinski, and Noam Nisan. The Communication Complexity of Local Search. Technical report, arXiv, 2018. arXiv:1804.02676.
5 Yakov Babichenko and Aviad Rubinstein. Communication Complexity of Approximate Nash Equilibria. In Proceedings of the 49th Symposium on Theory of Computing (STOC), pages 878-889. ACM, 2017. doi:10.1145/3055399. 3055407.
6 Paul Beame, Stephen Cook, Jeff Edmonds, Russell Impagliazzo, and Toniann Pitassi. The Relative Complexity of NP Search Problems. Journal of Computer and System Sciences, 57(1):3-19, 1998. doi:10.1006/jcss.1998.1575.
7 Paul Beame, Toniann Pitassi, and Nathan Segerlind. Lower Bounds for Lovász-Schrijver Systems and Beyond Follow from Multiparty Communication Complexity. SIAM Journal on Computing, 37(3):845-869, 2007. doi:10.1137/060654645.
8 Paul Beame and Søren Riis. More on the Relative Strength of Counting Principles. In Proceedings of the DIMACS Workshop on Proof Complexity and Feasible Arithmetics, volume 39, pages 13-35, 1998.
9 Aleksandrs Belovs, Gábor Ivanyos, Youming Qiao, Miklos Santha, and Siyi Yang. On the Polynomial Parity Argument Complexity of the Combinatorial Nullstellensatz. In Proceedings of the 32nd Computational Complexity Conference (CCC), volume 79, pages 30:1-30:24. Schloss Dagstuhl, 2017. doi:10.4230/LIPIcs.CCC.2017.30.
10 Eli Ben-Sasson and Avi Wigderson. Short Proofs Are Narrow—Resolution Made Simple. Journal of the ACM, 48(2):149-169, 2001. doi:10.1145/375827. 375835.
11 Andrei Bulatov. A Dichotomy Theorem for Nonuniform CSPs. In Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS), pages 319-330, 2017. doi: 10.1109/FOCS. 2017. 37.

12 Joshua Buresh-Oppenheim and Tsuyoshi Morioka. Relativized NP search problems and propositional proof systems. In Proceedings of the 19th Conference on Computational Complexity (CCC), pages 54-67, 2004. doi:10.1109/CCC.2004.1313795.
13 Sam Buss, Dima Grigoriev, Russell Impagliazzo, and Toniann Pitassi. Linear Gaps between Degrees for the Polynomial Calculus Modulo Distinct Primes. Journal of Computer and System Sciences, 62(2):267-289, 2001. doi:10.1006/jcss.2000.1726.
14 Siu Man Chan. Just a Pebble Game. In Proceedings of the 28th Conference on Computational Complexity (CCC), pages 133-143, 2013. doi:10.1109/CCC.2013.22.
15 Siu Man Chan and Aaron Potechin. Tight Bounds for Monotone Switching Networks via Fourier Analysis. Theory of Computing, 10(15):389-419, 2014. doi:10.4086/toc. 2014. v010a015.
16 Stephen Cook, Yuval Filmus, and Dai Tri Man Lê. The Complexity of the Comparator Circuit Value Problem. ACM Transactions on Computation Theory, 6(4):15:1-15:44, 2014. doi:10.1145/2635822.
17 Constantinos Daskalakis and Christos Papadimitriou. Continuous Local Search. In Proceedings of the 22nd Symposium on Discrete Algorithms (SODA), pages 790-804. SIAM, 2011. URL: http://dl.acm.org/citation.cfm?id=2133036.2133098.

18 Susanna de Rezende, Jakob Nordström, and Marc Vinyals. How Limited Interaction Hinders Real Communication (and What It Means for Proof and Circuit Complexity). In Proceedings of the $5^{7}$ th Symposium on Foundations of Computer Science (FOCS), pages 295-304. IEEE, 2016. doi:10.1109/FOCS. 2016.40.
19 John Fearnley, Spencer Gordon, Ruta Mehta, and Rahul Savani. End of Potential Line. Technical report, arXiv, 2018. arXiv:1804.03450.
20 Tomás Feder and Moshe Vardi. The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory. SIAM Journal on Computing, 28(1):57-104, 1998. doi:10.1137/S0097539794266766.

21 Anna Gál. A Characterization of Span Program Size and Improved Lower Bounds for Monotone Span Programs. Computational Complexity, 10(4):277-296, 2001. doi:10.1007/ s000370100001.
22 Anat Ganor and Karthik C. S. Communication Complexity of Correlated Equilibrium with Small Support. In Proceedings of the 22nd International Conference on Randomization and Computation ( $R$ ANDOM), volume 116, pages 12:1-12:16. Schloss Dagstuhl, 2018. doi: 10.4230/LIPIcs.APPROX-RANDOM. 2018.12.

23 Ankit Garg, Mika Göös, Pritish Kamath, and Dmitry Sokolov. Monotone Circuit Lower Bounds from Resolution. In Proceedings of the 50th Symposium on Theory of Computing (STOC), pages 902-911. ACM, 2018. doi:10.1145/3188745.3188838.
24 Mika Göös, Pritish Kamath, Robert Robere, and Dmitry Sokolov. Adventures in Monotone Complexity and TFNP. Technical Report TR18-163, Electronic Colloquium on Computational Complexity (ECCC), 2018. URL: https://eccc.weizmann.ac.il/report/2018/ 163/.
25 Mika Göös and Toniann Pitassi. Communication Lower Bounds via Critical Block Sensitivity. In Proceedings of the 46 th Symposium on Theory of Computing (STOC), pages 847-856. ACM, 2014. doi:10.1145/2591796. 2591838.
26 Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic Communication vs. Partition Number. In Proceedings of the 56th Symposium on Foundations of Computer Science (FOCS), pages 1077-1088. IEEE, 2015. doi:10.1109/FOCS. 2015.70.
27 Mika Göös, Toniann Pitassi, and Thomas Watson. The Landscape of Communication Complexity Classes. In Proceedings of the 43 rd International Colloquium on Automata, Languages, and Programming (ICALP), pages 86:1-86:15. Schloss Dagstuhl, 2016. doi: 10.4230/LIPIcs.ICALP. 2016.86.

28 Mika Göös and Aviad Rubinstein. Near-Optimal Communication Lower Bounds for Approximate Nash Equilibria. In Proceedings of the 59th Symposium on Foundations of Computer Science (FOCS), 2018. To appear. arXiv:1805.06387.
29 Michelangelo Grigni and Michael Sipser. Monotone Complexity. In Proceedings of the London Mathematical Society Symposium on Boolean Function Complexity, pages 57-75. Cambridge University Press, 1992. URL: http://dl.acm.org/citation.cfm?id=167687. 167706.

30 Bernd Halstenberg and Rüdiger Reischuk. Relations Between Communication Complexity Classes. Journal of Computer and System Sciences, 41(3):402-429, 1990. doi:10.1016/ 0022-0000 (90) 90027-I.
31 Alexandros Hollender and Paul Goldberg. The Complexity of Multi-source Variants of the End-of-Line Problem, and the Concise Mutilated Chessboard. Technical report, Electronic Colloquium on Computational Complexity (ECCC), 2018. URL: https://eccc.weizmann. ac.il/report/2018/120/.
32 Pavel Hubáček and Eylon Yogev. Hardness of Continuous Local Search: Query Complexity and Cryptographic Lower Bounds. In Proceedings of the 28th Symposium on Discrete Algorithms (SODA), pages 1352-1371, 2017. doi:10.1137/1.9781611974782.88.
33 Trinh Huynh and Jakob Nordström. On the Virtue of Succinct Proofs: Amplifying Communication Complexity Hardness to Time-Space Trade-Offs in Proof Complexity. In Proceedings of the 44 th Symposium on Theory of Computing (STOC), pages 233-248. ACM, 2012. doi:10.1145/2213977.2214000.

34 Russell Impagliazzo, Toniann Pitassi, and Alasdair Urquhart. Upper and lower bounds for tree-like cutting planes proofs. In Proceedings of the 9th Symposium on Logic in Computer Science (LICS), pages 220-228. IEEE, 1994. doi:10.1109/LICS.1994.316069.

35 David Johnson, Christos Papadimitriou, and Mihalis Yannakakis. How easy is local search? Journal of Computer and System Sciences, 37(1):79-100, 1988. doi:10.1016/ 0022-0000 (88) 90046-3.
36 Stasys Jukna. Boolean Function Complexity: Advances and Frontiers, volume 27 of Algorithms and Combinatorics. Springer, 2012.
37 Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require superlogarithmic depth. In Proceedings of the 20th Symposium on Theory of Computing (STOC), pages 539-550. ACM, 1988. doi:10.1145/62212.62265.
38 Mauricio Karchmer and Avi Wigderson. On span programs. In Proceedings of the 8th Structure in Complexity Theory Conference, pages 102-111, 1993. doi:10.1109/SCT. 1993. 336536.

39 Eyal Kushilevitz and Noam Nisan. Communication Complexity. Cambridge University Press, 1997.
40 László Lovász. On determinants, matchings, and random algorithms. In Proceedings of the 2nd Conference on Fundamentals of Computation Theory (FCT), pages 565-574, 1979.
41 Ernst Mayr and Ashok Subramanian. The complexity of circuit value and network stability. Journal of Computer and System Sciences, 44(2):302-323, 1992. doi:10.1016/ 0022-0000 (92) 90024-D.
42 Nimrod Megiddo and Christos Papadimitriou. On total functions, existence theorems and computational complexity. Theoretical Computer Science, 81(2):317-324, 1991. doi:10. 1016/0304-3975 (91) 90200-L.
43 Tsuyoshi Morioka. Logical Approaches to the Complexity of Search Problems: Proof Complexity, Quantified Propositional Calculus, and Bounded Arithmetic. PhD thesis, University of Toronto, 2005.
44 Ketan Mulmuley. A fast parallel algorithm to compute the rank of a matrix over an arbitrary field. Combinatorica, 7(1):101-104, 1987. doi:10.1007/BF02579205.
45 Igor Oliveira. Unconditional Lower Bounds in Complexity Theory. PhD thesis, Columbia University, 2015. doi:10.7916/D8ZP45KT.
46 Christos Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. Journal of Computer and System Sciences, 48(3):498-532, 1994. doi: 10.1016/S0022-0000(05)80063-7.

47 Toniann Pitassi and Robert Robere. Strongly Exponential Lower Bounds for Monotone Computation. In Proceedings of the 49th Symposium on Theory of Computing (STOC), pages 1246-1255. ACM, 2017. doi:10.1145/3055399. 3055478.
48 Toniann Pitassi and Robert Robere. Lifting Nullstellensatz to Monotone Span Programs over Any Field. In Proceedings of the 50th Symposium on Theory of Computing (STOC), pages 1207-1219. ACM, 2018. doi:10.1145/3188745.3188914.
49 Aaron Potechin. Bounds on Monotone Switching Networks for Directed Connectivity. Journal of the ACM, 64(4):29:1-29:48, 2017. doi:10.1145/3080520.
50 Pavel Pudlák. On extracting computations from propositional proofs (a survey). In Proceedings of the 30th Foundations of Software Technology and Theoretical Computer Science (FSTTCS), volume 8, pages 30-41. Schloss Dagstuhl, 2010. doi:10.4230/LIPIcs.FSTTCS. 2010. 30.

51 Pavel Pudlák and Jiří Sgall. Algebraic models of computation and interpolation for algebraic proof systems. DIMACS Series in Discrete Mathematics and Theoretical Computer Science, 39:279-295, 1998. doi:10.1090/dimacs/039.
52 Ran Raz and Pierre McKenzie. Separation of the Monotone NC Hierarchy. Combinatorica, 19(3):403-435, 1999. doi:10.1007/s004930050062.

53 Alexander Razborov. Lower bounds on monotone complexity of the logical permanent. Mathematical notes of the Academy of Sciences of the USSR, 37(6):485-493, 1985. doi: 10.1007/BF01157687.

54 Alexander Razborov. Lower bounds on the monotone complexity of some Boolean functions. Doklady Akademii Nauk USSR, 285:798-801, 1985.
55 Alexander Razborov. Unprovability of lower bounds on circuit size in certain fragments of bounded arithmetic. Izvestiya of the RAN, pages 201-224, 1995.
56 Robert Robere, Toniann Pitassi, Benjamin Rossman, and Stephen Cook. Exponential Lower Bounds for Monotone Span Programs. In Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS), pages 406-415. IEEE, 2016. doi:10.1109/ FOCS.2016.51.
57 Tim Roughgarden. Complexity Theory, Game Theory, and Economics. Technical report, arXiv, 2018. arXiv:1801.00734.
58 Tim Roughgarden and Omri Weinstein. On the Communication Complexity of Approximate Fixed Points. In Proceedings of the 57th Symposium on Foundations of Computer Science (FOCS), pages 229-238. IEEE, 2016. doi:10.1109/FOCS.2016.32.
59 Thomas Schaefer. The Complexity of Satisfiability Problems. In Proceedings of the 10th Symposium on Theory of Computing (STOC), pages 216-226. ACM, 1978. doi:10.1145/ 800133.804350.

60 Dmitry Sokolov. Dag-Like Communication and Its Applications. In Proceedings of the 12th Computer Science Symposium in Russia (CSR), pages 294-307. Springer, 2017. doi: 10.1007/978-3-319-58747-9_26.

61 Éva Tardos. The gap between monotone and non-monotone circuit complexity is exponential. Combinatorica, 8(1):141-142, 1988. doi:10.1007/BF02122563.
62 Alasdair Urquhart. Hard examples for resolution. Journal of the ACM, 34(1):209-219, 1987. doi:10.1145/7531. 8928.

63 Dmitriy Zhuk. A Proof of CSP Dichotomy Conjecture. In Proceedings of the 58th Symposium on Foundations of Computer Science (FOCS), pages 331-342, 2017. doi: 10.1109/FOCS. 2017. 38 .

## A Appendix: TFNP Class Definitions

For each TFNP subclass there is a canonical definition of its communication or query analogue: we simply let communication protocols or decision trees (rather than circuits) implicitly define the objects that appear in the original Turing machine definition. Each communication class $C^{\text {cc }}$ (resp. query class $\mathrm{C}^{\mathrm{dt}}$ ) is defined via a C-protocol (resp. C-decision tree) that solves a two-party search problem $S \subseteq\{0,1\}^{n / 2} \times\{0,1\}^{n / 2} \times \mathcal{O}$ (resp. $S \subseteq\{0,1\}^{n} \times \mathcal{O}$ ). The class $\mathrm{C}^{\text {cc }}$ (resp. $\mathrm{C}^{\mathrm{dt}}$ ) is then defined as the set of all $n$-bit search problems $S$ that admit a polylog $(n)$-cost C-protocol (resp. C-decision tree). We only define the communication analogues below with the understanding that a query version can be obtained by replacing mentions of a protocol $\Pi_{v}(x, y)$ by a decision tree $\mathcal{T}_{v}(x)$; the cost of a C-decision tree is defined as $\max _{v, x} \#$ (queries made by $\mathcal{T}_{v}(x)$ ). In what follows, sink means out-degree 0 , and source means in-degree 0 .

- Definition 11. (FP)

Syntax: $\Pi$ is a (deterministic) protocol outputting values in $\mathcal{O}$.
Object: n/a
Correctness: $\Pi(x, y) \in S(x, y)$.
Cost: $|\Pi|:=$ communication cost of $\Pi$.

Definition 12. (EoML)
Syntax: $V$ is a vertex set with a distinguished vertex $v^{*} \in V$. For each $v \in V: o_{v} \in \mathcal{O}$ and $\Pi_{v}$ is a protocol outputting a tuple $\left(s_{v}(x, y), p_{v}(x, y), \ell_{v}(x, y)\right) \in V \times V \times \mathbb{Z}$.
Object: Dag $G_{x, y}=(V, E)$ where $(v, u) \in E$ iff $s_{v}(x, y)=u, p_{u}(x, y)=v, \ell_{v}(x, y)>\ell_{u}(x, y)$.
Correctness: If $v^{*}$ is a sink or non-source in $G_{x, y}$, then $o_{v^{*}} \in S(x, y)$.
If $v \neq v^{*}$ is a sink or source in $G_{x, y}$, then $o_{v} \in S(x, y)$.
Cost: $\log |V|+\max _{v}\left|\Pi_{v}\right|$.

- Definition 13. (SoML)

Syntax: Same as in Definition 12.
Object: Same as in Definition 12.
Correctness: If $v^{*}$ is a sink or non-source in $G_{x, y}$, then $o_{v^{*}} \in S(x, y)$.
If $v \neq v^{*}$ is a $\operatorname{sink}$ in $G_{x, y}$, then $o_{v} \in S(x, y)$.
Cost: $\log |V|+\max _{v}\left|\Pi_{v}\right|$.

- Definition 14. (PLS)

Syntax: $V$ is a vertex set. For each $v \in V: o_{v} \in \mathcal{O}$ and $\Pi_{v}$ is a protocol outputting a pair $\left(s_{v}(x, y), \ell_{v}(x, y)\right) \in V \times \mathbb{Z}$.
Object: Dag $G_{x, y}=(V, E)$ where $(v, u) \in E$ iff $s_{v}(x, y)=u$ and $\ell_{v}(x, y)>\ell_{u}(x, y)$.
Correctness: If $v$ is a sink in $G_{x, y}$, then $o_{v} \in S(x, y)$.
Cost: $\log |V|+\max _{v}\left|\Pi_{v}\right|$.

- Definition 15. (PPAD)

Syntax: $V$ is a vertex set with a distinguished vertex $v^{*} \in V$. For each $v \in V: o_{v} \in \mathcal{O}$ and $\Pi_{v}$ is a protocol outputting a pair $\left(s_{v}(x, y), p_{v}(x, y)\right) \in V \times V$.
Object: Digraph $G_{x, y}=(V, E)$ where $(v, u) \in E$ iff $s_{v}(x, y)=u$ and $p_{u}(x, y)=v$.
Correctness: If $v^{*}$ is a sink or non-source in $G_{x, y}$, then $o_{v^{*}} \in S(x, y)$.
If $v \neq v^{*}$ is a sink or source in $G_{x, y}$, then $o_{v} \in S(x, y)$.
Cost: $\log |V|+\max _{v}\left|\Pi_{v}\right|$.
Definition 16. (PPADS)
Syntax: Same as in Definition 15.
Object: Same as in Definition 15.
Correctness: If $v^{*}$ is a sink or non-source in $G_{x, y}$, then $o_{v^{*}} \in S(x, y)$.
If $v \neq v^{*}$ is a $\operatorname{sink}$ in $G_{x, y}$, then $o_{v} \in S(x, y)$.
Cost: $\log |V|+\max _{v}\left|\Pi_{v}\right|$.
Definition 17. (PPA)
Syntax: $V$ is a vertex set with a distinguished vertex $v^{*} \in V$. For each $v \in V: o_{v} \in \mathcal{O}$ and $\Pi_{v}$ is a protocol outputting a subset $\Pi_{v}(x, y) \subseteq V$ of size at most 2 .
Object: Undirected graph $G_{x, y}=(V, E)$ where $\{v, u\} \in E$ iff $v \in \Pi_{u}(x, y)$ and $u \in \Pi_{v}(x, y)$.
Correctness: If $v^{*}$ has degree $\neq 1$ in $G_{x, y}$, then $o_{v^{*}} \in S(x, y)$.
If $v \neq v^{*}$ has degree $\neq 2$ in $G_{x, y}$, then $o_{v} \in S(x, y)$.
Cost: $\log |V|+\max _{v}\left|\Pi_{v}\right|$.

- Definition 18. (PPP)

Syntax: $V$ is a vertex set with a distinguished vertex $v^{*} \in V$. For each unordered pair $\{v, u\} \in\binom{V}{2}: o_{\{v, u\}} \in \mathcal{O}$. For each $v \in V: \Pi_{v}$ is a protocol outputting values in $V-v^{*}$.
Object: Bipartite graph $G_{x, y}=\left(V \times\left(V-v^{*}\right), E\right)$ where $(v, w) \in E$ iff $\Pi_{v}(x, y)=w$.
Correctness: If $(v, w)$ and $(u, w), v \neq u$, are edges in $G_{x, y}$, then $o_{\{v, u\}} \in S(x, y)$.
Cost: $\log |V|+\max _{v}\left|\Pi_{v}\right|$.


[^0]:    1 Work done while at Harvard University; supported by the Michael O. Rabin Postdoctoral Fellowship.
    2 Supported in parts by NSF grants CCF-1650733, CCF-1733808, and IIS-1741137.
    ${ }^{3}$ Work done while at University of Toronto; supported by NSERC.
    4 Supported by the Swedish Research Council grant 2016-00782, Knut and Alice Wallenberg Foundation grants KAW 2016.0066 and KAW 2016.0433.

[^1]:    5 That is, for any input $(x, y)$, every accepting computation of the nondeterministic protocol outputs a feasible solution $o \in \mathcal{O}$ for $(x, y)$. An alternative, more restrictive definition of TFNP ${ }^{\text {cc }}$ (which is closer to how the classical class is defined) is to require that there is an efficient deterministic protocol that on input $(x o, y)$ decides whether $o \in \mathcal{O}$ is feasible for $(x, y)$. In this paper we stick with the stronger (and simpler) definition for convenience. All results hold equally well under the more restrictive definition.

